

Practical Language: Its Meaning and Use

by

Nathan A. Charlow

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Philosophy)
in The University of Michigan
2011

Doctoral Committee:

Professor Allan F. Gibbard, Chair
Professor Richmond H. Thomason
Assistant Professor Ezra Russell Keshet
Assistant Professor Eric Peter Swanson
Professor Paul H. Portner, Georgetown University

© Nathan A. Charlow 2011

All Rights Reserved

For P&N

ACKNOWLEDGEMENTS

It was a great surprise to have to finish this dissertation so soon. It was also a huge task. I want to thank so many people for their help with it. I will, I'm sure, forget a few of you.

I especially want to thank Allan Gibbard for taking this project on. I am sure he was a bit skeptical about why a philosopher of language would wish to be supervised by an ethicist. In reading this, I think he (and you) will see why it made such good sense.

My brother Simon introduced me to hardcore formal semantics and kept me in touch with the field during the years when Michigan lacked a bona fide semanticist. He, more than anyone else, helped me to become a linguist in addition to a philosopher.

Eric Swanson has been my mentor since I took the plunge into philosophical linguistics in his 2006 seminar. He has read and given comments on nearly everything I have written since then. I owe him a great deal.

Rich Thomason taught me modal logic. It has served me better than any other tool I acquired during graduate school. Rich also pressed me early and often with questions about the conceptual foundations of what I was doing with imperatives and Expressivism. I hope I've been able to answer some of those questions in what follows.

I learned an incredible amount from Thony Gillies and Peter Ludlow. Thony introduced me to Paul Portner's work on imperatives, which is the foil for so much of what follows here, and he supervised my Candidacy Dossier, for which I eventually won the Charles L. Stevenson Prize. Peter taught me *Heim and Kratzer*, and was the first person I knew to try and give Allan's Expressivism its linguistic due.

Mark van Roojen and Al Casullo are pretty much wholly responsible for my going to graduate school. I'm glad to call them my friends.

For sundry inspiration, advice, and discussion, I'm grateful to Chris Barker, Gordon Belot, Billy Dunaway, Kai von Fintel, Benj Hellie, Lina Jansson, Ezra Keshet, Sam Liao, Clayton Littlejohn, Sarah Moss, Dilip Ninan, David Plunkett, Paul Portner, Peter Railton, Craig Roberts, Anders Schoubye, Alex Silk, Will Starr, Andreas Stokke, Sergio Tenenbaum, Dustin Tucker, Jonathan Weisberg, and Jessica Wilson.

My Michigan friends are the best. Lina Jansson is a great person, a great philosopher, and was an ideal flatmate for the past two years. Jason Konek turned down a record deal with Akon to be my roommate for my first two years in Michigan. Sam Liao and I have shared many good times, in bars, in restaurants, in New York, Bloomington, and New Haven. Dave Wiens helped me remember how awful I am at basketball. Alex Plakias has

impeccable taste in Nineties Hip-Hop. Eduardo García-Ramírez's friendship was especially important to me last year.

There is no way that I would have found a job—even secured an interview—without Linda Shultes' tireless help. Judith Beck has been unfailingly warm and helpful. Molly Mahony kept my spirits up during the job market. Thank you all.

I will keep things informal for people, things, and cats non-academic. *zilla. VT. Jingjing. Zingerman's, which still makes the best Americano in Ann Arbor. Rori Cat, my beloved three-fanged, barbarically yawping hillbilly =^_-=

Lots of love to C. *Qui nolet fieri desidiosus, amet.*

Finally, thanks Mom. Thanks Dad. I love you both.

TABLE OF CONTENTS

DEDICATION	ii
ACKNOWLEDGEMENTS	iii
LIST OF APPENDICES	viii
ABSTRACT	ix
CHAPTER	
I. Introduction	1
1.1 Interest	1
1.2 Overview	2
1.3 Synopsis	4
1.3.1 Identifying a Conventional Use	4
1.3.2 Embeddings	6
1.3.3 Applications: Expressivism and Conditionals	7
II. Dynamic Accounts and the Problematics of Conventional Force	10
2.1 Masters, Slaves, and Permissibility Spheres	11
2.2 A Less Idealized Model	13
2.3 Imperatives and Obligation-Imposition	15
2.4 Non-Command Uses	18
2.4.1 The Data	18
2.4.2 Interpretation and Explanation	20
2.4.3 A Problem: Instruction Imperatives	22
2.4.4 A Ninan-Inspired Reply?	23
2.4.5 Varieties of Permission	26
2.5 Recap	30
III. Imperatives and Necessitation	32
3.1 Indirect Speech Acts: A Primer in Defeasibility	33
3.2 Conventions as Defaults	37
3.2.1 Convention and Cancellation	38
3.2.2 The Mechanics of Interpretation	41

3.2.3	Data for Declaratives and Interrogatives	44
3.3	Imperatives Express Necessitation	45
3.3.1	Conventional Use Data	46
3.3.2	Necessitation Data	48
3.4	Computing ISAs	50
3.4.1	ISAs in General	50
3.4.2	Voiding	53
3.4.3	Binding	55
3.4.4	Permission as an ISA	57
3.4.5	Instruction as an ISA	62
3.5	Recap	66

IV. A Dynamic Theory of Imperative Meaning 67

4.1	A “Standard” Account	69
4.1.1	Motivations	70
4.1.2	Modeling the Standard Notion of a Speech Act	76
4.2	Information, Issues, and Preferences	81
4.2.1	Information States	81
4.2.2	Integrating Issues	82
4.2.3	Integrating Preferences	85
4.2.4	A Compositional System	87
4.3	Direction as Necessitation (as Coming to Recognize a Requirement)	93
4.3.1	Imperative Presuppositions	93
4.3.2	Coming to Recognize a Requirement	98
4.3.3	Instituting Requirements	99
4.3.4	<i>Conservative</i> Preference Update	101
4.3.5	Prima Facie and Inconsistent Directives	107
4.3.6	Instituting Recognition	109
4.3.7	Requirement Recognition as Decision	113
4.3.8	Strong Permission as Recognized Permission	115
4.3.9	Practical Questions and Preferential Issues	117
4.3.10	Summing Up	119
4.4	Against the Standard Account	121
4.4.1	Bad Arguments Against the Standard Account	121
4.4.2	The Challenge of Embedded Imperatives	129
4.4.3	Regular Operations	131
4.4.4	Conjunction: Contrived Examples	133
4.4.5	Conjunction: Sequencing Designation	135
4.4.6	Disjunction, Permission Content, and the Ross Paradox	137
4.4.7	Quantifiers?	151
4.5	Conditional Imperatives	162
4.5.1	A Standard Account (Wide-Scoping)	163
4.5.2	Conditional Speech Acts (Narrow-Scoping)	166
4.5.3	Challenging the Dynamic Program: The Modal Analysis	169
4.5.4	Against the Modal Analysis	173
4.5.5	A New Account	178
4.5.6	The Semantics of Conditional Requirements	187

4.5.7	Conditional Imperatives as Fundamental	193
4.5.8	Upshot	195
V.	Concluding Remarks	197
5.1	Taking Stock	197
5.2	Editorial Digression	197
5.3	Meaning, Force, Semantics, and Logic	199
APPENDICES	201
BIBLIOGRAPHY	271

LIST OF APPENDICES

Appendix

A.	Meaning for Expressivists	202
B.	Conditional Utility, Preference, Obligations, and Imperatives	234

ABSTRACT

Practical Language: Its Meaning and Use

by

Nathan A. Charlow

Chair: Allan F. Gibbard

I demonstrate that a “speech act” theory of meaning for imperatives is—contra a dominant position in philosophy and linguistics—theoretically desirable. A speech act-theoretic account of the meaning of an imperative $!\phi$ is characterized, broadly, by the following claims.

LINGUISTIC MEANING AS USE

$!\phi$'s meaning is a matter of the speech act an utterance of it conventionally functions to express—what a speaker conventionally uses it to do (its *conventional discourse function*, CDF).

IMPERATIVE USE AS PRACTICAL

$!\phi$'s CDF is to *express* a practical (non-representational) state of mind—one concerning an agent's preferences and plans, rather than her beliefs.

Opposed to speech act accounts is a preponderance of views which deny that a sentence's linguistic meaning is a matter of what speech act it is used to perform, or its CDF. On such accounts, meaning is, instead, a matter of “static” properties of the sentence—e.g., how it depicts the world as being (or, more neutrally, the properties of a model-theoretic object with which the semantic value of the sentence co-varies). On one version of a static account, an imperative *shut the window!* might, for instance, depict the world as being such that the window must be shut.

Static accounts are traditionally motivated against speech act-theoretic accounts by appeal to supposedly irremediable explanatory deficiencies in the latter. Whatever a static account loses in saying (prima facie counterintuitively) that an imperative conventionally represents, or expresses a picture of the world, is said to be offset by its ability to explain a variety of phenomena for which speech act-theoretic accounts are said to lack good

explanations (even, in many cases, the bare ability to offer something that might meet basic criteria on what a good explanation should be like).

I aim to turn the tables on static accounts. I do this by showing that speech act accounts are capable of giving explanations of phenomena which fans of static accounts have alleged them unable to give. Indeed, for a variety of absolutely fundamental phenomena having to do with the conventional meaning of imperatives (and other types of practical language), speech act accounts provide natural and theoretically satisfying explanations, where a representational account provides none.

CHAPTER I

Introduction

This dissertation develops a largely new understanding of a central dimension of human communication: the meaning and function of *practical language*. Language is practical (as opposed to descriptive or inquisitive) if it has, as a matter of its conventional meaning, a direct bearing on what we want and plan to do. Some central cases of practical language are imperatives, directives, deontic or obligation modals, and value-laden (evaluative and/or moral) language.

- (1.1) Pass the salt! (*imperative*)
- (1.2) I ask that you keep quiet. (*directive/request*)
- (1.3) a. Employees must wash their hands. (*strong deontic modal*)
b. Non-employees should, too. (*weak deontic modal*) (von Stechow & Iatridou 2009)
- (1.4) Chocolate is better than vanilla. (*evaluative claim*)
- (1.5) It's wrong to raise animals for meat. (*moral claim*)

In this introduction, I will try to describe, briefly, a few reasons that both philosophers and linguists might be interested in the question of the meaning and function of practical language. I will then give a broad overview of the project, as pursued in this dissertation. The project encompasses both the theory of the meaning and function of practical language that I develop in the main body of the dissertation, as well as the applications of that theory that I pursue in two lengthy appendices.

1.1 Interest

Practical language is of significant theoretical interest, both for the study of language and its role in human communication, and for philosophical (especially ethical and meta-ethical) theorizing.

Imperatives, especially, are recognized by linguists as one of three major sentence-types in human languages (declaratives and interrogatives are the others). Declaratives and inter-

rogatives have deservedly received an enormous amount of attention from theoreticians of the major paradigm in linguistic semantics—model-theoretic semantics in the Montagovian tradition (Montague 1973). Imperatives, however, have been comparatively neglected by theoreticians of this paradigm. Extant treatments cover only a very specific type of imperative (and are unsatisfactory, I shall argue, even with respect to that circumscribed domain). Imperatives, then, represent a major blind-spot in our understanding of natural language and the dynamics of human communication. This dissertation will attempt to correct this, by developing a systematic (and empirically sophisticated) theory of the *meaning* of a wide variety of imperatives in natural language, the *uses* to which agents put such imperatives in influencing their interlocutors' behavior, and the *relationship between* these.

The project's interest, however, far outstrips the value of extending the Montagovian paradigm's coverage to include imperatives. A major theme of my dissertation is that imperatives comprise a kind of language particularly amenable to *dynamic analysis*—a distinctive, and relatively new, approach to theorizing about linguistic meaning, which privileges dynamic properties of a sentence (what it's used *to do*) over its static properties (how it *depicts* the world to be, or, more neutrally, the set of objects assigned to it, relative to a model, by an interpretation function), and on which I will have much more to say below. In developing a dynamic analysis of imperative language, and demonstrating its explanatory power, the dissertation will aim to indirectly provide a largely new vindication of the dynamic program in theorizing about linguistic meaning.

The study of practical language is also important for its bearing on questions that are fundamental for theorizing about ethics, morality, and their associated discourses. Understanding practical communication—the communicative profile/function of practical language—enriches our understanding of its subject-matter, in a novel and surprising way. Normative questions of central interest to philosophers—the nature of moral obligation, whether to admit the existence of normative *facts*, saying what normative discourse could possibly be *for*, if not for exchanging factual information about the world—are, I argue, illuminated by studying how people use practical language. (I shall flag such applications, as appropriate.) In all cases, we will notice a common methodological lesson. Namely, progress on intransigent philosophical debates is made by looking at (and availing ourselves of the best empirical work on) how “ordinary people” talk about those subjects.

1.2 Overview

This dissertation has three main chapters and two appendices. The main chapters develop a systematic and wide-ranging theory of the meaning of various types of imperative clauses in natural language. The appendices address questions like (i) how this sort of account might be extended to normative practical language, (ii) what sort of bearing this has on philosophical theorizing about the nature and purpose of normative discourse.

The main chapters demonstrate that a “speech act” theory of meaning for imperatives

is—contra a dominant position in both philosophy and linguistics—theoretically desirable. A speech act-theoretic account of imperative meaning is characterized, broadly, by the following two claims.¹

LINGUISTIC MEANING AS USE

An imperative's meaning is a matter of the speech act an utterance of that imperative conventionally functions to express—what a speaker conventionally uses it to do (its *conventional discourse function*, hereafter "CDF").

IMPERATIVE USE AS PRACTICAL

The CDF of an imperative is not to state or represent facts about the world, or to express the beliefs of someone who utters it. Instead, it is to *express* practical (non-representational) states of mind—states of mind concerning an agent's preferences and plans, rather than her beliefs—and invite an agent's interlocutors to adopt corresponding states of mind.

Opposed to speech act accounts is a preponderance of views which deny that a sentence's linguistic meaning is ever a matter of what speech act it is used to perform, or its CDF. On such accounts, meaning is, instead, a matter of "*static*" properties of the sentence—e.g., how it depicts the world as being (or, more neutrally, the properties of a model-theoretic object² with which the semantic value of the sentence co-varies). On one version of a static account of imperative meaning, an imperative *shut the window!* might, for instance, depict the world as being such that the window must be shut.

As discussed at length in this dissertation, static accounts are motivated against speech act-theoretic accounts by appeal to supposedly irremediable explanatory deficiencies in the latter. Whatever a static account loses in saying (prima facie counterintuitively) that an imperative conventionally represents, or expresses a picture of the world, is said to be offset by its ability to explain a variety of phenomena for which speech act-theoretic accounts are claimed to lack satisfactory explanations (or, in many cases, the bare ability to offer something that might meet basic criteria on what a satisfactory explanation should be like).

What unites the chapters of this dissertation (and the disparate discussions contained therein), then, is the following overarching aim. I will aim to turn the tables on static accounts. I will do this by showing, not only that speech act accounts are capable of giving explanations of phenomena which fans of static accounts have alleged them unable to give, but also that, for a variety of absolutely fundamental phenomena having to do with

1. The exemplary reference, from my perspective, is [Portner \(2007\)](#); more references can be found in the chapter itself. As we will see, our account of imperatives has much in common with (and draws much inspiration from) the views in philosophical meta-ethics known as Expressivism ([Gibbard 1990, 2003](#)) or Quasi-Realism ([Blackburn 1984, 1988](#)). Viewing imperative use as practical, rather than representational, in nature is not, strictly speaking, a requirement of viewing imperative meaning as a matter of use. But I know of no speech act-theoretic or dynamic account of imperatives that is inclined to give them a representational use. For more on this, see [Appendix A](#).

2. Perhaps, but not necessarily, truth. For extensive discussion of other options, see my [\(2009\)](#).

the conventional meaning of imperatives, speech act accounts can provide natural and theoretically satisfying explanations, where a representational account can provide none.

1.3 Synopsis

A variety of tasks (each with a host of its own subtasks, and some of those with a host of their own subtasks) will go into achieving this aim. Here I will try to give an overall blueprint of the argument as I develop it through the following chapters. I will not try to be too detailed. I will just try to give an impression of some of the specific problems on which I'll be focused, and a general sense of how I will propose to go about solving them.

1.3.1 Identifying a Conventional Use

Chapter II introduces versions of some classic speech act accounts of imperative meaning. The version on which we focus first is due to David Lewis' *Convention* (Lewis 1969). This account holds that the CDF of an imperative is to *command* (where commanding is understood, roughly, as the creation of a new obligation on the utterances addressee).

Variability. Chapter II then presents a worry for this sort of account of imperative meaning—one that can be generalized so that it becomes a foundational worry for any speech act account of imperative meaning. The problem, in the case of Lewis's account, is that imperatives can be put to an impressive variety of uses: not just commands, but permissions, suggestions/advice, instructions. . . (see esp. Schwager 2006; Kaufmann & Schwager to appear).

(1.6) It's ok! Have an apple! (*permission-granting*)

(1.7) Talk to your advisor (to improve your research) (*advice*)

(1.8) Take Broadway (to get to 14th) (*instruction*)

Of course, someone who permits her addressee to have an apple, by issuing the permission imperative (1.6), does not thereby command her to have an apple. Likewise for advice and instruction imperatives. More generally (and therefore more worryingly for speech act accounts of imperative meaning as such), the wide variety of potential uses for imperatives seems to warrant pessimism about analyzing their meaning in terms of CDFs at all. A sentence's context-independent linguistic meaning, it is generally agreed, is a kind of skeleton that gets fleshed out on an occasion of use, by a context of utterance. The result of fleshing out a skeletal meaning of a sentence, for instance, relative to an occasion of use, is that sentence's meaning, on that occasion of use. It would seem that the context-independent meaning of an imperative, if it is to be given in terms of uses or speech acts, must be very thin indeed to accommodate the variety of meanings an imperative can

have, on a given occasion of use. So thin, indeed, that it would not seem to be capable of explaining much at all of interest to the theoretician.

Defaults. In Chapter III, I will develop a new kind of speech act account—one sufficiently flexible to cover all of these cases (and more besides). Imperatives, I argue, do have a CDF. It is, moreover, a thick CDF—one which I’ll gloss informally here as “requirement-generating” or “necessitating”, and one that is capable, as we shall see throughout the dissertation, of doing a serious amount of explanatory heavy-lifting.

This broad idea is not too different from that which animates Lewis (1969) (although in its specific implementation, it is, as we will see, considerably different). Which should lead us to wonder: how can a thick CDF—one sufficiently thick to have real explanatory traction—be sufficiently flexible to account for permission, advice, and instruction imperatives? What sets our view apart from Lewis’s here is a distinct account of linguistic interpretation. The CDF of an imperative is, I shall argue, best understood as an interpretive *default*. To flesh out what, precisely, is meant by this, I will develop an account of interpretive defaults on which cases like (1.6) and (1.8) (and, more generally, cases in which no requirement is generated by an imperative’s utterance) are explained as cases where information about the communicative goals of the speaker *overrides* the default. I develop a formal model—one that draws on state-of-the-art uses of interpretive defaults in the literature on Segmented Discourse Representation Theory—that generates testable (and correct) predictions about when, and how, the goals of a speaker at a context successfully override an imperative’s CDF.

Update. In Chapter IV, I shift to the task of giving a substantive account of what, exactly, the CDF of an imperative should be understood as consisting in. I will develop an account on which speech acts, generally, are understood as proposals for modifying the cognitive states of one’s interlocutors. Different clause-types in natural language (declarative, interrogative, imperative, and so on) are conventionally associated with different kinds of cognitive proposals. Declaratives, for instance, conventionally function to express assertions or resolve issues. Interrogatives conventionally function to introduce issues.³ Imperatives, I will argue, conventionally function to propose that an agent come to *recognize a requirement*.

This is broadly sympathetic to the Lewisian account of imperative meaning, but, as we will see, the requirement-recognition account, as I develop it, has significantly better explanatory coverage. I will argue for understanding the state of recognizing a requirement that ϕ as a complex state of mind, consisting roughly in (i) wondering whether or not to see to it that ϕ (treating whether or not to see to it that ϕ as a practical issue), (ii) resolving the practical issue in favor of ϕ . The major task here, insofar as we are trying to give a formal model of the relevant cognitive states (as well as formal definitions of acceptable

3. Such views about declaratives and interrogatives figure prominently in the literature on speech act theory and more recent work on clause-type. References are given in the chapter proper.

transitions between the relevant cognitive states), is to construct a theoretical apparatus for representing both (i) what it is to treat something as a practical issue, (ii) what it is to resolve a practical issue in favor of one of its answers. We will spend some effort developing such apparatus. The result of this effort will be a precise and detailed account of how an agent who accepts the speech act conventionally expressed by an imperative can be expected to update her cognitive state accordingly.

The recognitional part of this theoretical apparatus (and the notion of a non-informational issue, more generally) is, as far as I am aware, novel to this account.⁴ The apparatus is, I suggest, apt for application to diverse phenomena like (i) modeling the update conventionally proposed by a grant of permission, (ii) distinguishing between strong permission and merely implicit permission, (iii) explaining the Ross Paradox and related paradoxes of free choice permission, (iv) giving anti-representationalists (or “non-cognitivist”) of all stripes a way to think of questions and issues that does not presuppose that any resolution of an issue must be factual in nature.

1.3.2 Embeddings

We also, as I will emphasize in Chapter IV, find imperatives *embedded* in environments where they apparently have no commanding function. Even allowing that an imperative sentence Q commands some state of affairs (call it S), a *conditional* imperative—a sentence of the form (if P)(Q), like (1.9)—generally doesn’t. Similar points seem to hold for other environments embedding imperatives.

(1.9) If you’re cold, shut the window! (*conditional imperative*)

(1.10) Mail the letter, or give it to me! (*disjunctive imperative*)

(1.11) Someone bring me their chair! (*existentially quantified imperative*)

There is a general problem here, one which has been received comparatively little attention in recent work on the conventional meaning of imperatives. The problem is simply that of giving an account of the meaning of complex environments embedding imperatives.⁵ This is moreover no mere engineering problem. The very viability of a speech act-theoretic account rests on its ability to furnish an account of the meanings of such environments.

Worryingly, on what I will call the *Standard Account of Speech Acts*—embraced by all major 20th Century speech act theorists—it is difficult to identify any such speech act for a variety of such environments. The worry is a simple one. On the Standard Account a fundamental distinction is made between *illocutionary force* and *representations of states of*

4. Although it has a clear precursor in Gibbard (1990, 2003).

5. Apparently contra the predictions of the speech act account, (1.9) does *not* command that you shut the window (rather, it commands you do this *if* you’re cold), while (1.10) does *not* command that you mail the letter. This is analogous to the Frege-Geach problem (Searle 1962; Geach 1965; Green 2000) and is addressed indirectly in Chapter IV.

affairs. A speech act results from applying some force to a state of affairs. A command to shut the window, e.g., is the result of applying a specific force (*commanding* force) to a state of affairs (the state of affairs where the addressee shuts the window). But, contra the Standard Account, it is, I'll argue, fairly easy to see that the speech act expressed by complex imperatives cannot generally be represented as a force-content complex.

For speech act accounts to be viable for complex imperatives, then, we must jettison the Standard Account and broaden our understanding of (i) the types of speech act an utterance can express, (ii) what kinds of speech acts should be regarded as explanatorily fundamental (i.e., the building blocks out of which a domain of meanings for a speech act-theoretic account of meaning can be constructed). I am, as I will emphasize in Chapter IV, hardly the first (not even the first among friends of speech act-theoretic accounts of linguistic meaning) to make this point against the Standard Account. However, many of the arguments given for this point are, we will see, unconvincing, and are not very often accompanied by compelling alternative visions for how a speech act-theoretic account of linguistic meaning should be developed.

To give a bit of a flavor for the account we will end up developing, I will argue for representing (1.9) as proposing to introduce a *complex planning state* in an agent—one represented, very roughly, by sequentially pairing facts (relevant contingencies, like the addressee being cold) with planned outcomes (that the addressee shut the window). Borrowing from dynamic logic and the logic of computing, I will construct an appropriately rich range of meanings \mathcal{M} for a variety of environments embedding imperatives. The ingredients to this construction will be a family of *basic speech acts* (represented as a kind of program—an instruction for updating a cognitive state), together with a set of *recursive operations* on programs (sequencing, restriction, and so forth).

These devices make possible a speech act-theoretic account of the full range of complex imperatives. This account has the fullest empirical coverage of any of which I am aware. The only accounts with potentially comparable coverage—modal accounts of the imperative—are, I will argue, actually explanatorily inert. The account we develop in Chapter IV, then, represents the fullest account of the meaning of imperatives in natural language in the contemporary literature. It also represents a novel direction in speech act theory, especially as it pertains to the theory of meaning for natural language.

1.3.3 Applications: Expressivism and Conditionals

The appendices sketch different ways of applying this sort of account of meaning to problems and puzzles outside of the theory of meaning proper.

Meaning for Expressivists. Appendix A argues that an account of imperative meaning along the lines defended in the prior chapters can form the basis for an account of the meaning of normative sentences—one that is largely true to the motivations that have

animated Expressivism in meta-ethics. In particular, it shows how to transform an account of the *speech act* conventionally expressed by a practical (imperative or normative) sentence into a non speech act-theoretic *semantic theory* for that sentence. I argue that such a theory is capable of explaining, in a theoretically satisfying way, a host of uncontroversial logical and inferential properties of imperative and normative sentences, in a way that skirts traditional objections to Expressivist accounts of these properties. The result is something many meta-ethicists have thought chimerical: an Expressivist semantic theory for normative language as explanatorily powerful as its Cognitivist counterparts. This appendix ends with a sustained methodological discussion. The objections to Expressivist accounts of practical language (and Expressivist accounts of language, generally) arise, I argue, from an exaggerated understanding of Expressivism's methodological commitments—one, remarkably, embraced by *both expressivists and their critics*. Revising this understanding opens the door to Expressivist analysis of a much wider range of discourse than has been generally thought possible or desirable.

Conditional Utility, Preference, Obligations, and Imperatives. Appendix B explores the notion of conditional utility, arguing that it is philosophically important, not just for decision theorists interested in the correct formulation of the Law of Expected Utility, but also for philosophers interested in the meaning of natural language constructions that function to *express conditional preferences*, including deontic conditionals (*if you want A, you should see to it that A*) and conditional imperatives (*if you want A, see to it that A*). Building on the analysis of conditional imperatives that I developed in Section 4.5, I argue that the proper analysis of indicative deontic and imperative conditionals is in terms of what is planned, desired, or preferred, given monotonic changes to an agent's information.

Implementing this conception of conditional preference in a semantic analysis of indicative deontic conditionals turns out to be incompatible with a host of popular philosophical accounts of their meaning. The only analysis that does a satisfactory job with simple deontic conditionals like *if you want A, you should see to it that A* is an account that blends (i) a Context-Shifty account of indicative antecedents with (ii) an Expressivistic treatment of their deontic consequents (both elements of which emerge from the discussion of conditional imperatives in Section 4.5). Non-Expressivistic analyses of indicative conditionals encode a fundamental misunderstanding of what it is to be *best, given some condition*.

Entertaining a deontic conditional, I argue, is a matter of *consulting one's preferences*, on update with its antecedent. This, importantly, is distinct from consulting one's beliefs about one's preferences, on update with the antecedent. Correspondingly, deciding that a deontic conditional is the case is a matter of ending up in a cognitive state characterized by a specific, suppositional preference. This, I argue, is distinct both from (i) deciding *that* one has such a suppositional preference, and (ii) *coming to have* such a suppositional preference. Any account that conflates a decision to accept a deontic conditional with

either (i) or (ii) is incorrect as an account of their cognitive force; the former because it misconstrues the subject matter of such a decision, the latter because it makes it mysterious how a decision to accept a deontic conditional can *resolve a question* for the agent (namely, the issue of *whether* the deontic conditional is the case). On our account of the cognitive force of deontic conditionals, it is possible to entertain whether a deontic conditional is the case. Nevertheless, deciding that a deontic conditional is the case is not a matter of deciding that some body of preferences has some property; it is, rather, a matter of responding to a *successful query* of one's suppositional preferences.

The analytical gains of this dissertation are thus two-fold. We gain a new understanding of a relatively under-studied dimension of human discourse. We also show that (and, more importantly, *how*) an influential, but embattled, orientation to theorizing about linguistic meaning can ground serious, state-of-the-art analysis of human discourse.

CHAPTER II

Dynamic Accounts and the Problematics of Conventional Force

Perhaps surprisingly, sentences of the imperative clause-type (hereafter typically just “imperatives”) have to come to occupy an important place in the debate between proponents of dynamic and static approaches to theorizing about linguistic meaning. To explain this, it will be useful to have a rough and ready division of theoretical space. (The picture sketched here is quite general and crude. It will be sharpened progressively as this chapter unfolds.)

A *static account* of imperative meaning, very roughly, is one which understands the meanings of imperative clauses as traditional, “forceless” model-theoretic entities (propositions, properties, sets of model-theoretic constructs, and the like). A *dynamic account* of imperative meaning, on the other hand, is one which understands the meanings of imperative clauses as active entities (operations on a discourse, update instructions, context change potentials, and the like). For static accounts, saying what an imperative means is a matter of associating it with a model-theoretic object (or a logical form expressing a model-theoretic object) with the appropriate logical profile. For dynamic accounts, it is, very roughly, a matter of saying what an imperative is conventionally used to do. This is not to say, for dynamic accounts, that there is no static dimension of imperative meaning, or that imperatives have no interesting logical profile. It is to say that, for dynamic accounts, the static dimension of imperative meaning is theoretically peripheral; insofar as that dimension includes theoretically interesting facts about imperative meaning (e.g., facts about imperatives’ logical profile), those facts will typically be recast as dynamic facts, and given an explanation within the dynamic theory.¹

This chapter tries to do two things. First, we will describe two accounts of imperative meaning in the dynamic vein: the classic language-game account of David Lewis (Section 2.1), and the more recent, and very influential, account of Paul Portner (Section 2.2). While these accounts are not the only game in town, they are the most prominent in the philosophical and linguistic literature, and will serve as good foils for presenting the core features of dynamic accounts of imperative meaning, as well as outlining their difficulties. Second, we will describe a small set of data that any account of imperative meaning should

1. For more on this, see Section 5.3 and Appendix A.

explain. We will ultimately argue that a dynamic account of imperative meaning, in the vein of either of the accounts sketched in Sections 2.1 and 2.2, is not particularly well-suited to explaining these data. Indeed, these data are *prima facie* problematic for any account of imperative meaning which assigns imperative clauses a conventionalized type of use.

2.1 Masters, Slaves, and Permissibility Spheres

The classic account of imperative meaning is due to David Lewis (“A Problem about Permission,” 1979a).² Lewis describes an idealized language game between two players, MASTER and SLAVE (we shall ignore the KIBITZER), in which MASTER uses utterances of imperative and permission sentences (sentences of respective forms $\lceil !\phi \rceil$ and $\lceil \text{;}\phi \rceil$) to restrict and expand SLAVE’s options for acting (I will henceforth mostly ignore permission sentences). In the game, MASTER’s utterance of an imperative $!\phi$ is conventionally interpreted as expressing a specific kind of speech act—one that Lewis dubs *commanding*.³ Commanding that ϕ is a matter of incrementing (intersecting) a sphere of worlds (the *permissibility sphere*, conceived as the set of accessible worlds that SLAVE can permissibly actualize) with the proposition that ϕ . In summary form:

- For each world $w \in W$, there is an *accessibility sphere* $A_w \subseteq W$ representing the states accessible from w in conjunction with some action of SLAVE.
- For each world $w \in W$, there is a *permissibility sphere* $P_w \subseteq A_w$ representing the w -accessible states that SLAVE can permissibly realize, given the commands of MASTER in force at w .
- There are three *moves* available to MASTER: *command*, *permit*, and *assert*.
- SLAVE’s goal is to realize only permissible states. MASTER’s goal is to control SLAVE’s actions, by modifying the sphere of permissibility.

Evolution of the permissibility sphere. Lewis models the move associated with MASTER’s utterances of an imperative $!\phi$ as a restriction of the permissibility sphere to states satisfying ϕ .⁴ Formally:

2. Important predecessors of Lewis’s account are Stenius (1967); Lewis (1969); Chellas (1971), although Lewis (1979a) is an elegant synthesis of the key insights of each of these references. Important predecessors for the notion that meaning of clauses of the imperative “mood” could be explicated in terms of a conventionally expressed speech act (e.g., requesting or commanding) include Searle (1969); Austin (1975).

3. There are some quirky aspects of Lewis’ account. The semantic value of an imperative $!\phi$, for Lewis, is a *modal proposition*: the set of worlds w such that ϕ is required by the permissibility sphere at w . An imperative varies in its interpretation depending on whether it is MASTER or SLAVE that is using it; if it is MASTER, the utterance is interpreted as a command, and the permissibility sphere adjusts to make MASTER’s utterance true; if it is SLAVE, the utterance is interpreted as an assertion. Restricting our focus to cases of action-guiding uses, Lewis endorses something rather like a conventionalized connection between an imperative and the speech act of commanding.

4. Modeling the move of permission turns out to be fraught with trouble (hence the “problem about permission” referenced in the title of Lewis 1979a).

Definition 2.1. The **update potential** for an utterance by MASTER of the form $\lceil !\phi \rceil$ (notation: $[\!|\phi|]$) is a function mapping the permissibility sphere at a world w into $P_w \cap \llbracket \phi \rrbracket$. More formally, $[\!|\phi|] = \lambda w. \lambda P. P_w \cap \llbracket \phi \rrbracket$.

We write the update of w with MASTER's utterance of $!\phi$ in standard postfix notation: $w[\!|\phi|]$.

The conventional role of MASTER uttering $!\phi$ at w is, in an intuitive and familiar sense, to make it *obligatory* for SLAVE to bring about a state that realizes ϕ : since ϕ is true at every $w[\!|\phi|]$ -permissible state (i.e., $P_{w[\!|\phi|]} \subseteq \llbracket \phi \rrbracket$), there is no $w[\!|\phi|]$ -permissible state at which ϕ is false. If SLAVE is to realize a $w[\!|\phi|]$ -permissible state, SLAVE *must* realize a state in which ϕ . All of this is echoed in the familiar accessibility-relation semantics for modals expressing deontic necessity, on which it is deontically necessary that ϕ at w iff all worlds in A_w that are deontically accessible from w (i.e., all worlds permissibly realizable from w) satisfy ϕ . The Lewisian permissibility sphere at w is functionally equivalent to the domain of quantification induced by a deontic accessibility relation at w .

The **basic idea** of Lewis' account is that the conventional role of an imperative utterance is to add information to a parameter of the context that determines which courses of action the addressee (SLAVE) can permissibly select for realization. To this basic idea, Lewis adds a number of idealizing assumptions—all of dubious motivation, insofar as we are interested in phenomena that occur in actual conversations. The most important of these idealizations is the assumption that there is no essential distinction between the permissibility sphere (equivalently, the domain of quantification or deontic accessibility relation) and the parameter of the context that an imperative utterance targets for update.

It is a familiar point that this assumption makes sense only if we further stipulate that authorities may *never* issue imperatives whose prejacent is incompatible with the prior permissibility sphere: i.e., for all ϕ and w , if an authority utters $!\phi$ at w , $\llbracket \phi \rrbracket \cap P_w \neq \emptyset$. Adopting this idealization means that the model will fail to provide a satisfactory representation of any of the following kinds of cases.

- **Prima facie directives.** Cases where imperatives express *Prima Facie* Duties (Ross 1930)—general guidelines, things which are defeasibly (“tend to be”) our duties, to be followed when possible and when no other *Prima Facie* duties conflict, and which count against actions which fail to uphold them, without thereby *determining* those actions to be impermissible. Suppose a parent tells a child to (i) keep her promises, (ii) never harm others. If, on some occasion w , the child *must* break a promise to avoid harming someone, then there is no accessible world where the child satisfies all of the parent's instructions ($P_w = \emptyset$). Lewis' model predicts (incorrectly) that there is no world the child can permissibly realize.
- **Inconsistent authority.** Cases where an imperative is issued which is inconsistent with some prior imperative (or sequence of imperatives). Consider a case with *multiple authorities*, from (Kratzer 1991b: 642): one stipulates “that owners of goats are liable for damage their animals inflict on flowers and vegetables. . . [another] that owners of goats are not liable for damage caused by their animals.” Because these

imperatives are inconsistent, Lewis' model predicts that no world can permissibly be realized. That is incorrect; it is permissible to pay one's taxes, even if such "inconsistent" directives have been issued.

This basic idea can (and ought to) be divorced from the idealization that there is no essential distinction between the permissibility sphere and the parameter of the context that an imperative utterance targets for update. Paul Portner has recently developed just such an account. I describe it in the next section.

2.2 A Less Idealized Model

The source of the trouble here is the conflation of the *ideal* (satisfaction of *everything* the authority has commanded) with the *good enough* (satisfaction of *enough* of what the authority has commanded). Because Lewis' model includes no device for recording separately the content of the authority's utterances, it lacks the resources to determine which courses of action are good enough (with respect to all of the things the authority has commanded) in the event that none are ideal. The natural fix is to introduce such a device: we understand an imperative utterance as a proposal for adding the content of the direction to a body of norms, which subsequently functions to determine a permissibility sphere.⁵

This is basically analogous to the "To-Do List" of Portner (2004a, 2007): a body of norms that the addressee of an imperative utterance is committed, in some nebulous sense, to fulfilling, and which imperative utterances update via sequential addition of their prejacent. Here—discarding Lewis' arbitrary restriction to single-authority cases, and introducing Kaplanian contexts of utterance into the model—is a summary of Portner's account. We understand a context c as determining a speaker S_c , addressee A_c , and To-Do List function T_c from individuals into sets of propositions. Then,

Definition 2.2. $[\!|\phi|]$ is a function mapping a context c into an updated context c' such that $T_{c',A_c} = T_{c,A_c} \cup \{[\!|\phi|]^c\}$ and c' and c are otherwise identical.

What is the relationship between the To-Do List for an individual (a set of imperative prejacent that she is committed, *pro tanto*, to realizing) and what she must and may do? On the dominant picture,⁶ due mostly to Kratzer (1977, 1981, 1991b) (but see also Lewis 1981), a set of propositions \mathcal{P} characterizes a space of \mathcal{P} -permissible actions indirectly, by *inducing a partial ordering* on worlds. Portner suggests that we understand the relationship the To-Do List and what an individual must and may do along the lines of this model.

5. This is not necessarily to say that, in the final analysis, the relevant parameter must be understood as a set of propositions. Such a set may (indeed, certainly will) prove too crude for modeling certain phenomena. It's just to say that, in the final analysis, the parameter must be sufficiently structured to characterize a set of propositions (or something reasonably close thereto).

6. Dominant, in any case, among formal semanticists. Deontic logicians (particularly those working in the STIT and Preference Logic traditions) have worked to develop different, better-motivated, foundations for deontic logic, but the issues I'm interested in here can mostly be approached in the more familiar framework of Kratzer.

- The To-Do List functions as an *ordering source*:⁷ it determines a unique preorder on possibilities.
- What an individual must and may do is a function of the properties of possibilities occupying a sufficiently good place in the preorder determined by her To-Do List. Roughly, if all such possibilities share a property ϕ , the individual must see to it that ϕ ; if at least one satisfies ϕ , she may see to it that ϕ .

How does the To-Do List manage to determine a unique preorder on possibilities? According to Kratzer (1977, 1981, 1991b):

Definition 2.3. A world w is **at least as good as** a world v with respect to a set of propositions $\mathcal{P} \subseteq 2^W$ (notation: $w \leq_{\mathcal{P}} v$) iff $\{p \in \mathcal{P} : v \in p\} \subseteq \{p \in \mathcal{P} : w \in p\}$.

Some notational conventions: $w <_{\mathcal{P}} v$ iff $w \leq_{\mathcal{P}} v$ and $v \not\leq_{\mathcal{P}} w$, while $w =_{\mathcal{P}} v$ iff $w \leq_{\mathcal{P}} v$ and $v \leq_{\mathcal{P}} w$. The guiding intuition here, one suspects, is that a world is absolutely better with respect to a body of imperatives \mathcal{P} , *ceteris paribus*, the more things in \mathcal{P} it satisfies. As Kratzer notes, the relation $\leq_{\mathcal{P}}$ is reflexive and transitive, but it is not necessarily connected. A world w and a world v are *incomparable* with respect to \mathcal{P} (i.e., $w \not\leq_{\mathcal{P}} v$ and $v \not\leq_{\mathcal{P}} w$) in, for instance, a situation where w and v satisfy disjoint, non-empty subsets of \mathcal{P} (and, more generally, just in case, for some $p \in \mathcal{P}$, $w \in p$ and $v \notin p$, and, for some $q \in \mathcal{P}$, $v \in q$ and $w \notin q$).

Rankings on worlds are naturally tied to an analysis of the modal notions *must* and *may* (here we follow Kratzer 1981, 1991b). Given a finite domain of worlds, there will necessarily be worlds that are *closest* to the ideal established by \mathcal{P} . These are the worlds that are *minimal* with respect to the relation $\leq_{\mathcal{P}}$: worlds w such that for any v , if $v \leq_{\mathcal{P}} w$, then $w \leq_{\mathcal{P}} v$ (informally, worlds that it is impossible to strictly improve upon).

Proof. Suppose that $\neg \exists w : \forall v : v \leq_{\mathcal{P}} w \Rightarrow w \leq_{\mathcal{P}} v$. Then $\forall w : \exists v : v <_{\mathcal{P}} w$. The infinite cardinality of the domain of worlds follows immediately from the transitivity, irreflexivity, and antisymmetry of $<_{\mathcal{P}}$. \square

In cases with a finite domain, the desirable conditions on the modal notions *must* and *may* are pretty clear: *must* quantifies universally over $\leq_{\mathcal{P}}$ -minimal worlds, while *may* quantifies existentially over $\leq_{\mathcal{P}}$ -minimal worlds (and, hence, is the dual of *must*).⁸

Definition 2.4. $\llbracket \text{must}_a(\phi) \rrbracket^{c,w} = 1$ iff $\min(\mathcal{A}_{c, \leq_{\mathcal{P}_c(w)}(\llbracket a \rrbracket^c)}) \subseteq \llbracket \phi \rrbracket^c$

7. This simplifies somewhat; according to Portner (2007), the To-Do List is typically a *subset* of the ordering source relevant for evaluating a statement of necessity or possibility.

8. The assumption of a finite domain entails the Limit Assumption of Lewis (1973). Kratzer's official analysis of *must* and *may* avoids the Limit Assumption, but at the cost of considerably complicating the analysis. Assuming a finite domain is a common idealization, and the issue of the domain's cardinality is completely orthogonal to our concerns. So we will assume a finite domain.

Three notes about this definition: (i) for convenience, we refer to the set of $\leq_{\mathcal{P}}$ -minimal worlds in a set of worlds X (i.e., $\{u \in X : \forall v \in X : v \leq_{\mathcal{P}} u \Rightarrow u \leq_{\mathcal{P}} v\}$) using the abbreviated notation $\mathbf{min}(X, \leq_{\mathcal{P}})$; (ii) to introduce the requisite world- and individual-relativity of obligation, we treat \mathcal{P} as a function from worlds and individuals into sets of propositions; (iii) \mathcal{A}_c functions as a *modal base*, a function determining a domain of relevant worlds, basically equivalent, in the cases in which we're interested, to a Lewisian accessibility sphere.

2.3 Imperatives and Obligation-Imposition

Portner (2007: 359) writes, “Without any concern for formal theories of discourse, one could accept the following description of language use: In some cases, the utterance of an imperative imposes an obligation on the addressee.” Lewis’ model, following in the august footsteps Austin (1975); Searle (1969), treats the obligation-imposing function of imperative utterance as a *conventionalized* property of the utterance: the conventional function of an imperative utterance $!\phi$ by MASTER is to make it the case that SLAVE must see to it that ϕ . But Portner’s model is more circumspect: addition of the proposition that ϕ to an ordering source \mathcal{P} does not, in general, make it the case that the $\leq_{\mathcal{P} \cup \{\llbracket \phi \rrbracket\}}$ -minimal accessible worlds will satisfy ϕ (and, so, does not, in general, make it the case that utterance of an imperative $!\phi$ imposes an obligation to realize ϕ on the addressee).⁹

Of course, the addition of the proposition that ϕ to an ordering source \mathcal{P} will *tend to make it the case* that the $\leq_{\mathcal{P} \cup \{\llbracket \phi \rrbracket\}}$ -minimal accessible worlds will satisfy ϕ . Specifically, it will do this *whenever* there are both ϕ and $\neg\phi$ worlds that are:

- i. \mathcal{P} -comparable to each other, and
- ii. Among the best accessible worlds prior to update with the imperative, i.e.:
 - a. ϕ is antecedently permitted
 - b. $\neg\phi$ is antecedently permitted

Proof. Let \mathcal{P} be an ordering source, \mathcal{A} a set of accessible worlds, $w \in \mathbf{min}(\mathcal{A}, \leq_{\mathcal{P}}) \cap \llbracket \phi \rrbracket$, and $v \in \mathbf{min}(\mathcal{A}, \leq_{\mathcal{P}}) \cap \llbracket \neg\phi \rrbracket$, and suppose w and v are \mathcal{P} -comparable (i.e., $w \leq_{\mathcal{P}} v$ or $v \leq_{\mathcal{P}} w$). Then, by the definition of \mathbf{min} , $v \leq_{\mathcal{P}} w \Leftrightarrow w \leq_{\mathcal{P}} v$. So $w =_{\mathcal{P}} v$. By Definition 2.3, $\{p \in \mathcal{P} : w \in p\} = \{p \in \mathcal{P} : v \in p\}$. Then, since $w \in \llbracket \phi \rrbracket$ and $v \notin \llbracket \phi \rrbracket$, by Definition 2.3, $w <_{\mathcal{P} \cup \{\llbracket \phi \rrbracket\}} v$. And since $w \in \mathbf{min}(\mathcal{A}, \leq_{\mathcal{P}})$, $w \in \mathbf{min}(\mathcal{A}, \leq_{\mathcal{P} \cup \{\llbracket \phi \rrbracket\}})$. So $\mathbf{min}(\mathcal{A}, \leq_{\mathcal{P} \cup \{\llbracket \phi \rrbracket\}}) \subseteq \llbracket \phi \rrbracket$. \square

What of cases where the assumptions of this proof fail—cases where there are, among the \mathcal{P} -minimal accessible worlds, either (i*) no ϕ and $\neg\phi$ worlds that are \mathcal{P} -comparable to

9. Indeed, as we see below, Portner (2010) claims that this yields a satisfying treatment of permission-uses of imperatives. I don’t think that is right; an appreciation of the diversity of uses of the imperative actually leads to a less circumspect (i.e., more Lewisian) understanding of the conventional function of imperative utterances. See Chapter III for discussion.

each other, (ii.a*) no ϕ -worlds (ϕ is antecedently forbidden), or (ii.b*) no $\neg\phi$ -worlds (ϕ is antecedently required)? Let's take these in order.

Notice that, in the event of (i*), that addition of $\llbracket\phi\rrbracket$ to the ordering source *cannot* make it the case that ϕ is required.

Proof. If, among the \mathcal{P} -minimal worlds, no ϕ and $\neg\phi$ worlds are \mathcal{P} -comparable to each other, then, for any $w \in \mathbf{min}(\mathcal{A}, \leq_{\mathcal{P}}) \cap \llbracket\phi\rrbracket$ and $v \in \mathbf{min}(\mathcal{A}, \leq_{\mathcal{P}}) \cap \llbracket\neg\phi\rrbracket$, $\{p \in \mathcal{P} : w \in p\} \not\subseteq$ or $\not\supseteq \{p \in \mathcal{P} : v \in p\}$. So, $\{p \in \mathcal{P} \cup \{\llbracket\phi\rrbracket\} : w \in p\} \not\subseteq$ or $\not\supseteq \{p \in \mathcal{P} \cup \{\llbracket\phi\rrbracket\} : v \in p\}$. So $w \not\leq_{\mathcal{P} \cup \{\llbracket\phi\rrbracket\}} v$ and $v \not\leq_{\mathcal{P} \cup \{\llbracket\phi\rrbracket\}} w$. Since $w \in \mathbf{min}(\mathcal{A}, \leq_{\mathcal{P} \cup \{\llbracket\phi\rrbracket\}})$, $v \in \mathbf{min}(\mathcal{A}, \leq_{\mathcal{P} \cup \{\llbracket\phi\rrbracket\}})$. \square

Informally, incomparabilities among worlds *cannot be remedied by addition* to the ordering source; if two incomparable worlds w and v are among the best worlds, then addition to the ordering source of a proposition that w , but not v , satisfies cannot “improve” w 's position relative to v in the updated ordering.

If this seems unintuitive, notice that, in Kratzer's framework, comparability between worlds obtains only when one of those worlds satisfies all of the propositions that the other does (and perhaps more besides). This is a consequence of the fact that, in this framework, there is a (reasonable) idealization away from considerations of relative priority (e.g., cardinal or ordinal utility) among elements in the ordering source—the only way of assessing the goodness of a possibility, relative to an ordering source, is by reference to the notion that a given world can be strictly improved, with respect to an ordering source, by expanding the set of propositions in the ordering source that it satisfies. So questions like *how many further propositions must a world satisfy to offset its failure to satisfy a given element of the ordering source?* do not strictly make sense in this setup. Indeed, any device for inducing an ordering on worlds where neither satisfies all of the propositions that the other does will tacitly presuppose some sort of weighting mechanism on elements of the ordering source. For instance, the simplest such device (ranking by the sheer quantity of propositions satisfied) tacitly presupposes that elements of the ordering source are *equivaluable*. If it is reasonable to idealize away from considerations of relative priority, we must tolerate its consequences in cases where (i*) holds.¹⁰

Cases where (i*) holds are the most delicate (because the predictions are *prima facie* counterintuitive). Cases where (ii.a*) and (ii.b*) hold can be dealt with more easily. Case (ii.b*) is trivial: when ϕ is antecedently required, an utterance of $!\phi$ will not *introduce* an obligation to see to it that ϕ , for that obligation is *already* in force.

For case (ii.a*), it is controversial how to model the effect of uttering an imperative $!\phi$

10. And, whether or not it is reasonable, natural language semantics for modals with evaluative meaning has taken it mostly for granted. I will follow that tradition here, with the proviso that this is an idealization that future work should work to relax. (There are, I note, some interesting exceptions within this tradition—e.g., Kratzer's own, but ultimately undeveloped, suggestion that considerations of relative priority might be semantically encoded by making use of *ordered sequences* of ordering sources. See Kratzer 1981 and von Stechow & Iatridou 2005, 2008 for discussion. Within Kratzer's semantics, however, ordered sequences of ordering sources are too clumsy to mimic full-blown ordinal/cardinal utility scales. I discuss this in other work.)

when ϕ is antecedently forbidden (perhaps because $\neg\phi$ is already an element of the ordering source). Indeed, there is some reason to think such an utterance should not *always* produce a new obligation. Portner (2010) argues that in typical such cases, imperative utterances function to generate permissions (by nullifying earlier commands), rather than obligations. For instance, *take tomorrow off!* may function both to create an obligation to take the next day off and to generate a permission to take the next day off. A discourse-initial utterance of this imperative would typically perform the former function, while certain non-discourse-initial utterances, as in (2.1), would typically perform the latter function (2010: 10).

(2.1) Monday carry rocks! Tuesday carry rocks! And Wednesday carry rocks! [...Tuesday comes along.] Take tomorrow off!

Contra the Lewisian account, an utterance of (2.1) intuitively does not function to establish an *obligation* to take the next day off. Rather, it functions to nullify an earlier command, to *allow* the addressee to take the next day off.

(2.2) A: (...) Take tomorrow off!
 B: No, thank you. I prefer to work.
 A: Alright then.

Portner (2010) has an almost immediate account of this: the last imperative in the sequence introduces a *direct inconsistency* into the ordering source (i.e., makes it the case that, for some p and q in the ordering source, $p \cap q = \emptyset$). It follows from Definition 2.3 that worlds in which the addressee takes Wednesday off (p -worlds, for short) are incomparable to worlds in which she carries rocks on Wednesday (q -worlds). For reasons we have already seen, this in turn will tend to make it the case that (a) the best worlds will satisfy one or the other of p or q , (b) the set of best worlds will be a subset of neither p nor q .

Similar cases—in which the obligation-imposing tendency of imperative utterance is defeated by the presence of antecedent conflicting obligations—are easily constructed for contingently conflicting *prima facie* directives and multiple inconsistent authorities. Portner’s account handles both kinds of case with relative ease. The upshot: that Portner’s account predicts reliable obligation-introduction only in a limited range of contexts—paradigmatically, contexts in which conditions (i) and (ii) are satisfied—is not, as such, any strike against that account. Indeed, Portner seems to have an attractive—and, for all we have seen here, *appropriately flexible*—account of the pre-theoretical intuition that, typically, “the utterance of an imperative imposes an obligation on the addressee.”¹¹

11. To repeat: I don’t accept this account of permission imperatives, nor do I actually think that Portner has the right account of the conventional force of an imperative utterance. For further discussion, see Sections III and 4.3.2.

2.4 Non-Command Uses

Section 2.3 argued that Portner’s account is well-suited to explaining a range of non-command uses of imperatives. But there is some reason to think that range is not quite wide enough. Indeed, Schwager (2005); Kaufmann & Schwager (to appear) object to dynamic accounts—accounts which understand the meanings of imperative clauses in terms of a conventionalized use, discourse role, or speech act—on the grounds that they fail to predict the *wide variability in functional potential* of utterances of imperatives:

[E]ven if we could find a non-ad hoc way of associating linguistic objects with speech acts, we would still face the problem that there is no straightforward common core to the broad range of speech acts that imperatives are used for. They easily express not only commands or requests, but also wishes, advice, curses, etc. Therefore their semantic interpretation cannot strictly determine the speech act they are used for (Kaufmann & Schwager to appear).

In this section and the subsequent chapter, I will present the problematic data, and argue, contra Schwager and Kaufmann, that it represents no problem for a dynamic account of imperative meaning as such (although it does, I will argue, represent a deep problem for Portner’s specific version of that account).

2.4.1 The Data

To illustrate the wide variability in use of imperative clauses, Schwager and Kaufmann fix on cases like the following (these examples are taken from Schwager 2005):

(2.3) It starts at 9, but come earlier, if you wish. (*permission*; cf. Hamblin 1987)

(2.4) Fine, don’t come then. (*concession*)

The salient point about (2.3) and (2.4), according to Schwager and Kaufmann, is that, despite being *syntactically imperative*, neither, intuitively, expresses a command or request (in the sense that neither, on the most natural interpretation, seems to impose any sort of *obligation* on the addressee). (2.3) issues a permission to arrive early, while (2.4) issues a permission—in a wounded, perhaps insincere manner—not to come at all. To examples like these, we might add (2.5)-(2.14), all of which share the salient property of (2.3) and (2.4).

(2.5) Have a piece of fruit. (*invitation*; Portner 2007: 355)

(2.6) Talk to your advisor more often. (*suggestion/advice*; Portner 2007: 355)

(2.7) A: How can I get to Union Square?

B: Take Broadway to 14th. (*instruction*)

(2.8) Let me never hear her voice again! (*impersonal exhortation*)

(2.9) Get well soon! (*good wish*; Wilson & Sperber 1988)

- (2.10) Go on, throw it. Just you dare. (*threat/dare*; Wilson & Sperber 1988)
- (2.11) Rot in hell! (*curse*)
- (2.12) Aw, fuck me! (*exclamation*)
- (2.13) Stick it in your arse! (*expression of contempt*)
- (2.14) Go fuck yourself! (*expression of contempt*)

Some brief descriptive remarks about these examples, by way of stage-setting. Invitations like (2.5) are most naturally interpreted as nullifications of actual or assumed prohibitions—i.e., as a type of permission.

- (2.15) A: It's okay! Have a piece of fruit!
 B: No, thank you.
 A: Alright then.
- (2.16) A: It's okay! Have a piece of fruit!
 B: Do I have to?
 A: No, of course you don't.

Suggestions/pieces of advice/instructions like (2.6) and (2.7) are most naturally interpreted as giving the addressee advice about what is desirable, in view of what the addressee wants or a stipulated goal; such advice is not obligation-imposing.

- (2.17) A: I really want to make more progress on this paper.
 B: Talk to your advisor more often.
 A: I don't think that's a good idea. She hates me.
 B: Alright then.
- (2.18) A: I want to get to Union Square
 B: Take Broadway to 14th (for example).
 A: Can I take 4th Avenue instead?
 B: That will work too.

Impersonal exhortations, good wishes, threats/dares, curses, exclamatives, and expressions of contempt are all typically used without any *practical force* (commanding, permissive, suggestive, etc.) at all: they function to express a conative attitude of the speaker: they have (and are intended to have) no effect *per se* on the intentions of the addressee.¹²

- (2.19) A: Rot in hell!
 ??B: I'd really rather not.

12. The "*per se*" is a necessary qualification. In some cases, such uses may affect the intentions of the addressee *indirectly*, by informing her of the speaker's attitude toward some state of affairs over which she has control, in a context where she has some reason to care about the speaker's attitude toward that state of affairs.

While interesting, such constructions strike me as (in some nebulous sense) *non-literal* or *non-serious* uses of the imperative. The central data for any account of imperative meaning is comprised by literal or serious uses of imperative clauses; I will therefore ignore such constructions in the remainder of the paper (see [Wilson & Sperber 1988](#) for a less hand-wavy treatment). Interestingly, however, we will in [Section 2.4.2](#) describe cases of apparently literal uses of the imperative that have no practical force.

2.4.2 Interpretation and Explanation

From the discussion of [Section 2.3](#), it is clear that the fact that imperatives, as a clause-type, “express not only commands or requests,” but also speech acts such as permissions, suggestions, instructions, exhortations, curses, exclamations, and expressions of negative attitudes, is no strike, *as such*, against a dynamic account like Portner’s. A command is a specific kind of speech act: one which imposes a (deontic) obligation on the addressee. Portner’s account does not, we have seen, have imperatives conventionally expressing commands. Rather, the conventional function of an imperative utterance $!\phi$ is to add $[\phi]$ to the addressee’s To-Do List. Some, *but not all*, additions of a proposition $[\phi]$ to a To-Do List result in making it obligatory that ϕ . So, on Portner’s account, and consistent with the data described in [Section 2.4.1](#), some, but not all, imperative utterances constitute commands; it is a contingent fact about the context c that an imperative utterance at c constitutes a command/request.

There is, however, more to be said. Portner owes us an explanation of the multiplicity of functions an imperative can serve. We have seen cases in which an imperative utterance’s *nullification of a prior command* (by way of establishing a permission) can be explained by appeal to its conventional function, as Portner understands it (i.e., addition to the To-Do List). But it remains to be seen whether the various kinds of speech acts mentioned in [Section 2.4.1](#) can be subsumed under To-Do List addition—whether, in fact, there is a “common core to the broad range of speech acts that imperatives are used for” ([Kaufmann & Schwager to appear](#)).

[Portner \(2007, 2010\)](#) suggests that (i) literal imperative utterances may vary in the *speaker’s grounds* for making the utterance (roughly, the speaker’s communicative intention), (ii) *this variation explains* the variety of uses to which imperatives can be put, (iii) different grounds determine *what kind of ordering source* is selected for update. Contextual information about the speaker’s grounds determines the interpretation of the speaker’s utterance at a context. The account is something like the following (where ‘ $a \Rightarrow b \Rightarrow c$ ’ means that speaker’s grounds a explain interpretation b , which in turn explains which ordering source c is selected for update) ([Portner 2010](#): 6–7):

- To help speaker achieve a further goal in a situation where she has authority over addressee \Rightarrow *Command* \Rightarrow Deontic

- To help speaker achieve a further goal in a situation where she lacks authority over addressee \Rightarrow *Request* \Rightarrow Deontic
- To help addressee achieve a goal \Rightarrow *Suggestion/advice* \Rightarrow Bouletic/Teleological¹³
- To help addressee achieve a goal in a situation where speaker is viewed as an epistemic superior \Rightarrow *Instruction* \Rightarrow Bouletic/Teleological
- To help addressee achieve a goal in a situation where speaker has authority to prevent the addressee from doing so \Rightarrow *Invitation* \Rightarrow Bouletic/Teleological

The key idea is that variation in interpretation is a function of variation in the grounds a speaker can have for issuing an imperative. Variation in interpretation is, in turn, modeled as variation in either:

- **Type 1.** The *kind of ordering source* that is selected for update.
- **Type 2.** Context-dependent *facts about the speaker's understanding of her relation to the ordering source* that is selected for update (as embodied in her grounds for issuing the imperative).

TYPE-(I) variation distinguishes commands and requests, on the one hand, from suggestions, advice, instructions, and invitations, on the other. TYPE-(II) variation cuts things more finely. Commands are distinguished from requests, not in virtue of the kind of ordering source they select for update (both select deontic ordering sources), but in virtue of the speaker's understanding of her ability to effect additions to the ordering source. Similarly, instructions and invitations both select a bouletic ordering source for update; however, instructions involve an assumption by speaker that she is an epistemic authority about how the addressee can realize her desires, while invitations involve an assumption by speaker that she has the authority to prevent addressee from realizing her desires.

Here is an informal sketch of Portner's account of dialogue (2.17). B's ground for issuing the imperative is to help A realize her desires. So B's utterance is interpreted as a suggestion, and selects a bouletic ordering source for update: B's utterance is interpreted as a proposal to update the portion of A's To-Do List having to do with what A wants (a proposal, roughly, for A to add the proposition that A talk to her advisor more often, to her desires). But A has special knowledge, in view of which such a modification appears to be a bad idea. So, *A rejects B's suggestion* (a consequence of the fact that addressee, rather than speaker, is the authority with respect to her desires). Given that, by stipulation, B's grounds for giving her suggestion are to help A realize her desires (together with the further information that A's advisor hates her), B is *prima facie* committed to withdrawing her proposal. And that is exactly what she does. Most of this can be carried over to the explanation of dialogue (2.18).

13. Bouletic and teleological ordering sources are supposed to be determined by an agent's desires and goals, respectively. For my purposes, this is a distinction without a difference.

2.4.3 A Problem: Instruction Imperatives

So far, so intuitive. But consider this modification of dialogue (2.18):

- (2.20) A: How does one get to Union Square from here?
B: Take Broadway to 14th (for example).

In response to A's question, B offers an *instruction*. Worryingly, however, A can ask this question even in the absence of any desire to get to Union Square (e.g., out of mere curiosity). More damningly, B's ground for issuing the imperative obviously need not have any sort of practical content: supposing it is common ground that A is merely curious about how to get to Union Square from here—but has no interest in doing so herself—B utters the imperative without intending any effect on the goals/intentions of the addressee (or any other individual). B's ground for issuing the imperative is, rather, to *inform* A that a way to get to Union Square from here is by taking Broadway to 14th. B's utterance is, in other words, interpreted as resolving the issue raised by A's question; it is interpreted *as an assertion*.

Although dialogue (2.20) is an especially clear illustration of this phenomenon, it can also be seen to hold for dialogue (2.18). Assume simply that B's ground for issuing the imperative¹⁴ is simply to inform A how she can get to Union Square. By stipulation, B has no interest in helping A achieve her goal (or, indeed, in affecting A's intentions in any manner whatever). Her sole interest is in being a cooperative interlocutor (hence, in resolving

14. Paul Portner (pc) has expressed a worry about whether this is actually an imperative, rather than elliptical for the declarative *you take Broadway to 14th*. The problem here, as Portner notes, is that the imperative and second-person present-indicative inflections of English *take* are morphologically identical. But consider a different example.

- (2.21) A: How can I get her attention?
B: Hmm. Be cool as a cucumber. She'll notice that.

The example is interesting because imperative and second-person present-indicative inflections of English *be* (*be* and *are*, respectively) are not morphologically identical. *Be* cannot be second-person present-indicative.

Portner worries further that B's reply is either elliptical for *you should be cool as a cucumber* or, else, has directive force. I.e., he worries that there is no reading for B's reply on which the first sentence is typed as an imperative, but fails to have directive force.

Regarding the first worry, imperatives can plausibly subordinate result clauses, even if the resulting meaning can be paraphrased as using a necessity modal.

- (2.22) Pour obtenir son attention, soyez insouciant. (Soyez insouciant pour obtenir son attention.)
[Lit: *To get her attention, be insouciant.*]
(2.23) Pour obtenir son attention, vous devriez être insouciant.
[Lit: *To get her attention, you should be insouciant.*]

These sentences seem to be very close in meaning, but the former does not seem to be *elliptical* for the latter. *Soyez* is an inflected imperative, while *être* is an uninflected infinitive.

Regarding the second worry, notice that linguistic acquiescence is licensed for a clearly directive interpretation of the imperative.

- (2.24) A: How can I get her attention?
B: Please just be cool as a cucumber. She'll notice that.
A: Okay, fine.

the issue raised by A's question). A's imperative utterance in (2.18) is clearly an efficient way for A to realize her stipulated communicative intentions. To handle these dialogues, Portner must in each case identify an ordering source that B's utterance targets for update (otherwise, make the case that this sort of utterance is somehow an anomalous/non-central case). But in either case no such ordering source seems to be available.

The point here generalizes to *instructive uses* of imperatives (uses of the sort paradigmatically tokened in recipes, assembly manuals, user guides, etc.). It seems plausible that the typical function of such uses is to inform an agent how to achieve some specified end, rather than to propose the prejacent for addition to any agent's goals.¹⁵ In some cases—specifically, cases in which the agent has the specified end as one of her goals—the provision of information will often and naturally lead an agent to add the prejacent to her goals.¹⁶ But such cases are special; the effect on the addressee's plans is not, *per se*, part of the speech act of instruction. So instructive uses of imperatives cannot generally be modeled as proposals for addition to an addressee's goals. A Portner-style account of instructive uses of imperatives fails.

2.4.4 A Ninan-Inspired Reply?

Here is a possible line of reply. Portner (2007: 366) considers this case (from Ninan 2005):

(2.26) The Pope must relax his stance on contraception.

(2.26) is worrying for Portner because Portner he regards unembedded strong necessity modals (e.g., *must*) as having a conventionalized obligation-imposing (hence, ordering source-modifying) function. This is a key element in Ninan (2005)'s explanation (endorsed by Portner) of the following two classes of phenomena:

Phenomenon 1. Infelicity of conjoining imperatives/strong necessity modals to sentences expressing speaker's expectation that the imperative/strong necessity modal will be dis-

But linguistic assent and denial are available for A as a reply to B's utterance.

(2.25) A: How can I get her attention?
B: Be cool as a cucumber.
A: No, you're wrong about that. She's not very observant.

Linguistic acquiescence, as I will suggest below, is a characteristic linguistic hallmark of directive interpretation. Linguistic assent and denial are characteristic linguistic hallmarks of issue-resolving interpretations.

15. Schwager is in broad agreement. See her (2005: 5).

16. Of course, in cases like (2.20) where the specified means is not a *necessary* means—cases where the speaker is informing the addressee of one way, among others, to achieve the specified goal—it seems the addition of the prejacent to an ordering source will often be blocked. Portner's account predicts that this constitutes a *rejection* by the addressee of the speaker's proposal. That prediction is *prima facie* incorrect.

obeyed.

- (2.27) a. You should go to church, although I know that you won't.
b. #Go to church, although I know that you won't.
c. #You must go to church, although I know that you won't.

Phenomenon 2. Obligatorily epistemic interpretation of strong necessity modals over past-tense auxiliaries (+infelicity of past-tense imperative prejacent).

- (2.28) a. Sam must have gone to church. (*obligatorily epistemic*)
b. Sam should have gone to church. (*possibly deontic*)
c. #See to it that you went to Church. (*epistemic reading unavailable, hence odd*)

Ninan and Portner propose to explain these contrasts by appeal to the following facts (very roughly described here):

- Strong necessity modals like *must* have, like imperatives, a conventionalized To-Do List-modifying function.
- Proposing addition of a proposition that ϕ to a To-Do List is felicitous only when the speaker thinks it possible that the addressee will see to it that ϕ (perhaps because violating this condition manifests incoherence on the part of the speaker).

A *prima facie* problem with this account of the contrasts in (2.27) and (2.28) is that (2.26) is not typically uttered with the intention of imposing an obligation on the Pope. What ordering source might a typical utterance of (2.26) select for update?

According to Portner, (2.26) “is most plausibly understood with a teleological meaning: in view of the goal of protecting women’s health in poorer countries. . . or some such conversational background” (2007: 366). It selects a teleological ordering source for update. Similarly, perhaps, for (2.20). Perhaps there is a *very* specialized teleological ordering source—the ordering source invoked by sentences of the form *to get to Union Square from here, one must/should. . .*, given by $\{\llbracket \text{one gets to Union square from here} \rrbracket\}$ —that B’s utterance selects for update. On this suggestion, the result of updating with B’s utterance would be an updated ordering source: $\{\llbracket \text{one gets to Union square from here} \rrbracket, \llbracket \text{one takes Broadway to 14th} \rrbracket\}$. Against appearances, imperative instructions can be understood as proposals for addition to a teleological ordering source.

Whether or not this rescues Ninan and Portner’s account of (2.26)—I do not think it does, but won’t argue that here—there are good reasons not to like this analysis for imperative instructions. There is a basic intuition that the force of B’s utterance is *issue-resolving* (hence, *informative*), as well as both *non-directive* and *non-practical*. And this intuition is supported by suggestive linguistic data. Notice, for instance, that linguistic assent and denial are

clearly licensed as replies to imperative instructions in dialogue.

(2.29) A: How does one get to Union Square from here?

B: Take Broadway to 14th (for example).

A: Actually, that's wrong. Broadway is closed today. (*or*: You know, you're right.

I'd forgotten that you can get to Union Square via Broadway.)

Linguistic assent and denial are *relational* speech acts: they establish a specific relation between their agent and *prior assertions* in a discourse; they are, we might say, anaphoric to prior assertions (see esp. Asher & Lascarides 2001, 2003). By contrast, the sorts of relational speech acts we would expect to be licensed if some sort of addition to an ordering source had been proposed are not licensed.

(2.30) A: How does one get to Union Square from here?

B: Take Broadway to 14th (for example).

#A: Sure, okay.

Notice that Portner cannot explain this oddness by appeal to the claims that (i) the relational speech act expressed by *sure, okay* (which we might roughly dub "linguistic acquiescence") requires a command for its "antecedent," (ii) B's utterance proposes [*you take Broadway to 14th*] for addition to a *teleological* (as opposed to deontic) ordering source (hence, is not a command). This is because Portner's analysis of imperative suggestions/advice (Section 2.4.2) has them adding their prejacent to a teleological ordering source as well. But linguistic acquiescence *is* generally licensed as a reply to imperative suggestions/advice.

(2.31) A: I really want to make more progress on this paper.

B: Talk to your advisor more often.

A: Sure, okay.

Notice that the same points hold for strong modals in similar contexts.

(2.32) A: How does one get to Union Square from here?

B: You must take Broadway to 14th.

A: Actually, that's wrong. Broadway is closed today.

(2.33) A: How does one get to Union Square from here?

B: You must take Broadway to 14th.

#A: Sure, okay.

(2.34) A: I really want to make more progress on this paper.

B: You must talk to your advisor more often.

A: Sure, okay.

These data are surprising, on an account like Portner’s. They support an issue-resolving (assertoric) interpretation of B’s replies to A’s query about how to get to Union Square (whether B makes use of an imperative or strong modal), while also failing to support ordering source-updating interpretations of the same. Notice also that they make trouble for the Ninan-Portner explanation of the contrasts in (2.27) and (2.28).

- (2.35) A: How do I get to Union Square from here?
 #B: You must take Broadway to 14th, although I know that you won’t.
 #B: Take Broadway to 14th, although I know that you won’t.
- (2.36) A: How did I get here?
 B: You must have taken Broadway to 14th. (*obligatorily epistemic*)
 (or) #B: Take Broadway to 14th. (*epistemic reading unavailable, hence odd*)

Issue-resolving interpretations of imperatives and strong modals manifest the same interactions with expressions of speaker expectations and discourse about the past as ordering source-updating interpretations of the same. It is reasonable to expect these interactions to have a common explanation. Since issue-resolving interpretations of imperatives cannot be understood as proposals for addition to a To-Do List, the correct explanation of the contrasts in (2.27) and (2.28) should not be given in terms of imperatives/strong modals being interpreted as proposals for addition to a To-Do List.

2.4.5 Varieties of Permission

A similar objection arises for permission (and invitation) imperatives: the data tell against understanding them as cases of ordering source-addition. We addressed Portner’s (2010) treatment of permission imperatives in Section 2.3, but it is worth expanding on it here. Portner treats the distinction between permission and advice imperatives as TYPE-(II) variation: permission imperatives are distinguished from, e.g., advice imperatives by the fact that “a speaker imposes a requirement in order to *overcome the reluctance* of the addressee to undertake an action in his or her interests.” There is a pragmatic/sociological component of this analysis that I do not wish to dispute: the notion that permission imperatives are used by speakers, in part, to overcome addressee reluctance to pursue their prejacent. But an analysis in terms of ordering source-addition (glossed by Portner as “impos[ing] a requirement”—a bit roughly since, we’ve seen, addition to an ordering source does not generally result in the imposition of a new requirement) is not sufficiently general to cover all relevant cases.

Let \mathcal{P} be the ordering source that a permission imperative targets for update. As seen in Section 2.3, the Portner analysis succeeds in predicting the characteristic force of a permission imperative $!\phi$ —making it permitted that ϕ , i.e., that some of the $\leq_{\mathcal{P} \cup \{\llbracket \phi \rrbracket\}}$ -minimal worlds are ϕ -worlds—where there is, roughly, a standing prohibition on ϕ (i.e., $\llbracket \neg \phi \rrbracket \in \mathcal{P}$).

However, an imperative $!\phi$ can have a permission interpretation in contexts lacking such standing prohibitions. An addressee’s reluctance to realize ϕ (hence, also, a speaker’s ground for issuing a permission imperative) can be a function of *many things* besides a standing prohibition on ϕ . Portner (2010) proposes to handle such cases by suggesting that a standing prohibition is “accommodated”: $\llbracket\neg\phi\rrbracket$ is added to \mathcal{P} as a prerequisite to the interpretation of $!\phi$ and the addition of $\llbracket\phi\rrbracket$ to \mathcal{P} , so that the result of updating with a permission imperative in such a context is, not $\mathcal{P} \cup \{\llbracket\phi\rrbracket\}$, but rather $\mathcal{P} \cup \{\llbracket\neg\phi\rrbracket\} \cup \{\llbracket\phi\rrbracket\}$. Permission imperatives *presuppose*, in other words, the presence of a standing prohibition.

Positing the existence of such a presupposition is, however, *ad hoc* and does not appear to be supported by the data. For instance:

(2.37) A: It’s okay. Have an apple.

#B: Hey, wait a minute! I haven’t been told I *can’t* have an apple.

B’s response is odd because A’s grounds for issuing the imperative need not be to counter-balance a standing prohibition. A may just, for instance, be granting a right to B, where one did not exist previously. Notice that, if (i) the prior context contains no standing counter-vailing prohibition, and (ii) there is, indeed, no reason to suppose that one is added as a prerequisite to the interpretation of the permission imperative, Portner’s account predicts—apparently contrary to fact—that the typical effect of a permission imperative $!\phi$ (addition of $\llbracket\phi\rrbracket$ to a selected ordering source) is to make it *required*, with respect to the ordering source, that ϕ (recall the discussion of Section 2.3).

To buttress this general argument, I will describe two sorts of functions that permission imperatives can serve (and which all of the standard accounts have difficulty handling):

- Establishment of a STRONG PERMISSION
- Grant of a RIGHT

Strong Permission. The notion of strong permission is familiar from deontic logic (Anderson 1966 is the classic discussion). While its precise analysis is controversial, it can be intuitively glossed as follows. Strong permission that ϕ is what obtains when ϕ is *expressly*, as opposed to merely implicitly, permitted. The mere absence of a prohibition on ϕ (i.e., a weak, or implicit, permission that ϕ) does not mean that it is expressly permitted that ϕ (although, of course, express permission that ϕ implies absence of prohibition on ϕ).

What sort of role might the establishment of a strong permission play in a discourse? Here is one (among many others). Suppose it is antecedently weakly permitted for B to realize ϕ (because the authority A has failed to prohibit ϕ), but that B is nevertheless uncertain about whether the speaker wants that ϕ (and B has a strong interest in staying on

A's good side). Establishing a strong permission may function to alleviate this uncertainty.

(2.38) (*B looks unsurely at the fruit plate*)

A: It's okay. Have an apple!

(or) A: It's okay. You may take an apple!

The permission imperative and the performative *may* are, it seems, functionally equivalent in contexts like this: (i) they function to introduce a strong permission into the discourse; (ii) neither requires a standing prohibition be in force to be felicitous.¹⁷

So there is good reason to that strong permission imperatives are not *interpreted as* proposals for addition to a selected ordering source. Nevertheless, as was the case with instruction imperatives, we find that strong permission imperatives display the same interaction with "I know you won't" continuations:

(2.39) #It's okay. Have an apple, although I know you won't.

We have good reason to think the infelicity of (2.39) does not result from the imperative clause proposing addition to an ordering source. (2.39)'s infelicity seems to be due to something other than proposing addition to an ordering source, despite not thinking it possible that the addressee will realize the prejacent. What? Roughly, I will suggest that it is due simply to the fact that some imperatives conventionally, but defeasibly, express a kind of directive, or requirement-generating, speech act. Even when requirement-defeating interpretations are defeated by context, the fact that they are conventionalized means that they are not coherently cancelable by the speaker. (We will greatly expand on this explanation in Chapter III.)

In general, the weakest description of the function of a performative *may* is to introduce a strong permission into the discourse: all performative *may*'s function to introduce a strong permission; some function to do more besides (e.g., granting a right; see discussion below). Giving a complete account of the discourse role played by strong permission imperatives—an account of what, precisely, the establishment of a strong permission in a discourse consists in—is well beyond the scope of this dissertation (although we will have something to say in this direction in Sections 3.4.4 and 4.3.8). What is clear is that strong permission imperatives (like performative *may*'s) do not, contra Portner, require an antecedent countervailing prohibition in order to receive their intended, strong-permission-creating interpretations. Indeed, strong permission imperatives are paradigmatically licensed in cases where (i) a corresponding weak, or implicit, permission is in force, (ii) a speaker wishes to strengthen this weak permission into an explicit permission. In such contexts, treating permission

17. Portner replies (pc): "In this case, B would have *don't take an apple* on her [To-Do List], not as an explicit prohibition, but just out of caution." I don't think that is plausible. Being cautious about something doesn't mean planning, even provisionally, *not* to do it. It could just mean that one won't expressly plan *to* do it until one has more information. These are decidedly different states of mind.

imperatives as proposals for addition to a selected ordering source wrongly predicts that they typically establish requirements, as opposed to strong permissions.

Grant of right. Permissions can also function as *restrictions on future commands*: a speaker can issue a permission imperative (and a permission, generally) as a provision against future restrictions on the addressee's options.¹⁸ Suppose there has been a severe blizzard, and employee A calls into work to see whether she must come in that day. Suppose further that A's schedule is irregular—there is no standing expectation that she show up to work on a given weekday. There are a series of people in the company that can decide whether A must come in on any given day (B and C among them), but it is understood that whoever A talks to first cannot be overridden by any of the others. Consider the following dialogue.

- (2.40) A: Am I needed at work today?
B: No. Stay home.
(later, C, swamped, calls to see why A is not in the office that day)
C: We need you at work today! Get to the office immediately!
A: But B said it was okay to stay at home today!
C: Oh! My bad!

Intuitively, B's permission imperative in (2.40) has a two-fold function: it establishes a strong permission (for A to stay at home today), as well as a right (for A to resist certain future impositions on her plans, namely, those inconsistent with her staying at home today). Replacing the permission imperative with a performative *may* would not affect this. And neither the permission imperative nor the performative *may* requires the presence of a countervailing standing prohibition in order to have the intended, right-creating interpretations.

I do not want to get too deep into the pragmatic and social dimensions of grants of rights here (not least because this is a question best suited for sociolinguists and social psychologists, rather than philosophers). Suffice it to say that when the rules of the relevant discourse allow an authority to grant a right (and that grant is accepted by its subject), the rules of the discourse entail that later attempts of authorities to constrain the subject's options must be compatible with the subject's right. On the other hand, if it is known that the rules of the discourse allow C to revoke an earlier permission by B (in which case B's permission does not count as a grant of right at all), we observe different facts about which

18. For similar thoughts, see Veltman (2008); Charlow (2009). Veltman treats update on imperatives that are inconsistent with established rights as leading to the broken context \emptyset . This approach to permission is to be contrasted with, e.g., that of Lewis (1979a); van Rooy (2000) (call it the 'Lewis approach'). On the Lewis approach, permissive force consists solely in a one-off *expansion of allowable options* for the addressee, thus leaving it completely open to the authority to re-restrict the addressee's options at a later time. Both approaches to permission are necessary. The Lewis approach works well for cases where the speaker's intent is simply to reverse a standing prohibition. But prohibition-reversing permissions represent only one kind of permission, among many others.

relational speech acts are licensed. In such a context, A's objection fails to be licensed (except as an expression of petulance), while the following response becomes available to C.

(2.41) C: What does that matter? I'm the boss around here, not B.

Finally, as was the case with strong permission imperatives, we find that right-granting imperatives display the same interaction with "I know that you won't" continuations:

(2.42) A: Am I needed at work today?

#B: No, stay home, although I know you won't.

As before, such infelicities demand a different explanation than the sort suggested by Portner and Ninan. We will eventually tackle this below.

2.5 Recap

To recap, in Section 2.3, we argued that a Portner-style understanding of the conventional role of imperatives—on which imperatives rigidly express a general type of operation on ordering sources (addition of the prejacent), which in turn subsumes a wide range of speech acts—is sufficiently flexible to handle a range of imperatives' core (but, nevertheless, non-obligation-imposing) uses. But it is not, we have seen, flexible *enough*: there is a range of apparently core uses of an imperative of the form $!\phi$ which cannot be explained as either variation in the kind of ordering source to which $[\phi]$ is added (Type 1), or variation in the speaker's understanding of her relation to the ordering source to which $[\phi]$ is added (Type 2). Instruction and permission interpretation of imperatives cannot, in a wide range of cases, be modeled as ordering source-addition at all.

On Lewis' account, (i) an imperative $!\phi$ conventionally commands, and (ii) commanding is modeled as intersection of $[\phi]$ with the prior permissibility sphere. By any reasonable measure, Portner's account *strictly improves* on Lewis'; Lewis fails to accommodate instruction and permission interpretations (and plenty more besides). Note that Portner's account represents a *generalization* of Lewis'. As proved in Section 2.3, addition of $[\phi]$ to an ordering source \mathcal{P} constitutes something very much like commanding in Lewis' sense (restriction of the permissibility sphere to ϕ -worlds) in a class of contexts in which there are antecedently both ϕ and $\neg\phi$ -worlds that are \mathcal{P} -comparable to each other, and both ϕ and $\neg\phi$ are antecedently permitted; in such contexts, the updated ordering source will characterize a set of permissible worlds (a permissibility sphere), all of which satisfy ϕ .

So it might be thought that the proper reaction to the data is to *further generalize* the account: to identify a more general kind of speech act for imperatives to conventionally express—one that subsumes, not only commands, advice, and suggestions, but instructions, grants of rights, establishment of strong permissions, too.

This is the wrong way to react to the data. Unlike commanding, advice, and suggestions

(which, I will argue, may all be subsumed under the heading *necessitation with respect to a contextually salient ordering source*), such a class of speech acts does not plausibly comprise anything like a natural kind. So while it is, of course, possible to identify a disjunctive speech act-type that “subsumes” the speech acts of this class—the speech act of commanding, *or* advising, *or* suggesting, *or* instructing, *or* . . .—it is doubtful that this sort of *ad hoc* speech act will figure prominently in a properly explanatory account of the meaning of imperative clauses in natural language. An *ad hoc* kind of this type will ultimately explain nothing at all. Indeed, it is for reasons such as these that Kaufmann & Schwager to appear seem to object to any analysis of imperative meaning which has imperatives conventionally contributing a speech act to their interpretation at a context. If further generalizing the speech act for imperatives to conventionally express is a prerequisite of a descriptively adequate analysis in terms of conventionalized speech acts, any descriptively adequate analysis in terms of conventionalized speech acts will fail to be explanatorily adequate.

CHAPTER III

Imperatives and Necessitation

In this chapter, we will pursue a different tack. The lay of the argument will be this. We will identify a *core speech act* (which we will give the handy name “necessitation”—necessitation, namely, with respect to a salient, practical ordering source) which imperative utterances, as a matter of their conventional meaning, “express” (in a sense to be precisified). We will identify a class of *core uses* of imperatives on which they contribute the core speech act of necessitation to their final interpretation; this will include command uses, advice uses, suggestive uses, and some (but not all) permission uses. We will also identify a class of plausibly (although disputably) non-core uses of imperatives, on which they contribute something other than the core speech act of necessitation to their final interpretation; this class arguably includes at least some permission uses of imperatives, as well as central instruction uses. We will argue for assimilating uses of this class roughly to the class of so-called indirect speech acts (although there are some differences between this class and the class of indirect speech acts, as they are usually understood). All of this will be situated within a particular understanding of linguistic convention. On this understanding, imperative clauses may have a conventionalized necessitating use, but nevertheless fail to be interpreted, by an addressee, as expressing that use, on a given occasion of use. Still, the conventionalized use is plausibly seen to *contribute* in a precise and predictable way to the interpretation of an imperative utterance on such an occasion of use.¹

As we will see, the fruits of the necessitation analysis of the conventional use of im-

1. Portner (pc) suggests that, if my objection is that his theory only has trouble with *non-core* uses of the imperative, this is a good, rather than bad, result for his theory of imperatives. But that does not get my objection quite right. My objection is that Portner’s theory does not challenge the notion that linguistic conventions are infeasible interpretive facts. This notion has it that saying an imperative conventionally adds its prejacent to an ordering source is to say that, on *all literal uses of an imperative*, it adds its prejacent to an ordering source. Portner plausibly sees no need to challenge this notion because he thinks that, e.g., permission imperatives can be represented as proposals to add their prejacent to an ordering source. But this strategy is not generally effective: it works neither for permission imperatives nor for instruction imperatives. In a slogan, then, Portner’s notion of a core use for the imperative is *much too broad*. There are a variety of uses of imperatives on which they arguably do not add to an ordering source or necessitate their prejacent. The aim of this chapter is to develop an account which explains these uses, not as parametric variation in a single parametrized use, rather as non-standard uses, computable in a regular and predictable way from the imperatives’ core use, together with facts about the context of utterance.

peratives are many, and we will be focused on developing them for most of the rest of this dissertation. First, and foremost, this analysis gives a satisfyingly unified account of the wide range of imperative uses, when an imperative occurs either syntactically unembedded or semantically or discursively unembedded (for convenience, we will refer to all such occurrences using the syntactic idiom “unembedded”). Second, as we will see in detail in Section 4.4, the analysis suggests a constructive template for theorizing about the meaning of imperatives occurring under arbitrary syntactic, semantic, and discursive environments and conditions. The template leads to an account with wide empirical coverage—much wider, and with fewer methodological difficulties, than any competitor account. I will argue that this evidence, considered together, indicates that a plausible *theory of meaning* for imperatives (and related constructions—e.g. modals with some sort of practical normative import) can be grounded in and based on a theory of their conventional use.²

3.1 Indirect Speech Acts: A Primer on Defeasibility

As I said above, we will here opt for a different generalization of Lewis’ account than Portner’s—one that does *not* model permission and instruction interpretations of imperatives as variation in a parametrized single use for imperatives, but which instead treats such uses roughly on the model of (3.1).

(3.1) Can you pass the salt?

As a matter of clause-type, (3.1) is an *interrogative*. Interrogatives, it is widely agreed, paradigmatically function to introduce issues into a discourse (i.e., to perform the speech act of *questioning* or *interrogating*). But (3.1) is paradigmatically not used to introduce any issue; it is used to make a request.

There are two ways to react to this fact:

- Deny that an account of the conventional meaning of interrogatives should encode their typical issue-introducing function.
- Offer a non-semantic explanation for the request interpretation of (3.1). (Or, less popularly, offer a semantic explanation for the request interpretation, but nevertheless retain the encoding of the issue-introducing function of the interrogative in its semantics.)

As a sociological fact, accounts of the meaning of interrogative clauses do not take cases like (3.1) to pose a problem for encoding their issue-introducing function in their representations of interrogative meanings. Indeed, all major treatments of which I am aware—static and dynamic alike—analyze interrogative meanings, statically, in terms of *sets of alternative answers* and, dynamically, in terms of *introducing issues* into a discourse (by, for instance,

2. The ramifications of this general result for linguistic theorizing about meaning in the Dynamic and Expressivistic traditions will continue be developed and pursued in the remainder of this dissertation.

partitioning the context set into worlds compatible with the question's various alternative answers) (e.g. Groenendijk & Stokhof 1984; Roberts 1996; Groenendijk & Roelofsen 2009).

But such an interpretation, of course, is inappropriate for representing the typical interpretation of a polar interrogative like (3.1): (3.1) typically functions, not to introduce an issue into a discourse, but instead to request, of a salient addressee, that she pass the salt. This creates a *prima facie* problem for the notion that (3.1) conventionally expresses a polar question, with mutually exclusive and exhaustive alternative answers (i.e., *yes* and *no*). For one thing, a speaker who utters (3.1) and intends its typical interpretation can be expected to know the answer to this question in advance; she could not properly intend this interpretation if she did not assume her addressee were capable of passing the salt. For another, an addressee who knows the speaker intends (3.1)'s typical interpretation cannot felicitously respond simply *yes, I can* or *no, I can't*.

- (3.2) A: Can you pass the salt?
(B replies, pleasantly, but without lifting a finger)
#B: Yes, I can. [#No, I can't.]

If a speaker is known to be making a request, a felicitous reply must somehow involve action, in speech or otherwise, to address the speaker's request. A *yes, I can* response is typically appropriate only when accompanied by an effort by the addressee to fulfill the request, as in (3.3); a *no, I can't* response is typically appropriate only when accompanied by some sort of explanation of the addressee's unwillingness to pass the salt, as in (3.4).

- (3.3) A: Can you pass the salt?
(B reaches for the salt)
✓B: Yeah, sure thing!

- (3.4) A: Can you pass the salt?
(B gives A an icy glare)
✓B: No, I think you've had enough salt today.

All this is just to say that a speaker who utters (3.1) is not typically interpreted, by an addressee, as expressing a polar question, instead as expressing a request.³ So if, as per the philosophical and linguistic consensus, (3.1) does conventionally express a question (or set of alternative answers), then the sense in which it does so cannot be one which requires any literal utterance of (3.1) to be interpreted, by an addressee, as expressing a polar question.

In the spirit of charity to the philosophical and linguistic consensus, then, it seems worthwhile to search for a sense of ϕ conventionally expresses P which does not require

3. The technical idioms available to me at this point are a bit too clumsy to state the point precisely. There is, we'll see, a sense in which such a speaker is interpreted as expressing such a question. There is also a sense in which she is not. One useful property of the apparatus developed here is that it will allow us to distinguish precisely between these two senses.

that any literal utterance of ϕ be interpreted, by an addressee, as expressing P . One such sense makes use of Searle (1975a)'s notion of the *indirect speech act*—roughly a speech act performed by a speaker via the performance of another speech act. In the case of (3.1), indirect speech act analyses (hereafter ISAs) will hold that:

- The speaker, in a sense, both makes a request and asks a question.
- The speaker's asking the question has (what I will call) *indirect interpretational significance*: the fact that the speaker's utterance expresses a question, together with certain facts about the context of utterance (e.g., the triviality of the issue expressed by the question relative to the context-relevant information), generate the request interpretation.
- The speaker's asking the question has *merely* indirect interpretational significance. Although the speaker's utterance asks a question (in the sense that she makes an utterance that expresses, as a matter of its conventional meaning, a question), she does not intend to communicate a question (i.e., to actually introduce an issue into the discourse) through her linguistic performance. This, together with certain facts about the context of utterance, causes her not to be *interpreted as* asking a question.

I admit the theoretical idioms deployed here are not entirely natural. What, e.g., does it mean to ask a question (and have it be common ground, in view of the conversational participants' presumed linguistic competence, that one is asking a question), but fail to be interpreted as asking a question? Perhaps there is something to say here, but I will just encourage the reader not to rely on the idioms too heavily. It suffices to say that, on the view being considered here, (3.1) *contributes a question* to the interpretive hopper—the thing that takes a conventionalized speech act and facts about a context and spits out something we can gloss, roughly, as what speech act a speaker communicates with her utterance at that context—and what the interpretive hopper typically *spits out* is a request.

More specifically, the sort of ISA I want to sketch for (3.1) has the following sort of informal structure. Let c be a typical context at which (3.1) is uttered, S_c be c 's speaker, and A_c be S_c 's addressee. A request interpretation for S_c 's utterance is generated, and a question interpretation vitiated, by means of the following chain of conversational reasoning.

- (3.1), in virtue of its clause-type, contributes a question to its interpretation at c .
- S_c is presupposed at c to know the answer to this question; the issue expressed is *trivial* relative to the c -relevant information. It is inferred that S_c does not intend to be interpreted as proffering a question. So S_c is not interpreted, in the final instance, as proffering a question.

[I use *proffer* as a relatively neutral label for whatever speech act(s) a speaker is, in the final instance, interpreted as expressing. When the question is expressed by a speaker, without being proffered, I will say that it is *voided*.]

- It is known that S_c intends to do *something* with her linguistic performance; her linguistic performance is, after all, intentional. What does S_c intend to do? The

fact that (3.1) contributes a question whose “subject matter,” loosely speaking, is the state of affairs in which A_c sees to it that S_c has the salt (together with some background knowledge of S_c ’s psychological state) generates the conclusion that S_c desires the salt, and would like A_c to realize the salient state of affairs.⁴

It’s important to note that the conclusion that a speaker does not intend to be interpreted as proffering a question, generated in Step (ii) of our reconstruction, can be arrived at in multiple ways. If, for instance, it is common ground that a speaker *does not care* about the answer to a question, even if she is known not to know the answer, she will, *ceteris paribus*, not be interpreted as asking a question. So the explanation developed here will extend to variations of (3.1) that do not make use of ability constructions, e.g., (3.8) and (3.9).

(3.8) Would you like to pass the salt?

(3.9) How *ever* am I going to get that salt?

It is also important to note that the voiding of a speech act does not mean that the speech act is *unavailable* as the antecedent for relational or anaphoric speech acts. It is, for instance, generally infelicitous to assent to a request using the sorts of replies that would be appropriate for a polar (*yes/no*) question.

(3.10) A: Please pass the salt.

#B: Yes. [#No.]

4. For stories in a similar vein, see [Gordon & Lakoff \(1975\)](#); [Searle \(1975a\)](#). As with most attempts to reconstruct supposed conversational reasoning, there is strong Just So Story flavor to this explanation (cf. [Morgan 1975](#)). A better explanation is that the regularity of contexts in which such implicatures are generated has led to the existence of a well-established linguistic convention in English (and some Romance languages, e.g., French) on which INTERROGATIVE + CONDITIONAL ASPECT constructions function to express *polite requests* (cf. [Searle 1975a](#)).

(3.5) Pourrais-tu me passer le sel?

[Lit. *would you be able to pass me the salt?*]

(3.6) Voudrais-tu me passer le sel?

[Lit. *would you like to pass me the salt?*]

In the helpful idiom of [Morgan \(1975\)](#); [Asher & Lascarides \(2001\)](#), the request interpretation of such constructions, initially generated as a conversational implicature, is eventually “short-circuited” by being lexicalized. As [Asher & Lascarides](#) note, this also explains the *relative ease* of getting the request interpretation for such constructions, when compared to semantically equivalent paraphrases like (3.7):

(3.7) Do you have the ability to pass the salt?

If the request interpretation of (3.1) were a genuine implicature, such an interpretation should be readily available for its paraphrase (3.7), since “Paraphrases typically have the same calculable implicatures” ([Asher & Lascarides 2001](#): 184).

Significantly, this sort of view opens up the possibility for a different explanation of both (i) how (3.1)’s request interpretation is triggered, (ii) how (3.1)’s question interpretation is voided. Regarding (i), if (3.1)’s request interpretation is conventionalized, S_c will express, simply by uttering (3.1), a request that A_c pass the salt. Regarding (ii), in making this request, S_c will *presuppose* that A_c is able to pass the salt; generally, one cannot rationally request that A do X unless one presupposes that A is able to do X . Hence a question interpretation of (3.1) is unavailable (as one of its answers will be known to be presupposed by S_c). This suffices to void the interrogative force of (3.1).

Not so for request interpretations of questions. As illustrated by (3.2), (3.3), and (3.4), a question act can serve as the antecedent for anaphoric speech acts (hence licensing *yes/no* replies) without the speaker of that question being interpreted by her audience as proffering that question; as Asher & Lascarides put it, though the interrogative utterance proffers a request, it nevertheless bears some of the “linguistic hallmarks” of a question (2001: 185). As (3.3) and (3.4) show, such replies are licensed in the case of (3.1). But, as (3.2) shows, a respondent generally will be regarded as uncooperative if she interprets the speaker of (3.1) as *proffering* a question (for, if she were interpreted as proffering a question, one would expect “bare” *yes/no* replies to be issue-resolving, hence at least *partially* responsive to the speaker’s linguistic performance, hence pragmatically licensed⁵). (Indeed, these facts together constitute solid, independent, evidence for the conclusion that the question act is contributed, without being proffered, by the speaker’s utterance; that is to say, they constitute evidence that the question act is voided by the context.)

An analysis in this vein is, I think, well-suited to accounting for a conventionalized link between clause-type and use in the case of interrogatives. Working off that assumption, I will develop the outlines of a somewhat more formal account of this analysis in the next section. After that, I gesture at how both the informal and formal versions of the account of interrogatives can be extended to accounting for the wide range in functional potential for imperative clauses. If, as I’ll argue, an adaptation of this sort of view to the wide range of uses for imperatives is suited to predicting the relevant data, this will vindicate the necessitation analysis of the conventional role of imperatives against both (i) advocates of further generalizing the speech act expressible by an imperative to subsume instructions, grants of permission, etc., and (ii) critics, like Schwager (2006); Kaufmann & Schwager (to appear), of any analysis that attempts to assign imperative clauses a conventional discourse role.

3.2 Conventions as Defaults

In this section, I want to give the outlines of a system that is sufficiently flexible to represent the range of speech acts a clause of a certain type might be used to express, but which also leaves room for speech acts to be associated, as a matter of linguistic convention, with typed clauses. It is important for the reader to understand that my goals in this section are modest: I want to sketch a system which is able to accommodate such interpretive facts, by way of responding to the notion that these cannot be interpretive *facts*, because it is not possible to construct a system of interpretive representation that accommodates them.

The structure of the account I will be pursuing here is plotted as follows.⁶ We will use

5. Relevant here is the pragmatic licensing of partial answers in response to question acts. On this topic, see, e.g., Groenendijk & Stokhof (1984); Groenendijk & Roelofsen (2009).

6. The account presented here draws heavily on Asher & Lascarides (2001, 2003), in particular their emphasis on “speech act-related goals.” One central difference will be that Asher & Lascarides regard semantic type as determining speech act type; in the case of an indirect speech act, an incompatible *semantic type* is inferred via

\Rightarrow and $>$ indicate, respectively, *monotonic* and *non-monotonic* inferences within the theory of interpretation (or, if you prefer, monotonic and non-monotonic conditionals linking descriptions of utterances to facts about interpretations):

Clause Type, Semantic Type, Speaker Goals, and Interpretation

CLAUSE TYPE	\Rightarrow	SEMANTIC TYPE
SEMANTIC TYPE	$>$	SPEAKER'S GOALS
SPEAKER'S GOALS	\Rightarrow	USE/SPEECH ACT TYPE

Here is some notation to make the presentation more precise. Let ϕ be any clause-typed sentence and c be any context of utterance. Then:

- $\llbracket \phi \rrbracket^c := \phi$'s model-theoretic denotation⁷ at c
- $\phi \Rightarrow \mathcal{I} := \phi$ has interpretation \mathcal{I} at any context whatever
- $\phi >_c \mathcal{I} := \phi$ has interpretation \mathcal{I} at context c
- $\phi > \mathcal{I} := \mathcal{I}$ is a *conventional interpretation* for ϕ

I will usually understand \Rightarrow and $>$ as relations holding between utterance types and speech act types (or, if you like, descriptions of speech acts types); $>_c$ will be a relation holding between utterance tokens and speech act tokens. (I will sometimes speak as if these relations relate sentences to interpretations, but this is just an economical way of referring to utterances of these sentences by a speaker, either on an occasion of use, or qua type.) \mathcal{I} will count as a conventionalized interpretive feature of ϕ just if, roughly, at any *normal* or *default* context c , $\phi >_c \mathcal{I}$. Claims of the form $\phi > \mathcal{I}$ will be termed *interpretive defaults*.

3.2.1 Convention and Cancellation

What is the force of saying that some fact or other is an interpretive default for a clause? Here I don't have all that much to say (I will have somewhat more in Section 3.2.2). I will not be attempting to state any sort of semantics for non-monotonic conditionals here, instead favoring suggestive, informal descriptions of their inferential behavior (e.g., as defaults). Here are some gestures in that direction. In one standard formalization, such defaults are styled as *defeasible axioms* for an interpretive theory (Asher & Lascarides 2001, 2003). Defaults are distinguished from calculable implicatures, which are not lexicalized and do not have the status of interpretive axioms, but which arise when lexical information interacts with knowledge concerning the specific context, general norms of conversation,

Gricean reasoning, thus allowing the inference of an incompatible speech act type (see, e.g., Asher & Lascarides 2001: 195). This strikes me as a rather radical view of the semantics/pragmatics interface. (As we saw above, Portner 2007, 2010 also gives speaker goals—or, as he terms them, “grounds for issuing an imperative”—a prominent role in explaining interpretive variability.)

7. Model-theoretic denotations are assumed here to be constructions out of possible worlds models for a first-order language. We will have more on the model theory below.

and the psychologies of one's interlocutors, and which can typically be readily manipulated by a speaker. In particular, calculable implicatures admit of *cancellation* (see esp. Grice 1989; Sadock 1978).

Interpretive defaults, in the intended sense, although overridable, do *not* admit of cancellation; non-cancelability will be treated here as the characteristic *linguistic hallmark* of an interpretive default. What I will mean by this is that it is not possible to *coherently* attempt to override an interpretive default for a sentence ϕ in the course of uttering ϕ . The lack of coherence typically manifests as linguistic infelicity as in (3.11), while, conversely, coherent cancellation typically manifests as linguistic felicity as in (3.12).

(3.11) #He didn't quit smoking. Indeed, he never smoked at all.

(3.12) ✓Some hipsters smoke. Indeed, all do.

It might be tempting (especially when focusing on presuppositional examples of conventional implicature) to define cancelability *semantically*, so that an implicature ψ of ϕ is cancelable just if $\{\phi, \neg\psi\}$ is *consistent*. Clearly this will not do in general, but especially for my purposes. I take, for instance, declaratives, tokened unembedded and in a sincere discourse, to conventionally express assertions (and assertions to constitutively involve acceptance of the asserted proposition by the speaker); in short, if ϕ is declarative, an utterance of ϕ conventionally implicates the speaker's acceptance of ϕ . But if cancelability is defined purely in terms of semantic inconsistency, this cannot be possible. The following sentence, though odd, is nevertheless consistent (indeed, *Moore Paradoxical*).

(3.13) #It's raining, but I don't accept that it's raining.

We thus need a broader notion of cancelability to capture the sense in which the implication that a speaker who utters ϕ (and thereby expresses an assertion that ϕ) accepts that ϕ is non-cancelable. In place of the semantic notion, I suggest a *rationality-centric* understanding of coherence and cancelability—one that invokes rational relationships between speech acts, rather than between propositions.

COHERENCE

A sequence of speech acts $\langle \alpha_1, \dots, \alpha_n \rangle$ is coherent at c iff updating c with α_1 , then . . . then α_n does not lead to an irrational context.⁸

NON-CANCELABILITY

An implicated speech act β is non-cancelable with respect to a speech act α at c iff, for any γ such that $\langle \beta, \gamma \rangle$ is incoherent at c , $\langle \alpha, \gamma \rangle$ is incoherent at c .

8. The use of this sort of notion of coherence to handle sentences like (3.13) has some affinities with Thony Gillies' analysis of Moore's Paradox (Gillies 2001) (and more generally with Dynamic Semantic treatments of discourse coherence). We will have more to say in the way of characterizing the rationality of a context (or, more generally, of a cognitive state) in later sections.

The idea here is just that a speech act β is a non-cancelable implication of a speech act α just if any way of attempting to cancel β leads to incoherence with α . My major claim about the nature of linguistic convention (to be supported indirectly by the viability of a theory of conventional interpretation for imperatives built around it) can now be expressed as follows:

CONVENTIONAL INTERPRETATIONS ARE NECESSARILY NON-CANCELABLE

If $\phi > \beta$, then for any context c and speech act α such that $\phi >_c \alpha$, β is non-cancelable with respect to α at c .

In other words, if ϕ conventionally expresses β , then any attempt to cancel β at any context whatever will lead to incoherence at that context. Notice that this does not require, contra Schwager (2005); Kaufmann & Schwager (to appear), that ϕ “strictly determine” its interpretation, i.e., at any context whatever. For a sentence ϕ to conventionally express a speech act β , it is required only that β be non-cancelable with respect to ϕ ’s actual interpretation—that an interpretation that is assigned to ϕ at a context be coherent with β . So, for instance, holding that an interrogative like (3.1) conventionally expresses a question does not require that (3.1) be *interpreted as* expressing a question in any context whatever. It requires only that any use to which (3.1) be put be compatible with (i.e., not necessarily non-cancelable with respect to) the question it conventionally expresses.⁹

Notice that the welcome results of a semantic understanding of cancelability are captured as special cases, in view of our conventionalizing the association of proposition-denoting utterances (declaratives) with the speech act of assertion.

CANCELABILITY OF IMPLICATED PROPOSITIONS

An implicated proposition ϕ is cancelable with respect to a speech act α at c only if $\langle \alpha, \triangleright \neg \phi \rangle$ is coherent at c (where \triangleright represents assertion).

Whenever, e.g., $\{\phi, \neg\psi\}$ is inconsistent, $\langle \triangleright\phi, \triangleright\neg\psi \rangle$ will be incoherent; hence, if ϕ implicates ψ , ψ will be a non-cancelable implication of a speaker’s assertion that ϕ . Notice, though, that the *speaker’s acceptance* of ϕ will also count as a non-cancelable implication of her assertion that ϕ , on the plausible assumption that a context that is updated with both (i) a speaker’s assertion that ϕ , (ii) the assertion that the speaker does not accept that ϕ , is irrational. Updating with a speaker’s assertion that ϕ constitutively involves representing the speaker as accepting that ϕ . Updating further with the assertion that the speaker does not accept that ϕ leads to irrationality. Hence, both the proposition that the speaker accepts ϕ and the assertion whose content is that the speaker accepts ϕ are non-cancelable implications of the speaker’s assertion that ϕ .

9. If you’re thinking *didn’t he just suggest that these speech acts are incompatible?*, I’ll address this a bit below. The answer, to preview, is that the question act and the request act expressed by (3.1) are compatible. It is the presuppositional act expressed by the request that is incompatible with the question act. But this presuppositional act is not part of the *linguistic* interpretation of (3.1), as such. It is contributed, rather, by the interpreter’s theory of mind.

3.2.2 The Mechanics of Interpretation

With this in mind, let us return to the view of linguistic conventions sketched in this section's introduction, and try to flesh out a rough theory of interpretation that is consonant with it.¹⁰ First, some notation to make our lives easier here:

- Lowercase letters $u, v, x, y, z \dots$ designate utterances, with $speaker(u)$ the speaker of u , $addressee(u)$ the addressee of u , and $context(u)$ the context of utterance of u .
- If an utterance u consists of the sentence ϕ being uttered at c , $\llbracket \phi \rrbracket^c := \llbracket u \rrbracket$.
- $u : \tau$ indicates that u is of clause-type τ , while $\llbracket u \rrbracket : \tau$ indicates that u 's denotation is of type τ . So, for instance:
 - $u : ! := u$ is imperative
 - $u : ? := u$ is interrogative
 - $u : \triangleright := u$ is declarative
 - $\llbracket u \rrbracket : s$ (where s is the type of subsets of W) $:= u$ denotes a proposition

To summarize the view of linguistic conventions on offer, then:

- A clause's type (declarative, interrogative, imperative) *determines* the semantic type of its model-theoretic denotation (proposition, partition, ...). Thus, for any context c and semantic type τ , if $\llbracket \phi \rrbracket^c : \tau$, then for any context c' , $\phi \Rightarrow \llbracket \phi \rrbracket^{c'} : \tau$ and $\phi > \llbracket \phi \rrbracket^{c'} : \tau$.
- A clause's semantic type *normally determines* the category of its use (assertion, question, or necessitation).
- Different uses are individuated in terms of properties of the speaker's communicative goals or intentions.¹¹ As such, facts about communicative intentions determine the type the speech act proffered by an utterance.

It is a *consequence* of this picture that there is a conventional, but nevertheless defeasible, relationship between clause type and use type: since clause type determines semantic type, semantic type typically determines speaker goals, and speaker goals determine use type, it follows that clause type typically determines use type.

Interpreters may (and it is well-known that they do) rely on interpretive defaults such as these—not all interpretive inferences are usefully represented as the application of “general” rules (rules covering every subclass of a class of contexts) to specific cases, and interpretation, in any case, must often take place in the absence of any explicit cues for

10. The theory developed in this section (and its associated formalism) follow closely that of Asher & Lascarides (2001, 2003), though there are some important differences. Most significantly, as I will emphasize below, we do *not* assume that an utterance's semantic denotation determines use (a view which requires that semantic denotations be re-computed “on the fly” in cases where an utterance receives a non-default interpretation).

11. This is a view about the metaphysics of speech acts that I happen to endorse, on its strongest understanding, for independent reasons. Dealing adequately with this would take us too far afield. In the text, I will generally rely on the weakest understanding of this claim that is needed for the relevant argument.

applying such rules (see esp. Asher & Lascarides 2003). For some contexts, an interpretation \mathcal{I} for a clause ϕ is generated by deploying a conditional of the form $\phi > \mathcal{I}$, rather than one of the form $\phi \Rightarrow \mathcal{I}$. Crucially, for our purposes, adding information about speaker goals may block the inference from clause type to use type. If normal goals for uses of ϕ are given by G (and G yields interpretation I), but the speaker's *actual* goals at c are non-normal (and described by G'), we have: $\phi > G$; $(G \wedge \phi) \Rightarrow I$; therefore, $\phi > I$; but $(G' \wedge \phi) \not> I$ and $\phi \not> I$. It is important to note the role of the non-monotonicity of $>$ in securing this result: from $\phi > I$ it does not follow that $(G' \wedge \phi) > I$. An interpreter can, relying on a default of the form $\phi > \mathcal{I}$, generate interpretation \mathcal{I} for ϕ . Enriching the interpreter's information about the context (in particular, its information concerning the communicative goals of the speaker) can, however, generate a different interpretation for ϕ .

As noted above, some defaults might be styled as defeasible axioms for an interpretive theory. Whether or not we want to dub *logical consequences* of the axioms 'defaults' is mostly a matter of preference. I have been supposing that we do (indeed, I have styled something a default interpretation for ϕ just if it is part of ϕ 's conventional interpretation). But that is strictly optional. We might, for instance, introduce a monotonic axiom linking clause type to semantic type and a non-monotonic default linking semantic type to speaker goals. (Here, τ_{IMP} serves as a stand-in for whatever type of semantic denotation imperatives end up being assigned, sentences of the form $\mathcal{I}_\alpha \phi$ express that α intends that ϕ , and $R(u)$ designates the state of affairs enjoined by an imperative u .)

INDIRECT CLAUSE TYPE-INTERPRETATION LINK

$$\begin{aligned} u : ! &\Rightarrow \llbracket u \rrbracket : \tau_{\text{IMP}} \\ \llbracket u \rrbracket : \tau_{\text{IMP}} &> \mathcal{I}_{\text{speaker}(u)} \mathcal{I}_{\text{addressee}(u)} R(u) \\ \therefore u : ! &> \mathcal{I}_{\text{speaker}(u)} \mathcal{I}_{\text{addressee}(u)} R(u) \end{aligned}$$

Writing these down as axioms will mean that, in normal contexts, interpreters exploit semantic information to determine the speaker's goals; in the case of imperatives, that the speaker intends the addressee to realize the state of affairs her imperative enjoins. From this it *follows* that clause type typically determines speaker intentions (as indicated by the \therefore) but this need *not* be written down as an interpretive axiom. An alternative formalization might have it that semantic information, as such, is explanatorily inert in generating speech act interpretations; clause-type settles semantic value, and is also exploited in reasoning about the goals a speaker has in making an utterance in default contexts.

DIRECT CLAUSE TYPE-INTERPRETATION LINK

$$\begin{aligned} u : ! &\Rightarrow \llbracket u \rrbracket : \tau_{\text{IMP}} \\ u : ! &> \mathcal{I}_{\text{speaker}(u)} \mathcal{I}_{\text{addressee}(u)} R(u) \end{aligned}$$

Although *we* will not at this point make much of the differences between these alternatives, it will be useful to note that they offer a handy way of distinguishing different

linguistic methodologies that are often run together. The general sort of view articulated by Portner (esp. in his 2004a) would be best represented within the former axiomatization (although note that Portner himself does not go in for the view of linguistic convention we are developing here). According to Portner, imperatives’ characteristic conversational role is not assigned at the level of semantic interpretation (or even interpretation of the clause, as such), but rather as a result of deploying a general conversational norm linking utterances of a specific semantic type (proposition, partition, etc.) with specific types of conversational roles (assertion, issue-introduction, etc.). Accounts of imperatives within the “Dynamic Semantic” tradition—views which assign conversational roles (“update potentials”) as a matter of clausal interpretation (e.g. Veltman 2008; Starr 2011)—will tend to endorse sentences linking clause-types to conversational roles as axioms, while views closer to Portner’s will endorse them as derived consequences of other axioms.

We can classify things a bit more clearly if we add to our notation devices for distinguishing properly linguistic axioms—axioms derived from the grammar proper, the set of which we’ll denote as \mathcal{G} —from axioms concerning general norms of conversation and the psychologies of one’s interlocutors—the sorts of axioms that are “relevant to cognitive modelling,” the set of which we’ll denote as \mathcal{M} (Asher & Lascarides 2001: 209). A Portner-style view would be best characterized as endorsing the following claims (here, to be clear, \Rightarrow and \vdash represent monotonic and non-monotonic consequence relations within the linguist’s theory, rather than the theory of interpretation exploited by the interpreter):

$$\begin{aligned} \mathcal{G}, u : ! &\Rightarrow \llbracket u \rrbracket : \tau_{\text{IMP}} \\ \mathcal{M}, \llbracket u \rrbracket : \tau_{\text{IMP}} &\vdash \mathcal{I}_{\text{speaker}(u)} \mathcal{I}_{\text{addressee}(u)} R(u) \\ \therefore \mathcal{G}, \mathcal{M}, u : ! &\vdash \mathcal{I}_{\text{speaker}(u)} \mathcal{I}_{\text{addressee}(u)} R(u) \end{aligned}$$

Whereas a “Dynamic Semantic” theory will tend to endorse claims like the following:

$$\mathcal{G}, u : ! \vdash \mathcal{I}_{\text{speaker}(u)} \mathcal{I}_{\text{addressee}(u)} R(u)$$

The “Dynamic Semantic” theory represents conventional conversational role as a properly grammaticalized convention, while the Portner-style theory represents it as non-grammaticalized (but nevertheless conventionalized) convention.

Whether or not there are *empirical* differences to decide between these points of view—in Section 5.3, I will express skepticism that there are (contra, e.g., Starr 2011)—philosophical and/or methodological reasons are quite frequently given for preferring one theoretical orientation to the other. So it is a good idea to have some means of distinguishing them. Indeed, such distinctions will be a major theme of the chapter on Expressivism.

That ends our discussion, brief though it was, of the mechanics of interpretation, on the picture of linguistic convention we have endorsed. We will have more to say on the mechanics of interpretation—in particular, in cases where indirect speech act interpretations must be computed—in later sections.

3.2.3 Data for Declaratives and Interrogatives

To close this section, some evidence that assertion and questioning are, in fact, conventionally associated with clauses of the declarative and interrogative types, respectively. Unsurprisingly, the evidence takes the form of non-cancelability data. Specifically, we will consider *abnormal* contexts—contexts in which the relevant sentences lack their conventional interpretations. In such contexts, we find that attempts to cancel the speech act conventionally associated with the clause lead to some sort of irrationality in the context.

Suppose A wishes to signal to B the arrival of their compatriot C. But A and B are being closely watched by D, who will punish them if A attempts to relay any such signal. So A and B agree in advance that A will remark on the shininess of D's shoes when C is to arrive. Clearly, the force of (3.14) in this context is to signal to B the arrival of C, rather than to assert that D's shoes are shiny today. (Let's stipulate that D's shoes are actually quite dull, so that it is clear that the intended effect of this remark is not to inform B of a feature of D's shoes.)

(3.14) D's shoes are shiny today!

So, though (3.14) is normally used to assert that D's shoes are shiny today, in this context it is evidently used to do something else. Nevertheless, note the inability of A to append continuations that express speech acts—speech acts we would expect to be entirely appropriate, given the actual dullness of D's shoes—that are incoherent with this assertion.

(3.15) #D's shoes are shiny today, and they are not shiny today.

(3.16) #D's shoes are shiny today, but I don't accept that D's shoes are shiny today.

I take it that the hashmarks indicate incoherence or uninterpretability. Even in a context such as this, where it is arguably common ground that the declarative (3.14) is not even being used *sincerely*, the assertion typically expressed by (3.14) is non-cancelable with respect to the speech act that it is, in this context, used to express. Cases like this for declaratives can be multiplied ad infinitum.

Similar data can be found for interrogatives. Although (3.1) is typically used to make a request, it is *typed* as an interrogative, and, as such, I claim, conventionally expresses a question. We predict that the question act conventionally expressed by (3.1) is non-cancellable with respect to its request interpretation. And that prediction seems to be borne out in (3.17) and (3.18):

(3.17) #I know you can pass the salt, but could you pass the salt?

(3.18) #I'm not interested in whether you can pass the salt, but could you pass the salt?

Example (3.17) might seem problematic, in view of our Gricean explanation in Section

3.1 of how the request interpretation of (3.1) is generated. In that explanation, the speaker’s lack of interest in the answer to the question expressed by (3.1) plays a central role; indeed, the speaker is typically *presupposed* to know the answer to this question. We might say that (3.1) “carries” the presupposition that the speaker knows the answer to the question expressed. In uttering (3.1), then, a speaker typically is interpreted as expressing two speech acts:

- The act of presupposing that her addressee is physically capable of passing the salt.
- The act of requesting, of the addressee, that she pass the salt.

But this is just what (3.17) typically does. So why, on our account, is (3.1) admissible, even when (3.17) is not? The only relevant difference is that (3.17) expresses the presuppositional act as a matter of convention (in particular, the lexical properties of the factive verb ‘know’), while (3.1) expresses the presuppositional act *contingently*—specifically, in contexts where the question act conventionally expressed is voided. We’ll account for this difference by saying that a presuppositional act $\partial\psi$ is properly part of an utterance ϕ ’s *interpretation* only when it is contributed by ϕ itself, i.e., only when it is part of the conventional meaning of ϕ .¹² In other words, $\phi >_c \partial\psi$ only if $\phi > \partial\psi$. Whereas the presupposition in typical utterances of (3.1) is contributed by *background knowledge* about the context of utterance for (3.1); it has nothing to do with features of the utterance itself. On this picture, we will *predict* that, when at some context c an utterance of ϕ is interpreted as presupposing ψ ($\phi >_c \partial\psi$), and $\partial\psi$ is incoherent with the speech act conventionally expressed by ϕ , ϕ will be inadmissible at c if $\phi > \partial\psi$. Typical utterances of (3.1) will be admissible because the presupposition is not assigned by the theory of linguistic interpretation proper, but some other theory (e.g., the interpreter’s theory of mind or norms of interaction or whatever). Non-cancelability involves the demand to maintain coherence *within* the theory of linguistic interpretation.

3.3 Imperatives Express Necessitation

The account of bare imperatives (simple imperatives of the form $!\phi$) I will be defending in the remainder of this essay can be summarized as follows:

- Imperatives semantically express characteristic functions of sets of bodies of preferences (equivalently, functions from bodies of preferences into $\{0, 1\}$). In particular, the semantic value of an imperative clause $!\phi$ is the characteristic function of the set of bodies of preferences according to which ϕ is necessary.

12. This way is a tad stipulative (and probably a bit over-general). Still, I don’t think there is any doubt that the manner in which (3.17) expresses the presupposition in question is different from the manner in which (3.1) expresses it. Whatever that difference consists in, we will be able to construct a principled theory of cancelability around it.

- Imperatives conventionally express the speech act of *necessitation*: an utterance of an imperative $!\phi$ conventionally proposes to alter a relevant body of preferences so that ϕ comes to be necessitated by that body of preferences. Hence, there is a conventionalized relationship between an *utterance* of $!\phi$ and an *attempt* to make it the case that ϕ is necessitated by a relevant body of preferences.
- Imperatives can be properly used by a speaker who does not want it to be the case that ϕ comes to be necessitated by a relevant body of preferences—whose goal in uttering $!\phi$ is not to cause ϕ to be necessitated by a relevant body of preferences. In such cases, the correct interpretation for the utterance can be explained as an indirect speech act.

We will postpone a sustained discussion of imperative semantics here (although see Section 5.3 for some discussion). In the remainder of this section we will focus our attention on the latter two bullets—in particular, on developing a model for representing necessitation acts and showing how that model can accommodate non-normal uses of imperative clauses. (In Section 4.4, we will show how this model can be extended to account for a range of data—having primarily to do with *embedded* imperatives—that make trouble for all of the major competitor accounts.)

Supporting a necessitation account of the conventional use of imperatives will involve two things:

- Adducing data to show that imperatives do have a conventionalized use.
- Adducing data to show that conventionalized use can be understood in terms of necessitation with respect to a relevant body of preferences.

3.3.1 Conventional Use Data

With respect to the first task, it is useful to recall earlier examples (some of which are reproduced here) which show that both instruction and permission interpretations of imperatives bear the sorts of linguistic hallmarks that are characteristic of command interpretations of imperatives, even though it is difficult to represent these interpretations as proposals for modification of a salient ordering source.

- (3.19) A: How do I get to Union Square from here?
 #B: You must take Broadway to 14th, although I know that you won't.
 #B: Take Broadway to 14th, although I know that you won't.
- (3.20) #It's okay. Have an apple, although I know you won't.
- (3.21) A: Am I needed at work today?
 #B: No, stay home, although I know you won't.

Recall that the Portner-Ninan explanation of these infelicities is that one cannot, rationally, attempt to create a preference or intention while simultaneously expecting that that

preference or intention to go unfulfilled or to fail to be formed at all. This explanation misses the mark, I argued, because no such *attempt* is plausibly conveyed by any of the given imperative utterances. A speaker's aim in uttering the instruction imperative in (3.19) cannot be understood as preference-modification; her aim is to resolve an issue raised by her addressee. Similarly for the permission imperative (3.20); the speaker's aim is to establish a strong permission, rather than modify her addressee's preferences. And similarly for the permission imperative (3.21); the speaker's aim is to establish a right, rather than modify her addressee's preferences.

I propose to explain these infelicities as stemming from the non-cancelability of whatever speech act is conventionally expressed by imperative clauses. That's to say, I propose to explain these infelicities by appeal to a conventionalized use for imperative clauses. "I know you won't" continuations are inadmissible in these cases because (i) the speech act conventionally expressed by the imperative is *present* (although voided), (ii) the speech act the speaker is interpreted as expressing with her utterance of an imperative is incoherent with the speech act conventionally expressed by the imperative. Again, lest there be any air of mystery about what it might mean for a speech act to be present without the speaker being interpreted as proffering that speech act, I mean only that the imperative clause contributes this speech act to the interpretive hopper (even when what the interpretive hopper spits out is an instruction or permission). This speech act—whatever its identity—remains *detectable* under non-normal interpretations. As Asher & Lascarides suggest, this model of linguistic interpretation is "ampliative not destructive," in that a voided speech act nevertheless "remains in context for further predication" (2001: 193).

That is not to say that *all* the interpretive phenomena customarily associated with that speech act will remain detectable. (If that were the case, the imperative would not lack its normal interpretation.) It is enough for the argument I'm making here that some of these phenomena are preserved. Nor is it to say that detectable phenomena will be constant across cases. Indeed, we find that *please* (and similar linguistic signals of politeness) is admissible under permission interpretations, but inadmissible under instruction interpretations.

(3.22) Please, have an apple!

(3.23) Veuillez entrer!

[Lit: *Please enter!*]

(3.24)??Please, take Broadway to 14th.

(3.25)??Please, whip the egg whites before folding them into the custard.

(*As part of a cookbook recipe for soufflé*)

(3.26)??Pour regarder les images, veuillez appuyer sur la touche étoile.

[Lit: ??*To view the images, please press the star button.*]

This asymmetry should not, I think, be construed as evidence against the hypothesis that

imperative clauses conventionally express some type of speech act. Rather, in light of the availability of *please* on an evidently permissive interpretation of the imperative (and, e.g., the unavailability of *please* on a performative *may*; cf. 3.27), it should instead be construed as further evidence *for* the hypothesis.

(3.27) #Please, you may have an apple!

Conventionalizing the speech act expressed by the imperative will allow us to explain the availability of the *please* adjunct even in cases where the imperative cannot reasonably be interpreted as proffering any sort of speech act to which *please* might “attach.”¹³ Whatever the explanation of the unavailability of the *please* adjunct for instruction imperatives, conventionally linking imperative clauses to some speech act or other will aid in explaining its availability for permission imperatives.

3.3.2 Necessitation Data

Although we have yet to give a model of what necessitation with respect to a body of preferences involves, we can take it on faith that all of the following speech acts will be incoherent, in an intuitive sense, with a proposal to make ϕ necessary relative to a relevant body of preferences P .

- Proposing to make $\neg\phi$ P -permitted
- Proposing to make $\neg\phi$ P -necessary

We can also take it on faith that performative uses of the modals *may* and *must* count as proposals to make their complements permitted and necessary, respectively, relative to some salient body of preferences. Concretely, when spoken by a boss to her employee, (3.28) makes it permitted that the employee take the following day off, while (3.29) makes it necessary or required that the employee take the following day off.

(3.28) You may take tomorrow off.

(3.29) You must take tomorrow off.

It follows that, if imperatives conventionally express the speech act of necessitation, we will predict that performative uses of modals that, uncontroversially, express speech acts that are incoherent with a specific act of necessitation, are inadmissible in a context in which an imperative conventionally expressing that necessitation act is uttered.

13. For a related point, offered as an argument for conventionalizing the request interpretation of (3.1), see Searle (1975a); Morgan (1978); Asher & Lascarides (2001).

That prediction seems (unsurprisingly) to be a good one when command-expressing uses of the imperative are considered.

(3.30) #Come in to work tomorrow. By the way, you may take tomorrow off.

(3.31) #Come in to work tomorrow. By the way, you must [should] take tomorrow off.

I say unsurprising because, if any use of an imperative is accurately modeled using the rubric of necessitation, it will be the command use.

The real test comes when we consider uses of the imperative that are evidently not interpreted as expressing necessitation, i.e., permission and instruction uses. Here too, the data seem to square with our predictions. First, permission uses.

(3.32) a. #Take tomorrow off, but you may also come in tomorrow.

b. #You may come in tomorrow, but take tomorrow off!

(3.33) #Take tomorrow off, but you must [should] come in tomorrow.

The infelicity of (3.33) is, of course, unsurprising (since, even if the permission imperative did not conventionally express a necessitation act, we would expect the permission it is interpreted as expressing to be incoherent with the force of a contrary performative *must*). The inadmissibility of (3.32a) and (3.32b) is rather more surprising. Why, if permission imperatives simply express permission—if imperatives do not conventionally contribute a necessitation act to their interpretations—would these sequences be infelicitous? Permitting an employee to take the following day off is, of course, not incoherent with permitting the employee to come in.

A similar story goes for instruction uses.

(3.34) a. #Take Broadway to 14th. You may, however, take 5th Avenue.

b. #You may take 5th Avenue to 14th, but take Broadway to 14th!

(3.35) #Take Broadway to 14th. You must, however, take 5th Avenue.

These data are more ambiguous, as the infelicity might seem explicable without appeal to the conventional expression of a necessitation act. If, as I've argued, the instruction imperative in, e.g., (3.34a) communicates an assertion whose content is something like *to get to Union Square, you must take Broadway to 14th*, it is not surprising that following this with "you may, however, take 5th Avenue" leads to incoherence. For the latter is plausibly interpreted as expressing something *inconsistent* with the assertion communicated by the instruction imperative, namely, that you may take 5th Avenue to get to Union Square.

It is, however, possible to control for this. Notice that an instruction interpretation of the imperative is available even when it is understood that there are many ways to get to Union Square. In such cases, the instruction imperative can be paraphrased using a *weak*

necessity modal, i.e., as *to get to Union Square, you should take Broadway to 14th*. Suppose we are in such a context. Notice that the infelicity remains, even when we replace the *may* construction with something that clearly serves to establish a permission for the addressee to take 5th Avenue.

(3.36) #Take Broadway to 14th. But you may also take 5th Avenue, instead.

The weak necessity interpretation of the instruction imperative conveys something like *to get to Union Square using the least effort, you must take Broadway to 14th*.¹⁴ This certainly seems to be coherent with the most obvious interpretation of the second clause, which establishes a permission to ignore the advice proffered by the instruction imperative; one can (and often does) give advice while granting permission for the addressee to ignore it, as in (3.37).

(3.37) You should take Broadway to 14th. But you may also take 5th Avenue, instead.

Even in a context primed to generate the weak necessitation of the instruction imperative, the infelicity of (3.36) remains. Once again, positing a conventionalized necessitation act for the instruction imperative explains the infelicity. The necessitation act, since conventionalized, is non-cancelable (although voidable). Since the second clause involves an attempt to cancel a non-cancelable necessitation act, we predict, apparently correctly, the inadmissibility of (3.36).

3.4 Computing ISAs

This section lays out the explanatory payoff of the ideas developed above. First, I summarize the Asher & Lascarides (2001) method of computing ISAs. Then I'll explain why the method is actually not sufficiently general to cover all ISAs. I will develop an alternative, and then show how it yields a satisfying treatment of the various kinds of permission imperatives, as well as instruction imperatives.

3.4.1 ISAs in General

The central idea in Asher & Lascarides (2001)'s treatment of ISAs is that interpretation in such cases involves the assignment of a complex semantic type (what they term a "dot-type") to an utterance. Dot-types are constructed out of ordinary types by use of the following rule:

CONSTRUCTION OF DOT-TYPES

If τ_1 and τ_2 are types, then $\tau_1 \bullet \tau_2$ is a type

14. On this interpretation of weak necessity, see von Fintel & Iatridou (2008).

An utterance may be assigned a dot-type in interpretation either by convention (i.e., as a function of the information in the grammar \mathcal{G}), or by inference (2001: 211-12).

INFERRING DOT-TYPES

$u >_{\text{context}(u)} \tau \bullet \tau'$ if, for some non-actual utterance u' :

- i. $u >_{\text{context}(u)} \mathcal{I}_{\text{speaker}(u)}\phi$
- ii. $u : \tau$
- iii. $\forall v (v : \tau \not\prec \mathcal{I}_{\text{speaker}(v)}\phi)$
- iv. $u' > \mathcal{I}_{\text{speaker}(u')}\phi$
- v. $u' : \tau'$

This definition states that u is interpreted in its context of utterance as having a complex type $\tau \bullet \tau'$ if (i) u 's speaker is interpreted as intending ϕ , (ii) u is of type τ , (iii) normally, speakers of utterances of type τ do not intend ϕ , (iv/v) normally, speakers of an utterance u' of another type τ' do intend ϕ . In other words, an interpreter infers a dot-type for an utterance u as a kind of *accommodation strategy*: typically speakers of u don't intend ϕ , but this speaker does, and is assumed to be rational, so the speaker must be interpreted as performing another speech act of type τ' , in addition to one of type τ . Notice that this is recognizably similar to the Gricean reasoning invoked to explain the request interpretation of (3.1).¹⁵

There are *many* details left out of this presentation, but this summarizes the central bits. And even at this level of detail, we can identify two major flaws with this treatment.

- It conflates the semantic value of an utterance with its interpretation at a context. This is an odd picture of the semantics/pragmatics interface, especially so if one is a fan of a classic picture of semantic content, on which an utterance's semantic content supervenes on its LF (with values of some open parameters perhaps fixed by context), while its LF in turn supervenes on its *syntactic form*. An utterance's syntactic form is generally thought to be an input to interpretation, rather than one of its products.

15. One might think there is a worry of *circularity* for this type of account. Namely, meanings are individuated in terms of intentions: to know the speaker's intention is to know the utterance's meaning. But the result of deploying the INFERRING DOT-TYPES rule is the utterance's meaning. So to compute the utterance's meaning, the speaker has to *already know* its meaning. In reply, it is reading what I've said a bit too literally to see these rules as psychologically realized by interpreters, or to think that interpretations arise simply by deploying these rules. As a matter of psychology, linguistic interpretation will involve at least two distinct cognitive modules that messily interface with one another over time. Their *diachronic inter-relations are constrained by the sorts of rules I describe here*, without, I think, actually being *guided by* them. In reality, in cases of indirect speech act interpretations, I would say that an agent deploys a theory of cognitive modeling (that takes facts about the type of a clause as input) to arrive at an initial intention-attribution. The intention-attribution is seen somehow not to mesh with what the agent knows about the context, and so a new intention is inferred. This intention is subsequently associated with a cognitive update by the interpreter's linguistic theory. Although I often speak loosely, I want to emphasize that it is not in any way my goal to model this psychological process with the series of rules I will be articulating here. I only want to articulate a set of plausible constraints (which have the form of defaults) on how an agent's theory of mind and her theory of linguistic interpretation relate to one another.

- More worryingly, it offers no account of how default interpretations can be *voided*. The constituent types τ_1 and τ_2 of a dot-type $\tau_1 \bullet \tau_2$ are on “equal footing” so to speak: assigning an utterance an interpretation of type $\tau_1 \bullet \tau_2$ means interpreting the utterance as proffering both a speech act of type τ_1 as well as a speech act of type τ_2 . But, as argued in Section 3.1, there are many cases in which, although a speech act of τ_1 is in some sense conventionally contributed by the utterance, it is not part of the utterance’s final interpretation. Indeed, there are many cases in which a speaker would have to be represented as *irrational* if she were interpreted as proffering speech acts of both type τ_1 and type τ_2 with her utterance.¹⁶

The account we’ve been developing above—distinguishing systematically, as it does, between semantic denotation and interpretation—is already well-suited to remedying the first flaw.

Remedying the second requires a somewhat subtler picture of interpretation: we must recognize different “stages” of interpretation, rather than running together the stage at which conventional meanings are assigned with the stage at which a final interpretation is assigned.¹⁷ The following grid summarizes roughly what I have in mind.

Stages of Interpretation

- Stage 1: Assignment of Logical Form
- Stage 2: Preliminary Interpretation
- Stage 3: Contextual Contribution
- Stage 4: Final Interpretation

The class of types expressible at the preliminary stage will be the same as the class of types expressible at the final stage. That’s to say, preliminary interpretations are not *indexical*; they don’t require contextual supplementation to express something at a context of utterance.¹⁸ But preliminary interpretations can, I suggest, interact with contexts to generate different final interpretations, in at least the following ways. Let τ be the type of a preliminary interpretation and τ' the type of a final interpretation.

- **Supplementation.** Gricean considerations motivate a final interpretation on which $\tau' = \tau \bullet \tau^*$, for some τ^* .

16. Asher & Lascarides (2001) themselves recognize a point in this neighborhood. They note that interrogatives conventionally express the intention that the speaker come to believe an answer to the interrogative (207), but also recognize that attributing this intention *fails to make sense* of the intentions of a speaker who utters (3.1) (214). It is puzzling why they do not take this to motivate accounting for voiding the question interpretation (in addition to generating the request interpretation).

17. The language of “stages” of interpretation brings to mind Christopher Potts’ work on non-semantic “dimensions” of interpretation on which expressive content and conventional implicature content are computed (Potts 2005, 2007). But that work’s utility as a precedent for the picture I want to pursue here is limited. I am interested in developing a speech act theory of meaning (for imperatives, in particular) that explains *how* conventionally assigned interpretations interact with contexts to generate actual interpretations. Since both expressive content and conventional implicatures are, for Potts, assigned at the conventional meaning stage, they will not provide a useful analogy.

18. Notice, by the way, that we have suppressed the role of context in determining logical forms, interpreting variables, etc. This is just for simplicity’s sake.

- **Voiding.** Gricean considerations motivate a final interpretation on which $\tau' \neq \tau$ is a simple type, rather than a dot type partly composed of τ .
- **Binding.** If τ is the type of a function ($\tau : \tau_1 \mapsto \tau_2$), and context supplies material of type τ_1 , this generates a final interpretation on which $\tau' = \tau_2$.

The Asher & Lascarides account is suited only to explaining the interpretations of ISAs of the sort generated by Supplementation. That is a good start—these are the paradigm of a speaker performing one speech act *by virtue of* performing another, and it seems to cover a large portion of the examples discussed in the ISA literature, in particular, most examples of performatives via assertions.

(3.38) I command you to shut the door.

(3.39) I ask you again what you were doing yesterday evening.

(3.40) I hereby invite you to my party tomorrow.

But there are a wide range of closely related phenomena that cannot be fit into this mold.

3.4.2 Voiding

Here, roughly, is the chain of reasoning by which simple final types might be inferred from either complex or simple initial types.

VOIDING TYPES

$u > \tau, u >_{\text{context}(u)} \tau'$, and $\tau \neq \tau'$ if, for some non-actual utterance u' :

- i. $u >_{\text{context}(u)} \mathcal{I}_{\text{speaker}(u)}\phi$
- ii. $u : \tau$
- iii. $\forall v(v : \tau > \mathcal{I}_{\text{speaker}(v)}\psi)$
- iv. $u' > \mathcal{I}_{\text{speaker}(u')}\phi$
- v. $u' : \tau'$
- vi. $\forall x((\mathcal{I}_x\phi \wedge \mathcal{I}_x\psi) \Rightarrow \mathcal{IRR}_x)$

The key conditions here are (iii) and (vi). (iii) says that normally an utterance of type τ would be accompanied with an intention ψ . (vi) says that the intentions expressed by ϕ and ψ are rationally incompatible. Notice this chain of interpretive reasoning is, in one sense, *non-ampliative* (i.e., destructive). An utterance's initial interpretation can *fail to be proffered* as part of its final interpretation, if constructing the final interpretation means inferring a speech act of rationally incompatible type.

It is important to notice that speech acts of types τ and τ' are not *themselves* assumed to be incompatible. The voiding is *not* done as the result of being interpreted as proffering a speech act that is incompatible with the utterance's conventional interpretation (recall from Section 3.2.1 that the hallmark of conventionally expressing a speech act is that it cannot be

cancelled in this way). It is done at the level of *intention-attribution*: a speaker who proffers speech acts of types τ and τ' is represented as having rationally incompatible intentions. The question expressed by (3.1) is not *cancelled* by the request that it expresses (*please pass the salt*). An interpreter (and a context) could easily accept/integrate both of these speech acts without lapsing into irrationality. To contrast, neither an interpreter nor a context can accommodate the speech acts expressed by, e.g., (3.13), repeated here.

(3.41) #It's raining, but I don't accept that it's raining.

Updating with the left conjunct involves representing the speaker as accepting that it's raining; adjusting one's information in response to an assertion that ϕ constitutively involves representing the speaker as accepting that ϕ ; representing the speaker as insincere *pre-empts* acceptance of the speaker's assertion. Updating with the right conjunct involves representing the speaker as not accepting that it's raining. So updating with both involves representing a *contradiction*.

This is not the case with the question and request acts expressed, at different levels of interpretation, by (3.1). Those *acts* are compatible, in that both can be incorporated into an information state without lapsing into irrationality. (Incorporating an issue raised by a speaker constitutively involves representing the speaker as raising an issue, but this does not constitutively involve representing the speaker as *uncertain* about the issue.) But a speaker who is interpreted as expressing both would have to be interpreted as irrational: she would be interpreted as having an intention (to find out whether the addressee can pass the salt) that, in making the request, *she* takes to already be realized. This, I claim, is what triggers a re-interpretation on which (i) the speaker is not interpreted as intending the goal associated with the conventionalized speech act, (ii) the speaker is interpreted as expressing another speech act that would achieve the goal she is interpreted as having. This is, as before, an accommodation strategy, but it differs significantly from the one outlined for Supplementation cases.

Endorsing an interpretive principle of this sort requires modifying the sort of explanation of the "linguistic hallmark" data in Section 3.3 that the Asher & Lascarides account would have led us to expect. On that account, given an utterance u that conventionally expresses a speech act of type τ , u will bear, *ceteris paribus*, the linguistic hallmarks of a speech act of type τ because u is interpreted as expressing τ , even when u 's interpretation is not of type τ (for in this case, u 's interpretation will be a dot-type, one of whose constituents is τ). We must say something different, since we hold that speech acts are voidable, and, as such, when voided, are not part of the utterance's final interpretation. What explains linguistic hallmark data for us is simply that certain speech acts are conventionalized, in virtue of clause-types, and conventionalized speech acts, though voidable, are non-cancelable. This is a somewhat less ambitious explanation than that of Asher & Lascarides; indeed, it does not attempt to improve on the explanation that we suggested in Section 3.3. But it

is an explanation nonetheless; conventional meanings can explain hallmark data without being part of the final interpretation of an utterance.

3.4.3 Binding

Binding happens when an utterance expresses a functional type $\langle \tau_1, \tau_2 \rangle$ (i.e., some function $f : \tau_1 \mapsto \tau_2$) and something of type τ_1 is supplied by context, so that the utterance ultimately expresses something of type τ_2 . We have not discussed binding at all to this point. Indeed, we have not even motivated its existence.

One precedent for recognizing something very much like Binding is the Variable-Free Semantic treatment of syntactic items that are traditionally represented using variables, e.g., pronouns and ellipsis sites (see esp. [Jacobson 1999](#)). On Variable-Free treatments, the free pronoun in a sentence like (3.42) is not represented (as is standard) as a free variable that receives an interpretation relative to an assignment function supplied by context.

(3.42) He went swimming

Instead, the pronoun is treated as semantically vacuous (in this case, as expressing the identity function $\lambda x.x$), and combines with the semantic value of the predicate *went swimming* ($\lambda x.x$ *went swimming*) via function-composition, so that the entire sentence expresses simply the property $\lambda x.x$ *went swimming*. To contrast, on a variable treatment of the free pronoun, (3.42) is interpretable only relative to an assignment function; relative to such a function, it expresses a proposition. In order to recover a proposition from (3.42) on the Variable-Free treatment, an individual must be supplied. But even without such an individual, the sentence receives a complete semantic interpretation: it expresses a property.

Something similar can, I'll argue, happen *at the level of speech acts* (although it is possible, as I will show below, to style this as a purely semantic phenomenon). I think it is plausible, following [Swanson \(2006, 2008\)](#),¹⁹ that many kinds of language can be understood as expressing properties of cognitive states—equivalently, characteristic functions of acceptable cognitive states (or contexts, or what have you). To express such a function in discourse is, generally but also very roughly, to express a proposal that the cognitive state of one's addressee(s) (or the context, or what have you) come to satisfy the property expressed by the utterance. But such a proposal can be voided if either the speaker or the discourse supplies an argument to the function conventionally expressed by the speaker's utterance. So, for instance, modal claims like the following typically express a function that maps a cognitive state S to 1 iff S satisfies some property or other—in the case of (3.43) the property of recognizing the possibility that Obama is reelected, in the case of (3.44) the property of idealizing the possibility that Obama is reelected.

(3.43) Barack Obama might be reelected.

19. See also my later chapter "Meaning for Expressivists," as well as [Charlow \(2010b, 2011a\)](#).

(3.44) Barack Obama should be reelected.

Equivalently, we can think of (3.43) as expressing a function mapping S to the proposition that S recognizes the possibility that Obama is reelected, (3.44) as expressing a function mapping S to the proposition that S idealizes the possibility that Obama is reelected. So, if \mathcal{I} is the type of cognitive states and \mathcal{P} the type of propositions, (3.43) and (3.44) express each express some function $f : \mathcal{I} \mapsto \mathcal{P}$, a type conventionally associated with the speech act of making a cognitive proposal (or, as Swanson 2008 styles it, giving “cognitive advice”).

But, when an argument to this function—something of type \mathcal{I} —is made salient (in the right way), we will predict that these sentences ought to express propositions rather than characteristic functions of cognitive states. And, indeed, that is precisely what we seem to observe.²⁰ Suppose Tom is a big Obama fan—he both thinks it possible and desirable that Obama be reelected. Uncontroversially, I think, (3.45) and (3.46), express propositions whose contents concern Tom’s cognitive state (specifically, features of his beliefs and his desires).

(3.45) In view of what Tom thinks, Barack Obama might be reelected.

(3.46) In view of what Tom thinks desirable, Barack Obama should be reelected.

Rejecting (3.45) means rejecting an *assertion about* Tom’s cognitive state, rather than a piece of advice about which possibilities one should recognize. Similarly, rejecting (3.46) means rejecting an assertion about Tom’s preferences, rather a piece of advice about what one ought to prefer. The view we are developing here immediately predicts these interpretations, if they are treated as expressing characteristic functions of cognitive states (and if the “in view of” adjuncts are interpreted as supplying an argument to this function).²¹

Similarly if it is common ground that we are speaking from the “point of view” of some cognitive state or other (as often happens in cases of free indirect discourse and modal subordination; see esp. Roberts 1989), we will predict that context can, in certain cases, supply the argument to an utterance conventionally expressing a characteristic function of cognitive states, so that, in context, the utterance ultimately expresses a proposition (assertion), rather than a piece of doxastic or preferential advice. That’s to say: in a context where it is common ground that we are talking about what Tom thinks and prefers, (3.43) has an equivalent interpretation to (3.45) and (3.44) to (3.46). The difference between these pairs lies in the *stage of interpretation* at which the argument to the relevant function is

20. I don’t have the space to develop a sophisticated compositional treatment of these sentences here, so what follows will be fairly rough, but still plausible and definite enough to give the general flavor of what I have in mind.

21. For a broadly similar take, from an “Expressivist” point of view, see Gibbard (1990, 2003). Although I don’t want to harp on this, this sort of treatment represents somewhat of a descriptive improvement over standard treatments, which understand sentences like (3.43) to express propositions (relative to contextually fixed ordering sources), and must treat “in view of” adjuncts, somewhat implausibly, as “monstrous” (in the Kaplanian sense) devices for *shifting* the values of these contextually fixed ordering sources.

supplied. Plausibly, with (3.45) and (3.46), the argument is supplied at Stage 2 (so that, in the end, this is best understood as a properly semantic phenomenon, rather than a speech act-level phenomenon, since the argument is specified explicitly in the syntax), whereas with the equivalent interpretations of (3.43) and (3.44), the argument is more plausibly supplied at Stage 3. Though these phenomena are obviously very similar, I will be taking the latter as my model for the phenomenon informally described above as Binding.

It is, for all we have said here, a completely open question whether the argument supplied in Binding cases should be represented as being supplied to the characteristic function conventionally expressed by the modal sentence, or, instead, to the speech act conventionally contributed by it. I tend to prefer the former way, owing to an armchair commitment of mine to a certain view about the mechanics of interpretation, and I will be developing an account in that vein below. But nothing really turns on this, and, for all I will say, we might just as easily develop an account in the latter vein.

There is much more to say by way of description of these analyses, but the discussion so far should give a good, if rough, sense of the ideas we will be pursuing to account for “unconventional” readings of imperatives. To preview, permission readings will generally be accounted for as cases of Voiding, instruction readings as cases of Binding. We will go into more detail about Voiding and Binding, respectively, as needed below.

3.4.4 Permission as an ISA

Permissions, as I suggested in Section 2.4.5, come in a variety of forms:

- Reversals of standing prohibitions
- Establishment of strong permissions
- Establishment of rights

Accounts, like those of Lewis (1979a) and Portner (2007, 2010), which rigidly associate imperative clauses of the form $!\phi$ with a specific operation on a context—in Lewis’ case, making ϕ required, in Portner’s case, adding $[[\phi]]$ to the addressee’s To-Do List—do not seem sufficiently flexible to handle all of these varieties. Lewis’ account handles none of them. Portner’s is designed to only handle the first.

Against permission as necessitation. Claim: if I am right that an imperative of the form $!\phi$ conventionally express a necessitation act (a proposal to make ϕ necessary relative to a salient body of preferences), then *all* of these varieties of permission will have to be treated as ISAs. This idea goes against one prominent line of theorizing about permission, according to which permission imperatives function, in Portner’s words, to “[o]vercome the reluctance of the addressee to undertake an action in [her] own interests” (Portner

2010: 8). As I said above, I think this is right, at least for some versions of permission imperatives. But some theorists maintain, in my view erroneously, that (i) such speech acts can be understood in terms of necessitation or preference-addition, (ii) all interpretations of permission imperatives can be understood in terms of such speech acts. Such “permissions” are actually claimed to be rather akin to advice, in that the speaker is recommending the addressee modify her desires/intentions so that ϕ becomes desired (required in view of her desires) or intended (required in view of her intentions) (see, e.g., Wilson & Sperber 1988; Schwager 2005). The felt difference in force of such “permissions” is explained as a variation in the speaker’s authority with respect to the relevant body of preferences; in the case of her own preferences, the addressee is presumed to be the ultimate authority, and can accept/reject the speaker’s recommendation basically at her discretion.

As already noted (Section 2.4.5), this is wrong on two counts. First, not all attempts to address an addressee’s reluctance to take an action in her own interests can be understood in terms of advice. Some should be understood in terms of establishment of a strong permission, and others in terms of grants of rights; neither strong permission nor grant of right is properly understood in terms of advice. There are also, as Portner argues, cases of reversals of standing prohibitions where “the permission-granting sentence is accepted, but its content is not required” (2010: 8). Second, not all permissions can be understood in terms of attempts to address an addressee’s reluctance. Strong permissions and rights are rather *thin* notions: one can of course establish a strong permission or a right without, e.g., thinking the exercise of this permission or right is in the interests of the addressee at all.

Here is how things stand at this point. We have reason to think that utterances of imperative clauses conventionally express necessitation acts. Permission uses of imperatives are interpreted, at least in some cases, as expressing speech acts that cannot be represented as necessitation acts. It follows, then, that such permission uses *must* be modeled as ISAs. In what follows, I will briefly outline ISA accounts of permission uses using the Voiding template of Section 3.4.2.

Strong permission. Let’s first recall example (2.38), repeated here, in which an imperative utterance seems plausibly to be used to establish a strong permission.

(3.47) (*B looks unsurely at the fruit plate*)

A: It’s okay. Have an apple!

As I noted in Section 2.4.5, strong permissions are paradigmatically licensed in cases where:

- A corresponding weak, or implicit, permission is in force
- A speaker wishes to strengthen this weak permission into an explicit permission

Strong permission interpretations of *imperatives* will, I claim, paradigmatically be licensed in cases where strong permissions are licensed.

Our template for voiding a conventionalized speech act states, roughly, that an utterance that conventionally expresses a speech act of type τ can be interpreted as expressing a speech act of type τ' when (i) the speaker is interpreted as having a communicative goal that is better explained by interpreting her as expressing a speech act of type τ' than by interpreting her as expressing one of type τ , (ii) it would be irrational for the speaker simultaneously to have the goals characteristically associated with speech acts of respective types τ and τ' .

The speaker of (3.47), A, is interpreted, I've argued, as intending to establish merely a strong permission for B (relative to a salient body of preferences, which we will gloss, simply, as *what A requests*). Nevertheless, her utterance conventionally expresses a necessitation act: the act of proposing that some salient body of preferences (here, it would seem, the desires of the addressee, given that A's imperative seems by default to express a *suggestion*) come to require the taking of an apple. The intention characteristically associated with utterances of imperatives is the successful performance of such a necessitation act—in this case, a suggestion.

Now, these acts—the permission and the suggestion—are strictly compatible. There is obviously nothing incoherent about a context (so nothing incoherent about generating a context) in which:

- In view of A's requests, B is strongly permitted, but not required, to take an apple.
- In view of B's desires, B is required to take an apple.

Nevertheless, there is something *irrational* about a speaker who has these states as *communicative goals*; who intends, through her utterance, to bring both states of affairs about. B is presumed to regard A as an authority; she tries, insofar as is reasonable, to conform her pursuit of her desires to A's requests. Requests function to exploit this authority: the point of a request that ϕ , addressed to B by A, is to get B to intend to see to it that ϕ (which requires that B desire to see to it that ϕ). So requesting, to B, that ϕ is a way for A to realize the communicative goal of B's desires requiring that she take an apple. There are, moreover, no other *communicative* means by which A can reach this goal. B's willingness to listen to what A requests—B's deference—is the only leverage A has over B's desires. For one of A's utterances to *bring it about* that B's desires require her to take an apple, A must request of B that she take an apple.²² Put in the simplest possible terms: if the speaker wants the addressee to want to have the apple, rationally, she should encourage the addressee to have the apple!

So, the goal characterizing the necessitation act conventionally associated with A's imperative is irrational, given A's goal to make it strongly permitted, but not required,

22. We're ignoring some special cases here. For instance, we're ignoring cases where A intends to bring about the desire by informing B that one of her existing goals is served by taking the apple (e.g., by highlighting or otherwise emphasizing desirable features of the apple). In this case, the imperative would not have the force of necessitation with respect to B's desires. Trivially, it would have the force of informing A that taking the apple is *already necessitated* by B's desires. A's utterance typically gives rise to a new desire, but this does not amount to necessitation. More on instruction imperatives in Section 3.4.5.

that B take an apple. Since A is interpreted as having the latter goal (and is presumed to be rational), it is inferred, as an accommodation strategy, that A does not intend the goal characterizing the necessitation act. The speaker is thus interpreted as expressing simply the speech act characterized by the remaining intention, i.e., the establishment of a strong permission. Something like this sort of Gricean accommodation (which instantiates the general sort of interpretive process outlined in Voiding) is, I submit, what explains the strong permission interpretation of (3.47).

Grant of Right. The melody is basically the same for grants of right. Grants of right like (2.40) (repeated here) are paradigmatically licensed in cases where a speaker wants to strengthen an extant, perhaps implicit, permission into a right—an entitlement to resist certain future proposals to restrict her available, permissible options.

(3.48) A: Am I needed at work today?

B: No. Stay home.

(*later, C, swamped, calls to see why A is not in the office that day*)

C: We need you at work today! Get to the office immediately!

A: But B said it was okay to stay at home today!

C: Oh! My bad!

In this dialogue, speaker B's imperative is interpreted as intending to establish a right for A to stay home. Still, I contend, B's utterance conventionally expresses a necessitation act: the act of proposing that some salient body of preferences come to require that A stay home. The story from this point is the same as before. There is nothing incoherent about a context accommodating both a necessitation act and a grant of right that preserves A's freedom to do as A chooses, so long as these acts address different bodies of preferences (perhaps, as before, *what A desires* and *what B requests*). Still, there is an irrationality associated with mingling these speech acts' characteristic intentions *in a single agent*. And, so, the necessitation act is voided.

There is, of course, more to say about the pragmatics of grants of right, and below we will have a bit to say about how they constrain a compositional theory of speech acts for imperatives—in particular, how an entitlement to resist contrary commands should be represented in a formal pragmatics for imperatives. For now, we are satisfied to predict that, whatever grants of right actually consist in, so long as they consist in part in permitting, the necessitation acts conventionally expressed by an imperative utterance will generally be voided when that utterance is used to grant a right.

Does this overgenerate? A worry about this treatment of permission imperatives is that it might predict permission interpretations when none are intuitively present. It is actually rather difficult to infer a permission interpretation when a speaker utters (3.49), *even when*

it assumed that the speaker does not wish to establish a requirement for the addressee to leave, but rather a mere permission for the addressee to leave.

(3.49) (*B looks unsurely at the door. A wants to let B know it's ok to leave, but without forcing the issue.*)

#A: Leave the room at once!

I have been assuming, I think plausibly, that the body of preferences with respect to which permission imperatives make something permitted can be glossed, roughly, as *what the speaker commands*. For now, we'll stick with that. I have also claimed that permission interpretations of imperative utterances arise when an intention to establish a permission with respect to this body of preferences rationally conflicts with the intention conventionally associated with the imperative utterance. The problem is that such a rational conflict is plausibly present in (3.49), when it is assumed that the speaker does not wish to establish a requirement for the addressee to leave. So we seem to predict that (3.49) has an easily accessible permission interpretation which it, intuitively, seems to lack.

Significantly, though, not just any rational conflict is sufficient for generating an interpretation along the rubric given in the Voiding schema. Indeed, there are, as I have been emphasizing, two ways for communicative intentions to conflict. Let g and g' be communicative intentions, and let α and α' be their associated speech acts. Those ways are:

- **Merely rational conflict.** g and g' conflict, but $\exists c : \langle \alpha, \alpha' \rangle$ is coherent at c
- **Thoroughgoing conflict.** g and g' conflict, and $\forall c : \langle \alpha, \alpha' \rangle$ is incoherent at c

Cases (3.47) and (3.48) involve merely rational conflict; they establish a permission relative to what the speaker commands, while conventionally expressing a necessitation act relative to what the addressee desires. These speech acts are coherent, even if their associated intentions cannot be rationally harbored in a single agent. Case (3.49) plausibly involves thoroughgoing conflict: given the obvious, emphatic *force* of the utterance, it is hard to see the necessitation act conventionally expressed as a suggestion, i.e., as necessitation with respect to B's desires.²³ So (3.49) expresses, as a matter of A's intentions, a mere permission for B to leave at once (so that B is permitted, but not required, to leave at once), as well as, as a matter of the imperative's conventional meaning, a requirement for B to leave at once. These speech acts seem to pertain, moreover, to the *same* body of preferences, namely, what A requests. But no single body of preferences can both not require and require the addressee to leave at once. So the speech acts expressed, on the most accessible reading of (3.49), are incoherent. So there is no accessible interpretation of (3.49) at all; on the most accessible reading, it amounts to an attempt to *cancel* the speech act conventionally expressed by the imperative.

23. Why is that? I don't have an account, exactly. Here I'm just describing a pretty clear fact. So, while there is probably something very interesting to be said here, I'm going to leave that to another time.

Far from overgenerating, then, our account predicts, correctly, that the most natural reading of A's utterance renders the utterance uninterpretable.

3.4.5 Instruction as an ISA

Instruction imperatives, as I suggested in Sections 2.4.3 and 2.4.4, are paradigmatically licensed in the presence of some salient goal.

(3.50) A: How does one get to Union Square from here?

B: Take Broadway to 14th.

(3.51) A: I want to get to Union Square from here.

B: Take Broadway to 14th.

Again, if I am right that imperatives conventionally express necessitation acts, then instruction imperatives will have to be interpreted as a kind of ISA: the utterance expresses, as a matter of its conventional meaning, a proposal to necessitate A's taking Broadway to 14th. It is, however, interpreted as resolving an issue raised by A: how to achieve the goal of getting to Union Square. Issue-resolving interpretations cannot, I argued, be understood in terms of necessitation. So we need a story of how the issue-resolving interpretation arises from the presence of the salient goal, together with the conventional meaning of the imperative.

Asher & Lascarides address the issue briefly, suggesting an analysis along the following lines (2001: 206). According to this analysis, there is a default axiom of the following form:

$$\text{INFERRING REQUEST-ELABORATION} \\ \langle \langle D, u, u' \rangle \wedge u' : ! \rangle > \text{R-Elab}(D, u, u')$$

The default should be read as follows: when utterance u' is uttered in reply to an utterance u that is part of an antecedent discourse D and u' is typed as an imperative, interpret u and u' as related by the R-Elab relation. Letting c be a context characterized by $\langle D, u, u' \rangle$ and assuming that $u' : !, \text{R-Elab}$ is cashed out as follows:

$$\text{R-Elab}(D, u, u') \text{ iff:} \\ \text{i. } u >_c \mathcal{I}_{\text{addressee}_c} \phi \\ \text{ii. } u' >_c \mathcal{I}_{\text{speaker}_c} \mathcal{I}_{\text{addressee}_c} R(u') \\ \text{iii. } R(u') \Rightarrow \phi$$

Informally, R-Elab holds between u and u' at c iff (i) the addressee of u' (the speaker of u) is interpreted as expressing an intention that ϕ , (ii) the speaker of u' (the addressee of u) is interpreted as expressing an intention that her addressee realize the state affairs enjoined by u' , (iii) if the addressee sees to it that that state of affairs is realized, her intention that ϕ will be realized. Even more informally, R-Elab holds between u and u' at c iff the

speaker of u' requests her addressee to realize some state of affairs that in turn realizes some earlier-expressed goal of the addressee.

This analysis will not do, for the simple reason that it is not an ISA analysis at all. The speech act associated with R-Elab is, as Asher & Lascarides themselves note, a kind of requesting! In particular, it is understood as a suggestion: the speaker of u' proposes that the speaker of u come to desire $R(u')$, as a way of achieving her intention that ϕ . But this makes a mess of the linguistic data we outlined in Section 2.4.4. Relational speech acts that would be licensed if the speaker were interpreted as expressing a request (e.g., acquiescence) are not licensed in the case of instruction imperatives at all. Relational speech acts that would not be licensed unless the speaker were interpreted as making an assertion are, however, licensed. The R-Elab analysis can explain neither of these facts.

Instead, I suggest that we explain these readings as resulting from a kind of Binding (Section 3.4.3). I cannot fully do justice to the account here, since I do not have a fully fleshed-out theory of meaning for imperatives on the table (see Chapter IV). But suffice it to say that imperatives, like any other kind of clause, can be thought of as semantically expressing characteristic functions of cognitive states.²⁴ In particular, I will suggest that an imperative of the form $!\phi$ characterize the property of a body of preferences necessitating ϕ , equivalently, a function that maps a body of preferences \mathcal{P} into the proposition that \mathcal{P} necessitates ϕ . When a specific body of preferences is contextually salient (in the right way), that body of preferences can be supplied as the argument to this function, yielding an interpretation of the imperative utterance on which it is understood as proffering a proposition (that that body of preferences necessitates ϕ), as well as an assertion whose content is that proposition.

Here, then, is an informal pass at the sort of account that an analysis in this vein will endorse for (3.50) and (3.51). A's utterance makes salient a very specific body of preferences: one containing just the goal²⁵ that A gets to Union Square from here.²⁶ B replies with an imperative, which expresses a function from bodies of preferences to propositions. The body of preferences made salient by A's utterance goes in as the argument for this function. The result is that a proposition is ultimately computed as the semantic value of the utterance. The default connection between utterances that express propositions and the speech act of

24. One enormous virtue of this sort of view, as I alluded to in Section 3.4.3 and will talk more about in Section 5.3 and Appendix A, is that it can be seen to *explain* the default connection between an utterance expressing this type of semantic value and the speech act of necessitation.

25. Or, perhaps, it does not just make salient this goal, but also a body of "secondary" or "higher-order" goals, concerning the length of the route, its safety, and so on. In cases such as these, our account, together with the sort of account of weak necessity modals given in von Stechow & Iatridou (2008), easily (and correctly!) predicts, that the imperative will be interpreted as making a weakly modal assertion (glossable, roughly, as *to get to Union Square, you ought to, although you don't have to, take Broadway to 14th*). That is because the function of cognitive states expressed by the imperative can be viewed as taking a *merged* body of primary and secondary goals, rather than just the primary goal of getting to Union Square, as its argument. I have a much more in-depth discussion of weak necessity modals in my (2011b).

26. There's an interesting story to tell about how asserting that one has a certain goal, as in (3.51), can make that goal's achievement an issue in the discourse. Space and relevance prevent me from telling it here.

assertion is then exploited, so that B is interpreted as asserting this proposition, namely, that the relevant body of preferences necessitates taking Broadway to 14th.

There is, the reader will notice, a certain (ultimately, inessential) view about the mechanics of interpretation built into this analysis: we assume (rather like the Variable-Free treatment of free pronouns we discussed in Section 3.4.3) that semantic values can be re-computed on the fly, when supplemented by contextually salient material. So, the major contextual contribution to the imperative’s non-standard (i.e., issue-resolving) interpretation occurs at a semantic stage of interpretation (rather than, as was the case with Voiding, at a “speech act-assigning” stage of interpretation) (although these both occur at Stage 3 of interpretation; cf. the Stages of Interpretation described in Section 3.4.1).

There is a sense in which this semantic re-interpretation can be represented as occurring at the “level” of speech acts, as, e.g., an assignment of the following sort of interpretation (which I’ll style I-Elab) to B’s utterance. Let \mathcal{D} be the type of bodies of preferences, \mathcal{P} the type of propositions, and Δ an actual body of preferences (and, as before, c the context characterized by $\langle D, u, u' \rangle$). Then,

INFERRING I-Elab

$$\langle \langle D, u, u' \rangle \wedge u >_c \Delta \wedge \llbracket u' \rrbracket : \langle \mathcal{D}, \mathcal{P} \rangle \rangle > \text{I-Elab}(D, u, u')$$

I-Elab(D, u, u') iff:

- i. $u >_c \Delta$
- ii. $\Delta : \mathcal{D}$
- iii. $\llbracket u' \rrbracket : \langle \mathcal{D}, \mathcal{P} \rangle$
- iv. $\llbracket u' \rrbracket_c = \llbracket u' \rrbracket(\Delta) : \mathcal{P}$

The latter definition makes use of some new notation: adding a context subscript to the value of $\llbracket \cdot \rrbracket$ for an utterance. This is just meant to designate the semantic value of the utterance at the context, accounting for contextual argument-filling (so, whenever context does not supply an argument in $\text{dom}(\llbracket u \rrbracket)$ to $\llbracket u \rrbracket$, it will be the case that $\llbracket u \rrbracket = \llbracket u \rrbracket_c$). What these definitions together say, then, is that I-Elab is a default interpretation for an imperative when (i) the prior utterance u makes salient something of type \mathcal{D} and (ii) u' semantically expresses a function of type $\langle \mathcal{D}, \mathcal{P} \rangle$. Inferring I-Elab means semantically reinterpreting u' , by supplying the contextually salient body of preferences to $\llbracket u' \rrbracket$.

While I think this is an attractive picture, with at least some precedent in the Variable-Free Semantic treatment of free pronouns, I don’t want to rest anything on the acceptability of this view of the semantics/pragmatics interface.²⁷ To that end, note that it is possible to state an equivalent principle of interpretation that does not rely on re-computation of

27. I will note that it is compatible with the spirit of the complaint registered in Section 3.4.1 about Asher & Lascarides’s view of that interface, which allows values of “open parameters” to be filled in by context. The Asher & Lascarides picture of the semantics/pragmatics interface is *much* more radical than this.

semantic value—that is to say, a principle which regards Binding as taking place at the “speech act-assigning” state of interpretation. We can do this by, for instance:

- Representing the act of necessitating that ϕ as a function from preferences into assertions (namely, that function that maps a body of preferences Δ to the assertion that Δ necessitates ϕ). (And, more generally, if an utterance characterizes a cognitive property P , the speech act expressed by that utterance can be understood as a function mapping a cognitive state to the assertion that that cognitive state satisfies P .)
- Understanding the force of this function differently, depending on the source of its argument. If the function gets its argument from context, an assertion act is tokened in the context. If the function gets its argument from a cognitive state, the cognitive state *self-directs* the assertion, and the cognitive state subsequently adjusts to make the assertion true.²⁸

There are a variety of ways to implement this sort of picture formally. I won’t pursue any here. It is enough, for our purposes, to know that the notion of Binding is sufficiently flexible to be implemented in any number of different ways. This picture gives rise to the following version of the I-Elab rule, I-Elab*. (We let \mathcal{A} be the type of assertions.)

INFERRING I-Elab*

$$\langle \langle D, u, u' \rangle \wedge u >_c \Delta \wedge u' : \langle \mathcal{D}, \mathcal{A} \rangle \rangle > \text{I-Elab}^*(D, u, u')$$

I-Elab*(D, u, u') iff:

- i. $u >_c \Delta$
- ii. $\Delta : \mathcal{D}$
- iii. $u' > \langle \mathcal{D}, \mathcal{A} \rangle$
- iv. $u' >_c \mathcal{A}$
- v. $u' >_c \mathcal{I}_{\text{speaker}_c} \text{Believe}(\text{addressee}_c, \Box_{\Delta} R(u'))$

According to these definitions, I-Elab* is a default interpretation for an utterance when (i) the prior utterance makes salient something of type \mathcal{D} and (ii) normally, the utterance would express a function of type $\langle \mathcal{D}, \mathcal{A} \rangle$. Inferring I-Elab* means interpreting u' as expressing something of type \mathcal{A} —specifically, the assertion determined by the fact that the speaker is

28. A worry: as already noted, I tend to think of *all* speech acts as expressing, in some sense, properties of an admissible cognitive state. But, given that view, the account stated here would seem to predict that, given a speech act α that characterizes cognitive property P , when context supplies a suitable cognitive state σ , α can be used to *assert* that σ has P . So, for instance, suppose that an assertion that it is raining is represented as a function mapping a cognitive state σ to the assertion that σ accepts that it is raining. When the relevant σ is supplied by context, we seem to predict, oddly, that “it is raining” can be used to assert that σ accepts that it is raining. I actually think this is a *welcome* prediction—it offers a natural account of the force of utterances in certain kinds of utterances in free indirect discourse (which do function, relative to the basic context, very much like assertions about the cognitive state from whose vantage one is speaking). What is needed to make the account less odd is a statement of restrictions on *when* contexts can supply the cognitive arguments to speech acts. In the case of assertions, those restrictions will be very stringent indeed.

inferred to intend that her addressee come to believe Δ requires $R(u')$ (the state of affairs enjoined by u').

Recall that strong necessity modals exhibit similar behavior to instruction imperatives when some goal or preference is made salient.

(3.52) A: How does one get to Union Square from here?

B: You must take Broadway to 14th.

The explanation developed to account for instruction interpretations of imperatives will extend to instruction interpretations of strong necessity modals, so long as we make the assumption (for which I will argue in a later chapter) that such modals, like imperatives, characterize properties of preference-states.²⁹

3.5 Recap

In this section, I have tried to argue that it is possible:

- To identify a single type of speech act—one that can plausibly be viewed to constitute a natural kind—for imperative clauses to express by default, namely, necessitation.
- To explain interpretations of utterances of imperative clauses according to which they express some speech act other than necessitation, while still retaining (indeed, making central use of) the notion that imperatives express, by default, necessitation acts.

All told, the account developed to this point represents, I think, much more than just a “possibility proof.” Indeed, in arguing that a conventionalized speech act analysis was possible, we developed an account which was able to predict both permission and instruction uses of imperatives (while also explaining the persistence of imperatives’ characteristic linguistic hallmarks under these non-standard interpretations). *All* of these explanations made central use of the hypothesized conventionalized necessitation acts; insofar as the explanations are plausible, this adds further plausibility to the claim that imperative clauses conventionally express necessitation acts.

29. For now, it will be helpful to note that this *follows* from the view that (i) strong necessity modals have the same kind of conventionalized force as imperatives (as argued in [Ninan 2005](#); [Portner 2007](#); see the discussion in Section 2.4.4), (ii) imperatives conventionally express necessitation with respect to a body of preferences.

CHAPTER IV

A Dynamic Theory of Imperative Meaning

We have been arguing for the claim that speech acts are the sorts of things that could be part of a theory of conventional meaning for a clause (imperative clauses, in particular). But it is one thing to argue this point, another thing entirely to actually state a speech act theory of conventional meaning for imperatives. On the standard, Montagovian understanding of methodology for the construction of such theories, developing such a theory will mean doing all of the following:

- Describing a regimented formal language in which “logical forms” for sentences of the relevant fragment of natural language can be stated. (Note that while the standard notion of “logical form” comes with substantive commitments regarding natural language syntax, I want to work with a less committal notion. Logical forms, as I will emphasize below, will be understood as just syntactic representations of meanings or interpretations which may, or may not, be mirrored by natural language syntax.)
- Giving a principled and precise characterization of how sentences in our regimented formal language can be mapped to interpretations (in this case, *specific* speech acts, as opposed to general kinds of speech acts).
- Stating precisely what speech acts *are*. That is to say, we must define a data-type for speech acts, and also show, ideally, (i) how things of this type might be constructed out of familiar model-theoretic entities, (ii) how an assignment of things of this type to sentences in our regimented formal language meets various conditions of adequacy on a theory of meaning.

Our task is complicated by the fact that the notion of a speech act is surrounded by a somewhat mysterious aura. Indeed, this fact is what has plausibly hampered past attempts of prominent speech act theorists to mimic the formal precision and rigor of theorists who theorize about meanings in terms of set-theoretic model-theoretic denotations (e.g., possible worlds propositions) that stand in familiar Boolean relations (e.g., \subseteq) to one another and which can be manipulated using familiar Boolean operations (\cap, \cup, \prime , etc.). Speech acts are, of course, generally understood as a kind of *action*, but this does nothing to dispel the mystery, if it is not accompanied with some sort of theory about what sorts of things

actions are, what (logically relevant) relations actions stand in to other actions, and what sorts of operations, if any, would be useful for a theorist to use to manipulate actions. The task of coming up with such a theory will be referred to here as the **modeling challenge**. Solving the modeling challenge, we'll see, is further complicated by the presence of certain conditions of adequacy on such a solution. The "model theory" for imperatives must give rise to an account of, e.g., their logico-semantic properties that meets the usual conditions of adequacy on such theories (e.g., characterizing a consequence relation, predicting logical inconsistencies, and so forth). Call this corollary to the modeling challenge the **logical challenge**.

Our job is also, we'll see, complicated by the fact that philosophical and linguistic theorizing about the properties of speech acts has traditionally proceeded in a decidedly "top-down" fashion. As a prelude to theorizing, a stable of speech acts is typically posited or presupposed, for reasons having more to do with general, philosophical commitments about, e.g., the function of discourse (e.g., coordination on *propositional attitudes*, and, even more restrictively, coordination on factual information about the world, i.e., beliefs) than with a specific attention to (i) the enormous variety of linguistic devices available to speakers to express speech acts, (ii) evidence about what these linguistic devices characteristically function to do with tokened in a discourse. As I will argue below, any application of this standard methodology will have substantial difficulty accounting for a variety of non-exotic uses of imperatives. Let's call the task of developing a "bottom-up" approach to theorizing about speech acts that can be seen to explain the relevant data the **empirical challenge**.

All of these challenges represent, in their own way, further challenges to theorizing about the meaning of imperatives in terms of a conventionalized use. Absent solutions to these challenges (or real steps toward such solutions), it is reasonable to doubt whether they could, in principle, be resolved at all. And, indeed, certain theorists have exploited the lack of clear solutions (particularly to the logical and empirical challenges) to argue against the viability of any account of meaning in terms of conventionalized speech acts (or, more generally, uses). Our overarching task in this chapter is to put these doubts to rest.

The plan of this chapter is as follows. First, we will try to develop a "model theory" for imperatives in the spirit of the standard, top-down methodology. The chosen framework will borrow heavily from so-called "Dynamic Semantic" accounts of linguistic meaning, on which sentences are interpreted as expressing a specific type of *deterministic program*—a *cognitive change potential*, i.e., a function from cognitive states into cognitive states. Indeed, we will assume a tight relationship between the notion of a speech act we have been invoking and the Dynamic Semantic notion of a cognitive change potential: speech acts are, I'll suggest, to be understood as proposals to update a cognitive state so that it comes to satisfy some property or other. We will then consider a variety of challenges to that methodology. Many, we'll argue—indeed, nearly all found in the literature—turn out to be spurious. Others, we'll argue, have teeth. So we will modify our model theory, in

a host of non-standard ways—in particular, by introducing a variety of operations for manipulating programs and constructing complex programs out of simpler ones—with the aim of accounting for these problems. Along the way, we will try to flesh out some of the ideas invoked, but not explained, in the prior section, among them: (ir)rationality of contexts, interpretive types (including higher-order interpretive types), necessitation with respect to a body of preferences, rights, and so on.

4.1 A “Standard” Account

What I’ll refer to as the **Standard Account** of speech acts has a venerable pedigree—one which dates at least to the distinctive notation of Frege’s *Begriffsschrift* (specifically, Frege’s use of an assertion stroke scoping over a sentence to indicate that the sentence was to be interpreted as being put forward as asserted). Standard Accounts are characterized by the following two commitments:¹

- There is a basic distinction between *illocutionary force* and *semantic content*. Contents, as such, have no dynamic properties—no *sui generis* discourse-level significance. They must be put forward *with* some illocutionary force or other to be interpreted as making a contribution to a discourse.
- A *speech act* is the result of pairing a force (assertion, querying, commanding, and so on) with a content. Forces can be understood as functions mapping contents into speech acts; assertion, for instance, takes a proposition p as input and yields an assertion that p as output.
- A speaker *performs* a speech act when she proffers a specific force-content complex in a conversation: when she makes an utterance that is, in its final interpretation, interpreted as expressing that speech act.

Contents (variously known as propositions, thoughts, phrastics, descriptors, sentence or propositional radicals) are, then, “essentially forceless”: one and the same content will be apt for being put forth with a variety of distinct forces, and, depending on the force with which it is paired in a specific speech act, can serve the communicative purposes of speakers with a diverse array of communicative intentions (Green 2000: 436). The vast majority of work in contemporary linguistic pragmatics,² takes the standard account as its starting point. That is not surprising. The standard account has many empirical successes. The

1. A sampling of classic references that endorse, more or less, this basic picture: Dummett (1973); Hare (1949, 1952, 1967, 1989); Lewis (1970); Searle (1969, 1975b); Searle & Vanderveken (1985); Stenius (1967); Wittgenstein (1958); Zimmerman (1980). Dynamic semantics in the tradition of Kamp (1981); Groenendijk & Stokhof (1991); Veltman (1996) denies the distinction between linguistic meaning and illocutionary force, opting instead to assimilate linguistic meaning to illocutionary force or update potential. This controversy is not germane to our purposes here (although we will have a bit to say about how to draw the content/force boundary in Sect. 5.3, for a more involved discussion see Charlow 2009). What is important to note is that Dynamic Semanticists typically deny the standard account in letter only, and not in spirit: speech acts in Dynamic Semantics, while assimilated to linguistic meaning, are still represented (or representable) as force-content pairs.

2. There are a few notable exceptions (e.g. Krifka 2001a,b, 2004; Starr 2011), on which more below.

canonical (I think it's fair to say) treatments of assertoric force (Stalnaker 1978), interrogative force (e.g., Hamblin 1973; Groenendijk & Stokhof 1984; Roberts 1996; Groenendijk 1999; Groenendijk & Roelofsen 2009), and directive force (Lewis 1979a; Portner 2004a, 2007) effectively *presuppose* it.³

Here is an example to get the hang of things. The sentences in (4.1), on their typical uses, clearly express distinct speech acts.

- (4.1) a. The window is shut.
b. Is the window shut?
c. Shut the window!

Here is what the Standard Account says about these sentences. Letting *W* be the proposition that the window is shut, (4.1a) typically expresses an assertion (that *W*), (4.1b) typically expresses a question (whether *W*), whereas (4.1c) typically expresses a command that the addressee see to it that some proposition is true (namely, the proposition that *W*). The sentences, and their associated speech acts, are thought to share a *common propositional core*: the proposition that the window is shut. The claim that (4.1a) typically expresses an assertion is understood as shorthand for the claim that a speaker of (4.1a) will normally utter (4.1a) as a way of proffering a speech act of the form *assert(W)*.

4.1.1 Motivations

The Standard Account has considerable *prima facie* plausibility (something to which its considerable empirical successes only add). It can also be (and often is) motivated by appeal to the following sorts of considerations.

Mood, Clause-Type, and Force. There is often claimed to be an intimate connection between a sentence's mood and its function in discourse (see, for instance, Hare 1949, 1952, 1967; Stenius 1967; Lewis 1970; Wilson & Sperber 1988; Portner 2004a). A sentence in the indicative mood typically is used to express assertions, imperative mood to express commands, and so on. But care is required to state this connection in a plausible way. The central claim cannot be that a sentence's grammatical or syntactic mood (as determined by the mood of the main verb of a sentence) typically determines its force. For, as Portner (2004a) (cf. Sadock & Zwicky 1985) notes, typical interrogative clauses are grammatically indicative (because their main verbs are typically grammatically indicative), but are typically used to ask questions, rather than make assertions.

3. To elaborate: according to Stalnaker (1978), assertion amounts to proffering a proposition with assertoric force, the effect of which is to propose the proposition for addition to the Common Ground. According to Roberts (1996), interrogation amounts to proffering a proposition with interrogative force, the effect of which is to add the proposition to the Question Under Discussion Stack, thereby introducing a partition on the context set (Groenendijk & Stokhof 1984). According to Lewis (1979a); Portner (2004a, 2007), commanding amounts to proffering a proposition with directive force, the effect of which is to add the proposition to the agent's practical commitments (her "To-Do List").

The relevant notion from the linguistic literature is, rather, as [Portner](#) notes, the distinct but related notion of *clause-type*. Well-formed sentences in natural language are conventionally associated with a unique clause-type. Certain clause-types (declarative, interrogative, imperative) exist in all known human languages. In all such languages, these clause-types appear to be closely associated with a unique type of force (assertion, querying, and commanding, respectively) (see esp. [Portner 2004a](#)). There is, naturally, disagreement about what the first association (between sentences and clause-types) involves. The different positions fall, broadly, into two camps:

- **Syntactic Typing.** Clauses are typed in virtue of their syntax (in virtue of containing a covert, possibly modal, force-operator, or in virtue of the syntactic form of the clause⁴).
- **Semantic Typing.** Clauses are typed in virtue of their semantics (type of semantic denotation, especially⁵).

I will try, as best I can, to bypass the disagreement here. For our purposes, it is important only that clauses *can be clause-typed* in a principled way. If there are syntactically realized force-operators, then clauses can plausibly be typed using those. If not, we can fall back on syntactic form or semantics. We do not need to settle which notion of clause-type is fundamental in order to make good use of it in our theorizing about imperatives.

What does this have to do with the Standard Account? Clause-type (as well as its close cousin mood) is, of course, understood as a property of an entire clause. How do we arrive at a clause-type representation of a clause? One natural thought, which seems to work well in basic cases, would be to (i) somehow “factor out” the relevant clause’s type, (ii) attach the factored out clause-type to the result of factoring out the clause-type. As the following representations of the sentences in (4.1) show, the result of applying this procedure will, in paradigm cases, be a type-less proposition.

- (4.2)
- a. dec(that the window is shut)
 - b. int(that the window is shut)
 - c. imp(that the window is shut)

Assuming (as we have argued at length above) that clause-types can be mapped onto forces in some principled way, and that declarative type typically maps to assertoric force, interrogative type to questioning force, and imperative type to commanding or requesting force, this picture can be used to *explain* the typical uses of the sentences in (4.1). (4.1b), for instance, typically is used to ask whether the proposition that the window is shut is true *because* its clause-type representation is of the form given in (4.2b).

4. Some exemplary references: [Han \(1998\)](#) develops the former sort of view for imperatives (see [von Stechow & Iatridou 2009](#) for an excellent, critical discussion). [Ginzburg & Sag \(2000\)](#) develop the latter sort of view for interrogatives.

5. As we discussed above, [Portner \(2004a\)](#) develops this sort of view for subject-less imperatives. Such clauses, he claims, denote *properties*, and it is possible to infer their conventional force from this fact.

A couple of natural generalizations get us to the Standard Account from here:

- Assume that for an arbitrary natural language sentence S , the result of factoring out S 's clause-type is a type-less proposition.
- Assume that the clause-type representation for S is of the form $\tau(\phi)$, where ϕ is the result of factoring out S 's clause-type, and τ is S 's clause-type.

The Standard Account follows if we assume, I think plausibly, that a necessary condition on speech act-hood is being the typical interpretation of some typed clause in natural language; acts that lack a canonical natural language expression do not deserve the title “speech acts” at all.

The resulting account has tremendous explanatory appeal, even for theorists who reject the notion of conventionalized speech acts. Everyone must acknowledge that, e.g., interrogatives are at least *typically* used to ask questions. Adopting an account along these lines gives the theorist a recipe for predicting which question an arbitrary interrogative typically expresses: identify the interrogative’s “propositional core” and label it with some device encoding its type, then map this representation into a speech act furnished by the Standard Account.

Consider, e.g., a conditional sentence like (4.3).

(4.3) If Mom turns on the air conditioning, shut the window.

At first pass, (4.3) might be thought to be problematic for the account we’re considering here. We know roughly how to give a clause-type representation for the antecedent and consequent separately—respectively, $\text{dec}(\text{that Mom turns on the air conditioning})$ and $\text{imp}(\text{that the window is shut})$. But it is not clear (i) what the canonical clause-type and speech act representations of (4.3) even ought to look like, (ii) how these two constituent clauses (typed respectively as declarative and imperative) would interact with a conditional operator to generate whatever clause-type is appropriate for (4.3), (iii) how this interaction could be accounted for as an instance of the factoring-out procedure we outlined above.⁶

A response is, however, available. As Schwager (2006), citing syntactic work by Bhatt & Pancheva (2006), notes, “[T]he clause-type of hypothetical conditionals is determined by the clause-type of the matrix clause [i.e., the consequent].” A conditional imperative like (4.3) is, in other words, typed *as an imperative*. Here, building on that insight, is a plausible story about how the interpretation is computed—one that generalizes easily to well-formed conditionals whose constituent sentences are of any type at all.

- The grammar recognizes that (4.3) is a subordinate-matrix clause structure.
- The grammar computes a clause-type for the matrix clause (in this case, imperative).

6. A worry in this vein is pressed in Searle (1962) and has come to be known as one version of the so-called “Frege-Geach Problem” for the view of linguistic meaning known as Expressivism. Much more on Expressivism in a later chapter.

- The grammar assigns the whole clause the type *imp*.
- A typeless proposition is constructed from (4.3) (here, plausibly, the proposition that if Mom turns on the air conditioning, you shut the window).
- Using the resulting clause-type representation, it is inferred that (4.3) typically expresses a request that if Mom turns on the air conditioning, you shut the window.

Far from having a *problem* with sentences like (4.3), the account under consideration seems to have a reasonably complete story of how its typical interpretation is computed. In the end, this case serves only to emphasize its explanatory power, rather than undermine it. It appears, then, that the Standard Account can be grounded in an elegant and powerful account of the relationship between clause-type and force.⁷

Implicitly, on this picture, the computation of clause-type is *global*, rather than *local* (as contrasted with the computation of a clause's semantic value, on which the semantic values of a syntactic tree's terminal nodes combine using combinatory procedures to generate the semantic value for the top-most projection). A syntactic tree for the clause is constructed, and a clause-type is assigned to its top-most projection. Constituent clauses of a complex clause can be typed, but not *qua* constituent clauses. Typed clauses can be embedded, as in the antecedent of (4.3), but clause-type itself does not embed; the clause-type representation of (4.3) has *imp* scoping over a type-less proposition, rather than *imp* scoping over a declarative. The restriction on the "embeddability" of clause-type (and other clause-level indicators of conventional force) can be traced to Frege (*Über Sinn und Bedeutung*); Searle (1962); Geach (1965) (for relevant discussion, see also Zimmerman 1980; Price 1994; Green 2000).

Speech Acts as Public Propositional Attitudes. It is a platitude that speech acts function to express, in some sense or other, the mental states of a speaker. From here it is a short step to the notion that speech acts function as public expressions of a speaker's *propositional attitudes*.

Propositional attitudes, as such, are traditionally individuated by their functional role, rather than their objects: each propositional attitude is a relation between a subject and a propositional content, but the relations themselves have different functional profiles (beliefs and desires have different "directions of fit," desires and wishes have different variations of the same direction of fit, and so on). Speech acts function to voice or express such attitudes: assertion is the expression of belief, interrogation is the expression of a species of uncertainty or wonder, commanding and requesting are the expression of some species of desire. On this sort of picture, it would seem that it is possible to give a satisfactory theoretical account of an arbitrary type of speech act by (i) associating it with some sort

7. There are a number of sleights of hand in this argument, which ultimately cause it to fail, but I will wait until later on (Section 4.5) to discuss them. For now, my main goal is to build as good a case for the Standard Account as I can manage.

conversational “propositional attitude”—a public stance that a speaker can take toward a propositional content, corresponding to the attitude that the speech-act functions to express, (ii) defining an appropriate functional role for each of these conversational attitudes.

Propositional attitudes are, of course, usually understood as relations between an agent and a content.⁸ It follows that an agent’s propositional attitudes can be represented by specifying, for each of her propositional attitudes, (i) a propositional attitude type, (ii) a content. And so, insofar as speech acts represent public expressions of a speaker’s propositional attitudes, we will expect that, for any speech act α , there is some propositional attitude A such that α can be represented by pairing (i) a “public” analogue of the attitude-type of which A is an instance and (ii) A ’s content. This entails the Standard Account.

What is involved in the public expression of a propositional attitude? Here is a natural idea. Public expressions of propositional attitudes function as proposals to coordinate on a given propositional attitude. Given an arbitrary speech act of the form $\Delta\phi$, is natural to think of the force operator Δ as determining a *parameter* π of a mental state on which the speech act designated by $\Delta\phi$ proposes coordination (and perhaps the general type of update operation to be performed on π), while the contribution of the propositional complement ϕ is to determine a *specific condition on* π that coordination in accordance with $\Delta\phi$ will enforce.⁹ *Genera* of illocutionary force are individuated by the parameter of the mental state they function to update. So, for instance, what unites different kinds of assertoric force (explanation, narration, and so on) is that they all involve update to the informational parameter of the state (informally, what is believed by the state). What distinguishes assertoric force, as a category, from action-guiding or directive force is the fact that action-guiding or directive force (and all of its subtypes) involves updating the planning or preferential parameter of the state. *Species* of a single genus of illocutionary force may be individuated by the general type of operation they perform on a single contextual parameter: perhaps permission and commanding are both kinds of action-guiding force (and, so, both involve updates to plans), but permission involves the relaxation of plans, while commanding involves restricting the amount of choice within a given plan. The functional role for a genus of illocutionary force is a matter of *which* parameter of the state it targets for update. The functional role for a species of illocutionary force is a matter of *which* parameter of the state it targets for update, as well as the general type of operation it tends to perform on that parameter.

When presented abstractly in this way, it may seem the Standard Account has little in the way of empirical content—that rubber hits road only once substantive accounts of (i) the

8. Those who are fans of coarse-grained contents often introduce something like a mode of presentation or sense into the mix. This would require a modification of the Standard Account, but not a major one: speech acts would be modeled as expressions of an attitude toward a content under some mode of presentation or other (or, just as good, a proposal for one’s interlocutors to come to accept that content under the relevant mode of presentation). For my purposes, this distinction is unimportant: I will eventually be concerned to argue against any account of speech acts on which they are understood as essentially expressing propositional attitudes.

9. Portner (2004a, 2007) is explicit about his commitment to this sort of idea.

nature and functional role of the relevant cognitive parameters, (ii) how different varieties of speech acts propose to update these parameters are on the table. That is the wrong way to regard the Standard Account. Certain, propositional attitude-ish understandings of the relevant cognitive parameters—an understanding which underlies all of the major extant treatments of the empirical phenomena—integrate well with the Standard Account (indeed, as already mentioned, they presuppose it). But other profiles of these parameters—the ones, I’ll argue, with the best empirical support—*do not*. It is an empirical question whether a Standard Account-friendly characterization of, e.g., the assertion-, interrogation-, or command-sensitive parameter of the context is correct. And the evidence, I’ll argue, is not friendly to the Standard Account.

Indeed, the place to push on this “Propositional Attitude” argument for the Standard Account occurs right at its outset. It is, indeed, a platitude that a speech act is an expression of a speaker’s state of mind. What is *not* a platitude is that the only attitudes for which there exist canonical linguistic modes of expression are *propositional* attitudes. Arguing against this, however, will require identifying sentences that canonically express some mental state that cannot be represented as a propositional attitude. That, however, is no mean feat. Indeed, if the relationship between clause-type and force sketched above is correct, this will probably prove to be impossible. For, if all clauses can be represented as typed clauses of the form $\tau(\phi)$ (with ϕ a type-less proposition and τ a clause-type), and if τ typically is associated with a given speech act-type, then τ will very probably be associable with some attitude-type or other. What type of attitude? An attitude toward a proposition (in this case, toward the proposition ϕ), i.e., a *propositional* attitude.

Semantics and Pragmatics. There is a distinctive picture of the semantics/pragmatics divide that seems to be involved with adopting the orientation of the Standard Account. Semantic theory (or, more broadly, the theory of conventional meaning) assigns objects to sentences (contents) that are suitable for being put forward with some force or other. Pragmatic theory subsequently associates these objects with discourse-level phenomena. How does pragmatic theory accomplish this? Plausibly, by taking the content assigned to a sentence by the semantic theory, together with further information about an utterance (e.g., its clause-type!), to determine the illocutionary force with which the content is proffered.

While it is true that many theorists of the Standard Account have argued for recognizing certain notions typically regarded as pragmatic as part of conventional linguistic meaning (some randomly selected examples are [Hare 1949, 1967](#); [Searle 1969, 1975a](#); [Potts 2005, 2007](#); [Dummett 1973](#); [Grice 1989](#)), all of these theorists agree in spirit with the division of semantic and pragmatics I’ve just sketched. The *semantically* interesting properties of a sentence—if not all the properties relevant to the sentence’s conventional meaning—can be represented with a simple propositional content. Whereas the pragmatically interesting properties of a sentence arise as the result of performing some kind of illocutionary operation on that

content. So, for instance, the semantically interesting properties of a sentence like (4.3) are exhausted by the proposition that if Mom turns on the air conditioning, the window is shut. The fact that (4.3) expresses a conditional request can be explained as a reflex of the facts that (i) (4.3) semantically expresses a conditional proposition, (ii) (4.3) is typed as an imperative. The typical force of (4.3) can thus be represented by performing the illocutionary operation typically associated with imperatives on the conditional content expressed by the sentence. Similarly, (4.4)'s expression of a conditional question can be explained as a reflex of the facts that (i) (4.4) semantically expresses a conditional proposition, (ii) (4.4) is typed as a question.

(4.4) If Mary doesn't come to the party, will you be upset?

This is, on its face, an explanatory, elegant, historically popular, and widely shared picture of the semantics/pragmatics and conventional meaning/pragmatics divides. It is, nevertheless, one that is opposed by many theorists (including not a few partisans of the Standard Account). Although many of their reasons for opposing it in the case of imperatives will be resisted here (see Section 4.4.1), we will ultimately side with these critics. To preview, it is actually impossible to reconcile this picture with the notion, developed at length in the prior section, that clause-type determines type of semantic denotation. In the case of imperatives, this sort of picture arguably fails to account for (i) many of the fundamental properties of the conventional force of *embedded* imperatives, conditional (Section 4.5) and otherwise (Section 4.4). And, although we will not get into this here, it is well-known that this sort of picture of the force of conditional questions fails to account for the pragmatically interesting properties of conditional questions (see, e.g., Isaacs & Rawlins 2008; Groenendijk & Roelofsen 2009).

So much, then, for motivating the Standard Account. In the next section, I will sketch one way of formalizing that account, by way of resolving the “modeling challenge” described in this section's introduction. On this picture, speech acts will be understood as instructions for updating an information state—a notion which can be implemented precisely using some basic techniques from Dynamic Semantics and the formal logic of programming.

4.1.2 Modeling the Standard Notion of a Speech Act

In this section, we will pursue a formal implementation of the claim that clauses conventionally express speech acts that takes the Standard Account as its starting point.

Speech acts, on the Standard Account, are the result of pairing a specific kind of illocutionary force and a specific content. Although there are various ways to formalize this basic idea, the formalization opted for here integrates smoothly with extant formalisms in dynamic modal logic and the logic of computing, and allows us to draw out commitments

of the Standard Account and formulate questions about speech acts in natural language in an exceptionally clear way.¹⁰ The ingredients to the formalization will be:

- A first-order **base language** \mathcal{L} , the language of propositions.
- A finite set of **speech act operators** $\mathcal{O} = \{\triangleright, ?, !, \dots\}$, serving to represent the various types of illocutionary force with which a speaker can proffer a proposition.
- A recursive definition a set \mathcal{S} of **speech act designators** (SADs) from \mathcal{L}, \mathcal{O} .
- A **semantics** for \mathcal{S} .

Note that we are using the same symbols for illocutionary operators as we have been using for typing clauses. This is a notational convenience, rather than a theoretical sleight of hand, and it actually allows us to stay neutral on an important question. Our ultimate aim here is to define a mapping from a regimented formal language (one suitable for representing the “logical forms” of the relevant fragment of natural language) into speech acts. Here, the notion of logical form is to be understood in an exceptionally thin sense: a sentence’s logical form, in the sense of this section, will be just a syntactic object that represents its conventional meaning (or perhaps some theoretically interesting aspects of its conventional meaning). But should the logical form of a sentence be understood, on the Standard Account, as a *typed clause*, or as a *speech act designator*? For a view like that of Portner (2004a), it might make sense to understand logical forms, in the relevant sense, in the latter way. For Portner, typed clauses determine conventionalized speech acts only by *mediation* of pragmatic principles associating ordinary semantic objects (propositions, properties, etc.) with forces. A typed clausal representation of some natural language sentence, per se, represents only one dimension of that sentence’s conventional meaning.

The Base Language \mathcal{L} . The base language is defined in the standard way: we begin with sets of variables $\mathcal{V} = \{x, y, z, \dots\}$, constants $\mathcal{C} = \{a, b, c, \dots\}$, relation-symbols of varying arity $\{P, Q, R, \dots\}$, connectives $\{\vee, \neg\}$, and the quantifier \exists .

Definition 4.1. \mathcal{L} is the smallest set such that:

- i. If $t_1, \dots, t_n \in \mathcal{V} \cup \mathcal{C}$ and P an n -ary relation-symbol, then $P(t_1, \dots, t_n) \in \mathcal{L}$
- ii. If $\phi \in \mathcal{L}$ then $\neg\phi \in \mathcal{L}$
- iii. If $\phi \in \mathcal{L}, \psi \in \mathcal{L}$, then $(\phi \vee \psi) \in \mathcal{L}$
- iv. If $\phi \in \mathcal{L}$ and x is a variable, then $\exists x\phi \in \mathcal{L}$

We stipulate that $(\phi \wedge \psi) := \neg(\neg\phi \vee \neg\psi)$ and $\forall x\phi := \neg\exists x\neg\phi$.

10. For an approach in a broadly similar vein, see Isaacs & Potts (2003); Potts (2003).

Speech Act Designators. The set \mathcal{S} is constructed in the expected way from \mathcal{L} and \mathcal{O} : any pairing of a force-operator and a first-order sentence represents a possible speech act, and nothing else represents a possible speech act. That is to say:

Definition 4.2. \mathcal{S} is the smallest set such that, if $\phi \in \mathcal{L}$ and $\Delta \in \mathcal{O}$, then $\Delta\phi \in \mathcal{S}$.

Definition 4.2 is a straightforward corollary of the Standard Account. Note that the definition of \mathcal{S} , unlike that of the base language \mathcal{L} , is *non-recursive*, i.e., no provision is made (nor, on the Standard Account, *should* any provision be made) for any method of building “complex” SADs from “atomic” SADs. So long as speech acts are represented as force-content pairs, and both \mathcal{L} and \mathcal{O} are sufficiently rich, any force-content pair of interest will be designated by some member of \mathcal{S} . Of course, some force-content pairs will fail to be designated by some member of \mathcal{S} , as defined. This is, on the Standard Account, an artifact of the *expressive deficiencies* of the first-order base language—its inability to express some propositions that we are able to proffer with various kinds of illocutionary force in natural language. The strategy for ameliorating this is to *enrich the base language*. The manner in which \mathcal{S} is defined will be left unchanged.

A major consequence of Definition 4.2 is that no SAD ever embeds another SAD.¹¹ This is desirable: on the Standard Account, speech acts are composed of a force and a force-less content; no speech act ever contains another as a proper part. This consequence is often given the status of a dictum: “force,” we are often told, “does not embed” (see esp. Green 2000).¹²

Semantics for \mathcal{L} . Since \mathcal{L} is the “base” language, we want for formulas of \mathcal{L} to express propositional contents. The simplest fully explicit way of doing this is (i) to treat propositions as sets of possible worlds (equivalently, functions from possible worlds to truth-values), (ii) to state a constant domain possible worlds semantics for \mathcal{L} , which has formulas of \mathcal{L} characterizing possible worlds propositions.

Definition 4.3. A **constant domain frame** is a pair $\mathcal{F} = \langle W, D \rangle$, with W a set of possible worlds, D a set of individuals.

Definition 4.4. A **constant domain model** built on a constant domain frame $\mathcal{F} = \langle W, D \rangle$

11. We might design a base language in which the SADs are treated as *terms*, in which case some SADs would, technically speaking, embed other SADs. We will suppose that our base language is not like this (indeed, given that we will be treating SADs as denoting update functions, rather than individuals, the stipulation that the base language is first-order will require this).

12. This captures one salient sense of that dictum. But there is another salient sense that it does not capture: the sense on which syntactic objects that express speech acts as a matter of their meaning (so-called “conventional force indicators,” e.g., representations of clause-type) cannot adjoin to clauses that are embedded at the level of logical form—that clausal force indicators must take widest scope at logical form. While these ways of cashing out the dictum about the non-embeddability of force are distinct, it is natural to treat them as a package. And, indeed, as we saw above, this does fall out of one picture of the motivations for the Standard Account—the global understanding of clause-typing.

is a triple $\mathcal{M} = \langle W, D, \llbracket \cdot \rrbracket \rangle$, where $\llbracket \cdot \rrbracket$ is an interpretation function built on \mathcal{F} meeting the conditions in Definition 4.6.

Definition 4.5. An **assignment** for a constant domain frame $\mathcal{F} = \langle W, D \rangle$ is any function $g : \mathcal{V} \mapsto D$. A **modified assignment** $g^{[x/d]}$ for \mathcal{F} is an assignment for \mathcal{F} such that $g^{[x/d]}(x) = d$ and $g^{[x/d]}(y) = g(y)$ for all $y \neq x$.

Definition 4.6. An **interpretation function** $\llbracket \cdot \rrbracket_{\mathcal{F},g}$ built on a constant domain frame $\mathcal{F} = \langle W, D \rangle$ with respect to an assignment g is any function satisfying the following conditions. Let $w, v \in W$.

- i. If $a \in \mathcal{C}$, $\llbracket a \rrbracket_{\mathcal{F},g,w} \in D$ and $\llbracket a \rrbracket_{\mathcal{F},g,w} = \llbracket a \rrbracket_{\mathcal{F},g,v}$
- ii. If $x \in \mathcal{V}$, $\llbracket x \rrbracket_{\mathcal{F},g,w} = g(x)$
- iii. If P is an n -ary relation symbol, $\llbracket P \rrbracket_{\mathcal{F},g,w} \subseteq D^n$

Definition 4.7. Let $\mathcal{F} = \langle W, D \rangle$ be a constant domain frame and $\mathcal{M} = \langle W, D, \llbracket \cdot \rrbracket \rangle$ be a constant domain model built on \mathcal{F} . The **satisfaction relation** for \mathcal{L} , \models , is defined as follows:

- i. $\mathcal{M}, w, g \models P(t_1, \dots, t_n)$ iff $\langle \llbracket t_1 \rrbracket_{\mathcal{F},g,w}, \dots, \llbracket t_n \rrbracket_{\mathcal{F},g,w} \rangle \in \llbracket P \rrbracket_{\mathcal{F},g,w}$
- ii. $\mathcal{M}, w, g \models \neg\phi$ iff $\mathcal{M}, w, g \not\models \phi$
- iii. $\mathcal{M}, w, g \models (\phi \vee \psi)$ iff $\mathcal{M}, w, g \models \phi$ or $\mathcal{M}, w, g \models \psi$
- iv. $\mathcal{M}, w, g \models \exists x\phi$ iff $\exists d \in D : \mathcal{M}, w, g^{[x/d]} \models \phi$

For all $\phi \in \mathcal{L}$, we stipulate that $\llbracket \phi \rrbracket_{\mathcal{M},g,w} = 1$ iff $\mathcal{M}, w, g \models \phi$

Definition 4.8. Finally, if $\phi \in \mathcal{L}$, the **proposition** $\llbracket \phi \rrbracket_{\mathcal{M},g}$ expressed by ϕ in a model $\mathcal{M} = \langle W, D, \llbracket \cdot \rrbracket \rangle$ with respect to assignment g is $\{w \in W : \llbracket \phi \rrbracket_{\mathcal{M},g,w} = 1\}$

On a Semantics for \mathcal{S} . So far, so standard. What remains is to layer a semantics for \mathcal{S} on top of the semantics for \mathcal{L} . The problem, as we noted in this section's introduction, is that the Standard Account, per se, does not provide very much guidance about how to proceed here. But well-known work in Dynamic Semantics, which models conventional linguistic meaning in terms of update or context-change potentials, gives us some notion of how to state the semantics for \mathcal{S} that is friendly to the Standard Account.

In Dynamic Semantics, the use of a sentence is given a fundamental role: the meaning of a sentence is identified with what it is used to do to an information state (Kamp 1981; Groenendijk & Stokhof 1991; Veltman 1996). Actions "deployed" on a cognitive state are understood as a kind of program: for each input cognitive state, an action performs some update or other on that cognitive state. Actions can thus be represented as a kind of program: a relation (specifically, a function) that relates an input cognitive state to an

output cognitive state. The meaning of a sentence, then, can be understood in terms of its *update potential*: the sort of update it is disposed to perform on a cognitive state (or context).

We can gain a lot of traction in our project of modeling the Standard Account if we adopt a theoretical identification between the Dynamic Semantic understanding of meanings (as update potentials, or actions on a cognitive state) and the traditional notion of a speech act. On this picture, a speech act, very simply, is an instruction for updating a *cognitive* state. Insofar as actions, in general, can be modeled as relations on states, speech acts will be treated as a genuine *kind of action* (on this notion of an action, see esp. [Segerberg 1992](#); [Harel et al. 2000](#)).¹³ What distinguishes speech acts from “normal” acts are:

- The *kinds of states* on which they “operate” (cognitive states of mind versus non-cognitive states of the world)
- The *success-conditions* for the act. A normal act typically (although not universally) succeeds just if it is executed by the actor and the actual world is among the input states for the action. A speech act, however, can typically still fail even when it is executed by the actor and the interlocutor’s cognitive state is among the input states of the speech act, if the interlocutor *rejects* it.

Note that in making use of the Dynamic Semantic understanding of meanings, we are not thereby taking on the Dynamic Semantic understanding of meanings. We are simply saying that speech acts can be modeled using the sorts of things that Dynamic Semantics uses to represent meanings, even if speech acts themselves do not exhaust linguistic meaning.¹⁴ Informally, the picture this give rise to is this:

- A force—the semantic denotation of a force operator—is represented as a function from propositions into functions from states into states.
- A speech act is represented as the result of applying a force to a proposition.

This is pretty much all that is needed to formalize the Standard Account, as such. And the theoretical commitments of the standard account do seem faithfully rendered in our chosen formalism. Theoretically, of course, this account is not yet very satisfying. We are lacking substantive accounts of the nature and functional role of the relevant cognitive parameters that speech acts update, as well as accounts of how specific kinds of speech acts actually update these parameters. These accounts are essential components of any theory

13. Rich Thomason has pointed out to me that I could be misread as intending to identify a speech act with an operation on a cognitive state—a cognitive program. As he notes, this would be to conflate illocutionary acts (like assertion) with perlocutionary acts (like persuasion) ([Austin 1975](#)), since executing an assertion would be a matter of executing the cognitive program on the state (and thereby be identified with getting the addressee to accept the assertion’s content). To be clear: I, like [Swanson \(2006\)](#), understand illocutionary acts as *proposals* for an addressee to execute a program on her cognitive state. Perlocutionary acts are, in this framework, understood in terms of the acceptance by an addressee of an illocutionary act proffered by a speaker. The *common core* to an illocutionary act and a perlocutionary act is an update function or cognitive program. Illocutionary acts propose update functions to an agent, while perlocutionary acts involve an agent’s acceptance of such proposals.

14. For further discussion, see especially Section 5.3 and Appendix A.

of communication—without them, we have no real understanding of the conversational function of speech acts, of the role that speech acts play in the communicative enterprise. There is, not surprisingly, no shortage of such accounts for the major categories of speech act (assertions, interrogations, and commands). The bulk of our attention in the remainder of this section will be focused on such accounts.

4.2 Information, Issues, and Preferences

In this section, I'll start out with a very simple model of cognitive states and update functions. Cognitive states will be initially understood as information states (bodies of information, with these being understood as sets of worlds). Update functions will be initially understood as programs for increasing the information available at a given cognitive state (i.e., as performing intersective operations on the given cognitive state). This model turns out to be far too simple, so we will identify problems and suggest amendments as needed. At the end of this section, we will have a compositional system, fully consonant with the Standard Account, that associates a precise type of update function to an arbitrary speech act designator in \mathcal{S} .

4.2.1 Information States

Here and throughout, σ will designate a cognitive state (implicitly, a cognitive state for an agent, or, more abstractly, the mutually shared “cognitive state” characterized by a context). In the most basic case, σ is typed as a set of worlds, i.e., as a purely informational state. Let W be the universe. (Here and throughout, the model supplying the model-theoretic objects will be left implicit unless otherwise required.) Then:

Definition 4.9. σ is an **information state** iff $\sigma \subseteq W$.

Speech act designators of the form $\triangleright\phi$ represent assertions. In line with our earlier idea that...

- Speech act operators determine (i) the parameter of the state to be updated, as well as (ii) the specific operation to be performed on the state (i.e., a function from contents to specific set-theoretic operations)
- The speech act operator's complement provides the argument to this function

...we will say that the speech act operator \triangleright determines that (i) an information state is selected for update, (ii) an intersective operation is performed on the information state. The complement of \triangleright , a sentence $\phi \in \mathcal{L}$, supplies the argument to this operation. The complex $\triangleright\phi$ thus expresses a function mapping an information state to the result of intersecting that information state with the proposition that ϕ . Recursively, letting $A = \{a_1, \dots, a_n\}$ the (finite) list of constants for the base language:

Definition 4.10. $[\cdot]$ is an interpretation function mapping elements of \mathcal{S} into speech acts. If $\Delta \in \mathcal{O}$ and $\phi \in \mathcal{L}$, we write $\sigma[\Delta\phi]$ for the result of updating σ with $[\Delta\phi]$, i.e., $\sigma[\Delta\phi] := [\Delta\phi](\sigma)$.

$$\begin{aligned}\sigma[\triangleright p] &= \sigma \cap \llbracket p \rrbracket \\ \sigma[\triangleright \neg\phi] &= \sigma - \sigma[\triangleright\phi] \\ \sigma[\triangleright(\phi \wedge \psi)] &= \sigma[\triangleright\phi][\triangleright\psi] \\ \sigma[\triangleright\forall x\phi] &= \sigma[\triangleright\text{sub}^{a_1/x}(\phi)] \dots [\triangleright\text{sub}^{a_n/x}(\phi)]\end{aligned}$$

Notice that, in view of the non-recursive nature of \mathcal{S} , and in view of the fact that, for any $\phi \in \mathcal{L}$, $\llbracket \phi \rrbracket \subseteq W$, this definition is actually equivalent to the following *non-recursive* definition.

Definition 4.11. For all $\phi \in \mathcal{L}$, $\sigma[\triangleright\phi] = \sigma \cap \llbracket \phi \rrbracket$

This is an important point, one to which we will return repeatedly. Assuming that the domain of speech acts has any sort of recursive structure (or, more generally, any sort of complexity) is *not essential* to formalizing an account of assertion that is consonant with the motivations of the Standard Account. The recursion of Definition 4.10 is inherited from the recursion inherent in defining the base language \mathcal{L} .

A case against the Standard Account that makes use of examples of complex speech acts will ultimately require identifying speech acts that are **essentially complex**, as opposed to merely superficially complex.

Definition 4.12. A speech act α is **essentially complex** iff for no $\Delta\phi \in \mathcal{S}$ is it the case that $\alpha = [\Delta\phi]$

Definition 4.13. A speech act α is at most **superficially complex** iff for some $\Delta\phi \in \mathcal{S}$, $\alpha = [\Delta\phi]$

Superficially complex speech acts are those that can be represented (but needn't be represented) as complex. For the theorist of the Standard Account can always argue that, for instance, a speech act represented as the composition of two more basic speech acts $[\Delta\phi]$ and $[\Delta\psi]$ —as $[\triangleright(\phi \wedge \psi)]$ is represented in Definition 4.10 as the composition of $[\triangleright\phi]$ and $[\triangleright\psi]$ —is *fundamentally* just the speech act $[\Delta(\phi \wedge \psi)]$. The fact that a speech act has a complex *representation* does not require that the speech act itself be complex.

4.2.2 Integrating Issues

Nothing about the Standard Account requires typing either cognitive states or the central cognitive parameters as unstructured *sets of worlds*, or interpreting addition to such parameters in terms of the restrictive operation \cap . This sort of commitment would, in fact, be problematic. One reason is that modeling interrogative force seems to demand more

structure in the dedicated parameter of the discourse that interrogative speech acts select for update.

Standard treatments of interrogative speech acts treat them as introducing a question or issue into a discourse (e.g. Groenendijk & Stokhof 1984; Roberts 1996; Groenendijk & Roelofsen 2009). A question under discussion is typically represented as *partitioning the information state* into the salient alternatives or possibilities (i.e., the complete answers to the question; cf. Hamblin 1973). The question *whether* ϕ partitions the information state according to its complete answers, so that, if $\sigma \subseteq W$ is the initial information state, the result of updating σ with *whether* ϕ is the partition $\{\sigma[\triangleright\phi], \sigma[\triangleright\neg\phi]\}$.

Interrogation is not representable as restriction of a set of worlds. It is essential to distinguish a state where two questions (*whether* ϕ and *whether* ψ) are introduced from one where a single question (*whether* $\phi \wedge \psi$) is introduced. Four possibilities are salient at the former state, while just two are salient at the latter. If $\sigma \subseteq W$ is the initial state, then the former state is represented by the following partition:

$$\{\sigma[\triangleright\phi][\triangleright\psi], \sigma[\triangleright\phi][\triangleright\neg\psi], \sigma[\triangleright\neg\phi][\triangleright\psi], \sigma[\triangleright\neg\phi][\triangleright\neg\psi]\}$$

The latter, however, is represented by the partition:

$$\{\sigma[\triangleright(\phi \wedge \psi)], \sigma[\triangleright\neg(\phi \wedge \psi)]\}$$

Reading the partition off a set of worlds that is incrementally restricted by successive interrogative acts means that a state where the questions *whether* ϕ and *whether* ψ are introduced cannot be distinguished from one where just the question *whether* $\phi \wedge \psi$ is introduced.

A better model might keep separate track of each issue introduced at a state (cf. Roberts 1996). There is a set Q_σ of issues introduced at a state σ , where $\llbracket\phi\rrbracket \in Q_\sigma$ iff *whether* ϕ is introduced at σ . Each element of Q_σ induces a partition on the information state associated with σ , where each cell of the partition represents a complete answer to all of the questions introduced at σ . A unified account of illocutionary force might, therefore, treat cognitive parameters as sets of *propositions*. In general, an illocutionary act with content $\llbracket\phi\rrbracket$ that targets some cognitive parameter for update will yield a state at which that parameter *contains* that content. Formally:

Definition 4.14. If $\phi \in \mathcal{L}$ and $\Delta \in \mathcal{O}$ is an illocutionary force-type targeting cognitive parameter \mathcal{P}_σ for update, then for all $\mathcal{P}'_\sigma \neq \mathcal{P}_\sigma$, $\mathcal{P}'_{\sigma[\Delta\phi]} = \mathcal{P}'_\sigma$, while $\mathcal{P}_{\sigma[\Delta\phi]} = \mathcal{P}_\sigma \cup \{\llbracket\phi\rrbracket\}$.

For characterizing the functional role for certain parameters (e.g., the parameter devoted to keeping track of questions introduced at a state), this extra structure would be indispensable. It is an open question whether the extra structure is essential for characterizing the functional

role for other parameters of the state.¹⁵ For present purposes, it is only important to note that any sort of account on which speech acts essentially involved adding a proposition to a selected parameter would be *fully consonant* with the spirit of the Standard Account; speech act operators (elements of \mathcal{O}), when interpreted, would specify some parameter or other for update, while their complements would specify the proposition that would be added to the selected parameter. Speech act designators of the form $\Delta\phi$ would be sufficient to represent the entire range of available speech acts.

A different, and somewhat more technically elegant, approach might follow Groenendijk & Stokhof (1984); Groenendijk & Roelofsen (2009) in treating information states themselves as partitions (and cognitive states as exhausted by information states); on this picture, an information state (and a cognitive state) would function solely to keep track of both live possibilities and relevant issues in some inquiry or other. On this picture, assuming that an information state σ is a partition, we will have the following definitions of $[\triangleright\phi]$ and $[?\phi]$:

Definition 4.15. For all $\phi \in \mathcal{L}$, $\sigma[\triangleright\phi] = \bigcup_{p \in \sigma} \{p \cap [\phi]\}$

Definition 4.16. For all $\phi \in \mathcal{L}$, $\sigma[?\phi] = \bigcup_{p \in \sigma} \{p[\triangleright\phi], p[\triangleright\neg\phi]\}$

So, updating with $\triangleright\phi$ involves restricting all cells p in the partition to the ϕ -possibilities in p . Updating with $?\phi$ involves partitioning each cell p into the ϕ -possibilities in p and the $\neg\phi$ -possibilities in p . This allows us to define some pragmatically useful notions like the following (cf. Yalcin forthcoming).

Definition 4.17. $p \subseteq W$ is a **live possibility** at σ iff $\exists q \in \sigma : p \cap q \neq \emptyset$

Definition 4.18. $p \subseteq W$ is a **salient possibility** at σ iff $\forall q \in \sigma : p \cap q = q$ or $p \cap q = \emptyset$

Live possibilities are compatible with the information at σ , while salient possibilities correspond to issues already introduced at σ .

Notice that questioning is in Definition 4.16 defined *in terms of assertion*. Indeed, questioning might be understood in this case as the act of raising alternative assertions. This in turn might be thought to involve accepting a modicum of complexity into the domain of speech acts—a complexity that is incompatible with the spirit of the Standard Account. Questions are built out of assertions, rather than corresponding to a “basic” illocutionary force-type.

This is wrong, for at least two reasons. First, the Standard Account says nothing about illocutionary force-types needing to be basic. Even if questions are “built out of” assertions, they might well still correspond directly to some kind of propositional attitude or

15. Of special relevance for our purposes is whether a Lewisian permissibility sphere (typed as a set of worlds) has sufficient structure to represent the functional role of the cognitive parameter that directive speech acts target for update—whether treating the parameter that a speech act $[!\phi]$ targets for update as a set of propositions is essential in the way that it is essential for the parameter devoted to keeping track of the issues under discussion. In the next section, we argue that there is good reason to think that it is.

other—one, perhaps, that is built out of beliefs. Entertaining the issue that ϕ appears to be a propositional attitude, but it could be understood as entertaining alternative beliefs (namely, the belief that ϕ and the belief that $\neg\phi$). So the propositional attitude of entertaining might not be cognitively basic, but it would still, for all that, be a propositional attitude. As we saw above, associating each illocutionary force-type with some specific propositional attitude or other is sufficient to establish the Standard Account.

Second, as the use of recursion in defining updates on assertions was inessential, so too the use of assertion in defining updates on questions is inessential. Indeed, because an arbitrary assertion is non-recursively definable (as in Definitions 4.11 and 4.15), we might just as well replace Definition 4.16 with the following.

Definition 4.19. For all $\phi \in \mathcal{L}$, $\sigma[?\phi] = \bigcup_{p \in \sigma} \{p \cap \llbracket \phi \rrbracket, p \cap \llbracket \neg\phi \rrbracket\}$

Even if the non-basic-ness of a speech act-type were in itself problematic, it would not be problematic in the case of questioning. For only *essential non-basic-ness*—the inability to define the speech act-type without invoking other speech act types—could fairly be regarded as problematic. If some speech act-type is merely non-essentially non-basic (as questioning, on our partition-inducing understanding, is), the theorist of the Standard Account might simply maintain that the basic representation of the speech act, although equivalent to some non-basic representation, is in fact the fundamental representation.

4.2.3 Integrating Preferences

In this section, we will consider several alternatives for understanding the speech act expressed by a speech act designator of the form $!\phi$ (which I will neutrally label *direction*). Let's first recapitulate the Lewisian account (cf. Section 2.1), adapting it to the model we have been developing in this section.

The Lewisian Account. Let $\sigma = \langle A_\sigma, P_\sigma \rangle$ be a cognitive state. Then:

- The *accessibility sphere* $A_\sigma \subseteq W$ represents states of the world that are possible for σ .
- The *permissibility sphere* $P_\sigma \subseteq A_\sigma$ represents the σ -accessible states that σ can permissibly realize.

Recall that Lewis (1979a) models the move of directing that ϕ as a restriction of the permissibility sphere to states satisfying ϕ :

Definition 4.20. For all $\phi \in \mathcal{L}$, $\sigma[!\phi] = \langle A_\sigma, P_\sigma \cap \llbracket \phi \rrbracket \rangle$

The model can be adapted to account for issues in the following way. We understand cognitive states as pairs of information states (understood as partitions) and permissibility spheres. Instead of accessibility spheres, which represent for Lewis something like world-states that are *physically available* to the agent, we will make use of something that is cognitive

in character—bodies of information (subsets of W), which are partitioned when the state comes to recognize an issue.

- The *information state* I_σ represents both the information available at σ as well as the issues recognized by σ . (As such, I_σ is a partition of the information available at σ .)
- The *permissibility sphere* $P_\sigma \subseteq \bigcup I_\sigma$ represents world-states that are permissibly realizable at σ .

Integrating the Lewisian picture with the picture sketched in previous sections gives us the following, more complete specification of $[\cdot]$.

Definition 4.21. For all $\phi \in \mathcal{L}$, $\sigma[\triangleright\phi] = \langle \bigcup_{p \in I_\sigma} \{p \cap \llbracket \phi \rrbracket\}, P_\sigma \rangle$

Definition 4.22. For all $\phi \in \mathcal{L}$, $\sigma[?\phi] = \langle \bigcup_{p \in I_\sigma} \{p[\triangleright\phi], p[\triangleright\neg\phi]\}, P_\sigma \rangle$

Definition 4.23. For all $\phi \in \mathcal{L}$, $\sigma[!\phi] = \langle I_\sigma, P_\sigma \cap \llbracket \phi \rrbracket \rangle$

This account of direction is, of course, *fully Standard*. The relevant speech act is modeled as the pairing of an illocutionary force (what in Lewis’s case we could accurately label *commanding*) with a propositional content. What especially bears noting in the context of our discussion here is the fact that Lewis:

- Separates out a dedicated *action-guiding* or *preferential* cognitive parameter (permissibility sphere) from the informational/world-tracking parameter (information state).
- Treats assertion and interrogation as matters of performing distinct operations (information-addition and issue-addition) on the world-tracking parameter.
- Treats direction as a matter of performing a restrictive operation on this dedicated, action-guiding parameter. Direction is not understood in terms of an informational upshot. Instead, direction has a *practical* upshot—it directly affects the preferences according to which the cognitive state ranks outcomes and constructs plans to realize the best outcomes.

The Portner Account. Recall that Portner (2004a, 2007, 2010) also models the move of directing that ϕ as modification of a dedicated, action-guiding parameter (as opposed to modification of an information state). But while Lewis treats the relevant modification as restriction of the permissibility sphere to ϕ -possibilities, Portner does not make use of a permissibility sphere at all. Instead, there is a *body of preferences*—what Portner calls the “To-Do List”—to which an imperative adds its preadjacent. Letting π_σ designate this body of preferences at σ , we have the following definition for $[\cdot]$.

Definition 4.24. For all $\phi \in \mathcal{L}$, $\sigma[!\phi] = \langle I_\sigma, \pi_\sigma \cup \llbracket \phi \rrbracket \rangle$

The attendant connection between this parameter and action is much more sophisticated than in Lewis' treatment. A body of preferences induces an ordering on possibilities in the manner of Definitions 4.25-4.31 (Kratzer 1977, 1981, 1991b). Let π be a body of preferences. Then:

Definition 4.25. A possibility w is **at least as good as** a possibility v with respect to π (notation: $w \leq_{\pi} v$) iff $\{p \in \pi : v \in p\} \subseteq \{p \in \pi : w \in p\}$.

Definition 4.26. A possibility w is **strictly better than** a possibility v with respect to π (notation: $w <_{\pi} v$) iff $w \leq_{\pi} v \wedge v \not\leq_{\pi} w$.

Correct decision-making, on this way of thinking, is a matter of choosing actions that realize states of affairs that are not strictly worse than other available (given one's information) outcomes. One chooses incorrectly when one's action realizes a state of affairs that is not satisfied by any of the good enough possibilities compatible with one's information.

Definition 4.27. w is **good enough** with respect to σ , notation: $w \in \mathbf{min}(\cup I_{\sigma}, \leq_{\pi_{\sigma}})$, iff $\neg \exists v \in \cup I_{\sigma} : v <_{\pi_{\sigma}} w$.

Definition 4.28. ϕ is **required** at σ iff $\mathbf{min}(\cup I_{\sigma}, \leq_{\pi_{\sigma}}) \subseteq \llbracket \phi \rrbracket$

Definition 4.29. ϕ is **allowed** at σ iff $\mathbf{min}(\cup I_{\sigma}, \leq_{\pi_{\sigma}}) \cap \llbracket \phi \rrbracket \neq \emptyset$

As noted in Chapter II, making use of bodies of preferences (as opposed to simple permissibility spheres) allows the account to represent the phenomenon of prima facie directions—directions which normally, but defeasibly, determine what is preferred—as well as the phenomenon of inconsistent directions—directions that are inconsistent with some prior direction or sequence of directions. And that, I argued, is all to the good. Fortunately, all of these virtues can be retained without retaining the account of conventional directive force endorsed by Portner. Portner endorses a preference-addition account of the conventional force of imperatives, which, I argued in Chapter II, is inadequate with respect to the data.

Instead, I argued in Chapter III, we have reason to prefer a *necessitation* account of conventional directive force. In Section 4.3, I will develop a formal implementation of this account within the formalization of the Standard Account we have been pursuing here. That turns out to be rather involved. So, before getting to that, I will try to gesture at how the ideas developed in this section might be implemented in a type-driven compositional “semantics” for the relevant metalanguages.

4.2.4 A Compositional System

In this section, we will introduce a type-driven compositional system for *computing* the speech act expressed by an arbitrary speech act designator in \mathcal{S} —one that is fully consonant with the spirit of the Standard Account, as we have been developing it in this section. The

types I have in mind here correspond roughly, although not exactly, to the usual types in intensional Montague Grammar (Lewis 1970; Montague 1973). The key point of difference comes in distinguishing “static” types—set-theoretic constructions out of the ingredients supplied by a constant domain frames (properties, propositions, etc.)—from “dynamic” types—the types corresponding to cognitive state-representations and update operations on cognitive state-representations. Static types correspond to the types of the semantic values assigned by the static interpretation function $[[\cdot]]$ to pieces of syntax in the first-order base language. Dynamic types, however, correspond to types of the semantic values assigned by the dynamic interpretation function $[\cdot]$ to pieces of syntax in the illocutionary metalanguage (i.e., elements of \mathcal{O} and \mathcal{S}).

Standard Montagovian systems (e.g. Heim & Kratzer 1998; von Stechow & Heim 2007) avail themselves of neither dynamic types nor illocutionary notions, and they do not attempt to extend the tools developed by the Generative Semantic paradigm to the realm of speech acts.¹⁶ Although the system proposed here represents a departure from this tradition, it is a natural departure for the theorist interested in theorizing about linguistic meaning in terms of illocutionary notions. It is widely agreed that linguistic meanings are, in some sense, computed compositionally (i.e., the meaning of a string ϕ is a function of the meanings of its subconstituents, together with a rule-system specifying permissible ways of combining subconstituent meanings). And so the speech act theorist of meaning requires a compositional system for assigning speech acts to the regimented “logical forms” that her theory assigns to utterances in natural language.

This section represents a first, no doubt very simple, attempt to define such a system. Much of it, although quite formal, is nevertheless quite straightforward and will be presented without much in the way of comment. At the end of the section, I will step back from the formalism and attempt to survey what we’ve come up with, and what sort of picture of natural language syntax and semantics, if any, it is committed to.

Static Types Relative to a Frame $\mathcal{F} = \langle W, D \rangle$ and Assignment g .

e is a type, and $\{x : x \text{ is of type } e\} = D$
 t is a type, and $\{x : x \text{ is of type } t\} = \{0, 1\}$
 s is a type, and $\{x : x \text{ is of type } s\} = W$
 If τ and τ' are types, then $\langle \tau, \tau' \rangle$ is a type
 Nothing else is a type.

If τ is a type, we will use D_τ to refer to the set of objects of type τ . In general, if something is of type $\langle \tau, \tau' \rangle$, then it is a function from $\{x : x \text{ is of type } \tau\}$ into $\{x : x \text{ is of type } \tau'\}$.

16. There are exceptions, e.g., Groenendijk & Stokhof (1989), but such systems bear little similarity (except, perhaps, in inspiration) to the speech act-centric compositional system proposed here.

Static Types for Vocabulary Items of \mathcal{L} . We start with an ordinary description of semantic types for vocabulary items of \mathcal{L} .

If $k \in \mathcal{C} \cup \mathcal{V}$, $\llbracket k \rrbracket_{\mathcal{F},g,w}$ is of type e .

If P is an n -ary relation symbol, $\llbracket P \rrbracket_{\mathcal{F},g,w}$ is of type $\langle e^n, t \rangle$.¹⁷

$\llbracket \neg \rrbracket_{\mathcal{F},g,w}$ is of type $\langle t, t \rangle$, and $\llbracket \forall \rrbracket_{\mathcal{F},g,w}$ is of type $\langle t, \langle t, t \rangle \rangle$

If $Q \in \{\forall, \exists\}$, $\llbracket Q \rrbracket_{\mathcal{F},g,w}$ is of type $\langle \langle e, t \rangle, \langle \langle e, t \rangle, t \rangle \rangle$.

Comment. Note that, if $\llbracket \gamma \rrbracket_{\mathcal{F},g,w}$ is of type τ , then $\llbracket \gamma \rrbracket_{\mathcal{F},g}$ (a “ τ -concept”) is of type $\langle s, \tau \rangle$. So n -ary relations map n -length sequences of individuals to truth-values, while n -ary relational concepts map worlds to n -ary relations. Quantifiers express relations between sets, while quantifier concepts map worlds to relations between sets. And so on.

Dynamic Types Relative to a Frame $\mathcal{F} = \langle W, D \rangle$. If we were working within a static system, of course, the above types would be all that would be needed in order to state a compositional semantics for the relevant metalanguage. But, because we are interested in representing meanings in terms of speech acts, a compositional system for computing meanings (i.e., a compositional semantics for the illocutionary metalanguage \mathcal{S}) will require typing speech act operators (elements of \mathcal{O}) so that speech act operators may combine with things of propositional type to generate update functions. Within the rubric of the Standard Account, this turns out to be pretty straightforward, and I will move through the system quickly, keeping comment to a minimum.

Cognitive States. Since we have no referential devices for cognitive states in either \mathcal{L} or \mathcal{S} , we’re not strictly required to type them. So states will be typed as needed, depending on what kinds of speech acts we are considering. Representing information-addition (assertion) would require only that states be typed as sets of possibilities. But representing issue-introduction would require that states be typed as partitions of sets of possibilities (since issue-introduction involves making an existing partition more fine-grained). Etc.

A simplifying tool that will reduce notational clutter: we will assume that cognitive state-representations supply (indeed, are identical to) variable assignment functions; a cognitive state supplies not only values for the various relevant representations of cognitive parameters, but also determines the meaning the state associates with a variable.¹⁸ Update potentials, then, will update states (variable assignments) by *modifying* the value the state assigns to some variable. Given our focus here, the important variables will be variables for parameters of the cognitive states; an update function representing assertion, for instance,

17. Note on notation: $\langle e^1, t \rangle$ abbreviates $\langle e, t \rangle$, $\langle e^2, t \rangle$ abbreviates $\langle e, \langle e, t \rangle \rangle$, etc.

18. This is an idea from the logic of computing (so-called Dynamic Logic), which treats programs as updates on variable assignments (see, e.g., Harel et al. 2000). It also figures heavily in first-order Dynamic Semantic systems, e.g., Dynamic Predicate Logic (Groenendijk & Stokhof 1991) and Discourse Representation Theory (Kamp 1981).

will modify the value that the cognitive state assigns to a variable for that cognitive state's information. I will not introduce these variables explicitly, preferring instead to work with the intuitive idea that a cognitive state-representation is just a list of representations of relevant cognitive parameters. The main purpose of this idea here is just to allow us to drop reference to variable assignment functions for reference simply to cognitive states in what follows.

Speech Act Operators. Speech act operators, on the Standard Account, combine with propositional radicals to yield speech act designators. Which is to say that they express functions from propositions into speech acts—or, on our analysis, functions from propositions and cognitive states into updated cognitive states. Letting c be the type of cognitive states, we thus have:

If $\Delta \in \mathcal{O}$, $[\Delta]_{\mathcal{F}}$ is of type $\langle c, \langle \langle s, t \rangle, c \rangle \rangle$ and $[\Delta]_{\mathcal{F}, \sigma}$ is of type $\langle \langle s, t \rangle, c \rangle$

Combinatory rules. In the Montagovian tradition, computation of semantic values for complex syntactic structures is type-driven. Specifically, it is driven by function-argument structure. Following Heim & Kratzer (1998); von Stechow & Heim (2007), we introduce two rules for computing semantic values for complex syntactic structures. We will use binary-branching, tree-like syntactic structures to represent the relevant syntactic relations, but nothing of any significance turns on this.

FUNCTIONAL APPLICATION (FA)

Given a branching node α with daughters β and γ , if $[\![\beta]\!]_{\mathcal{F}, \sigma, w}$ is a function whose domain includes $[\![\gamma]\!]_{\mathcal{F}, \sigma, w}$, then $[\![\alpha]\!]_{\mathcal{F}, \sigma, w} = [\![\beta]\!]_{\mathcal{F}, \sigma, w}([\![\gamma]\!]_{\mathcal{F}, \sigma, w})$.

$$\left[\begin{array}{c} \alpha \\ \beta \quad \gamma \end{array} \right]_{\mathcal{F}, \sigma, w} = \begin{cases} [\![\beta]\!]_{\mathcal{F}, \sigma, w}([\![\gamma]\!]_{\mathcal{F}, \sigma, w}), & \text{if } [\![\gamma]\!]_{\mathcal{F}, \sigma, w} \in \text{dom}([\![\beta]\!]_{\mathcal{F}, \sigma, w}) \\ \text{undefined} & \text{otherwise} \end{cases}$$


DYNAMIC FUNCTIONAL APPLICATION (DFA)

Given a branching node α with daughters β and γ , if $[\beta]_{\mathcal{F}, \sigma}$ is a function whose domain includes $[\![\gamma]\!]_{\mathcal{F}, \sigma}$, then $[\alpha]_{\mathcal{F}, \sigma} = [\beta]_{\mathcal{F}, \sigma}([\![\gamma]\!]_{\mathcal{F}, \sigma})$.

$$\left[\begin{array}{c} \alpha \\ \beta \quad \gamma \end{array} \right]_{\mathcal{F}, \sigma} = \begin{cases} [\beta]_{\mathcal{F}, \sigma}([\![\gamma]\!]_{\mathcal{F}, \sigma}), & \text{if } [\![\gamma]\!]_{\mathcal{F}, \sigma} \in \text{dom}([\beta]_{\mathcal{F}, \sigma}) \\ \text{undefined} & \text{otherwise} \end{cases}$$

Dynamic Functional Application. DFA is used for cases where one sister node—in particular, a sister node containing a speech act operator—is “looking to combine” with a function from worlds to extensions, rather than an extension at the world of evaluation (cf. the Intensional FA rule of von Stechow & Heim 2007: 8–9). Given a branching node α with daughters $\beta \in \mathcal{O}$ and $\gamma \in \mathcal{L}$, if β 's interpretation is of type $\langle s, \tau \rangle$ and γ 's extension of τ , we allow β and γ to combine by DFA, even though regular FA would be undefined for them.

Deriving Meanings of Appropriate Type. The system proposed above is designed to extend compositional “semantic” (or quasi-semantic) interpretation into the realm of speech acts and illocutionary force. It is easy to see that it does just that. Consider an SAD of the form $\lceil !walk(a) \rceil$, with the Standard Account-friendly syntactic representation below.



The computation of the semantic value for this structure proceeds as follows (reverting to our less explicit superscript notation for the sake of simplicity). We begin with a toy set of lexical entries:

- $[!]^\sigma = \lambda p \in D_{\langle s,t \rangle} . \sigma'$
- $\llbracket walk \rrbracket^{\sigma,w} = \lambda x \in D_e . x \text{ walks in } w$
- $\llbracket a \rrbracket^{\sigma,w} = a$

Notice that, simply by making use of the lexical entries, we know that:

- $[!]^\sigma \in D_{\langle \langle s,t \rangle, c \rangle}$
- $\llbracket a \rrbracket^{\sigma,w} \in \text{dom}(\llbracket walk \rrbracket^{\sigma,w})$, and $\llbracket walk \rrbracket^{\sigma,w}(\llbracket a \rrbracket^{\sigma,w}) \in D_t$
- $\lambda \tau w . \llbracket walk \rrbracket^{\sigma,w}(\llbracket a \rrbracket^{\sigma,w}) = 1 \in \text{dom}([!]^\sigma)$

The rest of the derivation is trivial.

- i. $\llbracket walk(a) \rrbracket^{\sigma,w} = \llbracket walk \rrbracket^{\sigma,w}(\llbracket a \rrbracket^{\sigma,w})$ (FA)
- ii. $\lceil !walk(a) \rceil^\sigma = [!]^\sigma(\llbracket walk(a) \rrbracket^\sigma)$ (DFA)
- iii. $= [!]^\sigma(\lambda \tau w . \llbracket walk \rrbracket^{\sigma,w}(\llbracket a \rrbracket^{\sigma,w}) = 1)$

Comment 1. It would be easy to prove, by induction on complexity of sentences of \mathcal{L} , that, for any $\phi \in \mathcal{L}$, $\Delta \in \mathcal{O}$, and σ , $[\Delta\phi]^\sigma$ is defined and of type c . I will omit that proof here.

Comment 2. Once a lexical entry for an arbitrary speech act operator $\Delta \in \mathcal{O}$ is determined, the speech act expressed by a speech act designator of the form $\Delta\phi$ is derived simply by:

- Selecting an arbitrary cognitive state-representation σ and computing $[\Delta\phi]^\sigma$
- λ -abstracting over the cognitive state-representation parameter. That is to say, $[\Delta\phi] = \lambda \sigma . [\Delta\phi]^\sigma = \lambda \sigma . \sigma[\Delta\phi]$

Since, given a suitable lexical entry for Δ , the result of this procedure is generally trivial, this means that most of the effort in developing a substantive speech act theory of meaning for imperative clauses will go into describing, precisely, *what a suitable lexical entry for '!' might look like*. And that, indeed, is where we will place our focus in the next major chunk of this essay: a substantive theory of interpretation for the clause-typing device '!'.

On Our Commitments to Natural Language Syntax and Semantics. Although we have treated speech acts as computable from regimented logical forms by the use of compositional semantic rules—as, I suggested, is the *sine qua non* of any speech act theory of conventional linguistic meaning—our proposal does not incur any real commitments regarding natural language syntax or semantics.

How can that be? We are, after all, making use of familiar binary-branching syntactic trees, and assuming that speech act operators “adjoin” to clausal complements in such trees. I have suggested, moreover, that these trees represent *regimented logical forms for typed clauses in natural language*. And so it would seem to follow, on a fairly standard understanding of the notion of *logical form*, that we recognize the presence of speech act operators at some level of syntactic or semantic representation for the relevant typed clauses. (This might, in turn, be thought to pose a problem for our account, since we have, by and large, been utterly disinterested in the sorts of arguments by which it might be established that speech act operators are syntactically realized.)

The reasoning here rests on a contested understanding on the notion of logical form—one that is optional, but not required, for an account of the sort developed in this section. A clause’s logical form, in the most general sense, is just a representation of its *meaning* in a regimented metalanguage. That is the sense in which mapping a typed clause to a speech act designator gives us a logical form for that clause (at least from the standpoint of a speech act theory of meaning). While a clause’s logical form does, I think, *supervene on* its syntactic form, supervenience in no way requires that syntactic devices used in the metalanguage in which logical forms are stated have syntactic correlates in natural language; it merely requires that, whatever the clause’s correct syntactic representation, that representation contain enough information to determine a type for that clause (and hence a “logical form” for that clause) (cf. our discussion of clause type in Section 4.1.1).

In particular, then, the proposal developed here is compatible with a range of views about the syntax and semantics of sentence mood, clause type, and related notions. For all we have said (and for all we will say), imperative clauses in natural language may contain indicators of their type (or even speech act operators!) that adjoin, in their true syntactic representations, to an un-typed sentence (cf. Han 1998). Or they may contain no such exotic syntactic material at all (cf. Portner 2004a). It is, I think, a real virtue of our theoretical framework that it does not require us to decide this question *a priori*.

4.3 Direction as Necessitation (as Coming to Recognize a Requirement)

What is a necessitation act? Within the Standard Account, necessitation acts must be assumed, as all speech acts are assumed, to be of the form $\Delta\phi$. The most natural idea for the force of such an act is that it consists in making it the case that ϕ must be realized by that state. And, within the update-theoretic account we have been pursuing here, this will mean that $[\!|\phi|]$ takes a cognitive state σ and makes it the case that ϕ must be realized with respect to the updated cognitive state. In other words, in a standard case involving update on a cognitive state σ with $!\phi$, we will have:

- It is not the case that σ must make it the case that ϕ
- It is the case that $\sigma[\!|\phi|]$ must make it the case that ϕ

I will have a lot to say in the remainder of this section on how best to interpret these two components of necessitation acts. The first item, I will suggest, can happily be represented as a presupposition, i.e., as a definedness condition on $[\!|\phi|]$. (Fortuitously, this will give us a chance to critique various other proposals regarding the presuppositional content of imperatives.) The second item can be represented as the inducement of a state to satisfy a property I will gloss here as *recognition of a requirement*. This notion of requirement-recognition can, I'll show, be fruitfully applied to several of the phenomena we described in Chapter II—for instance, the task of distinguishing strong permission from merely implicit permission, as well as the parallel task of distinguishing a notion I'll term *strong obligation* from merely implicit obligation. The upshot of all this will be a model in which the action-guiding parameter of a cognitive state has a structure similar, in many ways, to the structure of the information-tracking parameter. Not only does the action-guiding parameter determine a *ranking* of outcomes (and hence a set of admissible outcomes); it also for various purposes *designates* outcomes.

4.3.1 Imperative Presuppositions

There are a couple of proposals about the presuppositions (or, more generally, preconditions for interpretability) of imperatives that have been entertained in the literature. [Portner \(2010\)](#), as we saw in Section 2.4.5, tacitly assumes that permission imperatives presuppose the presence of a standing, contrary prohibition. Positing this sort of presupposition, we argued, was implausible. By way of introducing our own proposal about imperative presuppositions, we will first argue against an incompatible proposal, on which imperatives presuppose that the addressee is an authority, in some sense, about what the relevant body of preferences requires.

Presupposing authority. [Schwager \(2006\)](#); [Kaufmann & Schwager \(to appear\)](#) suggest that the characteristic (what they call) “performative” interpretation of imperative utter-

ances (what I've been neutrally calling the "typical" interpretation of such utterances) is *explained* by the fact that these utterances conventionally presuppose that the speaker of the imperative is an authority, in a very special sense, about what is required by the relevant body of preferences.¹⁹ The ordinary "performative" interpretation of imperatives requires explanation, on their view, because imperatives semantically express *modal propositions*. More specifically, on the **Modal Analysis of Imperatives** (hereafter MAI), an imperative of the form $!\phi$ expresses that ϕ is necessary with respect to a salient body of preferences \mathcal{P} —something that holds iff all of the \mathcal{P} -best worlds compatible with the relevant information are ϕ -worlds. However, imperatives cannot generally be interpreted as proffering this proposition, i.e., as asserting that ϕ is required (or, indeed, as asserting anything at all).

(4.5) A: Go to your room!
 #B: False!

According to Kaufmann and Schwager, what explains this fact are the special presuppositions of imperative clauses—presuppositions that together restrict the range of contexts in which imperatives can be interpreted to those in which the imperative must have a "performative" interpretation, in view of the presupposed authority of the speaker.

What exactly are these presuppositions? As mentioned, they together are meant to ensure that the speaker of the context is an authority about what the relevant body of preferences \mathcal{P} requires. Specifically, an imperative of the form $!\phi$ presupposes that:

- The relevant body of preferences has to do with *ranking* or *prioritizing actions* (cf. Portner 2007). It does not represent, e.g., considerations of epistemic value. Rather, it represents considerations that bear on the value of actions available to the relevant agent. Using the standard linguistic jargon, it must be a deontic, bouletic, or teleological ordering source.
- The speaker has "perfect knowledge" about what is required by the relevant body of preferences at the context. According to Schwager (2006), the speaker "utters a necessity proposition he cannot be mistaken about." (Schwager tries to unpack this formally, but this is *formale gratia formalis*—it adds nothing to the intuitive description and seems, if anything, to misrepresent its content. I will ignore it here.)
- The speaker takes ϕ to be possible with respect to the relevant information.
- Finally, the body of preferences characterized by the speaker's wishes does not require $\neg\phi$, i.e., ϕ is compatible with what the speaker wants.

This proposal is mysterious. First, no evidence for the existence of such presuppositions is adduced by Kaufmann and Schwager. To be clear, I do agree that an imperative $!\phi$

19. The view expressed in Schwager (2006) is partly retracted in Kaufmann & Schwager (to appear). I will only discuss the more recent view here.

plausibly does presuppose that ϕ is in some sense possible.

(4.6) #Count all of the natural numbers!

#There are no unicorns. Bring me a unicorn! (Starr 2011)

(Infelicity when the context does not allow the presupposition to be accommodated)

(4.7) If you're able to, please pass me that 50-pound drum.

(Local accommodation; cf. Heim 1983)

(4.8) A: Please pass me that 50-pound drum.

B: Hey wait a minute! I can't lift 50 pounds!

("Hey Wait a Minute!" Test; cf. von Stechow 2004)

And plausibly imperatives do bear on the ranking or prioritization of *actions* (as opposed to, e.g., states of mind not under the voluntary control of one's interlocutor).

(4.9) #Be hungry!

But I see no reason whatever to think that either:

- i. That a speaker generally has perfect knowledge about what *is* required, at the time of utterance, by the relevant body of preferences
- ii. That, even if true, such knowledge would in any way explain the performative interpretation of imperatives.

With respect to (i), a speaker can felicitously use an imperative of the form $!\phi$ even when she knows (or is uncertain) that the relevant body of preferences does *not* require ϕ . Consider, for instance, a person with the ability to dictate law. Each of the following utterances is available to such a person.

(4.10) I've discovered that the law does not forbid jaywalking. So, don't jaywalk.

(4.11) The law is complicated and I have no idea whether it forbids jaywalking. To remove any doubt, please don't jaywalk.

Indeed, if an imperative of the form $!\phi$ does, as I'm going to argue, presuppose that the given body of preferences *does not recognize* a requirement that ϕ , such uses, far from being anomalous as predicted by Kaufmann and Schwager, will comprise the "core" class of uses for imperatives.

This leads to point (ii). If the speaker of an imperative $!\phi$ is not mistaken about ϕ being required by the relevant body of preferences, then *there is no reason* for the speaker to actually utter the imperative, except to ensure that this already-existing requirement comes to be recognized by her interlocutor(s). But, of course, speakers can rationally utter imperatives even when this does not accurately describe their reason for making the utterance. Speakers frequently use imperatives to *establish new requirements*—requirements that they know not

to be in force, at the context of utterance, relative to the relevant body of preferences.²⁰

Could a modified, non-epistemic, understanding of the “authority” requirement do the trick? Instead of construing the speaker as an authority *about* what the relevant body of preferences does require, we might construe the speaker as having the authority *to modify* the relevant body of preferences. There is, I agree, ample linguistic evidence that imperatives, on their typical interpretations, carry this sort of authority presupposition.

(4.12) If I’m not overstepping my authority, please leave the premises immediately!

(4.13) A: Go to your room!

B: Hey, wait a minute! That’s not up to you! You’re not my mom!

But recognizing such a presupposition (as I am happy to do) is not exactly friendly to the MAI. It does not explain—nor is it clear how it *could* explain—why uttering a necessity proposition would force a *performative* interpretation of that proposition—i.e., an interpretation on which the very utterance of the proposition would suffice to make the proposition true. Indeed, why would uttering a necessity proposition have any effect on whether the necessity proposition was true, unless the utterance itself involved a *proposal to make the proposition necessary*, relative to the relevant body of preferences. Positing such a proposal is, so far as I can tell, the only way to explain the typical effect of an imperative utterance—i.e., the creation of a new requirement—in the presence of the presupposition that the speaker has the authority to modify the relevant body of preferences. But the existence of such a proposal, as part of the conventionalized interpretation of an imperative utterance, is exactly what the necessitation account I am defending here posits. So it looks like the MAI must embrace some version of the necessitation account. And, moreover, it looks like the necessitation account, together with the authority-to-necessitate presupposition, explains all there is to explain about the typical force of imperative utterances. But then it is unclear what explanatory work the *modal* part of the modal analysis is actually doing at all.²¹

20. I made this point in my (2010a), and I am unaware of any response to it from either Kaufmann or Schwager. I suspect they would say something along the following lines. The speaker needs to be certain about the body of preferences—typically *her* preferences, or, more generally, whatever body of preferences prompts her issuance of the relevant imperative—to which the body of preferences she aims to modify with her utterance responds. The latter body of preferences subsequently adjusts in order to make the speaker’s utterance true, in much the same way as Lewis (1979a) (which also endorses a modal semantics for imperatives) has the permissibility sphere adjusting, in deference to MASTER, to make MASTER’s utterance true. Two points in reply. First, this has no predictive advantage over the much simpler necessitation account (indeed, it amounts to conventionalization, via presupposition, of necessitation acts). Second, it is nothing like the account that Kaufmann and Schwager actually give. They commit themselves to the claim that the body of preferences with respect to which the speaker is an authority is the *very same* body of preferences with respect to which the imperative is interpreted. That claim cannot be true, given the data recorded in this section.

21. For similar thoughts, see Portner (2007). The MAI is claimed, by Kaufmann and Schwager, to yield intuitively correct logical forms for a variety of constructions—conditional imperatives, in particular—for which speech act accounts are claimed to lack any sort of analysis. This claim is false. As we’ll argue in Section 4.5, it assumes the Standard Account of speech acts.

Presupposing efficacy. Instead, I want to suggest that imperatives (like, in fact, all clause-types conventionally associated with some speech act or other) presuppose that the context will not prevent the speaker's utterance from being *effective* in realizing the speaker's communicative intentions/goals. They presuppose that the speaker is *rational* in proffering the relevant speech act at the context of utterance (since if there were no expectation that the context would admit the speech act, the speaker would have no reason to perform the speech act).

The goal of an assertion, for example, is to induce a belief in the addressee. And there is linguistic evidence to suggest that assertions presuppose that the context will not prevent an assertion from inducing a corresponding belief in the addressee. A speaker generally cannot felicitously make an assertion unless she clears contextual roadblocks to the addressee accepting what she asserts. And the positing of such a presupposition seems to be supported by the standard presupposition-detecting tests.

(4.14) A: I have no reason to believe anything you say.

#B: There's a fire in the foyer.

[B: *Please*, listen to me! There is a fire in the foyer.]

(4.15) If you will just listen to me, there is a fire in the foyer!

(4.16) A: There is a fire in the foyer!

B: Hey wait a minute! Why would I have any reason to believe you?

Imperatives, I have been arguing, are conventionally associated with necessitation acts: proposals to make their prejacent required, relative to some relevant body of preferences. So if there is in general a presupposition that the context will not prevent the realization of the intention associated with a speech act, we will expect imperatives typically to presuppose that the context will not prevent an imperative utterance of the form ! ϕ from necessitating that ϕ , relative to a salient body of preferences.

There a variety of ways that the intention conventionally associated with an imperative utterance could be thwarted by a context. First, the speaker might not have the authority to modify the relevant body of preferences (by way of necessitating the imperative's prejacent). So imperatives ought to presuppose that the speaker has such authority. As examples (4.12) and (4.13) illustrate, that prediction seems to be on target. Second, the intention might already be realized by the context. In general, an intention that ϕ is an attitude to *alter* the world from a state which realizes $\neg\phi$ to one which realizes ϕ .²² So the intention that ϕ can be thwarted—in this case, thwarted in the sense of constitutively involving a misrepresentation of the state of the world—if ϕ is already the case. Linguistic evidence attests to the existence

22. There is a *dispositional* sense of intention, on which one can intend to ϕ even when one thinks ϕ is the case. This sense involves being disposed to see to it that ϕ if ever ϕ should cease to be the case. But that, pretty clearly, is not the sort of intention that is involved in executing a speech act.

of such a presupposition.

(4.17) A: Hold on a minute. I need to shut the window.

#B: Shut the window!

(4.18) If Dad didn't already ask you to shut the window, please do so.

(4.19) A: Shut the window!

B: Ummm, Dad already asked me to that. Jeez!

This, I think it's fair to say, constitutes further evidence for the necessitation account. The necessitation account, together with an independently plausible view about the presupposition generally involved in the performance of any kind of speech act (namely, an efficacy presupposition) predicts an array of presuppositions that will be borne, in typical cases, by imperative utterances. All of these predictions seem to be borne out by standard linguistic tests for detecting presuppositions.

4.3.2 Coming to Recognize a Requirement

The other component of the necessitation act conventionally expressed by an imperative of the form $!\phi$, I suggested, was that an agent whose cognitive state is representable as $\sigma[!\phi]$ must make it the case that ϕ . To express this more precisely, it is useful (i) to define a modal language containing the modal operator *must*,²³ (ii) to state the conditions under which a cognitive state "satisfies" a sentence of the form $must(\phi)$, (iii) to state the update potential expressed by a sentence of the form $must(\phi)$. We will assume, perhaps slightly idiosyncratically, an equivalence between the following notions:

- A cognitive state σ **satisfies** $must(\phi)$, notation: $\sigma \models must(\phi)$
- The result of updating σ with $must(\phi)$ is just σ , i.e., $\sigma[must(\phi)] = \sigma$
- A cognitive state σ **recognizes that ϕ is required**

The first equivalence here is mostly stipulative (since "satisfaction" relative to a cognitive state is a technical notion): for σ to satisfy $must(\phi)$, $must(\phi)$ must not express any "cognitive advice" not already heeded by the state. (Nevertheless, this is how "satisfaction" in Dynamic Semantics is typically understood; see esp. [Veltman 1996](#).) The second equivalence is plausible: if $must(\phi)$ does not express any "cognitive advice" not already heeded by the state, then the state recognizes that ϕ is required, and vice versa. The main idea we will

23. This modal operator's name is chosen deliberately: there are close correspondences between the technical notion I will be exploring here and the sort of concept expressed by strong necessity modals like *must* in natural language. That said, there are sure to be differences. I agree fully with [Ninan \(2005\)](#); [Portner \(2007\)](#) that strong necessity modals have broadly the same sort of conventionalized performative force as imperatives (see Section 2.4.4). So if I am right in thinking that imperatives express necessitation acts, strong necessity modals will do so too. Some of the properties of necessitation acts will not, however, be reflected in the interpretation for the modal operator *must* I will be giving here. So in that respect our treatment here does not represent—nor is it meant to represent—any sort of final analysis of strong necessity constructions in natural language.

be developing in this section, then, is the following: when defined, $\sigma[!\phi]$ is a state that recognizes that ϕ is required, i.e., $\sigma[must(\phi)] = \sigma$.

What are the conditions under which a state recognizes a requirement that ϕ ? First, the state must in fact require that ϕ . And, second, the state must recognize this requirement. Just as a state σ does not believe ϕ when ϕ merely happens to be entailed by the possibilities compatible with what σ knows (cf. Yalcin 2007), σ does not recognize a requirement that ϕ when ϕ merely happens to be true throughout $\min(\cup I_\sigma, \leq \pi_\sigma)$. So, when defined, $\sigma[!\phi]$ will both (i) require that ϕ , (ii) recognize a requirement that ϕ . Defining the update potential for an imperative of the form $!\phi$ will therefore require accounts of:

- i. **Requirement Institution.** How, when σ does not require ϕ , σ is altered by $[\phi]$ so that it comes to require ϕ . (Notice that some states for which $[\phi]$ is defined may already require ϕ . In that case they will have to fail to recognize this requirement.)
- ii. **Recognition Institution.** How σ is altered so that it comes to recognize a requirement that ϕ (Notice that for no state for which $[\phi]$ is defined can that state already recognize a requirement that ϕ , in view of our understanding of the presuppositions of imperatives.)

We provide these in turn.

4.3.3 Instituting Requirements

For a state σ to require ϕ , it must be the case that all the good enough possibilities at σ satisfy ϕ , i.e., $\min(\cup I_\sigma, \leq \pi_\sigma) \subseteq \llbracket \phi \rrbracket$. Equivalently, but perhaps more illuminatingly, requiring ϕ consists in strictly preferring the state of affairs characterized by ϕ to the state of affairs characterized by its negation. Here are some definitions to illuminate this second gloss.²⁴

Definition 4.30. A state of affairs p is **at least as good as** a state of affairs q with respect to π (notation: $p \preceq_\pi q$) iff $\neg \exists w \in p : \exists v \in q : v <_\pi w$.

Definition 4.31. A state of affairs p is **strictly better than** a state of affairs q with respect to π (notation: $p \prec_\pi q$) iff $p \preceq_\pi q \wedge q \not\preceq_\pi p$.

Definition 4.32. A state of affairs p is **at least as good as** a state of affairs q relative to σ (notation: $p \preceq_\sigma q$) iff $\neg \exists w \in p \cap \cup I_\sigma : \exists v \in q \cap \cup I_\sigma : v <_\pi w$.

Definition 4.33. A state of affairs p is **strictly preferred to** a state of affairs q relative to σ (notation: $p \prec_\sigma q$) iff $p \preceq_\sigma q \wedge q \not\preceq_\sigma p$.

24. In that sense, the account of Starr (2011), which has an imperative $!\phi$ introducing a preference for the state of affairs that ϕ over the state of affairs that $\neg\phi$, also has imperatives instituting requirements. He does not, though, have them instituting recognition of requirements (which means his account is incomplete with respect to the data I'll go on to outline below). (Our accounts were, for the record, developed independently. I will discuss some of his proposals further below.)

These definitions are, with the possible exception of Definition 4.32, straightforward. What Definition 4.32 says is that a state of affairs p is at least as good as some other state of affairs q relative to a cognitive state σ iff there are no p -possibilities live at σ strictly worse than some q -possibility live at σ .

That in mind, I would like to suggest the following informal gloss on the proposal by which a requirement that ϕ is instituted in a cognitive state:

MODIFY YOUR PREFERENCES

Modify your preferences in the most conservative way so that ϕ is strictly preferred to $\neg\phi$ relative to your cognitive state.

There are several things to note about this gloss. First, notice that although requirements could always be instituted by modifying one's *information*, that is not what the imperative advises; instead, it recommends instituting a requirement by way of *preference* modification. That, I submit, is a brute fact (or at least a fact for which I don't have any explanation) about necessitation acts; necessitation acts institute requirements, not by way of information-modification, but rather by way of preference-modification.²⁵ If there is no reasonably sane preference modification by which the relevant requirement can be instituted, then the necessitation act must be rejected by the state. That explains why, e.g., (4.20) is unacceptable by one who values human life absolutely.

(4.20) Kill that prisoner at once!

The addressee *would* achieve a state in which the relevant requirement was instituted, were she to come to believe that she kills the prisoner (without actually doing so). For a variety of reasons, of course, this sort of *wishful thinking* response is never (or nearly never) available to the addressee. But, and here is the key point, even if it were available, wishful thinking would obviously not be a response that is consonant with the speaker's goals in uttering (4.20). Necessitation acts are proposals for one's interlocutors to modify their preferences, rather than to modify either their preferences or their beliefs, as they wish.

Second, the preference modification is required to be *conservative* with respect to one's pre-cognitive preferences. There are, in general, a variety of modified preference states that will generate the relevant requirement, relative to one's information. Generally, however, shifting to a radically modified preference state is not going to comport well with the speaker's communicative intentions. A speaker obviously has their communicative intentions *thwarted*, rather than realized, if her addressee forms the relevant preference in a

25. Surprisingly, Schwager (2006) seems to favor the opposite position. She writes that "an utterance of [an imperative] in the given context *constrains the Common Ground* by leaving in only worlds where the speaker has the respective preferences," i.e., preferences according to which the prejacent is required. In other words, an utterance of $!\phi$ amounts to an assertion that the speaker prefers that ϕ . And this assertion is somehow (it is never actually made clear how) supposed to generate the performative interpretation of the utterance. I hope, at this point in the essay, that my reasons for rejecting this sort of view are clear.

bizarre and unanticipated way.

(4.21) A: Please bring me a soda.

[B responds by adding a preference for bringing A a soda, while abandoning her preference for achieving her ends legally. B brings A a soda and immediately lets her know that she knocked over a convenience store to get it.]

A: What! Obviously that's not what I was trying to get you to do!

4.3.4 Conservative Preference Update

What are the conditions on the relevant preference modification being conservative? Here is one thing we can say: in general, the most conservative preference modification simply involves adding the imperative's prejacent to the relevant body of preferences, in the manner suggested by Portner (2004a, 2007, 2010). This will not *always* result in a body of preferences which requires the prejacent, but, in a large class of cases, it will. Namely, as we proved in Section 2.3, adding $\llbracket\phi\rrbracket$ to the relevant body of preferences π_σ will yield a state in which ϕ is required when, for some σ -live ϕ -possibility w and some σ -live $\neg\phi$ -possibility v :

- w and v are *comparable* with respect to \leq_{π_σ}
 $w \leq_{\pi_\sigma} v$ or $v \leq_{\pi_\sigma} w$
- Both w and v are *good enough* with respect to σ , i.e., both ϕ and $\neg\phi$ are permitted with respect to σ
 $w \in \mathbf{min}(\cup I_{\sigma, \leq_{\pi_\sigma}})$ and $v \in \mathbf{min}(\cup I_{\sigma, \leq_{\pi_\sigma}})$

That is to say, when these conditions are satisfied, and when $\sigma[!\phi]$ is defined, it will follow that adding $\llbracket\phi\rrbracket$ to σ 's body of preferences will yield a state in which ϕ is required, i.e., all the good enough possibilities, with respect to that state, are ϕ -possibilities:

$$\mathbf{min}(\cup I_{\sigma, \leq_{\pi_\sigma \cup \{\llbracket\phi\rrbracket\}}}) \subseteq \llbracket\phi\rrbracket$$

Such cases do *not*, however, exhaust the range of cases in which an imperative utterance can succeed in realizing its conventional communicative intention (i.e., necessitating its prejacent). As we noted in Section 2.3, addition of $\llbracket\phi\rrbracket$ to a body of preferences can fail to institute a requirement that ϕ when, in $\mathbf{min}(\cup I_{\sigma, \leq_{\pi_\sigma}})$, there are:

- i. No ϕ and $\neg\phi$ -possibilities comparable with respect to \leq_{π_σ}
- ii. No ϕ -worlds (ϕ is antecedently forbidden)

I argued in Section 2.3 that the failure of adding $\llbracket\phi\rrbracket$ to one's preferences to make it required that ϕ was not necessarily a problem for the Portner account, in view of the idealizations built into Kratzer's procedure for generating an ordering on possibilities from a body of preferences.

I also tentatively defended the predictions of Portner’s account in case (ii). My grounds for defending Portner there were that there are a range of cases in which (ii) holds, but in which predicting that the imperative utterance instituted requirements seems to be incorrect. For instance, there are cases like (2.1), repeated here as (4.22), where a contrary imperative can be used to reverse a standing prohibition, by way of establishing a permission.

(4.22) Monday carry rocks! Tuesday carry rocks! And Wednesday carry rocks! [. . . Tuesday comes along.] Take tomorrow off!

Far from predicting that (4.22) institutes a requirement to take the next day off, an account ought to predict that it institutes a permission to do so. And Portner’s (2010) account—on which the final imperative adds the proposition that the addressee takes the next day off to her preferences (and thereby makes her preferences indifferent between carrying rocks or taking the day off on Wednesday)—has them doing precisely that.

Nevertheless, while Portner does have an explanation of the permission interpretation of the final imperative in (4.22), I do not think he has identified the *correct* explanation for it—or, for that matter, for permission imperatives in general (see Sections 2.4.5 and 3.4.4). I, of course, would explain the permission interpretation as an ISA (see Section 3.4.4), rather than (as Portner would explain it) a result of assigning the imperative its conventional, normal interpretation.

Nor do I think Portner has the right account of the typical effect of imperatives that are uttered in the presence of standing, contrary prohibitions. According to Portner, such imperatives *generally* have a permission interpretation: when ϕ is antecedently prohibited, updating with $!\phi$ will generally make one **indifferent** between ϕ and $\neg\phi$. Instead, I want to suggest, that imperatives, even when uttered in the presence of a standing, contrary prohibition, will typically have their default, requirement-instituting interpretations. Authorities often *change their minds*. Different agents often have authority over the content of a single body of preferences, and these agents *may not always agree*. Consider the sequence in (4.22). It is easy to imagine the final imperative being used to institute a requirement for the addressee to take the next day off.

- **Change of mind.** Suppose it is known that MASTER has realized on Tuesday morning that there is no need for SLAVE’s services on Wednesday (and would like to avoid paying overtime). Then the final imperative proposes to make it required that SLAVE take the next day off (while nullifying the earlier command to carry rocks on Wednesday).
- **Conflicting authorities.** Imagine two forepersons, each with equal authority over setting SLAVE’s work schedule, one of whom would like SLAVE to carry rocks on Wednesday, the other of whom would like SLAVE to take Wednesday off. If the final imperative is uttered by the latter foreperson, it will normally receive a requirement-instituting interpretation.

To be clear, in neither case is it necessarily obvious what SLAVE ought to do in reaction to these utterances. In reaction to a change of mind or conflicting authorities, SLAVE can respond with (4.23) and (4.24), respectively.

(4.23) Hold on! Earlier you told me I had to carry rocks on Wednesday!

(4.24) Hold on! The other boss told me I had to carry rocks on Wednesday!

But I think it is fairly clear that, for SLAVE to **accept** the imperative, on its intended interpretation, SLAVE must in each case come to recognize a requirement that she take Wednesday off. Portner's account says, incorrectly, that accepting the imperative in these cases involves instituting indifference between taking Wednesday off and carrying rocks on Wednesday.

What does this have to do with conservative preference modification? Suppose that the addressee does accept an imperative $!\phi$ that conflicts with a prior accepted imperative $!\neg\phi$. Then the addressee will come to recognize a requirement that ϕ . But this change cannot be effected simply by adding $[\phi]$ to the relevant body of preferences: this will tend to bring about indifference between ϕ -possibilities and $\neg\phi$ -possibilities, and that is not what accepting the imperative involves. So the relevant preference modification in these cases *cannot be simple preference-addition*. What will the relevant modification look like? It will, of course, be a modification that:

- Yields a state recognizing a requirement that ϕ
- Is otherwise conservative with respect to the agent's prior preferences

But this does not tell us what sort of cognitive mechanics are involved in revising one's prior preferences in a way that blunts the standing prohibition on ϕ while preserving as much of the content of that body of preferences as possible.

How Much Should a Theory of Speech Acts Have to Say? I will take a stab at describing an update function that builds in a substantive account of conservative preference update. But, before I do that, I should register my doubt that this sort of account should be part of a theory of speech acts at all (or, therefore, part of a speech act theory of linguistic meaning).

A theory of speech acts, as I have been understanding it, is a theory about what sort of cognitive instruction is proffered by an utterance in natural language. But cognitive instructions, as I understand them, are relatively *thin* things. An assertion that ϕ can be accurately described as a cognitive instruction whose force is to propose to get the addressee to accept that ϕ . When ϕ conflicts with an agent's prior information, then, in order to fully accept the assertion, the agent cannot simply add $[\phi]$ to her stock of beliefs. She must revise her prior beliefs to eliminate the source of conflict. And, of course, not just any old revision will do; the revision must be conservative with respect to those priors. How rational agents revise priors conservatively is the topic of a vast amount of work in formal and

informal *epistemology*, rather than linguistics.²⁶ That seems to me like the correct division of theoretical labor. The job of speech act theory is to associate sentences with general, relatively skeletal, cognitive instructions. It is the job of epistemologists (and theorists interested in rational changes in attitudes, more generally) to give a substantive account of how rational agents can comply with these instructions.

In slogan form: speech act theory furnishes a theory of cognitive directives. Epistemology furnishes a substantive theory of diachronic compliance for cognitive directives. Nevertheless, I'll try to do some of the epistemological work here (if only so that we can represent speech acts as functions from cognitive states into cognitive states, rather than as functions from cognitive states *and* a substantive theory of rationality into cognitive states).

The Mechanics of Preference Revision. Here's one way to approach the question of how to associate the speech act of direction with an update *function*. Adding a proposition $\llbracket\phi\rrbracket$ to a body of preferences does not always suffice to make it required that ϕ . In such cases, is it possible to isolate a conflicting preference (or set of preferences that together generate a conflict), and revise them accordingly?

Here is a simple case. The prior body of preferences is given by $\pi_\sigma = \{\llbracket\neg\phi\rrbracket\}$. To fully accept the proposal conventionally expressed by $!\phi$, what does σ do? The answer seems obvious:

$$\pi_{\sigma[!\phi]} = \{\llbracket\phi\rrbracket\}$$

That is to say: the conflicting preference (for $\neg\phi$) is isolated, modified (in this case, annihilated), and $\llbracket\phi\rrbracket$ is added, thus yielding a body of preferences that strictly prefers ϕ .

Here is a not quite as simple case. The prior body of preferences is given by:

$$\pi_\sigma = \{\llbracket\neg\phi \vee \psi\rrbracket, \llbracket\neg\psi\rrbracket\}$$

This body of preferences does prefer $\neg\phi$, but only as a consequence of preferring two things that together entail $\neg\phi$. To fully accept the proposal conventionally expressed by $!\phi$, what does σ do? In this case, the prior body of preferences strictly prefers both $\neg\phi$ and $\neg\psi$ to their negations. The posterior body of preference must prefer ϕ . And so the prior body of preferences must be revised. Should the posterior body of preferences:

- Prefer ψ ?
- Prefer $\neg\psi$? or
- Prefer neither?

It should prefer neither. Why?

26. Incidentally, I've recently learned that how rational agents revise prior preferences conservatively is also the topic of a good deal of research in the AGM tradition. Unfortunately, I do not know this literature well enough to comment responsibly on it here.

- Suppose ψ is preferred, so that $\pi_{\sigma[\uparrow\phi]} = \{\llbracket\phi\rrbracket, \llbracket\neg\phi \vee \psi\rrbracket, \llbracket\psi\rrbracket\}$. This preserves the preference for $(\neg\phi \vee \psi)$, but at the expense of the prior preference for $\neg\psi$.
- Suppose $\neg\psi$ is preferred, so that $\pi_{\sigma[\uparrow\phi]} = \{\llbracket\phi\rrbracket, \llbracket\neg\psi\rrbracket\}$. This preserves the preference for $\neg\psi$, but at the expense of the prior preference for $(\neg\phi \vee \psi)$.

So updating one's preferences conservatively will require "moderating" *both* the strict preference for ψ and the strict preference for $\neg\psi$, so that a strict preference for ϕ is guaranteed, but "as much" of the content of these other strict preferences is preserved as possible.²⁷ How will this work?

One very intuitive idea is that, in instituting a preference for ϕ , that preference should *percolate* through the body of preferences—i.e., that each prior preference should be *strengthened* with a preference for ϕ . On this view, the most conservative way of revising the state's prior preferences yields:

$$\pi_{\sigma[\uparrow\phi]} = \{\llbracket\phi\rrbracket, \llbracket(\neg\phi \vee \psi)\rrbracket \cap \llbracket\phi\rrbracket, \llbracket\neg\psi\rrbracket \cap \llbracket\phi\rrbracket\}$$

And this is equivalent to:

$$\pi_{\sigma[\uparrow\phi]} = \{\llbracket\phi\rrbracket, \llbracket\psi\rrbracket \cap \llbracket\phi\rrbracket, \llbracket\neg\psi\rrbracket \cap \llbracket\phi\rrbracket\}$$

On this understanding of $\pi_{\sigma[\uparrow\phi]}$, $\pi_{\sigma[\uparrow\phi]}$ prefers ϕ , but, given that preference, prefers neither ψ nor $\neg\psi$, in that $(\phi \wedge \psi)$ -possibilities and $(\phi \wedge \neg\psi)$ -possibilities are incomparable with respect to it (since, for any $(\phi \wedge \psi)$ -possibility w and $(\phi \wedge \neg\psi)$ -possibility v , both $w \not\prec_{\pi_{\sigma[\uparrow\phi]}} v$ and $v \not\prec_{\pi_{\sigma[\uparrow\phi]}} w$). I do not, however, think this sort of revision can rightly claim to be conservative. In preferring ϕ , an agent does not thereby commit herself to pursuing no other thing she desires, unless it *also* realizes ϕ . Put somewhat roughly, a proposal to require ϕ , just in virtue of being accepted, should not thereby establish a preference for ϕ as *prior* to all of the agent's standing preferences; an agent who accepts a direction to shut the window does not behave incorrectly if she satisfies another of her preferences without shutting the window. But saying that the preference for ϕ to must percolate throughout the preference state amounts to just that.

Another idea is that the prior preferences should be *weakened*, so that, if χ is antecedently preferred relative to π_{σ} , then $(\chi \vee \phi)$ is preferred relative to $\pi_{\sigma[\uparrow\phi]}$. On this view, the most conservative way of revising the state's prior preferences yields:

$$\pi_{\sigma[\uparrow\phi]} = \{\llbracket\phi\rrbracket, \llbracket(\neg\phi \vee \psi)\rrbracket \cup \llbracket\phi\rrbracket, \llbracket\neg\psi\rrbracket \cup \llbracket\phi\rrbracket\}$$

27. This intuition trades on the fact that, as I noted above, there is, in the ordering-semantic framework, a reasonable idealization away from considerations of relative priority (e.g., cardinal or ordinal utility) among elements in the ordering source. Choosing one or the other of the state's strict preference for ψ or $\neg\psi$ (rather than moderating both) would assume a priority that is not representable within the framework.

And this is equivalent to:

$$\pi_{\sigma[!\phi]} = \{[\phi], W, [\neg\psi] \cup [\phi]\}$$

While this, intuitively, does involve *moderating* the prior strict preferences for both ψ and $\neg\psi$, it falters for a similar reason: in preferring ϕ , an agent does not thereby give up all of her other basic preferences (cf. the “Problem about Permission” in Lewis 1979a). But saying that the preference for ϕ must weaken any prior preference so that it is realized if ϕ amounts to just that.

Instead, I suggest (tentatively) that percolation should be *restricted* to minimal classes of propositions that are inconsistent with the thing necessitated—classes I’ll term **revision candidates**.

Definition 4.34. If π is a body of preferences, then $\rho \subseteq \pi$ is a **revision candidate** at π with respect to an imperative $!\phi$ iff:

- i. ρ is consistent (i.e., $\bigcap \rho \neq \emptyset$)
- ii. $\rho \cup \{[\phi]\}$ is inconsistent (i.e., $\bigcap \rho \cap [\phi] = \emptyset$)
- iii. ρ is minimal (i.e., $\neg \exists \rho' \subset \rho : \bigcap \rho' \cap [\phi] = \emptyset$)

Given this definition, a conservative update function expressed by an imperative $!\phi$ can be defined, intuitively, as that function which (i) adds $[\phi]$ to the relevant body of preferences, (ii) has the preference for ϕ percolating throughout all and only the revision candidates at that body of preferences, with respect to $!\phi$.

Definition 4.35. $p \in \pi_{\sigma[!\phi]}$ iff

- i. $p = [\phi]$, or
- ii. For some revision candidate $\rho \subseteq \pi_{\sigma}$ and for some $q \in \rho$, $p = q \cap [\phi]$, or
- iii. $p \in \pi_{\sigma}$ and for no revision candidate $\rho \subseteq \pi_{\sigma}$ is it the case that $p \in \rho$

On the current understanding of revision candidates, however, this definition unfortunately does not guarantee that ϕ will be required at a cognitive state that accepts $!\phi$. There are cases where *there are no revision candidates* with respect to $!\phi$, but nevertheless, the updated state’s information rules out satisfying ϕ together with an optimal amount of the updated state’s preferences. For example, an imperative like *shut the window* is strictly consistent with a prior preference like *there is good circulation in the room*, but, nevertheless, it is known that shutting the window is not practically compatible with maintaining good circulation in the room (perhaps because the room lacks non-window ventilation). Updating such a state in line with Definition 4.35 will lead to a state that treats possibilities where the window is shut and possibilities where there is good circulation in the room as incomparable. And that, I’ve argued, is incorrect—the state should treat possibilities where the window is shut as strictly preferable.

In order for this general sort of update procedure to institute the relevant requirement in a case like this, we must modify our definition of a revision candidate as follows.

Definition 4.36. If σ is a cognitive state, then $\rho \subseteq \pi_\sigma$ is a **revision candidate** at σ with respect to an imperative $!\phi$ iff:

- i. ρ is consistent
- ii. $\rho \cup \{\llbracket \phi \rrbracket\}$ is non-realizable at σ (i.e., $\bigcap \rho \cap \bigcup I_\sigma \cap \llbracket \phi \rrbracket = \emptyset$)
- iii. ρ is minimal (i.e., $\neg \exists \rho' \subset \rho : \bigcap \rho' \cap \bigcup I_\sigma \cap \llbracket \phi \rrbracket = \emptyset$)

This definition marks as a revision candidate any minimal subset of the prior body of preferences that cannot be jointly realized, *given what is known*, in conjunction with the imperative's prejacent. Applying it to the problem case above, it will isolate the set consisting of the prior preference for good circulation in the room, given that this preference is non-realizable, given what is known, in conjunction with a preference for shutting the window.

In conjunction with Definition 4.35, Definition 4.36 will generally guarantee that acceptance of an imperative will institute the relevant requirement. In the problem case above, the prior preference for good circulation will be strengthened, so that the state comes to prefer that the window be shut (and, even better than that, good circulation in a room where the window is shut). That is a welcome result. Not only does the resultant state require that the window be shut. It also sets the stage for action—e.g., in the way of introducing new sources of ventilation into the room—designed to achieve the strengthened goal. That is all to the good.

There is, of course, much more to say about this topic, but this counts, I think, a passable first stab.

4.3.5 Prima Facie and Inconsistent Directives

A critic of the necessitation account—specifically, of the requirement-instituting dimension that I've covered in this section—might argue that it loses some of the gains reaped from switching from unstructured permissibility spheres to order-inducing bodies of preferences. In particular, it might be argued that the necessitation account—which holds that accepting an imperative $!\phi$ involves instituting a requirement that ϕ —inherits the Lewisian account's troubles with prima facie and inconsistent directives (see Section 2.1). Recall the cases that made trouble for the Lewisian account:

- **Prima facie directives.** Suppose a parent tells a child to keep her promises. The parent seems not to mean to institute an all-purposes requirement that the child keep her promises—she merely means to suggest to the child that keeping her promises is a good thing, and that, when she can do so without comprising her other basic commitments, she is required to. In other words, she wants the child to

include promise-keeping among her list of things preferred (*ceteris paribus*), but does not want the child to regard promise-keeping as a requirement

- **Inconsistent authority.** Suppose one authority issues a string of imperatives that are jointly inconsistent with an imperative $!\phi$ issued by a different authority. The addressee decides that the appropriate reaction to these utterances is indifference between ϕ and $\neg\phi$ —neither authority is superior, and she would like to bracket their disagreement by way of pursuing the other things her preferences require. The necessitation account *rules out* this sort of update as an available response to the relevant utterances.

Notice, however, that the necessitation account can represent both the state of mind involved in *prima facie preference*, as well as the state of mind involved in *indifferent preference*. *Prima facie preference* for ϕ involves *ceteris paribus* preferring ϕ , which can be understood as a defeasible preference for ϕ —defeasible under some possible modification of the agent’s information.

Definition 4.37. ϕ is **prima facie preferred** relative to σ iff:

- $[\![\phi]\!] \in \pi_\sigma$
- $\exists I : \exists \Delta\phi_1, \dots, \Delta\phi_n : I = I_{\sigma[\Delta\phi_1] \dots [\Delta\phi_n]}$ and $\min(\cup I, \leq \pi_\sigma) \not\subseteq [\![\phi]\!]$

Indifference between ϕ and ψ consists in lack of strict preference between ϕ and ψ :

Definition 4.38. ϕ and ψ are **indifferently preferred** relative to σ iff $\phi \not\prec_\sigma \psi$ and $\psi \not\prec_\sigma \phi$.

What the necessitation account cannot accommodate is the notion that accepting an imperative of the form $!\phi$, on its conventional interpretation, can involve creating a merely *prima facie* or *indifferent preference*.²⁸ An update in response to an imperative $!\phi$ that generates a merely *prima facie preference* for ϕ , or indifference between ϕ and $\neg\phi$ *cannot*, on our analysis, represent acceptance of the speech act conventionally proffered by $!\phi$. But that, in itself, is not problematic at all. *Prima facie directives*—i.e., directives that are *interpreted as* proposals to establish merely *prima facie preferences*—can be handled as ISAs. Indifference as a reaction to inconsistent authority can be handled as an accommodation strategy—a reasonable reaction on the part of an agent who is unable to coherently accept a sequence of jointly inconsistent imperatives, whether offered by different authorities, or by a single authority.

28. The necessitation account can accommodate the notion that accepting an implicitly *conditional imperative* can involve creating a merely *prima facie preference* (in the sense of Definition 4.37). Conditional imperatives, on my analysis, will create conditional requirements. Conditional requirements, on my analysis, hold relative to a state iff the relevant unconditional requirement holds at the state, under antecedent-constrained incrementations of the state’s information. For more discussion of conditional imperatives and requirements in this vein, see Section 4.5 as well as the section of the dissertation on Imperative and Deontic Conditionals.

4.3.6 Instituting Recognition

In Section 4.3.2, I suggested that necessitation acts are constituted by the fact that, when defined for a cognitive state, they induce the state to recognize a requirement. We have talked at length about what the institution of a requirement in a cognitive state involves. We now need to talk about what the recognition of a requirement by a cognitive state involves.

Recognition vs. Mere Requirement. Why think necessitation acts involve recognition-institution, rather than just requirement-institution, or that it is necessary to complicate the account with a distinct notion of a state recognizing a requirement (as opposed to merely requiring something)? As a “psychological” matter, a state’s recognizing a requirement is clearly distinct from the state simply requiring the thing (much as a state’s recognizing that, given its information, it must be the case that ϕ is distinct from it simply being the case that, given the state’s information, it must be the case that ϕ). So *if* necessitation acts do, in fact, involve recognition-institution, then an attendant complication of our account will be appropriate.

And there is good reason to think that necessitation acts do involve recognition-institution, on top of mere requirement-institution. An imperative of the form $!\phi$, for instance, can be felicitously used when the relevant body of preferences already requires ϕ , but the state does not recognize that ϕ is required. (Likewise for its performative modal equivalent, *must*.) Consider the following sequence of instructions.

- (4.25) If you draw a green marble from the bag, return it to the bag. If you return a marble to the bag, draw another marble.
- (4.26) So, if you draw a green marble, return it to the bag and draw another marble.
(Or: So, if you draw a green marble, you must return it to the bag and draw another marble.)

Although we do not have an account of conditional imperatives on the table yet (for that, see Section 4.5), it is nevertheless fairly clear that any state that accepts the first two imperatives in (4.25) will also *require* that a green marble, if drawn, be returned to the bag and that another marble be drawn. Nevertheless, the force of the imperative (or the corresponding performative modal) in (4.26) is not generally *vacuous*.

Informational Recognition. Similarly, the force of an epistemic *must* that explicitly indicates a logical consequence of a body of information is not *vacuous*.

- (4.27) If Bob folded, Jim called. And if Jim called, Jim won. And if Jim won, then he’ll run around the house celebrating.
- (4.28) Oh, Bob folded? Then Jim must be running around the house celebrating.

Yalcin (forthcoming) proposes to account for recognitional properties of an information state using the *partition device* discussed in Section 4.2.2. Very roughly, Yalcin’s idea is that a state σ will satisfy an epistemic necessity claim of the form $must(\phi)$ (or recognize ϕ as a consequence of the other information at σ) only when the proposition that ϕ is *visible* at σ . A proposition that ϕ is, in turn, visible at σ just if the issue *whether ϕ* is salient at σ . And this, finally, holds just if updating σ with a question $?\phi$ introduces no new issue into σ .²⁹

Definition 4.39. A proposition that ϕ is **visible** at σ iff $\sigma[?\phi] = \sigma$

Propositions, on this model, can be entailed by an information state (the information state can, in Yalcin’s words, be “tacitly committed” to the proposition), without being recognized by the state. For a proposition to be “explicitly believed” (as opposed to merely tacitly or implicitly believed), it must both be entailed by the information state, and be visible to the information state.

Preferential Recognition as Practical Designation. Similarly, I will suggest that the *permissibility sphere* jointly characterized by (i) an information state and (ii) a body of preferences be partitioned. What sort of cognitive state does such a partition serve to represent? While informational partitions represent informational issues that are recognized by a cognitive state—an explicit “decision” to commit to one of two issues recognized by the state—practical partitions represent *designated alternatives*—an explicit decision to realize one of two practical alternatives recognized by the state.

Before formalizing this idea, let me say a bit more by way of informally motivating this device. *For any action at all*, it is either the case that an agent will perform that action, or that the agent will not perform that action. But it is not the case that, for any action at all, that the agent has deliberated about whether or not to perform that action. These distinctions give rise to familiar distinctions in the realm of action (see esp. Davidson 1980):

- **Intentional performance and refraining.** An agent intentionally performs (or intentionally refrains from performing) an action of the form $\delta\phi$ (read: seeing to it that ϕ) only if the agent performs the action and recognizes the alternatives *seeing to it that ϕ* and *not seeing to it that ϕ* in her deliberation.
- **Non-intentional performance and refraining.** An agent non-intentionally performs (or non-intentionally refrains from performing) an action of the form $\delta\phi$ only if the agent performs the action and fails to recognize the alternatives *seeing to it that ϕ* and *not seeing to it that ϕ* in her deliberation.

29. For this to make sense, we must keep separate track of the issues that a state has considered (even after one of the possible answers has been eliminated as a possibility); we cannot represent an information state in the manner of Definitions 4.22 and 4.18. On that representation, there is no distinction between what Yalcin (forthcoming) calls a “tacit commitment” of an information state and an “explicit belief,” since both a tacit commitment that ϕ and an explicit belief that ϕ are such that for all $q \in I_\sigma$: $[[\phi]] \cap q = q$ or $[[\phi]] \cap q = \emptyset$.

To use Davidson’s famous example, consider an agent who flips a light switch by way of illuminating a room, but thereby, unbeknownst to her, alerts a burglar watching her window from the outside to her presence in the apartment. It is fairly intuitive to say, of such an agent, that she intentionally illuminated the room (and intentionally refrained from failing to illuminate the room), but non-intentionally alerted the burglar to her presence in the apartment (and non-intentionally refrained from failing to alert the burglar to her presence).

I want to suggest that the central thing (or, at least, one central thing) that separates intentional from non-intentional performance is the presence or absence of the *recognition of a proposition and its negation as alternatives* in one’s practical deliberations.³⁰ I propose to represent this recognition, formally, as a partition supplied by the action-guiding portion of an agent’s cognitive state on the possibilities that are live for that agent.

(To note and postpone an objection, it might be thought that none of the applications of this device explored here actually require a partition of one’s “practical alternatives” that is independent of the partitioning of information discussed just above. In other words, practical “issues” might be thought reducible to ordinary informational issues. For instance, to recognize a practical alternative between two options might be thought to be reducible to being explicitly uncertain whether one will perform one option or the other. I will address this objection properly in Section 4.3.9.)

To implement the idea formally, we will first re-define a cognitive state as an ordered pair, one of whose elements is an information state (typed as before), the other of whose elements is an action-guiding state (typed as a pair of a set of preferences and a partition, representing the agent’s recognized practical alternatives).

Definition 4.40. $\Pi_\sigma = \langle \pi_\sigma, \omega_\sigma \rangle$ is an **action-guiding state** for σ iff π_σ is a body of preferences and ω_σ partitions W .

Definition 4.41. $\sigma = \langle I_\sigma, \Pi_\sigma \rangle$ is a **cognitive state** iff I_σ is an information state for σ and Π_σ is an action-guiding state for σ .

Together, an information state and an action-guiding state generate what I’ll term resolved permissibility and ability spheres.

Definition 4.42. \boxplus_σ is a **resolved permissibility sphere** for σ iff:

- i. $\cup \boxplus_\sigma = \mathbf{min}(\cup I_\sigma, \leq \pi_\sigma)$
- ii. $\boxplus_\sigma \subseteq \omega_\sigma$

30. With formulations like this, the crudeness of our account starts to become apparent. To talk properly about the sorts of phenomena involved in decision-making, we require a much richer language and model theory—a language that, minimally, allows us to express the distinction between action and non-action events and a model theory that, minimally, allows us to represent the branching structure of a decision problem. I will accept this crudeness, for now, as a heuristic simplification adopted for the sake of constructing a proper model. For further discussion, see Section 4.4.1.

Definition 4.43. \boxtimes_σ is a **resolved ability sphere** for σ iff:

- i. $\cup \boxtimes_\sigma \subseteq \cup I_\sigma$
- ii. $\boxtimes_\sigma \subseteq \omega_\sigma$

A resolved permissibility sphere represents the recognized practical alternatives that an agent regards as permissible, given the relevant body of preferences and the information available to her. Condition (i) of Definition 4.42 says that a possibility is part of a resolved permissibility sphere iff that possibility generates only permissible outcomes, while condition (ii) says that every recognized practical alternative induces a partition on the permissibility sphere. A resolved ability sphere, on the other hand, represents the recognized practical alternatives that an agent regards as available, given the information available to her. In what follows, I will adopt the idealizing assumption that an agent will designate a practical alternative only if she regards that alternative as permissible (i.e., I will work only with resolved permissibility spheres).

It will be useful to distinguish between *genuine* alternatives—alternatives where the agent explicitly considers or designates *both* some action and its “complement” (but rules out neither)—from *decisions*. (Genuine) designated alternatives correspond to partitions of a permissibility sphere, with each alternative being realized in at least one cell of the partition.

Definition 4.44. $\delta\phi$ is **designated** at σ iff $\overline{\delta\phi}$ is **designated** at σ iff $\forall p \in \omega_\sigma : p \cap \llbracket \phi \rrbracket = p$ or $p \cap \llbracket \phi \rrbracket = \emptyset$

Definition 4.45. $\{\delta\phi, \overline{\delta\phi}\}$ is a **designated alternative** at σ iff:

- i. $\delta\phi$ is designated at σ
- ii. $\exists p \in \boxplus_\sigma : p \cap \llbracket \phi \rrbracket = p$
- iii. $\exists p \in \boxplus_\sigma : p \cap \llbracket \phi \rrbracket = \emptyset$

Decisions also correspond to partitions of a permissibility sphere. But in the case of a decision for an action of the form $\delta\phi$, ϕ is realized throughout *every cell of the partition*. That is to say, $\delta\phi$ is designated, and its alternative $\overline{\delta\phi}$ eliminated.

Definition 4.46. $\delta\phi$ is **eliminated** at σ iff $\forall p \in \boxplus_\sigma : p \cap \llbracket \phi \rrbracket = \emptyset$

Definition 4.47. $\delta\phi$ is a **decision** at σ iff $\delta\phi$ is designated at σ and $\overline{\delta\phi}$ eliminated at σ

Complicating the action-guiding parameter of our cognitive state-representations was motivated by reference to notions from what we might term a folk-psychological action theory. These notions included: intentional versus non-intentional performance, explicit designation of a practical alternative, decision for a designated action, and elimination of a

designated action. Our interest here, of course, is not in action-theoretic questions, but rather questions concerning language whose “subject matter,” roughly speaking, is practical, or action-involving. My contention here is that the analysis of much such language can make good use of the apparatus developed here. In the next two sections, I will substantiate this claim. The recognition-instituting force of an imperative $!\phi$ is to be understood, I will suggest, in terms of making a state decided on $\delta\phi$. I’ll also argue that phenomena like strong permission (recall the discussion of Section 2.4.5), grants of right, and a notion I’ll term strong obligation can be helpfully understood in terms of the recognitional apparatus we’ve developed here.

4.3.7 Requirement Recognition as Decision

What distinguishes a merely implicit requirement of a body of preferences from a recognized requirement? Here is an idea:

Definition 4.48. ϕ is a **merely implicit requirement** at σ iff ϕ is required at σ , but $\delta\phi$ is non-designated at σ

Definition 4.49. ϕ is a **recognized requirement** at σ iff ϕ is required and $\delta\phi$ designated at σ

It follows from Definition 4.49 that ϕ is a recognized requirement at σ just if ϕ is a decision at σ . That is because ϕ is required at σ iff $\overline{\delta\phi}$ is eliminated at σ .

Proof. (\Rightarrow) Suppose $\mathbf{min}(\cup I_{\sigma}, \leq \pi_{\sigma}) \subseteq \llbracket \phi \rrbracket$, and suppose for reductio $\overline{\delta\phi}$ is not eliminated at σ . By Definition 4.46, $\exists p \in \boxplus_{\sigma} : p \cap \llbracket \phi \rrbracket' \neq \emptyset$. So $\exists w \in \cup \boxplus_{\sigma} : w \in \llbracket \phi \rrbracket'$. But since $\cup \boxplus_{\sigma} = \mathbf{min}(\cup I_{\sigma}, \leq \pi_{\sigma})$, $w \in \mathbf{min}(\cup I_{\sigma}, \leq \pi_{\sigma})$. So $\mathbf{min}(\cup I_{\sigma}, \leq \pi_{\sigma}) \not\subseteq \llbracket \phi \rrbracket$. Contradiction.

(\Leftarrow) Suppose $\overline{\delta\phi}$ is eliminated at σ , and suppose for reductio $\mathbf{min}(\cup I_{\sigma}, \leq \pi_{\sigma}) \not\subseteq \llbracket \phi \rrbracket$. By Definition 4.46, $\forall p \in \boxplus_{\sigma} : p \cap \llbracket \phi \rrbracket = \emptyset$. Since \boxplus_{σ} partitions $\mathbf{min}(\cup I_{\sigma}, \leq \pi_{\sigma})$, we know that $\neg \exists w \in \mathbf{min}(\cup I_{\sigma}, \leq \pi_{\sigma}) : w \in \llbracket \phi \rrbracket'$. But then $\mathbf{min}(\cup I_{\sigma}, \leq \pi_{\sigma}) \subseteq \llbracket \phi \rrbracket$. Contradiction. \square

On this way of understanding things, to recognize a requirement is to *explicitly eliminate* the alternative from one’s practical deliberations.

My proposal for imperatives is, at this point, pretty obvious. I suggest that an imperative of the form $!\phi$, in addition to making it the case that ϕ is required by the relevant body of preferences, serves to make $\delta\phi$ designated. For a state to recognize a requirement that ϕ is for that state to satisfy *must*(ϕ). The creation of requirements (by way of modification of the relevant body of preferences) is handled as suggested in Definitions 4.35 and 4.36. The designation of the imperative’s prejacent is handled simply by adding the prejacent to the list of actions designated by the state.

Definition 4.50. $\Pi_{\sigma[!\phi]} = \langle \pi_{\sigma[!\phi]}, \omega_{\sigma[!\phi]} \rangle$, where:

- i. $\pi_{\sigma[!\phi]}$ is as defined in Definitions 4.35 and 4.36

$$\text{ii. } \omega_{\sigma[\!|\phi]} = \bigcup_{p \in \omega_{\sigma}} \{p \cap \llbracket \phi \rrbracket, p \cap \llbracket \phi \rrbracket'\}$$

This proposal is open to a number of different glosses (all of which are, as far as I can tell, plausible with respect to the data). Here is a gloss that we will run with: an imperative of the form $!\phi$ raises $\{\delta\phi, \overline{\delta\phi}\}$ as a genuine alternative for a cognitive state (which we could gloss as *raising the question whether or not to see to it that ϕ*), while also eliminating $\overline{\delta\phi}$ as an option for the state. This gloss suggests a four-way distinction between various species of conventional interpretations of imperatives, along the following two dimensions.

- **Agenda-setting.** The state has not designated $\{\delta\phi, \overline{\delta\phi}\}$ as an alternative (\approx has not considered whether or not to see to it that ϕ)
- **Preference-setting.** The relevant body of preferences is indifferent about ϕ

In the following grid, ‘+’/‘-’ indicates that updating a state in accordance with the necessitation act expressed by the imperative is/is not X -setting (where the value of X is determined by the row on which the sign occurs). (X -setting means simply that the act functions to modify the X parameter of the state.)

Typing Necessitation Acts, According to Prior State

	N1	N2	N3	N4
AGENDA-SETTING	+	+	-	#
PREFERENCE-SETTING	+	-	+	#

In the case of N4 acts, the prior state *already recognizes* the relevant requirement, and so, in line with the understanding of imperative presuppositions developed in Section 4.3.1, we will predict (correctly, I’ve argued) that imperatives are unacceptable, on their conventional interpretations, in such a case, as indicated by the hashmark. Here are examples of cases in which the remaining types of necessitation acts do seem to be expressed.

- (4.29) (*A is giving B her itinerary for the day. B has no idea what to expect from A, and has no real preference for what she does that day, except for wanting to please A*)
 A: Walk to Queens and get me a sugar cookie!
- (4.30) (*A has uttered (4.25) and, so, knows the relevant body of preferences requires B to return any green marble she draws to the bag. A also knows that B fails to recognize this property of the body of preferences.*)
 A: If you draw a green marble, make sure to return it and draw another.
- (4.31) (*B and A are fighting. B is looking unsurely at the door, considering whether to leave.*)
 A: Leave the room at once!

In the first example, the imperative generates a requirement and designates the action that the requirement requires. In the second example, the requirement is already in place; the force of the imperative is merely to effect recognition of the requirement. And in the third

example, the action of leaving the room is already designated (as is the action of staying) at the prior state; the force of the imperative is to eliminate the designated alternative of staying, by way of causing the state to decide on leaving.

I cautiously conjecture that any *conventional* use of an imperative can (modulo complications arising from any deviations from the Standard Account) be fit into one of the types described in the above grid. It is not, alas, fruitful to test this conjecture by looking at earlier examples, since they generally lack enough information about the prior state to classify them.

4.3.8 Strong Permission as Recognized Permission

Recall our discussion in Section 2.4.5 of the following kinds of permissions:

- **Strong Permission.** Intuitively, ϕ is strongly permitted, relative to a body of preferences, just if ϕ is expressly permitted relative to that body of preferences.
- **Grant of Right.** Intuitively, ϕ is a right, relative to a body of preferences, iff (i) ϕ is expressly permitted relative to that body of preferences, (ii) attempts to modify that body of preferences so that it comes to require $\neg\phi$ may be legitimately resisted.

A grant of strong permission, recall, is the speech act paradigmatically expressed by a performative *may*.

(4.32) It's okay. You can take an apple.

A grant of right has a strong permission as a “proper part,” but licenses, in appropriate contexts, certain relational speech acts—specifically, speech acts that express rejection of a contrary request. Recall dialogue (2.40), repeated (slightly modified) here.

(4.33) A: Am I needed at work today?
 B: No. You can stay home.
(later, C, swamped, calls to see why A is not in the office that day)
 C: We need you at work today! Get to the office immediately!
 A: But B said it was okay to stay at home today!
 C: Oh! My bad!

My first suggestion here is simple. Strong permission is *recognized permission*. Recognized permission is weak permission plus designation. Weak permission plus designation is for the resolved permissibility sphere to be appropriately partitioned (cf. Definition 4.44).

Definition 4.51. ϕ is a **merely implicit permission** at σ iff ϕ is permitted ($\neg\phi$ is not required) but $\delta\phi$ is non-designated at σ

Definition 4.52. ϕ is a **recognized permission** at σ iff ϕ is permitted and $\delta\phi$ designated at σ

A strong permission for an action $\delta\phi$, in other words, obtains when a state, restricting its focus to possibilities that it can permissibly realize (by restricting its focus to a resolved permissibility, rather than ability, sphere), designates $\delta\phi$.

This in turn suggests the following understanding of the speech act in which a strong permission is granted. We propose roughly the following understanding for this speech act.

- First, following roughly Lewis (1979a); van Rooy (2000): strongly permitting ϕ involves weakly permitting ϕ , which in turn involves adding the “least reprehensible” ϕ -possibilities to the permissibility sphere. Adopting the proposal from Section 4.3.4, as weakening each member of each revision candidate for ϕ , by adding to each member of each revision candidate the available ϕ -possibilities that are least reprehensible with respect to the prior state.
- Second: strong permitting ϕ involves adding $\llbracket\phi\rrbracket$ to the set of recognized practical alternatives.

To implement this formally, we’ll first define revision candidates with respect to permissions. (Notice that revision candidates for permissions are defined exactly as revision candidates for imperatives, as in Definition 4.36.)

Definition 4.53. If σ is a cognitive state, then $\rho \subseteq \pi_\sigma$ is a **revision candidate** at σ with respect to a permission ϕ iff:

- i. ρ is consistent
- ii. $\rho \cup \{\llbracket\phi\rrbracket\}$ is non-realizable at σ (i.e., $\bigcap \rho \cap \bigcup I_\sigma \cap \llbracket\phi\rrbracket = \emptyset$)
- iii. ρ is minimal (i.e., $\neg\exists\rho' \subset \rho : \bigcap \rho' \cap \bigcup I_\sigma \cap \llbracket\phi\rrbracket = \emptyset$)

We then define the expected update potential for a permission of the form ϕ .

Definition 4.54. $\Pi_{\sigma[\phi]} = \langle \pi_{\sigma[\phi]}, \omega_{\sigma[\phi]} \rangle$, where

- i. $p \in \pi_{\sigma[\phi]}$ iff:
 - a. For some revision candidate $\rho \subseteq \pi_\sigma$ and for some $q \in \rho$, $p = q \cup \mathbf{min}(\bigcup I_\sigma \cap \llbracket\phi\rrbracket, \leq_{\pi_\sigma})$, or
 - b. $p \in \pi_\sigma$ and for no revision candidate $\rho \subseteq \pi_\sigma$ is it the case that $p \in \rho$
- ii. $\omega_{\sigma[\phi]} = \bigcup_{p \in \omega_\sigma} \{p \cap \llbracket\phi\rrbracket, p \cap \llbracket\phi\rrbracket'\}$

Notice that, in contrast to Definition 4.35, we do *not* add the permissions’s preajcent (i.e., $\llbracket\phi\rrbracket$) to the body of preferences. This would have the effect of making ϕ required in conjunction with clauses (i.a) and (i.b) of Definition 4.54. The force of clause (i.a) is to change any preference P that is a candidate for revision when updating with ϕ by *weakening* it, so that the preference is realized if either P is realized or a possibility among the antecedently best (a.k.a. “least reprehensible”) available ϕ -possibilities is realized.

Grants of right are, we suggested, a sub-type of strong permissions. They are strong permissions that also, when the discourse is right, generate an entitlement to reject any future request that is incompatible with the permission.³¹ This entitlement could be modeled very simply by (i) adding a parameter to our cognitive state representations to keep track of such entitlements (in the interest of keeping things simple, we will not do this here), (ii) having an necessitation act of the form $[\!|\phi]$ be undefined for states in which an entitlement to $\neg\phi$ is antecedently in force).

The ability to give an account of strong permissions and rights—as well as an account the speech acts by which these things are established—is another point in favor of the recognitional apparatus we have been developing. In the next section, I will consider a challenge to the apparatus that arises, not from any sort of deficiency with respect to representing the relevant phenomena, but rather from a question regarding the “ontology” of practical questions. This challenge will ultimately stem from skepticism about whether there are any essentially practical (as opposed to informational) questions at all.

4.3.9 Practical Questions and Preferential Issues

Apart from the question of whether action-guiding or preferential partitions are *suitable* devices for representing the state of mind of designating or explicitly considering an action, there is the question of whether they are *necessary*. For it might be thought that the states of mind in question have an *informational*, rather than essentially action-guiding, flavor; although the states of mind in question are certainly about action, perhaps they can be characterized as features of the information-tracking, rather than action-guiding, parameter of the state. Insofar as introducing recognitional devices into the action-guiding cognitive parameter is a departure from tradition, and involves a complication of the technical apparatus, there would be some reason to prefer doing things that way, were it viable. Let us call the strategy of doing things using characteristics of an information-state **Cognitivism**. According to the sort of Cognitivism I have in mind:

- For $\delta\phi$ to be *designated* at σ is just for the proposition that the agent sees to it that ϕ to be *visible* at σ , i.e., for the agent to explicitly consider the proposition that she sees to it that ϕ as possible.³²
- For $\delta\phi$ to be *eliminated* at σ is just for (i) $\delta\phi$ to be designated at σ , (ii) the agent to know that she will not see to it that ϕ (her information rules it out that she sees to it that ϕ).

31. This could be very usefully modeled using the apparatus developed in Chapter III. Asher & Lascarides (2001), indeed, make a point of emphasizing the ease with which specific sub-types of a general type of T speech acts can be understood as T acts that place additional constraints, beyond those generally associated with T acts as such, on a discourse (as, as I suggested in Section 2.4.5, grants of right place additional constraints, beyond those associated with grants of permission as such, on a discourse).

32. One might also say that, for $\delta\phi$ to be *designated* at σ is just for the proposition that the agent is required to see to it that ϕ to be *visible* at σ . I’ll ignore this alternative, since it makes no difference to my arguments here.

Cognitivism about designation and elimination allows us to be Cognitivist about designated alternatives as well as decisions.

- For $\{\delta\phi, \overline{\delta\phi}\}$ to be a *designated alternative* at σ is just for (i) $\delta\phi$ to be designated at σ (hence also for $\overline{\delta\phi}$ to be designated at σ), (ii) neither $\delta\phi$ nor $\overline{\delta\phi}$ to be eliminated at σ .
- Decision is defined as in Definition 4.47.

Cognitivism (about both intention and the action-guiding attitudes that structure practical reasoning) has a venerable pedigree in the action theory and moral psychology literature.³³ Nevertheless, it is by far the minority position in those fields.³⁴

Practical Questions. Apart from this rather blatant appeal to authority, there are natural worries about going Cognitivist about the action-guiding states of mind that figure in our discussion here (in particular, the state of mind of *recognizing a requirement*. To recognize a requirement (make a decision, etc.) is, I suggested, to:

- Entertain a *practical question* (a question of the form *whether or not to see to it that ϕ*).
- *Eliminate one of the answers* to that question.

What are practical questions? Here, first, are some examples.

(4.34) To be or not to be?

(Self-directed practical question)

(4.35) Hmm. . . what to have for lunch?

(Self-directed practical wh-question)

(4.36) A [to C]: Leave at once!

B [to C]: No, don't leave!

C: Hmm. . . give me a minute to think it over.

(Practical question arising from conflicting imperatives)

All of these cases intuitively involve an agent entertaining a practical question. Informally, I want to suggest—I hope very intuitively—that entertaining a practical question, in all of these cases, is a matter of considering distinct decisions as alternatives. Or, if you do not like Noncognitivist talk about decisions, I instead suggest that practical questions are questions that *have imperatives, rather than declaratives, as their answers*. The imperative *be!* (as well as the imperative *don't be!*) counts as an answer to (4.34); the imperative *have an egg salad sandwich!* (as well as the imperative *have tuna salad!*) counts as an answer to (4.35); and the imperative *don't leave!* (as well as the imperative *leave!*) counts as an answer to (4.36).

33. See, e.g., Harman (1976); Velleman (1989, 2000). And for relevant discussion, see Bratman (2009).

34. One big reason: it is very hard to get the content of the information just right so that it counts as an intention, rather than a *de se* belief about the sorts of things one will, perhaps in spite of oneself, do.

Somewhat more formally, we can think of the force of a practical question at a state σ as modifying σ 's action-guiding parameter as follows: take the union of the result of updating σ 's action-guiding parameter with all of the imperative answers to the practical question. And this, in turn, is equivalent to simply, for each answer to the practical question $!\phi_1, \dots, !\phi_n$, making it the case that $\{\delta\phi_i : 1 \leq i \leq n\}$ is a designated set of alternatives at σ (cf. Definition 4.45).

Here I want to run with the notion that imperatives constitute the answers to practical questions to make an intuitive case that practical questions are *non-informational* in character. On this way of thinking, settling on an answer to a practical question typically induces a change in one's planning dispositions, rather than a change in one's beliefs. Any change in an individual's beliefs that is effected by such a change is account for as a reflex of the agent's reflective awareness of this change in her planning dispositions.³⁵ Put a bit differently, practical questions are questions that concern what *to do*, rather than what *to believe*. As such, they partition the space of possible plans, rather than the space of possible worlds. Practical questions are the sorts of questions that paradigmatically arise when an agent is entertaining alternative actions, as opposed to alternative depictions of the world. Considering what to eat for lunch involves considering a question of how to plan. Answering this question (*have egg salad!*) means eliminating any plan in which the agent does not have egg salad for lunch from consideration.

The fact that the question can be glossed in declarative, normative language (*what should I eat for lunch?*), as can its answers (*I should have egg salad for lunch*), is not as relevant as it might at first seem. For one, deciding, by way of resolving this practical question, that you should have egg salad for lunch is intuitively not a matter of coming to think that you should have egg salad for lunch. It is, rather, a matter of planning to have egg salad for lunch. Of course, if you are rational, coming to think that you should have egg salad for lunch will generally *lead you to* form an intention to have egg salad for lunch. But, as I will suggest in a later chapter, it seems to me mistaken to conflate the state of mind that arises from *deciding the practical question* with the state of mind of *thinking that* one should have egg salad for lunch. For two, it is by no means required that we adopt an informational analysis of normative language. A Norm Expressivist like Gibbard (1990, 2003) might (indeed, in a sense *does*) defend a plan-partitioning analysis of normative questions.³⁶

4.3.10 Summing Up

Let us conclude this section with a summation of our proposal about the speech act that imperatives conventionally express. Imperatives, I argued, conventionally express

35. The case I make here is mostly intuitive in nature. In particular, I am leaning heavily on the sorts of intuitions that motivate the Noncognitivist view of normative judgment (see esp. Gibbard 1990, 2003). This means that the Cognitivist, who will have her own explanations of these intuitions, will not have any reason to change her mind after reading it. My aim is only to prompt a sympathetic reader to see some of the appeal in Noncognitivism.

36. For a case for Expressivist about normative language, see my "Meaning for Expressivists," Appendix A.

necessitation acts, which have both presuppositional and proffered components.

- Presuppositional: saliently, imperatives presuppose that the prior state does not satisfy $must(\phi)$.
- Proffered: imperatives propose to update the prior state, so that it comes to satisfy $must(\phi)$.

There are, I suggested, two conditions on a state satisfying $must(\phi)$.

- Requiring that ϕ .
- Designating $\delta\phi$.

Here, then, is the total picture of necessitation acts defended in this section.

- $\sigma[!\phi] = \left[\begin{array}{l} \langle I_\sigma, \Pi_{\sigma[!\phi]} \rangle, \text{ if } \sigma \not\models must(\phi) \\ \text{undefined otherwise} \end{array} \right]$ (Presuppositional)
- $\Pi_{\sigma[!\phi]} = \langle \pi_{\sigma[!\phi]}, \omega_{\sigma[!\phi]} \rangle$ (Proffered)
 - $p \in \pi_{\sigma[!\phi]}$ iff
 - i. $p = \llbracket \phi \rrbracket$, or
 - ii. For some revision candidate $\rho \subseteq \pi_\sigma$ and for some $q \in \rho$, $p = q \cap \llbracket \phi \rrbracket$, or
 - iii. $p \in \pi_\sigma$ and for no revision candidate $\rho \subseteq \pi_\sigma$ is it the case that $p \in \rho$
 - $\omega_{\sigma[!\phi]} = \bigcup_{p \in \omega_\sigma} \{p \cap \llbracket \phi \rrbracket, p \cap \llbracket \phi \rrbracket'\}$

While there is, I think, a great deal that is novel about this account, it is nevertheless rather traditional, in that it falls entirely within the paradigm supplied by the Standard Account of speech acts. Necessitation, to be sure, is a rather complicated sort of speech act. But there is nothing, as such, that is non-Standard about it. Nowhere have we made use of any representations of any speech acts that are not constructed by composing a force with a content (i.e., essentially complex speech acts—speech acts that *must* be represented as taking “narrow scope” with respect to combinatory operation on speech acts).³⁷ In fact, we have spent much effort on trying to say how precisely the force expressed by the imperative operator ‘!’ (which, in every case considered here, takes wide-scope with respect to a propositional radical) should be modeled. Nor have we abandoned the notion that speech acts correspond, in some sense, to propositional attitudes. Explicitly requiring a state of affairs is undoubtedly a propositional attitude—indeed, the very sort of propositional attitude that figures centrally in any agent’s reflective planning behavior.

37. There is, of course, an uninteresting sense in which necessitation acts are composed of (i) a presuppositional act, (ii) a requirement-instituting act, (iii) a recognition-instituting act. In this sense, *most* speech acts are composed of multiple component speech acts, since most speech acts involve at least a presuppositional component. (Assertions, for instance, presuppose that the thing asserted is relevant to resolving an issue that is conversationally salient.) But this does not represent a serious challenge to the Standard Account (since, *inter alia*, it does not require abandoning logical representations of speech acts on which a speech act operator takes scope over a propositional radical).

In the following section, we will consciously move beyond the Standard Account. There is an array of prosaic uses of imperatives which require a more sophisticated system of representation—uses that do not correspond to any non-gerrymandered propositional attitude, and which cannot be represented under the force-content rubric. We will tackle this now.

4.4 Against the Standard Account

What does a case against the Standard Account ultimately require? Broadly, it will require identifying speech acts have more *complexity*—perhaps, but not necessarily, recursive complexity—than can be represented using the composition of an illocutionary force and a propositional content. Recall from Section 4.2.1 that a case against the Standard Account along these lines will require identifying speech acts that are essentially complex, as opposed to merely superficially complex.

Definition 4.55. A speech act α is **essentially complex** iff $\neg\exists\Delta\phi \in \mathcal{S} : \alpha = [\Delta\phi]$

Definition 4.56. A speech act α is at most **superficially complex** iff $\exists\Delta\phi \in \mathcal{S} : \alpha = [\Delta\phi]$

In other words, the fact that a speech act has a complex *representation* does not require that the speech act *itself* be complex.

This section will break down into two parts. We will first identify some specious arguments against the Standard Account—specious because, as we will see, they adduce evidence of speech acts that are at most superficially complex (given reasonably expectable improvements in the expressive power of the base language). We will then identify a few non-specious arguments—arguments which seem to require complicating the apparatus developed in the prior section along several dimensions. Specifically, we will see some reason to introduce connective- and quantification-like operations on speech acts—operations that are required to outscope speech act designators in representations of the relevant utterances’ logical forms. We will also see reason to introduce a hierarchical sort of structure into our representations of cognitive states, as well as devices in our illocutionary metalanguage for selecting different levels of this hierarchy for update.

4.4.1 Bad Arguments Against the Standard Account

The fact that a speech act is not representable with a speech act designator of the form $\Delta\phi$, for some $\phi \in \mathcal{L}$ is not, in itself, a strike against the Standard Account. The base language \mathcal{L} was, recall, assumed to be an ordinary, non-modal, first-order language, with a fully standard possible worlds semantics. It is, however, well-known that representing certain phenomena, particularly concerning action, requires, at the very least, the introduction of some intensional notions. In this section, I will run through some illustrations of this. In

each case I discuss, we have the option of maintaining an account that is Standard in spirit, or developing a non-standard account. Since the phenomena *can* be accommodated within a Standard Account, I don't think that any of these cases pose a problem for the Standard Account.

There are different ways to conceptualize the objections posed by the sorts of examples that we will consider in this section.

- **Linguistic.** This objection has to do with the adequacy of a *theory of meaning* built on top of the Standard Account. Fleshing this out a bit, the objection notes that there are various types of necessitation acts with dedicated modes of expression in natural language. Some types of such acts are claimed to be unrepresentable on the Standard Account. There are, for instance, necessitation acts that concern ordered sequences of actions, conditional necessitation acts, etc. Since some imperatives seem to conventionally express necessitation acts concerning an ordered sequence of actions or conditional necessitation acts, it is claimed that a speech act theory of meaning in the vein of the Standard Account cannot be adequate with respect these types of imperatives.
- **Representational.** This objection has to do with the *adequacy of the cognitive apparatus* we have developed to implement the Standard Account's understanding of speech acts. This objection involves doubt about whether our representational apparatus can account for the cognitive phenomena of merely conditional acceptance of a directive, a direction's being *in force* or *binding* with respect to a state at a given time, the fulfillment of one direction *giving way* to a new direction (the latter being in force conditional on the fulfillment of another), and a direction *lapsing* (or failing to lapse) upon fulfillment.

Lest the reader be left scratching their head, I should note that I will tend to slip pretty seamlessly between these two ways of conceptualizing objections to the Standard Account. This might seem strange: if our focus is on language, why pay any attention to "representational" objections, as such? The answer is that, for any representational objection we identify, there will generally be an associated objection to a speech act theory of meaning (since the states of mind in question can generally be induced by linguistic means). Sometimes it is easier to talk about the cognitive dimensions of planning than about phenomena within the rather rarefied realm of speech acts (and so I will choose to do so). Ultimately, though, our subject-matter, here and throughout, is properly linguistic, rather than cognitive.

Agentives. Directions are *agentive*, which is to say, loosely speaking, that a direction typically involves a speaker's prescription of a specific relation between an agent and an action—the relation of the agent *making it the case* or *seeing to it* that some state of affairs is realized. Intuitively, an imperative of the form $!\phi$ is not satisfied when the agent takes some action that, as a matter of coincidence or side-effect, ends up making it the case that ϕ .

Suppose Ishmael is conversing with Ahab, and Ahab is trying to get Ishmael to join him

in his pursuit of Moby Dick. Ahab accordingly utters the request (4.37).³⁸

(4.37) Sail with me in search of Moby Dick!

Ishmael decides to sleep on it. In the morning, after being bunked with Queequeg (and becoming fast friends), Queequeg issues the request (4.38).

(4.38) Sail with me on the Pequod!

And so Ishmael decides to go with Queequeg rather than Ahab. Nevertheless, of course, it turns out that the Pequod is the ship captained by Ahab. Has Ahab's request been *satisfied*? Has Ahab's communicative intention been realized?

Intuitively, no to both questions. (Which is not, of course, to say that Ahab would have to be displeased at the fact that a desirable state of affairs came to pass; it is only perhaps to say that it would be appropriate for him to be.) Ahab proposed that Ishmael come to recognize a requirement for him to see to it that he sailed with Ahab in search of Moby Dick. But Ishmael did not come to recognize this requirement (or, if he did, he eliminated it as an option when he decided for sailing with Queequeg). Nor did he actually see to it that he sailed with Ahab in search of Moby Dick. He saw to something else entirely—something that just so happened to result in him sailing with Ahab in search of Moby Dick. In a slogan, then, a direction act is typically agentive in character; it prescribes a specific relation between an agent (typically the addressee) and an action, and attempts to get the agent to plan to establish such a relation between herself and that action. When the agent does not so plan, the communicative goals of the speaker have not been realized (even if the agent does plan to do something that will, as a side effect, result in her performing the relevant action, without seeing to it that she performs that action).

It might seem that the Standard representation of direction acts (on which they are the result of combining directive force with a propositional content) obscures the agentive character of direction (cf. Green 1997). Typical directions are not fiats—mere prescriptions that *some proposition be made true*. They are prescriptions that an agent (the addressee) see to it that she makes the proposition true. On this way of thinking, the imperative operator '!' should take at least two arguments, rather than just one; representing directions as combinations of directive force with propositional contents collapses them into fiats. Direction should rather be represented as prescribing a kind of relation between *an agent* and *an action*, along the lines below (where *a* is typically a constant denoting the addressee):

$a!\phi$

This sort of speech act designator would be read as something like 'a: see to it that ϕ !'. Fiats, on the other hand, would receive a standard representation ($!\phi$), and would be interpreted

38. The example and much of the ensuing discussion is inspired by Belnap, Jr. & Perloff (1988).

as prescriptions that some state of affairs be seen to (cf. [Kenny 1966](#)).

This criticism of the Standard Account is mistaken. For one, I have suggested that the result of updating a state with an imperative of the form $!\phi$ is a state *that is decided to see to it that ϕ* . In this sense, the agentic character of direction, although not represented explicitly, is built into the account as it stands. The account we have developed can account for the fact that Ahab's communicative intention fails to be realized by Ishmael's plans by noting that accepting Ahab's request would lead to Ishmael's planning state being decided on *sailing with Ahab in search of Moby Dick*. This, of course, is not the state that Ishmael ultimately ends up in.

More generally, though, it might be worried that the structure of action-guiding states, as it stands, does not accommodate the distinction between viewing a state of affairs p as *ceteris paribus desirable* and viewing a state of affairs p as *ceteris paribus to be brought about by the agent*. Both states of mind would seemingly have to be represented simply by having p included among the state's preferences (thereby conflating them). Analogously, it might be thought that the account cannot accommodate the distinction between fiats—directives that express a speaker's desire for the realization of some state of affairs—and genuine directions—directives that express a speaker's desire that the addressee see to it that the relevant state of affairs is realized.

This worry, too, is misplaced. We can draw a distinction between viewing a state of affairs as desirable and as to be brought about by the agent along the following lines:

- ϕ is σ -**desirable** iff $[\![\phi]\!] \in \pi_\sigma$
- ϕ is σ -**to-be-brought-about** iff the proposition that the agent sees to it that $\phi \in \pi_\sigma$

Similarly, the Standard Account can draw a distinction between mere fiats and actual directions along the following lines.

- $!\phi$ is a **fiat** for σ iff ϕ is non-agentic for σ
- $!\phi$ is a **direction** for σ iff ϕ is agentic for σ

Following [Belnap, Jr. & Perloff \(1988: 179\)](#), we'll define agentic sentences, roughly, along the following lines:

- ϕ is **agentic for σ** just if it is paraphrasable as σ *sees to it that ψ* (for some ψ)

Putting all of this together with our stated preference-modification analysis of conventional imperative force (on which any imperative $!\phi$, fiat or otherwise, serves to make a state recognize a requirement that ϕ), we have that:

- The canonical effect of updating σ with a fiat $!\phi$ is for ϕ to become σ -desirable
- The canonical effect of updating σ with a direction $!\phi$ is for ψ to become σ -to-be-brought-about (for some ψ)

These distinctions in hand, we could further develop this proposal by explicitly defining a formal language of agentives and stating a semantics for it.³⁹ To give a flavor of what would be involved here, standard semantic treatments for *sees to it that* constructions generally have an agent seeing to it that ϕ when the agent *executes an action* that, in some sense, *guarantees the truth* of ϕ at some future moment. Prescribing a relation between an agent and the action of seeing to it that ϕ , then, would amount to prescribing that the agent execute an action which guarantees ϕ 's truth. Issuing such a prescription would amount to proposing that the agent plan to execute an action which guarantees ϕ 's truth. (Again, since Ishmael neither executes, nor plans to execute, such an action, Ahab's practical and communicative goals are both clearly thwarted).

We will not pursue this sort of proposal here.⁴⁰ What is important for our purposes is that the relevant distinctions can be accommodated by complicating the syntax and semantics of the base language of propositions, rather than by introducing complexity into the realm of speech acts proper. Any account developed within this vein would be fully consonant with the spirit of the Standard Account.

Temporality and Conditionality. Here are some more central phenomena about direction we would like to be able to represent, but, as things stand, cannot. For sake of concreteness, focus on the example of someone attempting to follow a recipe for a soufflé.

- Directions to perform a series of actions in some sequence
 - (4.39) Whip the egg whites, then fold them into the custard!
 - (4.40) Fold the egg whites into the custard, after whipping them!
- Directions to perform an action or series of actions *conditional on the performance* of an action or sequence of actions. Representing how specific courses of action will impact an agent's later commitments.
 - (4.41) Fold the egg whites into the custard, if you've whipped them!
 - (4.42) Make a pudding, if you've failed!
- Directions that are binding *now* versus those binding at some later point in time
 - (4.43) Whip the egg whites!
 - (4.44) Put the soufflé into the oven!

39. The best semantics for constructions like *Bob saw to it that Mary was notified* involves branching-time models (and satisfaction conditions) that are vastly more complicated than the constant domain models of which we made use in Section 4.1.2 (Belnap, Jr. & Perloff 1988, 1992; Belnap, Jr. 1991). See Thomason (1981b,a) for foundational work on branching-time models.

40. I see no technical barrier to doing this. Indeed, I tried to do something like it, although for a propositional dynamic logic of action, in Charlow (2009). But the formal complexity associated with this sort of treatment is formidable, and I will avoid it.

- Directions that lapse upon fulfillment versus those that do not—*dischargeable* and *non-dischargeable* directions, respectively

(4.45) Hand me that mixing bowl!

(4.46) Don't set fire to your kitchen!

To capture some of these phenomena (the first and the second, in particular) plausibly requires the ability to reference other actions in the description of the action prescribed for the addressee (specifically, in the direction's depiction of the prior contingencies upon which the agent's performance of the action is conditioned). Insofar as the base language \mathcal{L} is not equipped to do this, this highlights another expressive deficiency of \mathcal{L} . But, at the risk of sounding like a broken record, this is a strike against the standard account only if there are *no propositions* such that sequenced directions and directions conditional on the performance of other actions can be understood as directions having those propositions as their content. But that seems far-fetched: a direction to whip the egg whites, then fold them into the custard can be understood as having the proposition *the addressee sees to it that she folds the egg whites into the custard, after she has seen to it that she has whipped them* as its content. Even assuming a language capable of expressing agentives, some sort of tense operator is required to formulate this proposition in the base language. But it is, for all that, a proposition.⁴¹

That response, though, is a little bit hasty. It is important to separate out two related, but distinct, phenomena here.

- *Preconditioned* (but unconditional) direction. The sort of direction usually expressed by directions to perform a series of actions in some sequence.
- *Conditional* (or contingent) direction. The sort of direction usually expressed by conditional imperatives.

Preconditioned directives are, in a sense, stronger than conditional directives. A directive to see to it that ϕ , preconditioned on ψ , does intuitively “include” a conditional direction to see to it that ϕ , if ψ . But it is not exhausted by this. Preconditioned directives have *genuine directive force*, in the sense that they conventionally propose to institute new recognized requirements for the agent. An agent who accepts (4.39) will recognize requirements both to whip the egg whites *and* fold the whipped egg whites into the custard; she fails, with respect to this directive, if she fails to fold the whipped egg whites into the custard.⁴² This

41. This is comfortable terrain for approaches within the propositional dynamic logic (PDL) tradition. In PDL treatments of action, representing actions of the form $\lceil \pi, \text{ then } \pi' \rceil$ is trivial (see, e.g., Segerberg 1990, 1992). For extensive discussion of a PDL treatment of imperatives and direction, see Charlow (2009). A brief word by way of comparing the STIT and PDL approaches: STIT theories seem to have it right that any semantics for a performance predicate or relation will rely on branching-time models. This is something which PDL treatments, as usually presented, tend to neglect; see, e.g., Segerberg 1992: 375. (Indeed, I know of no way to sensibly formalize the claim that an agent has performed a particular action at w —crucially different from the claim that an agent is able to perform a particular action at w —without branching-time models.)

42. Castañeda (1958) actually disagrees with this claim. For his objection and a reply, see Section 4.4.6.

is not, in general, the case with conditional directives. Conditional directives, on the other hand (and as I'll be stressing repeatedly in Section 4.5), propose to institute new recognized *conditional* requirements for the agent. Accepting (4.41) in no sense generates an obligation for the agent to fold the whipped egg whites into the custard.

To spoil the surprise somewhat, the reason it is important to draw this distinction is that conditionality, rather the phenomenon of preconditioning, is, I'll argue (Section 4.5), a source of the trouble for the Standard Account. The latter, but not the former, is dealt with easily, using the sort of base-language-enriching strategy we've been happily deploying.

Preconditioning. Preconditioned directions are paradigmatically, but not necessarily, expressed by directions to perform a series of actions in some sequence. Other cases, drawing on distinctions drawn in von Wright (1963), involve directions to perform actions *under certain descriptions of initial conditions*. So far, we have presupposed that actions are individuated by the kinds of states of affairs they bring about (somewhat more precisely, by the propositions that are true in each state of the world that might result from performing the action). von Wright, however, argues for recognizing a distinction between the following kinds of action.⁴³

- **Change-bringing:** *producing ϕ from ψ (destroying ψ for ϕ)*
- **Change-preventing:** *maintaining ϕ from ϕ (suppressing ψ for ϕ)*

The production of a state of affairs in which ϕ and the maintenance of a state of affairs in which ϕ cannot, of course, be individuated with respect to their end-states. Nevertheless, a direction to produce a state of affairs in which ϕ is *prima facie* distinct from a direction to maintain a state of affairs in which ϕ ; indeed, these directions differ precisely in the preconditions they set for seeing to it that ϕ .

To explicitly rehearse the objection to the Standard Account that is being suggested here, it might seem that the Standard representation of acts of direction cannot accommodate this distinction. Instead, certain commands seem to prescribe a certain relation between an agent, an input state of affairs (initial conditions), and an output state of affairs (the result of the agent's action). Indeed, Zarnić suggests exactly this argument, writing that "If we take the sentence-radical [i.e., the content of the direction] to be a description of a . . . situation, then change-expression semantics threatens such a view, since for a change expression the use of *two radicals* is allowed" (2003: 2; my emphasis).

43. I am presenting the views of von Wright (1963) as interpreted by Segerberg (1992); Zarnić (2003).

Predictably, I want to suggest that the blame here (if in fact there be any⁴⁴) lies with expressive deficiencies in the base language. Even an agentive language cannot express the proposition that an agent sees to it that ϕ is maintained, or that an agent sees to it that ϕ is produced, because it lacks devices for expressing notions having to do with change and its absence. But such devices can be added to the language, and it seems clear that it would be possible to give them a reasonable semantics using the tools introduced for interpreting the agentive language. (Intuitively, for example, producing ϕ from ψ would involve seeing to it that ϕ , plus the condition that the point from which the agent saw to it that ϕ satisfied ψ .)

There *is* a problem in the neighborhood of this one, however. It is difficult to represent the conditionality that is inherent in a preconditioned direction (because, we see in Section 4.5, the Standard Account, along with its attendant cognitive apparatus, lacks any satisfactory treatment of conditional speech acts). It is, in a rather rough way, clear how the preference-update model of direction would have to be revised to accommodate directions prescribing sequenced actions: we would have to provide for addition of a proposition describing the addressee's execution of a sequence of actions to the ordering source for that agent. But it is, I'll claim, unclear how it could be revised to accommodate the *conditionality* inherent in preconditional direction (and conditional direction more generally). The Standard Account only gets us so far, and it is not nearly far enough.

Looking Ahead. The phenomena about imperatives and direction surveyed in this section have the following characteristic in common: non-Standard accounts could, but need not, be stated for each. And so they represent no real challenge to the hegemony of the Standard Account. In the subsequent sections, however, I will outline a series of phenomena—having principally to do with quantification, conjunction, and the aforementioned conditionality—that seem to *require* non-Standard treatments. This, in turn, will be used to make trouble for the very idea of a speech act theory of meaning for imperatives. If there *are* no suitable speech act meanings for the relevant kinds of imperatives to conventionally express, then, since since those kinds of imperatives are obviously not *meaningless*, as a class, the meanings of such imperatives, it would seem, cannot be given in terms of a conventionalized speech act. And this—like the fact that imperatives can only defeasibly be identified with a natural kind of speech act (Chapter III)—calls into question the very idea of a speech act theory of meaning for imperatives (and, thereby, for all of natural language).

44. For what it is worth, my best guess is that, yes, directions to produce and directions to maintain do differ, but only in virtue of their *circumstantial-cum-informational presuppositions*. Directions to *produce* a state of affairs in which ϕ presuppose that $\neg\phi$ currently holds, while directions to *maintain* a state of affairs in which ϕ presuppose that ϕ currently holds. They are not otherwise distinct. If that is right, there is not even a prima facie threat to the standard account here, and no real need to revise the base language. I would also like to tentatively suggest that imperative utterances of the form $!\phi$, by and large, presuppose that $\neg\phi$ is currently the case. Special language typically needs to be exploited to express a desire that some state of affairs be maintained.

4.4.2 The Challenge of Embedded Imperatives

To begin, there are two data points about imperatives that I want to highlight:

- **Point 1.** Imperatives *obviously* embed in natural language, in any syntactically reasonable sense of that notion.
- **Point 2.** There is a conventionalized connection between various syntactic subtypes of imperatives and corresponding subtypes of directive force.

The following examples illustrate both points well.

(4.47) Stop the car and turn off the engine [Conjoined imperative]

Conjunctive direction. The sort of force usually associated with a conjoined imperative. An utterance with conjunctive directive force constitutes a proposal for the agent to recognize a corresponding conjunctive requirement (*you must stop the car and you must turn off the engine*).

(4.48) Leave no stone unturned!

(4.49) No one move! [Quantified imperative]

Quantified direction. The sort of force usually associated with universally-quantified-subject imperatives. An utterance with universally quantified directive force constitutes a proposal that, for each x in the relevant domain, the agent(s) come to recognize an obligation of the form $must(Px)$. So, for instance, (4.48) proposes that the agent come to recognize an obligation of the form $must\ not\ leave\ x : stone(x)\ unturned$, for all x in the relevant domain.

(4.50) If you go out, take an umbrella! [Conditional imperative]

Conditional direction. The sort of force usually associated with conditional imperatives. An utterance with conditional directive force constitutes a proposal for the agent to recognize a corresponding conditional requirement of the form $(if\ \phi)(must(\psi))$.

Of course, clauses embedding clauses of non-declarative clause-type \mathcal{J} (interrogative or imperative), are familiar trouble cases for typical theories of the meaning of clauses of type \mathcal{J} (most of which can be classified to some degree as speech act theories).⁴⁵ A rough way of expressing the problem is this. We generally expect logical forms for syntactically complex natural language sentences to be themselves syntactically complex. But there is no real syntactic complexity in the illocutionary metalanguage in which speech act theories within the Standard paradigm state logical forms—just a set of simple speech act designators. So there is some reason to worry that any such theory can provide a satisfactory theory of the

45. As before, I follow the relatively standard practice of typing subordinate-matrix clause structures according to the clause-type of the matrix clause (cf. Bhatt & Pancheva 2006). Even if unembedded \mathcal{J} -clauses are basic, then, a clause of the form $(if\ \phi)(?/\!\psi)$ will count as an \mathcal{J} -clause.

meaning of clauses embedding \mathcal{J} -clauses.

Let me be a bit more specific about the nature of the worry. Typical (i.e., speech act) theories of the meaning of \mathcal{J} -clauses within the paradigm of the Standard Account have characteristically had the following two properties:

UNEMBEDDED OCCURRENCES AS BASIC

Unembedded occurrences of \mathcal{J} -clauses (occurrences of the form $!\phi$ or $?\phi$) are taken to be basic, and a theory of the meaning of \mathcal{J} -clauses is constructed around a theory of their meaning.

DYNAMIC PHENOMENA ARE EXPLANATORILY CENTRAL

Dynamic phenomena (e.g., the conventional discourse function of \mathcal{J} -clauses) typically occupies an explanatorily central role.

Although such theories are not always purely dynamic (in that they often assign ordinary, “static” model-theoretic entities, rather than update potentials, as the semantic values of \mathcal{J} -clauses), the conventional discourse function of \mathcal{J} -clauses plays a regulative role: it constrains the *type* of model-theoretic entity that may be assigned as the semantic value of a \mathcal{J} -clause. Interrogatives, as we’ve seen, introduce issues into a discourse, so their static semantic values are typed as divisions of (typically, partitions on) the space of possibilities. Imperatives, perhaps, introduce obligations into a discourse, so their static semantic values are typed as properties of an obligation-determining parameter of the discourse. For speech act theories of meaning within the Standard paradigm, then, theorizing begins by describing the conventional discourse function of utterances of unembedded \mathcal{J} -clauses, and we assign static semantic values to unembedded \mathcal{J} -clauses in accordance with that function. Theorizing about embedded occurrences involves identifying an unembedded \mathcal{J} -clause that has the *same meaning* (expresses the same speech act) as the embedded \mathcal{J} -clause.

There are two commitments embodied in this methodology that it is important to separate.

- **Ontological.** Giving a theory of meaning for \mathcal{J} -clauses can be accomplished within the paradigm of the Standard Account. That is to say, for any \mathcal{J} -clause C , there is some unembedded \mathcal{J} -clause $\blacktriangleright \psi$ such that $[\blacktriangleright \psi]$ is the speech act conventionally expressed by C . In a slogan, there *are* no essentially complex speech acts.
- **Explanatory.** Perhaps less obviously: even if it we end up having to abandon the Standard Account in favor of recognizing some class of complex speech acts, it is nevertheless the case that the speech acts posited by the Standard Account—the speech acts paradigmatically expressed by simple, as opposed to complex, typed clauses—will still count as basic.

Notice that violating the ontological commitment does not require violating the explanatory commitment (whereas violating the latter entails violating the former). For various reasons, it might be thought that jettisoning the ontological commitment, while retaining

the explanatory commitment, preserves the spirit, if not the letter, of the Standard speech act theory of linguistic meaning. Essentially complex speech acts might be admitted, if they could be understood as *constructions* out of the building blocks supplied by the Standard Account.

Generalizing from the examples of embedded imperatives presented above, the necessitation act expressed by an imperative of any form at all can be understood as a proposal for the relevant agent to recognize a requirement of the corresponding form.⁴⁶ So such examples will variously express conjunctive, quantified, and conditional necessitation acts (respectively: proposals for the agent to recognize conjunctive, quantified, and conditional requirements). Herein lies the challenge for a speech act theory of linguistic meaning: given the generalization that an imperative of arbitrary form expresses a proposal for an agent to come to recognize a requirement of corresponding form, identify a speech act within the Standard paradigm that can be understood as being constituted by such a proposal. If such a speech act cannot be located, identify Standard speech acts out of which the relevant necessitation act can be understood to be composed.

In the remainder of this section, I will argue two things. First, in attempting to meet the described challenge, *neither the ontological nor explanatory commitments can be sustained*. (The case against the ontological commitment will rely on a novel analysis of conjunctive and quantified imperatives. The case against the explanatory commitment will rely on a novel analysis of conditional imperatives.) Second, although this dooms the prospects for a speech act theory of linguistic meaning within the vein of the Standard Account, *in no way does this doom the prospects for a speech act theory of linguistic meaning*, as such.

4.4.3 Regular Operations

Here is a rather technical, but nevertheless effective, strategy for approaching these questions. We have typed speech acts as update instructions, which in turn have been typed as (deterministic) programs. It is therefore natural, at least from the standpoint of the logic of programming, to ask whether there are any *regular operations* on programs which it would be essential to represent in our illocutionary metalanguage. Such operations are often used in programming-inspired systems devised to model action, e.g., the Propositional Dynamic Logic systems of Segerberg (1990, 1992). And so it is reasonable to wonder whether such operations would be helpful in analyzing the sorts of actions that agents perform through linguistic utterances.

For the sake of illustration, let's focus on two binary regular operations on programs, '+' and ';', with the following rough and intuitive syntax and semantics. When α and β

46. The evidence by which this generalization could be robustly established will, not surprisingly, mimic the sort of evidence by which the claim that ordinary imperatives, of the form $!\phi$, express proposals that the relevant agent recognize a requirement that ϕ . It wouldn't, I think, be edifying to run through this evidence here; the reader can satisfy herself that the generalization holds by adapting some of the tests described in Section 3.3.

are program expressions, $\lceil \alpha + \beta \rceil$ is a complex program expression, denoting a program that executes at least one of $[\alpha]$ or $[\beta]$. $\lceil \alpha; \beta \rceil$, however, is a complex expression denoting an update potential that executes $[\alpha]$ then $[\beta]$. Somewhat more formally, it is standard to treat ‘+’ as expressing union, ‘;’ as expressing relation composition:

Definition 4.57. $[\alpha + \beta] = [\alpha] \cup [\beta]$

Definition 4.58. $[\alpha; \beta] = [\alpha] \circ [\beta] = \{ \langle \sigma, \sigma' \rangle : \exists \sigma^* : \langle \sigma, \sigma^* \rangle \in [\alpha] \wedge \langle \sigma^*, \sigma' \rangle \in [\beta] \}$

What would constitute evidence for the idea that operations like these should figure in the illocutionary metalanguage? Suppose some speech act with a canonical (and superficially syntactically complex) linguistic expression did not seem to be helpfully analyzed using the class of speech acts posited by the Standard Account. We might approach the analysis of this speech act (and the sentence that canonically expressed it) as follows.

- Do a quick first-pass syntactic analysis of the relevant sentence, so that its syntactic complexity is made explicit.
- Associate the proper parts of this analysis (i.e., the component sentences) with speech act designators.
- Link these speech act designators with some sort of illocutionary analogue of the syntactic device used to link the component sentences in the syntactic gloss.
- Associate this illocutionary operation with a regular operation on programs.

If successful, this procedure would amount to a *prima facie* case for the use of regular operations in our illocutionary metalanguage. The case could then be strengthened with an argument that the update potential resulting from executing the regular operation on the component update potentials *could not* be expressed using any reasonable candidate speech act supplied by the Standard Account. This is exactly the sort of argument we will be deploying against the Standard Account (specifically, its ontological commitment) in the next couple of sections.

All this, of course, will be just a prelude to actually stating a formal system for representing essentially complex speech acts—something we’ll postpone until all the relevant data is “collected.” Such a system will require both the introduction of a variety of devices for building complex speech act designators, together with a formal syntax and semantics for those devices. Because of our method of attack, however—which shall make extensive informal use of the relevant formal devices—this will figure mostly as an afterthought.

Looking ahead, the subsequent section will implement roughly the procedure outlined here for an array of data, some of which has been entertained elsewhere, some of which is new. Our verdicts regarding the Standard Account will split about fifty-fifty. Most of the data entertained elsewhere, I argue, does not yield a forceful case against the Standard Account. The new data proposed here does better.

4.4.4 Conjunction: Contrived Examples

Contrived (I will explain what I mean by this below) examples of natural language sentences that seem to sequence speech acts are fairly easy to come by. Here we have a sentence sequencing a request and an assertion.

(4.51) Drink another beer and Bob's a loser. (Starr 2011)

Note that the intended (sequencing) reading of (4.51) is not the “pseudo-imperative” or “conditional conjunction” reading discussed at length in Culicover & Jackendoff (1997); Russell (2007); von Stechow & Iatridou (2009). That reading is given roughly by the following gloss:

(4.52) If you drink another beer, Bob's a loser.

Predicting this reading is, of course, an interesting problem, but irrelevant for our purposes here. If (4.52) and (4.51) share a logical form (as Culicover & Jackendoff 1997 suggest), then, regardless of whether *drink another beer* bears the syntactic hallmarks of an imperative (as von Stechow & Iatridou 2009 suggest), (4.51) will plausibly be typed as a declarative, and will conventionally express an assertion of a conditional content.

More interesting is the genuine sequencing reading: the reading on which (4.51) expresses both a request (for the agent to drink another beer) and an assertion (that Bob's a loser). On its face, it would seem that this reading should be given a illocutionary logical form of the form $!\phi; \triangleright\psi$, so that $\sigma[(4.51)]$ will be given by a state-representation of the form $\sigma([\!|\phi] \circ [\triangleright\psi])$.

But this is a contrived example. There are *ands*, and then there are *ands*. The coordinating conjunction *and* can, it is widely recognized, express two kinds of coordination, one syntactic, the other discourse-level. Genuine syntactic coordinations—as contrasted with the sort of “extended coordination” relation that is often argued to obtain between sequential, but nevertheless syntactically individuable, utterances in a discourse⁴⁷—display, for instance, a distinctive sort of intonational pattern illustrated in (4.53). Extended coordination relations (of the sort involved in elaboration, narration, or explanation) display the distinctive sort of intonational pattern illustrated in (4.54).


(4.53) Fido fetched the stick_↑ and_↓ returned it to his owner.

(4.54) Fido fetched the stick_↓. He_↑ returned it to his owner.

Notice that giving (4.51) the sort of intonation in (4.53) *forces* the conditional conjunction reading. Indeed, getting the sequencing reading seems to require the sort of intonation in (4.54) that is characteristically associated with extended coordination.

47. On this view of discourses, see, e.g., Kamp (1981); Heim (1982); Grosz et al. (1989); Larson (1990).

Of course, a defender of this objection to the Standard Account could spin these facts in a way favorable to her. On extended coordination models of discourse, really any string of sentences uttered in sequence can be represented using a single discourse structure, in which, if utterance u_i occurs prior to u_{i+1} , the complex utterance of u_i then u_{i+1} is assigned the following discourse structure (where \bullet expresses extended coordination).



And so there is a sense in which the Standard Account goes really, *really* astray. It does not take exotic examples to show that the Standard Account's understanding of speech acts is impoverished. *Any* sequence of utterances related by extended coordination will do.

I don't have a knockdown argument against this spin, but insofar as it represents an argument against the Standard Account, it is a bit defeatist. If there is theoretical reason to recognize extended discourse structures in the representation of utterance meaning, then I agree that there is theoretical reason for a speech act theorist of meaning to recognize non-Standard speech acts (since interpreting these structures will involve the computation of a complex speech act by means of the regular operation ' $;$ '). Nevertheless, there is no reason for thinking that the *basic units of discourse*—unsequenced utterances—should have a non-Standard analysis. And, so, while this response does represent some give in the Standard Account's ontological commitment, it does not involve any relaxation of the explanatory commitment. What's more, the attendant relaxation of the ontological commitment is, it's fair to say, minimal. For, on its face, there is nothing that can be done with representing utterances as nodes of interpreted discourse structures that could not be done with a representation on which utterances were *separately* interpreted, and coordination relations were supplied to link these interpretations, as a strategy for maximizing discourse coherence (cf. Asher & Lascarides 2003).

Starr (2011) also suggests that account for the contrast between (4.55) and (4.56) requires the recognition of speech act sequencing.

(4.55) Marry Pat and Hillary!

(4.56) Marry Pat and marry Hillary!

Intuitively, (4.55) requests polygamy, while (4.56) could represent simply a conflicting string of requests through which a speaker might express her deep indecision. But the Standard Account has only one available representation for both of these imperatives, in which the imperative operator ' $!$ ' takes wide scope over the proposition that you marry Pat and Hillary (which is, notice, equivalent to the proposition that you marry Pat and marry Hillary). To represent the contrast, Starr suggests, we need to represent the second imperative as using

something like speech act sequencing. Notice that, by way of response, it will not do to note that (4.55) has a reading on which it does not request polygamy. The important claim is that there are two distinct readings that require distinct logical representations, but that the Standard Account has the resources to represent only one.

In reply, I will just deny that there are two distinct readings here. There is one reading—the reading requesting the agent to marry both Pat and Hillary. A speaker can fail to understand or endorse what she is, as matter of fact, requesting in uttering (4.56). Nevertheless, what she requests is that the agent marry both Pat and Hillary. Starr avoids this simple response by stipulating that the sentences are being used to *report* what, say, two different parents have requested of the agent. Intuitively, he is right about the content of these reports: (4.55) reports that the parents agree on polygamy, while (4.56) reports that the parents disagree about whom the agent should marry. But these different reports are easily represented within the standard account: one reports that the parents request polygamy, the other reports that the parents disagree. There are no sequenced imperatives involved in representing either reading.⁴⁸

4.4.5 Conjunction: Sequencing Designation

There is a distinction in meaning between the following conjunctive imperatives.

(4.57) Lick and seal the envelope!

(4.58) Take a shower and buy a new shirt!

The difference, intuitively, is this: (4.57) requests a complex action of the agent (the action of licking and sealing the envelope), whereas (4.58) requests separate, non-complex actions of the agent (the action of taking a shower, and the action of buying a new shirt).

To be clear, I am *not* suggesting that this fact, by itself, warrants representing (4.58) using a speech act designator of the form $!\phi;!\psi$. Were that the sole difference between these imperatives, then we could chalk this difference in meaning up to an expressive deficiency in the base language (namely, the inability of the language to represent the distinction between genuine complex actions and contingently connected sequences of simple actions). There would be no problem for the Standard Account that could not be remedied by improving the expressive power of the base language.

But there are other differences in meaning between these conjunctive imperatives, due to the fact that they intuitively involve different ways of *designating actions* (in much the same way as a question of the form $?(\phi \wedge \psi)$ raises a different issue than a question $?\phi$ followed by a question $?\psi$; see Section 4.2.2). Recalling that necessitation, on our account, amounts to designation plus requirement, here is what I am claiming this difference in

48. As for the matter of distinguishing the state of mind induced by accepting the parents' requests in (4.55) (and thus coming to recognize that polygamy is required) from that induced by treating the requests as conflicting, see Section 4.3.5.

designation behavior consists in:

- (4.57) designates and requires an action of the form $\delta(p \wedge q)$
- (4.58) designates and requires two actions, each of the form δp

This difference can be observed in various places.

- **Simple and complex requirements.** (4.57) and (4.58) generate distinct kinds of requirement. (4.57) generates a single, complex requirement: to lick the envelope by way of sealing it.⁴⁹ (4.58) generates distinct, simple requirements, which the agent can consider separately, namely, the requirement to take a shower, and the requirement to buy a new shirt.
- **Partial and total compliance.** There is, then, a sense in which an agent who takes a shower without buying a new shirt can be said to have complied partially with the force of a speaker's utterance of (4.58); this speech act consists of distinct requests, only one of which the agent fails to realize. There is no sense in which an agent who licks the stamp without sealing the envelope can be said to have complied partially with the speech act expressed by an utterance of (4.57); this speech act requests a complex state of affairs, and the agent fails to realize this request.
- **Simple and complex permissions.** (4.57) and (4.58) generate distinct kinds of permission. An utterance of (4.57) does not, in any sense, confer a bare permission for the agent to lick the envelope. It confers a permission for the agent to lick the envelope *as part of the action of sealing the envelope*; it confers a single, complex permission (which does not allow the isolation of either of its constituents). To contrast, an utterance of (4.58) does confer a bare permission for the agent to take a shower. It also confers a bare permission for the agent to buy a new shirt. That's to say: it generates distinct, simple permissions.

I do not think there is any real barrier to assigning (4.57) a Standard Account-friendly logical form—something like $!(p \wedge q)$ —since the meaning of the imperative is exhausted by its expression of a single necessitation act with content $(p \wedge q)$. This satisfactorily accounts for each of: (i) the fact that it generates a single requirement of the form $(p \wedge q)$, (ii) the fact that failing to satisfy either p or q makes the agent wholly non-compliant with the speech act expressed by the speaker, (iii) the fact that it confers a single permission of the form $(p \wedge q)$.⁵⁰

Not so for (4.58). Giving it a logical form of the form $!(p \wedge q)$ obscures its separate designation of δp and δq —a phenomenon attested by the three data points we have noted. To account for separate designation, it seems that we are forced to recognize a non-Standard

49. Again, our base language lacks the expressive power to actually represent this sort of action, but, for our purposes, an action designator of the form $\delta(\phi \wedge \psi)$ will suffice.

50. One thing for which this logical form does not account is that a speaker of (4.57) is intuitively enjoining the action of *licking and then sealing the envelope*. But this is a symptom of deficiency in the base language—its inability to distinguish complex actions that result from performing actions in *conjunction* and complex actions that result from performing actions in *sequence*. For systems that do allow such distinctions, see Segerberg (1990); Charlow (2009).

logical form for (4.58), of the form $!p;!q$. This satisfactorily accounts for each of: (i) the fact that the imperative generates distinct, simple requirements, one of the form that p , another of the form that q , (ii) the fact that failing to comply with either of these requirements may land the agent in a state of merely *partial* non-compliance with the force of the speaker's utterance, (iii) the fact that it confers distinct, simple permissions.

It bears (re-)emphasizing that all of this is an upshot of the question-like behavior of designation acts. A question of the form $?(ϕ ∧ ψ)$ raises a different issue—namely, whether $(ϕ ∧ ψ)$ —than those raised by a question $?ϕ$ followed by a question $?ψ$ —namely, whether $ϕ$ and whether $ψ$. So, too, a designation act with content $(ϕ ∧ ψ)$ raises a different practical issue than separate designation acts with respective contents $ϕ$ and $ψ$. This fact, to emphasize, has nothing to do with the requirement-generating dimension of necessitation acts: a state requires $(ϕ ∧ ψ)$ if and only if it requires $ϕ$ and it requires $ψ$. With respect to their requirement-generating dimensions, a necessitation act with content $(ϕ ∧ ψ)$ is indistinguishable from separate necessitation acts with respective contents $ϕ$ and $ψ$. Not so with respect to their action-designating dimensions. This is what makes a speech act of the form $!(ϕ ∧ ψ)$ non-equivalent to one of the form $!ϕ;!ψ$, and what, in turn, allows us to account for the difference between imperatives like (4.57) and those like (4.58).

In other words, we've identified what appears to be an essentially complex speech act with a canonical means of expression in natural language. That represents the first basic challenge to the Standard Account endorsed here. It will not be the last.

4.4.6 Disjunction, Permission Content, and the Ross Paradox

It is fairly obvious that imperatives embed, in some sense, under disjunction.

(4.59) Post the letter or burn it! (Ross 1941)

It might be thought that (4.59) has a straightforward Standard-friendly representation of the form $!(ϕ ∨ ψ)$. And at first pass this seems plausible. But this plausibility more than likely rests on an intuitive understanding of the meaning of the '!' operator (I'll say why in a bit). But assuming an intuitive understanding is, strictly speaking, illicit. According to the account we have been defending, the force of an imperative of the form $!(ϕ ∨ ψ)$ is to make an agent recognize a requirement that $(ϕ ∨ ψ)$. But if that is indeed the force of (4.59), then we would expect that the force of an imperative of this form is exhausted by making $(ϕ ∨ ψ)$ designated and required.

This, though, does not seem to be right: (4.59) does not generally function to make just posting or burning the letter designated (hence expressly permitted; see Section 4.3.8). It functions to make posting the letter expressly permitted, and burning the letter expressly permitted.

Permission Content and the Ross Paradox. The well-known Ross Paradox (Ross 1941) grows out of the felt invalidity of the following inference schema:

(4.60) $!\phi / \therefore !(\phi \vee \psi)$
(Post the letter! doesn't entail Post the letter or burn it!)

The claimed invalidity of this inference (as well as its claimed paradoxicalness) is a complicated matter, in large part because the semantic relationship that is supposed to be captured by a validity relation for an imperative language (in contrast with the validity relation for a propositional language) is a deeply contested matter.⁵¹ I largely avoid the question here.

It is, however, important to avoid the presupposition, implicit in the standard formulation of the paradox, that the felt invalidity is an intuition *about* licensed and unlicensed inferences *in* the regimented metalanguage. This takes for granted something that must actually be argued—that $!(\phi \vee \psi)$ is an appropriate logical form for a natural language disjunctive imperative, and that an intuition about whether inferring a disjunctive imperative from a non-disjunctive imperative in natural language is licensed counts equally as an intuition about whether inferring something of the form $!(\phi \vee \psi)$ from something of the form $!\phi$ is licensed.⁵²

This presupposition often, I think, arises from a more general mistake about the permission content of imperatives. Consider the following inference, which Castañeda (1958: 43–4) claims to be invalid (see also Vranas 2010: note 13).

(4.61) See to it that: if you read the book, you see me. $!(\phi \supset \psi)$
 Read the book! $!\phi$
 So, come see me! $!\psi$

About this inference, Castañeda (1958) notes:

[A] teacher who [issues the premises of 4.61] has not thereby ordered or told his student to come to see him, *regardless of the student's reading of the book.*

The idea being deployed here seems to be that meaning of any imperative $!\chi$ is accurately represented as *see to it that χ , regardless of whether χ' , for any χ' consistent with χ* . Notice that the *order χ regardless of χ'* locution expresses that both $(\chi \wedge \chi')$ and $(\chi \wedge \neg\chi')$ are permitted (when possible), and at least one required. On this picture of permission content, it's natural to think that an imperative of the form $!\psi$ does not follow from imperatives $!\phi$ and $!(\phi \supset \psi)$: the latter jointly prohibit a state of affairs that the former permits, namely, $(\neg\phi \wedge \psi)$.

How does this idea lead to the presupposition that the Ross Paradox is a puzzle about regimented imperatives of the form $!(\phi \vee \psi)$, rather than about natural language proper? On this idea, an imperative of the form $!(\phi \vee \psi)$ requests $(\phi \vee \psi)$ regardless of ϕ and

51. But ultimately, I think, a terminological matter. See my (2009) for discussion.

52. It is hard to believe that so many deontic logicians and semanticists take this for granted. But they do.

regardless of ψ , so that all of $(\phi \wedge \psi)$, $(\neg\phi \wedge \psi)$, and $(\phi \wedge \neg\psi)$ are permitted, if possible. It is, then, natural to think that an imperative of the form $!(\phi \vee \psi)$ does not follow from an imperative of the form $!\phi$: the latter prohibits a state of affairs that the former permits, namely, $(\neg\phi \wedge \psi)$.

The picture of permission content involves what seems to me to be a basic conceptual confusion: the conflation of failure to expressly forbid with express permission. It is important to distinguish what an imperative weakly permits (namely, all of the things with which it is consistent) with what permissions it actually functions to generate. Express permission, on our analysis, is a function of *designation*, rather than mere consistency with the thing requested. There is no reason to think, of an imperative $!\phi$, that every state of affairs consistent with ϕ is designated. On our analysis, in fact, it is just ϕ itself that is designated.

The Ross Paradox: Neutrally Formulated. To avoid this presupposition, I will formulate the Ross Paradox solely as a puzzle about natural language disjunctive imperatives. For our purposes, a disjunctive imperative can be characterized roughly as follows: a natural language imperative is disjunctive if its surface syntax is of the form $!\phi$ or $!\psi$. The Ross Paradox then receives the following formulation:

THE ROSS PARADOX: NEUTRALLY FORMULATED

A disjunctive imperative of surface form $!\phi$ or $!\psi$ does not follow from an imperative of surface form $!\phi$.

As I suggest above, I will opt for explaining the felt invalidity of the inference by appeal to the *permissive content* of disjunctive imperatives, on their most natural interpretations: (4.59) does not follow, intuitively, from *post the letter!* because, although a state that requires that you post the letter also requires that you post or burn the letter, a state that requires that you post the letter does not permit you to burn the letter.⁵³ And (4.59), on its most natural reading, permits the agent to burn the letter. Implicitly, of course, I am relying on something like the following conception of when a sentence or speech act follows from another:⁵⁴

Definition 4.59. $\Delta_1\phi_1, \dots, \Delta_n\phi_n$ **dynamically entails** $\Delta_k\phi_k$ (notation: $\Delta_1\phi_1, \dots, \Delta_n\phi_n \models \Delta_k\phi_k$) iff, for any state σ , $\sigma[\Delta_1\phi_1] \dots [\Delta_n\phi_n] = \sigma[\Delta_1\phi_1] \dots [\Delta_n\phi_n][\Delta_k\phi_k]$.

So (4.59) does not follow from its first disjunct because updating a state with (4.59) is not redundant, conditional on updating the state with its first disjunct. Updating with (4.59), although it will not introduce any new requirements into the state, intuitively should modify a state so that it comes to permit burning the letter (something that has been forbidden at

53. This approach to the Ross Paradox is also endorsed by Åqvist (1965); Aloni (2007); Veltman (2008); Charlow (2009), among others.

54. This is just the standard notion of dynamic entailment, familiar from such classic works in Dynamic Semantics as Kamp (1981); Groenendijk & Stokhof (1991); Veltman (1996).

the prior state). Though the Ross Paradox is often treated simply as a failure of some sort of closure of requesting under logical consequence (Hansson 1990; Kenny 1966; Segerberg 1990), a satisfying treatment requires more, namely, *an account of the permissive force of disjunctive imperatives*.

More specifically, we want an account of permissive force that has disjunctive imperatives partaking, in some sense—whether via presupposition, implicature, or conventional content—in *free choice permissions* (as well an account of free choice permissions which predicts the relevant implications). The rough idea is that, in addition to expressing a requirement (to post or burn the letter), (4.59) expresses, in some sense, a free choice permission, of the sort indicated in (4.62a).

- (4.62) a. You can post or burn it (I don't care which)
 \Rightarrow You can burn it
 b. You can post or burn it (I'm unsure which)
 $\not\Rightarrow$ You can burn it

It bears noting that there is, so far anyway, nothing non-Standard here. Necessitation acts, we've assumed, involve designation acts, which we've understood in terms of explicit permissions. So, if the logical form of a disjunctive imperative *is* given as $!(\phi \vee \psi)$, then a disjunctive imperative expresses an explicit permission with content $(\phi \vee \psi)$. And so long as this explicit permission is analyzed as a free choice permission, à la (4.62a), the account as we've developed it would be, contrary to initial appearances, well-suited to explaining the typical force of a disjunctive imperative.

In what sense does a disjunctive imperative partake in a free choice permission? A little bit more carefully, in what sense does an imperative of surface form $!\phi$ or $!\psi$ both permit ϕ and permit ψ ? The familiar menu of options includes: by presupposition, implicature, or conventional encoding. We'll survey each option briefly (though we'll ultimately opt for the last).

Presupposition? Williams (1963) was perhaps ahead of his time in treating the relationship between an imperative of surface form $!\phi$ or $!\psi$ and the permission that ψ as presuppositional (he refers to this permission as a “permissive presupposition” of the disjunctive imperative). But is there any evidence for typing this permission as a presupposition?

Unfortunately, no. Embedding a disjunctive imperative in the consequent of a conditional destroys the permissive implication; unlike a presupposition, then, it fails to project.

- (4.63) If Bob insists, then post or burn the letter!
 $\not\Rightarrow$ You may burn the letter

And when the authority presupposition (Section 4.3.1) of the imperative is satisfied, challenging the permission implication of a disjunctive imperative with a “Hey, wait a minute!”

echo, as in (4.64), is infelicitous.

- (4.64) A: Post or burn the letter!
#B: Hey, wait a minute! I'm not allowed to burn the letter!

When the authority presupposition is satisfied, A's utterance serves to grant B permission to burn the letter. Interpreting the utterance does not require that the permission is already in place.

Implicature? What about typing the permissions of disjunctive imperatives as *quantity implicatures* (Grice 1989), as suggested by Bennett (1970); Hare (1967)?⁵⁵

Quantity implicatures, as we noted in Section 3.2.1, do not involve speaker commitment; they are felicitously cancellable (cf. 4.65), and merely withdrawing a quantity implicature fails to sanction withdrawal of the statement itself (cf. 4.66).

- (4.65) Some, but not all, swans are white.
(4.66) A: Some whales are mammals.
B: Actually all are.
A: I didn't say they weren't! (#You're right. I was wrong. It's not the case that some whales are mammals.)

Are permission implications of disjunctive commands cancellable? It would appear not.

- (4.67) #Post or burn the letter, but don't burn it!
(4.68) A: Use shears or a chainsaw to cut the branch.
B: I'm not allowed to use a chainsaw!
A: My mistake. Use the shears. (#Noted, but the instruction still stands: use shears or a chainsaw.)

(4.67) indicates that there is something fishy about conjoining an injunction to a disjunctive command, such that the former conflicts with the permission implication of the latter. That is unexpected, if the permission implication is a quantity implicature. (4.68) indicates that denial of a disjunctive command's permission implications sanctions (indeed, requires) the withdrawal of the disjunctive command.

The latter piece of data is especially telling; quantity implicatures, especially those

55. Typing these implications as quantity implicatures commits one to specifying a plausible, Gricean way of deriving them, involving some sort of non-informational (practical) analogue of Grice (1989)'s Maxim of Quantity. For sophisticated recent attempts to derive free-choice implications in roughly this way, see Fox (2006); Kratzer & Shimoyama (2002); Schulz (2003, 2005). (I don't, by the way, want to claim that a sophisticated fan of neo-Gricean treatments of free choice implicatures should be persuaded by what I say here. Such a fan can read me as simply taking it for granted that free choice implications are, in some sense, entailments rather than implicatures.)

arising from disjunctions, sometimes give rise to weirdness when they are cancelled.

(4.69)??He's in the shed or the field, but he's not in the field.

But rejecting a quantity implicature fails even to license (to say nothing of *requiring*) withdrawal of the original utterance.

(4.70) A: He's in the shed or the field.

B: Uh, he's not in the field.

A: Well, I didn't say that he was! (#I guess I was wrong.)

Tentatively, these data seem to suggest that the permission in question meets the non-cancelability condition on conventionalization articulated in Section 3.2.1.

Conventional content? In light of such data, it is reasonable to suppose that there is a conventionalized distinction between *choice-offering* and *alternative-presenting* disjunctive imperatives (the terminology is from Åqvist 1965). Choice-offering disjunctive imperatives conventionally express that the prejacent of each disjunct are permitted; they conventionally express free choice permissions. Alternative-presenting disjunctive imperatives do not. (Both, however, express that at least one of the prejacent is required.) We shall treat the permissions of choice-offering interpretations as a conventionally proffered element of whatever speech act is conventionally expressed by a choice-offering disjunctive imperative.⁵⁶

There is the further matter of *how to account for these permissions*, or, more precisely, what logical forms to assign choice-offering and alternative-presenting readings of disjunctive imperatives. Alternative-presenting readings are easy: they may be represented with speech act designators of the form $!(\phi \vee \psi)$. On the analysis we've been developing, this means that alternative-presenting disjunctive imperatives simply designate and make required a disjunctive state of affairs $(\phi \vee \psi)$.

Choice-offering readings are more complicated. We will consider the various options below. Surprisingly, though, if Åqvist's distinction between choice-offering and alternative-presenting readings of disjunctive imperatives in natural language is right, there is reason to suspect that the Ross Paradox is actually ill-formulated. We have intuitions about the environments introduced by natural language connectives *masquerading* as intuitions about the interpretation of a formal imperative language (which lacks such environments).⁵⁷

56. This in turn goes naturally with a semantic resolution of the paradox of free choice permission. For some options, see Aloni (2007); Geurts (2005); Mastop (2005); Veltman (2008); Zimmermann (2000).

57. Similar methodological remarks apply mutatis mutandis to the Ross Paradox for deontic necessity operators, the paradox of free choice permission, and related paradoxes. None of these are, I think, properly thought of as paradoxes about the interaction of modal or quasi-modal operators with \vee (although they are almost always formulated as such). They are all, rather, properly thought of as paradoxes about the interaction of natural language operators with natural language *or*.

Faithfully representing the meaning of natural language disjunctive imperatives, then, does require grappling with a Ross-like Paradox—just not as traditionally formulated (since we are not forced to locate the relevant ambiguity in a formal imperative language). That is to say: it does not require predicting that $!(\phi \vee \psi)$ expresses permissions that $!\phi$ fails to express. Instead, it requires:

- Supplying appropriate LFs for choice-offering and alternative-presenting readings of disjunctive imperatives in natural language
- Stating conditions on when the choice-offering LF is used in the analysis of a natural language disjunctive imperative (something on which we will not have a great deal to say here)

We'll now tackle how to supply LFs for choice-offering disjunctive imperatives.

Defeasibly Preferring Both? One *prima facie* attractive choice, defended both by [Portner \(2010\)](#); [Starr \(2011\)](#), is that updating with a choice-offering imperative of the form $!\phi$ or $!\psi$ involves hypothetically updating the prior state separately with $!\phi$ and $!\psi$, and obtaining the updated state by somehow *merging* these states. In both accounts, the force of the merging operation on the state's preferences is given by \cup . (In presenting their analyses, I assign choice-offering imperatives logical forms of the form $!\phi \vee !\psi$ since both analyses basically involve making use of the ordinary Dynamic Semantic understanding of \vee .)

$$\pi_{\sigma[!\phi \vee !\psi]} = \pi_{\sigma[!\phi]} \cup \pi_{\sigma[!\psi]}$$

For [Portner](#), this means that updating with an imperative $!\phi \vee !\psi$ leads the state to defeasibly prefer both ϕ and ψ . [Starr's](#) analysis, although implemented in a different formal system, implements basically the same idea: the state comes to strictly prefer ϕ -possibilities to $(\neg\phi \wedge \neg\psi)$ -possibilities, as well as to prefer ψ -possibilities to $(\neg\phi \wedge \neg\psi)$ possibilities, but to have no preference between ϕ and ψ -possibilities.

The similarity of these analyses results in their sharing some faults. Choice-offering imperatives simply do not institute preferences, even defeasible preferences, for each of their disjuncts. The agent is given a *choice*: she may either plan to realize one disjunct (choosing fish) or plan to realize the other (choosing fish), and must do at least one.

(4.71) Please choose fish or chicken

The sort of preference state instituted by the [Portner-Starr](#) analysis simply does not correctly represent this state of mind: the agent comes to prefer both fish *and* chicken to nothing (but neither fish nor chicken to each other), and, in the typical event that they are not jointly realizable (or their joint realization is prohibited), is forced, presumably with some reluctance, to choose only one. This is a state of mind more suited to representing preferential *dilemmas*

(and its attendant cognitive phenomena, like regret, which would not be suitable responses to being confronted with a choice-offering imperative), rather than preferential choices.

To bring out why this is so implausible, consider a choice-offering imperative that is non-typical, in that its disjuncts *are* jointly realizable and their joint realization is *not* prohibited.

(4.72) For your hangover, take Advil or take Alka Seltzer

On the [Portner-Starr](#) analysis, updating with (4.72) leads the agent to prefer Advil and to prefer Alka Seltzer. And, lo, nothing prevents or prohibits her from taking both! And so the agent plans to take both. In other words, when a choice-offering imperative's disjuncts are jointly realizable, the [Portner-Starr](#) analysis holds that they function to create commitments to pursue the realization of both disjuncts, i.e., they have the same force as the conjunctive imperative in (4.73).⁵⁸

(4.73) For your hangover, take both Advil and Alka Seltzer

Contrary to the predictions of the [Portner-Starr](#) analysis, choice-offering imperatives are, so far as I can tell, *never* put to this sort of use. This is not to say that an agent is irrational to respond to (4.72) by planning to take both. It is just to say that the imperative, given the context, clearly does not require taking both.

(4.74) A: For your hangover, take Advil or take Alka Seltzer

(*B takes just the Advil*)

#A: Hey, why didn't you take my advice?

[Portner \(2010\)](#) notes a fact about disjunctions in natural language that might seem to resolve this difficulty. Disjunctions, he claims (following work within the Hamblin tradition, e.g., [Alonso-Ovalle 2006](#); [Aloni 2007](#)), introduce alternatives. "Crucially, these alternatives are typically, and perhaps always, exclusive" (2010: 15). In other words, the claimed reading of (4.72) is actually not available at all. According to [Portner's](#) analysis, accepting (4.72) is a matter of adding the proposition that one takes Advil (but not Alka Seltzer) as well as the proposition that one takes Alka Seltzer (but not Advil) to one's preferences. Since these propositions are mutually exclusive, the resulting body of preferences, according to [Portner \(pc\)](#), exhibits the following features:

58. This isn't 100% fair to [Starr](#). According to him, (4.73) leads to a preference state in which taking both is strictly preferred to not taking both, while (4.72) leads to a state which prefers taking Advil to nothing, and also prefers taking Alka Seltzer to nothing. But, of course, if faced with a choice between satisfying one preference (e.g., for Advil) and satisfying two, an agent is rationally obligated to select the latter. (Indeed, this is the exact intuition that lies behind the standard conception of the relationship between bodies of preferences and preference orderings embodied in, e.g., [Definition 4.25](#).) That is incorrect—an agent that accepts (4.72) is in no sense thereby rationally obligated to take both.

- Possibilities where one takes Advil (but not Alka Seltzer) are incomparable to those where one takes Alka Seltzer (but not Advil)
- Possibilities where one takes both Advil and Alka Seltzer (or neither Advil nor Alka Seltzer) are strictly dispreferred to possibilities where one takes Advil (but not Alka Seltzer)
- Possibilities where one takes both Advil and Alka Seltzer (or neither Advil nor Alka Seltzer) are strictly dispreferred to possibilities where one takes Alka Seltzer (but not Advil)

In short, accepting (4.72) does not commit an agent to pursuing the realization of both disjuncts.

In reply, I first reiterate my objection that this is committed to a misrepresentation of the phenomenology associated with the acceptance of a disjunctive imperative. Treating possibilities where one takes Advil (but not Alka Seltzer) as incomparable to those where one takes Alka Seltzer (but not Advil) is symptomatic of *indecision* between two valuable, but incomparably so, alternatives (hence also associated with psychological phenomena like regret that are associated with the choice of one or other of incomparable alternatives, a reflex of the agent's inability to reconcile two incomparable alternatives with respect to a common measure of value). That is likely incorrect. Even allowing that the alternatives induced *are* exclusive, it would be better, *ceteris paribus*, to treat the state of mind induced by accepting (4.72) as a state in which the alternatives are treated, not as *individually valuable* and *jointly incomparable*, but rather as simply available as planning options for the agent (who is indifferent, *ceteris paribus*, between them). In typical cases, such a state will be decidedly indifferent between these two alternatives, rather than regard them as incommensurable. It is a misunderstanding of the history and representational function of preference orderings (and their associated choice functions) to assimilate incomparability or incommensurability between *p*- and *q*-possibilities to decided indifference between *p*- and *q*-possibilities.⁵⁹

Moreover, contrary to the predictions of an analysis of disjunctions as introducing exclusive alternatives, I also doubt that an agent who accepts a disjunctive imperative like (4.72) is to be faulted, whether relative to the speaker's communicative intentions or relative to her own preferences, if she takes *both* Advil and Alka Seltzer. Suppose an agent who has accepted (4.72) does indeed take both. It seems rather infelicitous for the speaker to object to the speaker's action as follows (unless, e.g., it is understood that the drugs, taken together, will have an unpleasant interaction):

(4.75)??Hey, you're not supposed to do that!

59. Indeed, it is a mistake very much like the conflation of incommensurability (lack of comparability with respect to a common measure of value) with indifference so well-known in the economic and decision-theoretic literature. For relevant discussion, see, e.g., Broome (2000).

This constitutes evidence that the typical interpretation of choice-offering disjunctive imperatives as introducing exclusive alternatives (hence, as conventionally introducing a preference state in which satisfying one of the disjuncts is generally strictly preferred to satisfying both) is a pragmatic phenomenon, rather than a reflex of the conventional meaning of disjunction. That is a problem for Portner’s analysis of choice-offering disjunctive imperatives. For the analysis to be sufficiently general, and for it to avoid the objection that an imperative like (4.73) can be used to introduce a conjunctive requirement, it must not just be the case that the alternatives introduced by such imperatives are “*typically, and perhaps always, exclusive.*” It must further be the case that the alternatives introduced by such imperatives are *always* exclusive, as a matter of the conventional meaning of disjunctions in natural language. That, as I think this example well illustrates, is implausible.⁶⁰

Underdetermination and Alternative States? Instead, I’ll suggest that choice-offering readings are ripe for representation using the regular operation ‘+’ (Section 4.4.3).⁶¹

$$[!\phi+!\psi] = [!\phi] \cup [!\psi]$$

In effect, the idea here is to introduce something like *disjunction at the level of speech acts*—to have speech acts scoping, in essence, under disjunction. And, indeed, the effect of updating a state with a speech act designator of the form $\Delta\phi + \Delta\psi$ can be glossed metalinguistically using disjunction: for any state σ for which $[\Delta\phi]$ and $[\Delta\psi]$ are defined, $\sigma[\Delta\phi + \Delta\psi]$ is either $\sigma[\Delta\phi]$ or $\sigma[\Delta\psi]$.

Of course, a cognitive state attempting to update with a speech act designator of the form $\Delta\phi + \Delta\psi$ cannot rest at this; the update expressed by $\Delta\phi + \Delta\psi$ is not strictly a function, but a *relation*. For any state σ for which both $[\Delta\phi]$ and $[\Delta\psi]$ are defined, it is typically the case that $\sigma[\Delta\phi] \neq \sigma[\Delta\psi]$, and so there is no unique σ' such that $\sigma' = \sigma[!\phi+!\psi]$. In other words, the output state is *underdetermined* by the speech act itself.

Definition 4.60. An update is **underdetermined** relative to a speech act α and input state σ iff $\exists\sigma' \exists\sigma'' : \langle\sigma, \sigma'\rangle \in \alpha$ and $\langle\sigma, \sigma''\rangle \in \alpha$ and $\sigma' \neq \sigma''$.

A state that accepts a speech act that underdetermines an update is confronted with a “choice” (either side of which is available, given the cognitive constraint proposed by $\Delta\phi + \Delta\psi$).⁶²

60. This is not altogether surprising, given a recent trend in theorizing about disjunctions on which they introduce alternatives, but on which the alternatives are not necessarily incompatible (cf. Groenendijk & Roelofsen 2009 and related work in the Inquisitive Semantics program).

61. For proposals that seem to be in a similar sort of vein (though it is hard to tell for sure, given how quickly they are presented) see Krifka (2004); Veltman (2008).

62. I used to think (Charlow 2010a) that this sort of underdetermination was problematic, as it meant introducing unwelcome *uncertainty* or *indeterminacy* into the realm of speech acts. An utterance whose force is to update with α or β would, I argued, *commit* the speaker to α or β , depending on which update was executed, but the speaker would not have any idea which. There is ample evidence that disjunctive speech acts in *this*

- Option 1: update with $!\phi$ (decide on $\delta\phi$)
- Option 2: update with $!\psi$ (decide on $\delta\psi$)

This proposal has a considerable amount of intuitive appeal: choice-offering disjunctive imperatives of the form $!\phi$ or $!\psi$ intuitively propose that the agent either decide on $\delta\phi$ or $\delta\psi$, while allowing either.

So far so good. But we are still in need of a story about what the permission-granting force of choice-offering disjunctive imperatives consists in. For, in typical cases (where the disjuncts are either not jointly realizable or their joint realization is prohibited), a state that actually executes an update in response to an imperative of the form $!\phi+!\psi$ eventually ends up either (i) deciding on $\delta\phi$ and deciding against $\delta\psi$ or (ii) deciding on $\delta\psi$ and deciding against $\delta\phi$. In case (i), only ϕ is permitted (because required). In case (ii), only ψ is permitted (because required).

Our story here is a relatively simple one. The sort of permission granted by choice-offering disjunctive imperatives consists in the *introduction of alternative allowed planning states*. What it is for ϕ to be permitted by a choice-offering disjunctive imperative of the form $!\phi$ or $!\psi$ is for its regimented representation (i.e., $!\phi+!\psi$) to allow a transition to a planning state in which $\delta\phi$ is decided on. Relative to an input state σ , an imperative of the form $!\phi+!\psi$ introduces the following set of states as *allowed alternatives*:⁶³

$$\{\sigma' : \sigma[!\phi+!\psi]\sigma'\}$$

More generally:

Definition 4.61. $\langle\sigma, \sigma'\rangle$ is an **allowed transition** relative to a speech act α iff $\langle\sigma, \sigma'\rangle \in \alpha$.

Definition 4.62. σ' is an **allowed alternative** relative to an input state σ and a speech act α iff $\langle\sigma, \sigma'\rangle$ is an allowed transition relative to α .

These notions are important, as they will allow us to represent the state of a cognitive state that has accepted a speech act that underdetermines an update, but has not decided which allowed transition to execute (i.e., has not decided which allowed alternative to select). Such a state is represented as a *superposition of allowed alternatives*.

sense are not expressible in natural language.

(4.76) #Post or burn the letter (I don't know which)

(4.77) #You may hereby post or burn the letter (I don't know which)

Oddly, Hamblin (1987); Hare (1967: 315–6); Krifka (2004) seem to be open to the claim that constructions like (4.76) and (4.77) are felicitous. I dispute that (cf. Merin 1991). The analysis developed here is designed to get around this sort of worry, by understanding underdetermination in such a way that it does not involve unwelcome uncertainty or indeterminacy.

63. Krifka (2004) briefly suggests a similar sort of proposal, which he glosses in terms of “alternative commitment states.” He is not, as far as I can tell, interested in using it to give an account of the Ross Paradox (although he may be interested in using it to give an account of the paradox of free choice permission). It is hard to say for sure, given the brevity of his notes.

Indeed, it is useful to think of all cognitive states as superpositions of allowed alternatives relative to a range of allowed cognitive actions. In general, such allowed alternatives bear the following relation to their input state: they are either *more informationally decided* or *more practically decided* than their input states.⁶⁴ We are, in other words, adopting the idealizing assumption, common to the sorts of theories from which our account draws inspiration, that cognitive transitions, whether in response to speech acts or spontaneous changes in attitude, are generally *monotonic* in nature. That is to say, for any cognitive action α , $\langle \sigma, \sigma' \rangle \in \alpha$ only if σ' is at least as decided as σ . Here are some formal definitions to fill out this informal skeleton.

Definition 4.63. σ' is at least as **informationally decided** as σ iff $\bigcup I_{\sigma'} \subseteq \bigcup I_{\sigma}$ and every possibility that is salient at σ (in the sense of Definition 4.18) is also salient at σ' .

Definition 4.64. σ' is at least as **practically decided** as σ iff $\bigcup \boxplus_{\sigma'} \subseteq \bigcup \boxplus_{\sigma}$ and $\forall p \in \omega_{\sigma} : \exists p' \in \omega_{\sigma'} : p \subseteq p'$.

Definition 4.65. σ' is at least as **decided** as σ iff σ' is at least as practically and informationally decided as σ .

It follows that making a state strictly more informationally decided means either recognizing a new issue as salient or no longer considering some possibility to be compatible with one's information. Making a state strictly more practically decided means either adding to that state's preferences (thereby shrinking the permissibility sphere) or designating new practical alternatives. Making a state strictly more decided means either making the state strictly more informationally or practically decided.

We'll not have anything general to say about what sorts of spontaneous changes in attitude are allowed (nor, therefore, about what the allowed alternatives relative to an input state and a range of allowed spontaneous changes in attitude) are like. We will simply assume that a space of allowed alternatives relative to an input state is given. We will confine ourselves to saying that one thing these allowed alternative will have in common is that they are strictly more decided than their inputs.⁶⁵ Even given such a quietist attitude, a new, and satisfyingly more general, picture of the nature of speech acts can be seen to emerge. *What speech acts do is restrict the space of allowed alternatives relative to an input state.*

Here is an informal sketch of the idea as it applies to choice-offering disjunctive imperatives (to be followed with a more careful presentation). Consider a typical state prior to an utterance of (4.71). This state has no specific plans about what to select for dinner: for all the agent has planned, fish, chicken, beef, and pork are all consistent with the agent's standing

64. Gibbard (2003)'s notion of a hyperdecided state has been central to shaping my thinking on these matters.

65. Certain things are fairly obvious. For instance, relative to a state in which $\{\delta\phi, \bar{\delta}\phi\}$ is a designated alternative, a spontaneous change in attitude whereby the agent comes to decide on $\delta\phi$ (thereby eliminating $\bar{\delta}\phi$) is allowed; agents are allowed to make practical decisions. Relative to a state in which $\delta\phi$ is non-designated, a spontaneous change in attitude whereby the agent comes to designate $\delta\phi$ (thereby raising the practical question whether or not to see to it that ϕ) is also allowed; agents are allowed to raise practical questions.

preferences (although none have been designated, and so none are explicitly permitted). That is to say, the state is a superposition of these strictly more decided states, relative to a space of allowed changes in attitude. How does such a state update in response to accepting the speech act expressed by (4.71)? It responds by eliminating elements of this superposition. Letting σ be the input state, the resulting state is a state in which the allowed alternatives are $\sigma[!fish]$ and $\sigma[!chicken]$ (rather than an alternative superposition where the allowed alternatives are $\sigma[!fish]$, $\sigma[!chicken]$, $\sigma[!beef]$ and $\sigma[!pork]$). What it means for (4.71) to explicitly permit both selecting fish and selecting chicken is that it makes these alternatives available; each is such that *there is an allowed alternative* in which they are decided on. What it means for (4.71) to explicitly require at least one is that there is no allowed alternative in which neither fish nor chicken is selected.

It would be cumbersome to revise our system to adequately represent this new notion of a state (as superposition), and we'll opt against it. For our purposes, anyway, we can get all the mileage of this idea by simply inductively generalizing the notion of a state and the notion of a speech act, as follows.⁶⁶

Definition 4.66. $\Sigma = \{\sigma\} = \{\langle I_\sigma, \Pi_\sigma \rangle\}$ is a **cognitive state** if I_σ is an information state for σ and Π_σ is an action-guiding state for σ . If Σ and Σ' are cognitive states, then $\Sigma \cup \Sigma'$ is a cognitive state. Nothing else is a cognitive state.

Definition 4.67. $\Sigma[\Delta\phi] = \bigcup_{\sigma \in \Sigma} \{\sigma' : \sigma[\Delta\phi]\sigma'\}$

The specific sense in which a choice-offering disjunctive imperative expresses a free choice permission is given, then, by the following definitions.

Definition 4.68. ϕ is **expressly permitted** relative to Σ iff $\exists \sigma' \in \Sigma : \phi$ is a recognized permission at σ' .⁶⁷ (Definition 4.52 is captured a special case—the case where $\Sigma = \{\sigma\}$.)

Definition 4.69. An imperative Φ expressly permits ϕ relative to Σ iff, if $\Sigma[\Phi]$ is defined, then ϕ is expressly permitted relative to Σ .

Definition 4.70. An imperative Φ expressly permits ϕ iff $\forall \Sigma : \text{if } \Sigma[\Phi] \text{ is defined, then } \phi \text{ is expressly permitted relative to } \Sigma$.

Other applications. Once we have the regular operation '+' (and, less abstractly, the notion of speech act disjunction or relationalization) at our disposal, its potential utility in the analysis of natural language disjunctions (and disjunction-like elements, like the

66. Cf. again Krifka (2004) (although, again, he does not ultimately seem to favor an approach in this vein).

67. There is, interestingly, a similar idea embodied in the static semantic analysis of Zimmermann (2000). Zimmermann treats disjunctions as "allowing" (in the sense of expressing \diamond for) both of their disjuncts. In the case where the disjuncts are themselves permissions (perhaps of the form $\diamond\phi$ and $\diamond\psi$), the fact that the disjunction permits both ϕ and ψ follows from the fact that $\diamond\diamond\phi$ (it is permitted that it is permitted that ϕ) entails $\diamond\phi$. Similarly, in our analysis, we have a state permitting that ϕ if there is an allowed transition that terminates in a state that permits ϕ .

negative polarity item *any*) is not hard to see. Free choice phenomena have, of course, been identified for natural language modals of all types.

- (4.78) a. It might be raining $\not\Rightarrow$ It might be raining or sunny
 b. It might be raining or sunny \Rightarrow It might be raining
- (4.79) a. You may have a pear $\not\Rightarrow$ You may have a pear or an apple
 b. You may have a pear or an apple \Rightarrow You may have a pear
- (4.80) a. You should post the letter $\not\Rightarrow$ You should post or burn the letter
 b. You should post or burn the letter \Rightarrow You may post the letter

Our explanations of all of these phenomena are, more or less, automatic. To illustrate, consider the free choice readings indicated (B.48). We'll analyze *it might be raining or sunny* using an LF something like $might(raining) + might(sunny)$. We'll treat epistemic possibility modals as conventionally expressing a property of a cognitive state—the property of treating the prejacent as a salient possibility (cf. Swanson 2006; Yalcin 2007). So $might(raining) + might(sunny)$ conventionally expresses the property a state Σ has just if $\exists\sigma : \sigma[might(raining)] \in \Sigma$ and $\exists\sigma : \sigma[might(sunny)] \in \Sigma$. With the following set of definitions, the phenomena in (B.48) are predicted.

Definition 4.71. ϕ is a **salient possibility** relative to Σ iff $\exists\sigma' \in \Sigma : \phi$ is a salient possibility (in the sense of Definition 4.18) at σ' .

Definition 4.72. Φ **treats as possible** ϕ iff $\forall\Sigma : \text{if } \Sigma[\Phi] \text{ is defined, then } \phi \text{ is a salient possibility relative to } \Sigma$.

What about practical necessity modals (and requirements generally)? An idea (one that we will develop extensively in the Appendices) is that these, like imperatives, express a property of a cognitive state—the property, roughly, of treating the prejacent as required. So $should(post) + should(burn)$ conventionally expresses the property a state Σ has just if $\exists\sigma : \sigma[should(post)] \in \Sigma$ and $\exists\sigma : \sigma[should(burn)] \in \Sigma$. With the following set of definitions, the phenomena in (4.80) are straightforwardly predicted.

Definition 4.73. ϕ is a **treated as permitted** relative to Σ iff ϕ is expressly permitted (in the sense of Definition 4.68) relative to Σ .

Definition 4.74. ϕ is a **treated as required** relative to Σ iff $\forall\sigma' \in \Sigma : \phi$ is a recognized requirement (in the sense of Definition 4.49) at σ' .

Definition 4.75. Φ **treats as permitted** ϕ iff $\forall\Sigma : \text{if } \Sigma[\Phi] \text{ is defined, then } \phi \text{ is treated as permitted relative to } \Sigma$.

It is easy to see how to extend this method to a general account of free choice phenomena. The sort of quantification generally expressed by permission or possibility sentences is of

the form $\exists\exists$; ϕ is expressly possible, in some sense, relative to Σ iff there is some $\sigma \in \Sigma$ relative to which ϕ is expressly possible. The sort of quantification generally expressed by requirement or necessity sentences is of the form $\forall\forall$; ϕ is expressly required, in some sense, relative to Σ iff, for every $\sigma \in \Sigma$, ϕ is expressly required relative to σ . Permission sentences, in general, express freedom of cognitive transition (and freedom relative to a given cognitive transition); requirement sentences, in general, express constraints in cognitive transition (and constraints relative to a given cognitive transition). With this apparatus and the regular operation ‘+’ at our disposal, an account of an arbitrary kind of modal free choice is all but automatic.

Summing up. Once again, we have identified what appears to be an essentially complex speech act—generically describable as a free choice possibility—with a canonical means of expression in natural language. And so we have gradually built a case for incorporating tools allowing the recursive construction of complex speech acts—namely, the regular operations ‘;’ and ‘+’—into our theory of speech acts.

4.4.7 Quantifiers?

In this section, we will briefly consider some *quantificational* analogues of these regular operations. Specifically, we will argue for the existence of quantification-like devices with canonical modes of expression in natural language, and which exhibit the following features:

- Essentially taking scope over speech act designators
- Binding variables occurring within a speech act designator
- Definable as quantification-like generalizations of the regular operations ‘;’ and ‘+’

Krifka on pair-list questions. Evidence for a kind of “quantified” illocutionary force is given in recent work by Krifka (2001a, 2004) on the conventional meaning of questions. Questions of the surface form in (4.81) admit of so-called narrow-scope, functional, and pair-list readings, indicated semi-formally in (4.81a), (4.81b), and (4.81c), respectively.⁶⁸

(4.81) Which x : all y : $R(x, y)$? (E.g., Which dish did each chef prepare?)


- a. $?x : \forall y : R(y, x)$ (Narrow-scope LF)
- b. $?f : \forall x : R(x, f(x))$ (Functional LF)
- c. $\forall x : ?y : R(x, y)$ (Pair-list LF)

68. Note: we’re now assuming that some speech-act operators can bind variables in open formulae. This doesn’t, in my mind, count as a deviation from the Standard Account. The Standard Account denies complexity in speech acts, but it need not deny complexity in speech act *operators*—e.g., the sort of complexity that would allow them to introduce and bind a variable in the proposition over which the operator takes scope.

The force of the narrow-scope reading is to query the identity of a dish (e.g. pasta) satisfying the property in question (being prepared by each chef). The force of the functional reading is to query the identity of a function f (e.g., x 's favorite dish) such that $f(x)$ satisfies the property in question. Nothing is mysterious about either type of interrogative force. A *single question* is intuitively asked, and interrogative force adjoins, in either case, to an open first-order formula.



Not so, it would seem, for the pair-list reading. As [Krifka](#) notes, the force the pair-list reading seems to be to query, *for each chef*, the identity of the dish that chef made.



It is natural to interpret this sort of quantification as a *generalized kind of speech act sequencing*. As a universal first-order formula $\forall x : \phi$ is equivalent to the conjunction of ϕ 's substitution instances (assuming a surjective map from constants to finitely many individuals), so the force of the pair-list reading of $\forall x : ?y : \phi$ can be intuitively regarded as equivalent to the “conjunction” of the substitution instances of $?y : \phi$ (i.e., the result of taking the substitution instances and sequencing them with the regular operation ‘;’) (cf. [Krifka 2004: 3-4](#)).

Formally, then, letting $A = \{a_1, \dots, a_n\}$ be a finite list of constants, the following interpretation seems reasonable for a speech act designator of the form $\forall x : ?y : R(x, y)$.

$$[\forall x : ?y : R(x, y)] = [\bigwedge_i \text{sub}^{a_i/x} (?y : R(x, y))] = [?yR(a_1, y)] \circ \dots \circ [?yR(a_n, y)]$$

In sum, pair-list readings provide good evidence for recognizing a quantification-like operation that can take scope over (and bind variables occurring within) speech act designators. It is, we’ve seen, natural to analyze such quantification as sequencing.

Quantifying into imperatives. Similar considerations yield a case for recognizing quantification into imperative sentences. Consider the imperatives in (4.82).

- (4.82) a. See to it that everyone takes shelter!
 b. Everyone_i take their_i seat!

There is an unmistakable difference in force here, which no extant account, so far as I know, analyzes correctly.⁶⁹ Portner (2004a,b), for instance, seems to blur these together, assigning the following sentence the indicated meaning at a context of utterance c .

$$(4.84) \llbracket \textit{everyone eat!} \rrbracket^c = \lambda w. \lambda x : x \in \textit{addr}(c). (\forall y : y \in x. y \textit{ eats in } w)$$

The indicated function is a function from worlds to sets of groups of individuals, namely, the set of groups of individuals such that everyone in the group eats in that world. As we know from our earlier discussions of Portner’s account of the conventional force of imperatives, this means that an addressee (or group of addressees) of *everyone eat!* updates by adding the proposition *that everyone in the group eats* to their preferences. But this collapses the meaning of the imperative *everyone eat!* into that of the imperative *see to it that everyone eats!*. And that is incorrect—an agent satisfies the request conventionally expressed by *everyone eat!* simply *by eating!* The agent need not see to it that everyone else in the relevant group eat; although this would be reasonable, given the obviously desired state of affairs (consisting in everyone’s eating), it is not obliged. Portner’s account incorrectly says otherwise.

Intuitively, I will suggest that the distinction between an imperative in the mold of (4.82a) and an imperative in the mold of (4.82b) consists in the following:

- (4.82a) requests of a specific agent a that a see to it that the truth of a proposition of the form $\forall x : P(x)$ is secured
- (4.82b) requests of every individual x that x secure the truth of the proposition that x takes x ’s seat. (This in turn is distinct from requesting of the audience, as a group, that it take its seat.)


As a rough first pass, the differences between these requests appear to be perspicuously representable as a scope-difference in the illocutionary metalanguage. In the former, the quantifier phrase is represented as scoping under the imperative operator, while in the latter, the quantifier phrase *takes scope over* the imperative operator.

69. There may be a preliminary worry about how the quantifier phrase *everyone* figures in the interpretation of (4.82b). For instance, we (or at least *I*) would not say that (4.83) conventionally functions to express a sequence of assertions, i.e., an assertion directed at each individual in the relevant domain.

(4.83) Everyone: your book is under your seat.

That’s because the quantifier phrase here is a vocative, and, so, does not actually bind any variable in the LF of (4.83). Vocatives are, following Portner (2004b), “separate” performatives, meaning that they do not contribute to the “at-issue” content of the utterance (see also Green 2000; Potts 2003). They may contribute to the “meaning” of an utterance in some sense, but should not be represented as part of the proffered, or at-issue, meaning of the utterance.

The quantifier phrase in (4.82b) is clearly not vocative. Notice that it *binds* a third-person pronoun, and so *must* be represented in the LF for the imperative. Notice further that it intuitively does seem to contribute to the at-issue content or meaning of the utterance. The speaker of (4.82b) says, *of every person*, that she must take her seat. The speaker of (4.83) does *not* say, of everyone, *that*. . . anything! Perhaps she does manage to communicate, of every person, that her book is under her seat; but that is not, strictly speaking, what she *says* with (4.83).



Notice that we are now construing the imperative operator, as we earlier construed the question operator, as taking an individual argument, corresponding to an addressee for an imperative. The principal use for this innovation is to allow the imperative operator to bind variables within its scope. More importantly, notice that, in the latter tree, we once again have quantification scoping over a speech-act operator. And, as before, letting $A = \{a_1, \dots, a_n\}$ be a finite list of relevant constants, an interpretation of the speech act quantifier in terms of sequencing seems to be reasonable.

$$[\forall x : !x : P(x)] = [\bigwedge_i \text{sub}^{a_i/x} (!x : P(x))] = [!a_1 : P(a_1)] \circ \dots \circ [!a_n : P(a_n)]$$

What is the force of a speech act of this sort on a member (call her a) of the agent's audience? The more straightforward answer seems to be:

- a updates on $!a : P(a)$
- For any agent $a' \neq a$ in the audience, a is unable to update on $!a : P(a)$

Both of these things follow naturally from the idea that a speech act operator of the form $!x$ expresses conventional imperative force *directed at* x , i.e., that the individual argument of the operator specifies an addressee.

An elegant way of representing this—and one that will allow us to give a natural treatment of imperatives occurring under *existential* quantifiers—is to think of a speaker's intended addressees at a context of utterance c as comprising a complex of individuals, A_c —which we will understand as a function from each addressee to her cognitive state at c —and to have speech acts updating the context by way of updating this function (and thereby updating individual cognitive states). Speech acts will thus, strictly speaking, express *context change potentials*, rather than mere update potentials (although because of the attendant formal complexity, we will only make use of CCPs when absolutely necessary). Context change potentials are understood to be *multi-agent generalizations* of update potentials; sequenced speech acts can modify the cognitive states of distinct agents in sequence—for instance, when the speech act is of the form $!a : P(a); !b : P(b)$ —or the cognitive state of a single agent in sequence—for instance, when the speech act is of the form $!a : P(a); !a : Q(a)$. Formally, the idea is this:

Definition 4.76. If c is a context, then A_c is an **addressee state** at c . A_c is a relation mapping each addressee at c to a cognitive state for the addressee at c . So if $A_{c,a}(\Sigma)$, Σ is a cognitive state for a at c .

Definition 4.77. For all i such that $A_{c,a}(\Sigma_i)$, a 's cognitive state at c (notation: $A_{c,a}$) is $\cup_i \Sigma_i$.

Definition 4.78. A **context change potential** is any function $|\cdot|$ such that:

- i. $A_{c|>\phi|}(a) = \{\sigma' : A_{c,a}[>\phi]\sigma'\}$
- ii. $A_{c|?\phi|}(a) = \{\sigma' : A_{c,a}[?\phi]\sigma'\}$
- iii. $A_{c|!a:\phi|}(a) = \{\sigma' : A_{c,a}[!\phi]\sigma'\}$ (When $b \neq a$, $A_{c|!a:\phi|}(b) = A_{c,b}$)
- iv. $A_{c|\alpha;\beta|} = A_{c(|\alpha|\circ|\beta|)}$
- v. $A_{c|\alpha+\beta|} = A_{c(|\alpha|\cup|\beta|)}$
- vi. $A_{c|\forall x:\alpha|} = A_{c|\bigwedge_{k_i: \llbracket k_i \rrbracket \in D_c} \text{sub}^{k_i/x}(\alpha)|}$ (D_c gives the c -relevant domain of quantification)

The key clause is (vi). This says a context that has been updated with a speech act of the form $\forall x : \alpha$ is the result of updating A_c (the addressee state at c) with $|\text{sub}^{k_1/x}(\alpha); \dots; \text{sub}^{k_n/x}(\alpha)|$, for all i such that $\llbracket k_i \rrbracket$ is in the context-relevant domain of quantification. So, for instance, if the quantified speech act is of the form $\forall x : !x : P(x)$ (and the domain of quantification ranges over members of the speaker's audience), the update to the context is of the expected form: $!a_1 : P(a_1); \dots; !a_n : P(a_n)$. And this, per clauses (iii) and (iv), has the force of updating each addressee a 's cognitive state so that it is decided on $\delta P(a)$. And that is exactly what the conventional force of this quantified imperative seems to consist in.

From this perspective, it's not at all surprising that Portner has misidentified the meaning of imperatives like (4.82b) (although he is by no means alone in this). Without regular operations at his disposal, Portner is forced to try to model the force of (4.82b) using the rubric of adding a proposition to an agent's body of preferences. Any account that makes use of only this rubric is doomed to conflate the meanings of imperatives like (4.82a) and (4.82b). More generally, any account that does not help itself to quantification-like operations outscoping speech act operators (and binding variables within the scope of speech act operators) will have a very hard time accounting for this distinction. (4.82b) conventionally expresses a sequence of speech acts that cannot be represented with a speech act designator of the sort allowed by the Standard Account.

Existential quantification? So far, we have seen evidence for thinking that quantification-like operators with universal force—operators that express speech act sequencing, for each individual in some relevant domain—have dedicated modes of expression in natural language. We have not seen evidence for thinking the same about quantification-like operators with *existential* force—operators that would express iterated applications of speech act disjunction (i.e., '+'), for each individual in some relevant domain.


Perhaps that is no accident. There does seem to be some evidence to be skeptical that existential quantification can scope over speech act designators. Krifka (2001a) notes that pair-list readings of question sentences are seemingly available only when the relevant quantifier is universal—i.e., representable as expressing sequencing of some series of speech acts that result from uniform substitution for some variable.⁷⁰

- (4.85) How did each linguist write her dissertation?
 a. ✓ $?m : \forall x : m$ is how x wrote x 's dissertation (Narrow-scope LF)
 b. ✓ $\forall x : ?m : m$ is how x wrote x 's dissertation (Pair-list LF)
- (4.86) How did two linguists write their dissertations?
 a. ✓ $?m : 2x : m$ is how x wrote x 's dissertation (Narrow-scope LF)
 b. # $2x : ?m : m$ is how x wrote x 's dissertation (Pair-list LF)
 c. # Craig and Mandy used Word. Paul and Maria used L^AT_EX.

To close our discussion of speech act quantification, I want to suggest a series of imperative examples that would be usefully analyzed by making use of an existential quantification-like operation scoping over speech act designators.⁷¹

Indefinite addressee. Consider (4.87), the proposed representation (making use of existential quantification into speech acts), and the expected interpretation for such an operation.

- (4.87) Someone_{*i*} bring me her_{*i*} chair!



$$[\exists x : !x : P(x)] = [\bigvee_i \text{sub}^{a_i/x} (!x : P(x))] = [!a_1 : P(a_1)] + \dots + [!a_n : P(a_n)]$$

I say this is the expected interpretation, because, as an existential first-order formula $\exists x : \phi$ is equivalent to the disjunction of ϕ 's substitution instances, so the force of a speech act of the form $\exists x : !x : \phi$ can be intuitively regarded as equivalent to the “disjunction” of the substitution instances of $!x : \phi$

Here are some worries for this proposal. Since a speaker of (4.87) typically addresses

70. For some further evidence, see Krifka (2004).

71. This, as well as my account of the Ross Paradox, represents a shift in position from my earlier work (Charlow 2009, 2010a,b). This shift is driven largely by my development of an apparatus that places *deliberative alternatives*—practical questions, questions about how to plan, questions about what to believe—front and center in the model of cognitive states invoked by my favored theory of speech acts.

an audience, (4.87)'s force might be thought to be universal, rather than existential—aptly represented, that's to say, with a speech act of the form $\forall x : !x : P(x)$, with P expressing the property of seeing to it that someone brings the speaker her chair, and \forall ranging over the audience. Or, less extravagantly, we might just assign (4.87) a Standard imperative LF, treating the addressee simply as the audience, considered as a *fusion* of its component individuals, and the content of the request that the group see to it that someone bring the speaker her chair. Either analysis predicts (correctly) that:

- The quantifier phrase *someone* is able to bind the third-person pronoun *her* (since, on either, the sentential complement of the imperative operator contains the quantifier phrase)
- Accepting (4.87) would not thereby require an agent to bring the speaker *her* chair
- Accepting (4.87), however, would mean that an agent who knows that her bringing the speaker a chair is the *only* way the speaker's request will be fulfilled is required to bring the speaker her chair

On the former proposal, this is because the agent has added the preference that she see to it that someone brings the speaker a chair to her preferences; but this obliges her to bring the speaker a chair only if that is the only way to see to it that someone brings the speaker a chair. On the latter proposal, this is because the agent is part of a group that has been requested to see to it that the speaker is provided a chair; this obliges her to bring the speaker a chair only if that is the only way for the group of which she is a part to fulfill the speaker's request of the group. If this is all there is to explain about the force of (4.87), quantifying existentially into speech acts is not required.

More worryingly, the suggested representation of (4.87) might be thought to mean that it actually does express an *indeterminate* speech act (in the sense of Footnote 62). Why? Suppose there are two addressees, a_1 and a_2 . Per definition 4.78, such a speech act will update a context either by leading a_1 to plan that $P(a_1)$, or leading a_2 to plan that $P(a_2)$. One operation or other will be done to the context, but the speaker has *no idea* which. This is a bizarre kind of speech act, and there is good evidence that such a speech act cannot be expressed in natural language.

(4.88) #Someone bring me their chair (I don't know who)

Notice, however, felicity of the following continuation.

(4.89) Someone bring me their chair (I don't care who)

And notice that this is the same sort of contrast that obtains with choice-offering imperatives, which we analyzed in terms of their introduction of planning alternatives for an agent.

(4.90) #For your hangover, take Advil or Alka Seltzer (I don't know which)

(4.91) For your hangover, take Advil or Alka Seltzer (I don't care which)

This, I think, strongly suggests a similar sort of analysis of (4.87) as choice-offering disjunctive imperatives (the difference lying in the fact that the former is said to introduce alternatives for a single agent, while the latter introduces alternatives instead *for a context*). In other words, we extend Definition (4.78) by adding a clause for existential quantification into speech acts.

$$A_{c|\exists x:\alpha|} = A_{c|\bigvee_{k:\llbracket k \rrbracket \in D_c} \text{sub}^{k/x}(\alpha)|}$$

According to this clause, a context that has been updated with a speech act of the form $\exists x : \alpha$ is the result of updating with $|\text{sub}^{k_1/x}(\alpha) + \dots + \text{sub}^{k_n/x}(\alpha)|$, for all i such that $\llbracket k_i \rrbracket$ is in the context-relevant domain of quantification. So, for instance, if the speech act is of the form $\exists x : !x : P(x)$ (and the domain of quantification ranges over members of the speaker's audience), the update to the context is of the form $!a_1 : P(a_1) + \dots + !a_n : P(a_n)|$. And this, per Definition 4.78 clauses (iii) and (v), means that $|\exists x : !x : P(x)|$ expresses an update relation, rather than an update function—a relation which $c|!a_1 : P(a_1)|$ satisfies just as well as $c|!a_2 : P(a_2)|$. Following Definitions 4.66 and 4.67, which provided, respectively, a notion of a cognitive state meant to incorporate cognitive alternatives, and update instructions for such cognitive states, we define a notion of addressee states meant to incorporate contextual alternatives, and update instructions for such addressee states:

Definition 4.79. Any set of relations relating individuals to cognitive states is an **addressee state**. If \mathcal{A} and \mathcal{A}' are addressee states, then $\mathcal{A} \cup \mathcal{A}'$ is an addressee state. Nothing else is an addressee state. (Notation: Roman lettering will be used for relations between individuals and cognitive states; Cal lettering will be used for sets of such relations.)

Definition 4.80. $\mathcal{A}|\alpha| = \bigcup_{A \in \mathcal{A}} \{A' : A|\alpha|A'\}$

Consider an addressee state given by $\{A\}$, where A is an addressee state in the sense of Definition 4.76—a relation between individuals and cognitive states. The result of updating this state with a speech act of the form $\exists x : !x : P(x)$ is, according to Definition 4.80, is

$$\{A' : A|\exists x : !x : P(x)|A'\} = \{A|!a_1 : P(a_1)|\} \cup \dots \cup \{A|!a_n : P(a_n)|\}$$

In other words, it is a set of functions from individuals to cognitive states, one of which has a_1 planning that $P(a_1)$, another of which has a_2 planning that $P(a_2)$, and so on. The force of the speech act expressed by (4.87), on this account, is to get a context to recognize a set of admissible alternatives such that, in each admissible alternative, an agent plans to bring the speaker their chair.

How does this analysis do with the data that was handled with relative ease by either of the Standard analyses? It allows the quantifier to bind the third-person pronoun *her* (in the same way that universal quantification into speech acts allowed such binding).

As for the other two predictions, there is not, strictly speaking, anything corresponding to an *agent's* acceptance of an imperative like (4.87) in this analysis. *Contexts* (to which agents bear some sort of participant relationship, as either speakers or addressees) accept or reject such imperatives, and any effect on the agent is derivative of the imperative's effect on the context. And this seems to me to be correct. Here are paradigmatic circumstances in which an agent might come to be the subject of an imperative like (4.87), and in which update with such an imperative would be expected to affect the agent's own planning behavior.

- The agent is part of a group (the audience).
- The agent recognizes her membership in this group as in some sense *normative* for her: features of group-level "planning" can properly impact her own planning, depending on the role the agent sees herself as occupying in the group.
- An agent with authority over group-level "planning" (the speaker) proposes to update group-level "plans".

Given an update with (4.87), such an agent can be expected to reason as follows. There are a number of possible ways of planning by agents in the audience that are allowed by the speaker's communicative aims, each corresponding to an alternative introduced by update with (4.87). So updating the context with (4.87) does not thereby require this agent to bring the speaker her chair. But given decisions which eliminate all other alternatives as possible evolutions of cognitive states of other audience members, the agent is ultimately obliged to bring the speaker her chair. And so the analysis seems fully capable of predicting the other two phenomena for which the Standard analyses suggested above offer explanations.

It might be wondered what sort of leg up this sort of analysis has on the Standard analysis which assigns (4.87) an LF of the form $\exists x : P(x)$, but which has update with this imperative modifying the preferences of the *group*, considered as some kind of complex individual, so that it comes to recognize a requirement of the form $\exists x : P(x)$. And similarly for an imperative like (4.82b) (which we analyzed in terms of universal quantification into an imperative); why can't this be assigned an LF of the form $\forall x : P(x)$, while understanding the addressee to be the group?⁷²

My response has several parts. First, there is, on the alternative analysis, no prediction of the relevant free choice-like phenomena exhibited in (4.88) and (4.89). From the fact that the group comes to recognize a requirement of the form $\exists x : P(x)$, it, of course, does not follow that the group recognizes a permission for each relevant x to realize $P(x)$. Given a

72. This is the only credible challenge that I see to the non-Standard analysis I'm proposing here for indefinite addressee imperatives. So long as our non-Standard analysis of universal addressee imperatives stands, considerations of theoretical unity alone would motivate an attempt to extend that analysis to indefinite addressee imperatives, presuming such an extension to be workable. (The analysis developed here for indefinite addressee imperatives, which extends our analysis of universal addressee imperatives, presumably shows such an extension is workable.) An objection, to be credible, must push group-addressee analyses not only of indefinite addressee imperatives, but universal addressee imperatives as well.

natural extension of the analysis of choice-offering imperatives developed in Section 4.4.6, our analysis can predict this with ease.

Second, this objection seems to get other data wrong. The suggested Standard-friendly analysis predicts that (4.87) and (4.92) are basically identical in force.

(4.92) Audience: see to it that one of you brings me their chair!

That seems incorrect. (4.92) sets the stage for group deliberation about how best to realize the speaker's request (seeing to it that one of them brings the speaker her chair). Less pompously, it sets the stage for *collective* action—action by the group that is being explicitly addressed. But this, intuitively, is not what (4.87) does at all. It sets the stage for *individual* action. An agent who reacts to (4.92) by bringing the speaker her chair presumes (perhaps presumptuously!) to act for the group. An agent who reacts to (4.87) by bringing the speaker her chair does not presume to act for the group. Rather (and exactly as our account predicts) she sees the alternative evolutions of (individual) planning states allowed by the speaker's request; she sees that her planning to bring the speaker her chair is allowed by the speaker's request; and so she plans accordingly.

A related worry: conversation is fundamentally a cooperative endeavor between agents. Group "agents" are a useful toy in metaphysics and the philosophy of mind and can serve genuinely useful modeling purposes (e.g., representing group knowledge, collective action, etc.). But speech acts, fundamentally, are proposals by (bona fide) agents (speakers) and directed toward (bona fide) agents (their interlocutors). In the absence of workable alternative accounts, it might be a useful fiction to treat quantified-addressee imperatives as updating the "plans" of a group (thereby indirectly influencing the actual plans of agents, depending on the relations in which they understood themselves to stand to the group). But we have an alternative model, which does not require adopting any such useful fiction.

Finally, objecting to the use of non-Standard speech acts on the grounds that there is a construal, however inelegant, on which they can be represented as non-complex is a strategy with vastly diminishing returns. We have already seen multiple constructions that are usefully analyzed by making use of essentially complex speech acts. At this point, if we have a workable analysis of a troubling construction that makes use of a complex speech act, there is no real presumption, stemming from the fact of its complexity, against using it.

Free choice 'any'. Finally, there is the fact that imperatives invoking the negative polarity item *any* have *obligatorily* free choice readings.

(4.93) Pick any card! \Rightarrow You can pick a spade (Aloni 2007)

(4.94) Anyone_i bring me their_i chair! \Rightarrow Bob can bring me his chair

Both sentences, I suggest, should be analyzed using the apparatus for representing choice-offering disjunctive imperatives developed above. Notice again the way these imperatives pattern with *I don't care wh-* and *I don't know wh-* continuations.

- (4.95) a. #Pick any card (I don't know which)
 b. Pick any card (I don't care which)
- (4.96) a. #Anyone bring me their chair (I don't know who)
 b. Anyone bring me their chair (I don't care who)

Unsurprisingly, then, I'll suggest the following analyses of these sentences.

(4.93) is given an LF of the form:

$$\exists x : !a : P(a, x)$$

Notice that the quantifier here does not bind the addressee argument of the imperative operator: the force of (4.93) is to introduce the addressee *a*'s taking of each card as an alternative for *a*.

(4.94) is given an LF of the form:

$$\exists x : !x : P(x)$$

Notice that the quantifier here does bind the addressee argument of the imperative operator: the force of (4.93) is in this sense basically identical to that of (4.87).

The interpretations of these speech act designators follows straightforwardly from the account we have developed in this section (thus illustrating the far-reaching explanatory power of this account). Among the many welcome features of this account, beyond its prediction of the relevant free-choice phenomena, is the fact that it preserves the semantically standard analysis of NPI *any* as expressing existential quantification. It also has the advantage of being adaptable to the sort of explanation of NPI *any*'s peculiar distributional behavior favored in such classic accounts as that of Kadmon & Landman (1993) (on which *any* conventionally induces widening of the domain for the existential quantifier expressed by *any*, but is *licensed* only when this domain-widening leads to a *strengthening* of the sentence's content along some salient dimension; cf. Aloni 2007). And this in turn explains why the choice-offering reading is obligatory: on alternative-presenting representations—representations of the form $!\exists x : P(x)$ —domain-widening makes both (4.94)'s required and permitted contents—each given by $\exists x : P(x)$ —strictly *weaker*. On choice-offering readings, domain-widening ends up generating more permissions; and, so, the choice-offering readings of both (4.93) and (4.94) are obligatory.⁷³

73. Aloni notes: “[D]omain widening [in the case of choice-offering *any*] does not make our statement stronger. However, it does not weaken the statement either. None of the wide or the narrow interpretations of sentence [4.93] entail each other. Domain widening in this context creates a new meaning. This, I would like to

4.5 Conditional Imperatives

More or less everything that we have said about the linguistic phenomena surrounding “ordinary” imperatives carries over, *mutatis mutandis*, to conditional imperatives (CIs). Most importantly, and we noted in Section 4.4.2, it is natural (and potentially very fruitful) to generalize the necessitation account of imperatives so that an imperative of arbitrary form expresses a proposal for an agent to come to recognize a requirement of corresponding form. Applying this to CIs will mean that CIs conventionally (but defeasibly) express *conditional necessitation acts*. This means only that utterances of CIs canonically function to make an agent *recognize a conditional requirement*. Accepting the speech act conventionally expressed by (4.97) means coming to recognize the requirement in (4.98).

(4.97) If it rains, shut the window!

(4.98) If it rains, you should shut the window

It is this aspect of CIs that proves most troublesome for the sort of account of imperatives that I’ve been pursuing (and, more generally, any account which attempts to give a speech act-theoretic account of the meaning of imperative clauses in natural language).

Giving such an account requires locating a speech act that can be seen to fill this theoretical role. If such a speech act within the mold of the Standard Account cannot be located, that is not, in itself, a big deal. As our examination of embedded imperatives suggests, we may then go looking for Standard Account-friendly speech acts out of which the relevant necessitation act can be understood to be composed. There is good reason to think, though, that this sort of methodology cannot get us anywhere with CIs. And that represents a fundamental challenge, not only to the Standard Account, but to speech act-theoretic accounts of conventional meaning more generally.

Things, of course, are not so bleak as this. But it will take a fairly radical departure from the Standard Account—one, however, that is not all *that* radical on the treatments of conditionals prevalent in model-theoretic and dynamic semantic analysis of conditionals—to see why.⁷⁴ According to that account, the sorts of speech acts favored by the Standard Account play no fundamental role in theorizing about the meaning of imperative clauses in natural language. Complex necessitation acts are fundamental, and the sorts of “simple” necessitation acts to which we’ve been appealing throughout this essay are, in fact, best understood as constructions out of such complex necessitation acts.

suggest, supplies enough reason for widening to occur” (2007: 90). Although we favor different analyses of choice-offering *any*, I concur.

74. The discussion here will of necessity be abbreviated. For a fuller treatment that picks up many of the themes introduced here, see Appendix B and my (2011b).

4.5.1 A Standard Account (Wide-Scoping)

Let's first consider a Standard Account-friendly treatment of CIs, on which CIs conventionally express proposals for an agent to come to recognize a requirement to make some material conditional true. According to this account (let us call it the Wide-Scoping Account, or WSA) the LF of a CI of surface form $(if \phi)(!\psi)$ is given by:

$$!(\phi \supset \psi)$$

It is obvious that the WSA, if it can be made to work, is friendly to the Standard Account. Not that this is what matters, at this point; it is not. The primary concern of the fan of the Standard Account is simply to establish a bulwark against the fundamental challenge outlined above. If there is a Standard-friendly treatment against CIs, so much the better for this position.

Information-sensitivity. There is not, to my knowledge, a single wide-scope account of conditionals embedding information-sensitive operators—epistemic modals, deontic modals, teleological modals, whatever—that is empirically plausible. And since imperatives are, as I will argue at length in this section, information-sensitive, it should be no surprise that a wide-scope account of CIs in the mold above is a non-starter.

In what sense, though, are imperatives information-sensitive? They are information-sensitive in roughly the same way as the practical requirements they function to establish. Just as (4.99) can be satisfied or not, relative to a cognitive state that has updated with (4.98), depending on the information an agent has at her disposal (specifically, her information about the weather), so too can (4.100) be satisfied or not (in the sense that the cognitive update it expresses can be redundant or not) depending on the information about the weather she has at her disposal.⁷⁵

(4.99) You should shut the window

(4.100) Shut the window!

In a slogan, CIs typically articulate constraints on planning relative to *contingencies*. Contingencies, roughly, are situations *relevant to planning*, hence specified in terms of properties of an agent's *information state* (rather than in terms of, say, non-informational properties of the world). When a cognitive state comes to satisfy the property characterized by a relevant contingency, the relevant constraint on planning is enforced.

In this vein, then, here is some concrete data, having to do with the information-

75. This sort of information-sensitivity is pretty weak indeed (so weak that it is shared by declarative consequents of indicative conditionals). There is a stronger sense in which imperatives and corresponding modalities are information-sensitive (in that the *evaluation* of both is sensitive, in different but closely related ways, to a relevant base of information). See our later discussion of deontic and imperative conditionals for more on this.

sensitivity of natural language imperatives, that an account of CIs should be able to predict. Suppose an agent has updated on (4.98), and suppose further that at the time of update the agent believes it is not raining. Now consider some later time at which the agent has learned that it is raining. How does such an agent plan in these sorts of circumstances? The answer seems obvious: the agent plans to shut the window! Here, I suggest, are the key ingredients an account of CIs must supply in order to make this sort of prediction.

- Updating on a CI like (4.98) should update the agent's plans *for the contingency in which* she learns that it is raining
- Learning that it is raining should cause the agent *to plan*, in her current contingency, as advised by the consequent of the CI. In other words, it should cause the agent to plan to shut the window

Contingency planning. The WSA, of course, is eliminativist about contingency planning. Contingency planning, according to the WSA (at least when paired with a rough analysis of the conventional force of unembedded imperatives, along the lines proposed above), is just planning to satisfy a material conditional in one's current contingency.

We should not, however, be eliminativists about contingency planning: planning to see to it that ψ in a ϕ -contingency is simply not reducible to planning to see to it that $(\phi \supset \psi)$. Notice, for instance, that the following CIs (to which I've appended their LFs, according to the WSA) intuitively place *consistent constraints* on agent's plans.

- (4.101) a. If it rains, bring only the umbrella. $:= !(R \supset (U \wedge \neg S))$
 b. If it doesn't, bring only the sunglasses. $:= !(\neg R \supset (\neg U \wedge S))$
 c. Bring both. $:= !(U \wedge S)$

That is to say: this collection of imperatives will typically express useful advice for an agent. (4.101a) tells the agent how to plan in a contingency where it rains. (4.101b) tells her how to plan in a contingency where it's sunny. And (4.101c) tells her how to plan in her current contingency, according to which the question of whether it will rain or not is not settled. But if, per the WSA, the CIs don't express genuine contingency plans—if they instruct an agent simply how to plan in the current situation (with her current information)—this collection of imperatives is provably inconsistent! Given that the agent brings both the umbrella and the sunglasses, the only way all three imperatives are satisfied is if it both rains and does not rain. In other words, the WSA predicts that these three imperatives together express a collection of advice that the agent cannot consistently incorporate into her plans.

Sophisticated informants often see this sort of argument and decide to backtrack. This collection of imperatives cannot, they decide, express coherent advice for an agent. The antecedents of (4.101a) and (4.101b) together cover all the possible contingencies. So, in accepting (4.101a) and (4.101b), an agent must plan either to bring only the umbrella, or to bring only the sunglasses. The fact that the agent does not do either is evidence that she

does not accept these CIs. What she instead accepts are CIs like the following:

- (4.102) a. If you learn it will rain, bring only the umbrella
b. If you learn it won't rain, bring only the sunglasses

And these two CIs can manifestly be accepted while planning to bring both (since an agent who accepts the advice to bring both manifestly does know not whether or not it will rain).

The Ramsey Reading. It is, shall we say, revealing that such revisionary backtracking typically happens only after the problem has been laid out. It is revealing because the *salient* reading of the CIs (4.101a) and (4.101b) does not seem to be the one that the sophisticated informant is pressuring us to endorse. It is, instead, what I will refer to as the **Ramsey Reading**. What, though, is that?

The Ramsey Reading of an indicative conditional (including, incidentally, imperative indicative conditionals) is the reading accessed by executing the Ramsey Procedure. The Ramsey Procedure for the evaluation of an indicative conditional “invites us to add the information carried by the antecedent to the contextually relevant stock of information. . . and check the fate of the consequent” (Gillies 2010: 27). On the Ramsey Reading of (4.101a), we add the information carried by its antecedent to the relevant stock of information, and subsequently evaluate the consequent (in this case, by planning, from the vantage point of this hypothetical information-state, to bring only the umbrella).

Of course, on the Ramsey Reading of (4.101a), it seems like this CI expresses sound advice. *Of course*—said from the vantage point of an information-state according to which it rains—I should bring the umbrella and leave the sunglasses. And similarly for (4.101b). And, of course, since it might either rain or not, I should take both the umbrella and my sunglasses. There is obviously nothing incoherent about the cognitive state of such an agent. And there is good reason to think that this is the sort of cognitive state that the imperatives in (4.101), on their salient readings, function to generate.⁷⁶

The WSA lacks the sorts of LFs which would be required to adequately represent the Ramsey Readings of CIs. Perhaps even more worryingly, the WSA seems to eschew the sort of understanding of cognitive states that would be required to define a suitable speech act for such CIs to express. CIs, I have suggested, update *contingency plans* (where contingencies are individuated in terms of the property of an agent's information state they express). The WSA has no need for contingency plans, except in an eliminativist sense, in which an agent who plans that $(\phi \supset \psi)$ plans that ψ in the contingency that ϕ .

76. More generally, the salient reading for any indicative conditional is *always*, I submit, the Ramsey Reading. The Ramsey Procedure will differ depending on the nature of the antecedent. Antecedents that themselves express propositions about the agent's information state do not, typically, induce an information state that represents this proposition; rather, they induce an information state that satisfies the property specified by that proposition. And so an ordinary speaker will typically assign (4.101a) and (4.102a) the same interpretation. For this reason, a fan of the Standard Account cannot say that (4.101a) should receive whatever analysis (4.102a) ultimately receives; for it is likely that the salient reading of (4.101a) will require a non-Standard analysis too.

A richer notion of contingency planning is needed. Neither the WSA nor the Standard Account give any guidance in developing one. Worse, there is strong reason to think that neither could consistently make use of such a notion, if one were to be developed.


4.5.2 Conditional Speech Acts (Narrow-Scoping)

Next, I want to consider a family of accounts that involve treating CIs as expressing a complex speech act—one that is constructible out of basic speech acts together with some sort of operation (a *conditionalizing* operation) on speech acts—and see whether it fares any better than the Wide-Scope account of CIs discussed above. These analyses share the core idea that CIs express a conditional speech act: $[\psi]$ is somehow conditioned on ϕ . They differ primarily in the notion of conditioning that they employ.

Conditioning on truth. On one notion of conditioning, the speech act expressed by the consequent of a CI is conditioned on the bare truth of the antecedent (see, perhaps, [Portner 2004a](#)). So, if it is raining outside, updating with (4.97) means recognizing a requirement to shut the window. But, if it is not, updating with (4.97) requires nothing of the agent.

Such an account clearly fails to supply either of the key ingredients an account of CIs must supply. First, it fails to explain how updating with (4.97) should update the agent's plans for the contingency in which she learns it is raining. CIs express constraints on the agent's contingency plans, where contingencies are individuated in terms of the property of an agent's information state they express. Bare facts about the world (e.g., the fact that it's raining) express no property of the agent's information state; they express a property of the world. And so such an account also fails to explain how learning it is raining should cause the agent who has updated with (4.97) to plan to shut the window. An agent learns that it is raining when (i) it is raining, (ii) she did not know that it was raining, (iii) she comes to acquire information that allows her to know that it is raining. The account in question has update with (4.97) leading the agent to plan to shut the window when (i) is the case. It offers no explanation of how learning it is raining can cause an agent to adjust her plans (for in such circumstances, no adjustment is possible at all).

Conditioning on assertion. A much more promising idea understands the conditioning as imposing some sort of *informational restriction* on update with the CI's imperative consequent (see, e.g., [Potts 2003](#); [Asher & Lascarides 2003](#); [Lascarides & Asher 2004](#)). A CI, on this way of thinking, expresses a relation between assertion acts (namely, the assertion conventionally expressed by the CI's antecedent) and the speech act conventionally expressed by the CI's consequent.



There are various ways of understanding this relation, but I will run with what I take to be the most plausible one here (since they all ultimately suffer from the same flaw). On that understanding, CIs still express conditioned updates. But the updates are conditioned, not on non-informational properties of the external world, but rather on properties of the agent’s information state. Specifically, a CI like (4.97) tells the agent to check whether her information state satisfies the property expressed by the CI’s antecedent. If it does, update accordingly (by planning to shut the window). If it does not, do nothing. More formally, according to this account the update potential conventionally expressed by a CI of surface form $(if \phi)(!\psi)$ is given by:

$$[!\psi] \cap \{ \langle \Sigma, \Sigma' \rangle : \Sigma[\triangleright\phi] = \Sigma \} = \begin{cases} \Sigma[!\psi], & \text{if } \Sigma[\triangleright\phi] = \Sigma \\ \Sigma & \text{otherwise} \end{cases}$$

The force of the antecedent is, as with the above proposal, to *restrict* the force of the consequent. (And, in general, the force of a speech act of the form $\alpha \rightarrow \beta$ will be to restrict the execution of β to contexts where α ’s execution obtains.) Significantly, though, the restriction, in this case, has to do with properties of the agent’s information state (whether or not she thinks it’s raining), rather than non-informational properties of the world. A CI will adjust an agent’s plans, on the condition that the information expressed by its antecedent is available.

As something of an added bonus, this is something like the sort of formal analysis we might expect to go along with the informal Ramsey Procedure for the evaluation of CIs. The Ramsey Procedure says to evaluate the imperative consequent (i.e., update accordingly) from the vantage point of an information state representing the information containing in the CI’s antecedent. And this, at first glance, seems to be just what the suggested analysis of CIs has the agent doing.

In this case, though, the first glance is from a bit of a misleading angle. Indeed, this cannot be the right analysis for CIs in natural language. Update with CIs leads an agent to update her contingency plans; in other words, update with CIs causes the agent to recognize conditional requirements. And this holds even when the agent’s starting information state is undecided about the truth or falsity of the CI’s declarative antecedent. Even in cases of an agent’s uncertainty about the indicative antecedent—cases, it must be said, that obtain *whenever* an indicative conditional is evaluable for that agent (cf. [Stalnaker 1975](#); [von Stechow 1998](#))—updating with a CI of surface form $(if \phi)(!\psi)$ causes the agent to update her contingency plans, by way of coming to recognize a corresponding conditional

requirement.

This, however, is *impossible* on the account under consideration. According to this account, the original context is returned when the information expressed by the antecedent is not available. That is to say:

$$\text{If } \Sigma[\triangleright\phi] \neq \Sigma, \text{ then } \Sigma[\triangleright\phi \rightarrow !\psi] = \Sigma$$

Of course, the Ramsey Procedure does say that, when evaluating an indicative conditional and your actual cognitive state satisfies the property expressed by its antecedent, an agent evaluates (updates on) the conditional simply by evaluating (updating on) the consequent. And this part of the Ramsey Procedure is captured perfectly by the account under consideration. But the Ramsey Procedure says more besides: it tells an agent how to evaluate an indicative in the (standard) situation where the agent's actual cognitive state does not satisfy the property expressed by its antecedent. In such circumstances, the Ramsey Procedure says to shift to a hypothetical cognitive state that does satisfy the property expressed by the antecedent, and evaluate the consequent with respect to this hypothetical state. In the case of CIs, specifically, the Ramsey Procedure says that to update on a CI is to construct a contingency plan: to evaluate the CI's imperative consequent with respect to the information state induced by supposing the CI's declarative antecedent. The account under consideration simply ignores this part of the Ramsey Procedure.

What is instead needed is an account on which learning that the antecedent of a previously accepted CI is true causes the agent to update her actual plans in the way specified by the CI's consequent. More formally, we are in need of an update potential for $\triangleright\phi \rightarrow !\psi$ on which the result of updating on $\triangleright\phi \rightarrow !\psi$ and $\triangleright\phi$ is a state which recognizes a requirement to see to it that ψ (i.e., a state in which further update with $!\psi$ would be redundant).

$$[\triangleright\phi \rightarrow !\psi] \circ [\triangleright\phi] = [\triangleright\phi \rightarrow !\psi] \circ [\triangleright\phi] \circ [!\psi]$$

Of course, this requires defining a value for $[\triangleright\phi \rightarrow !\psi]$ on which $[\triangleright\phi \rightarrow !\psi]$ is not provably incapable of updating a state Σ when $\Sigma[\triangleright\phi] \neq \Sigma$. And if that were all that were required, an account of CIs which conditioned update on an imperative consequent on update with the antecedent might be workable after all. Unfortunately, much more is required than this.

What a speech act theory of meaning would need. A speech act theory of meaning for CIs also seems to be in need of:

- A way of indexing an agent's plans to *non-actual contingencies*—to possible future information states of the agent.⁷⁷
- An analysis of the antecedents of CIs on which they involve selecting, in typical

77. On the theme of information-sensitivity in planning and preferences, see the later chapters on deontic and imperative conditionals, as well as Charlow (2009: §4.10).

cases, an agent's plans relative to non-actual contingencies for update.

I see no way to build this sort of account given the tools available to the Standard Account, *even when* regular operations like ';' and '+' (and the sort of intersective operation mooted above) are thrown into the mix. The speech acts that the Standard Account regards as basic treat updates on an agent's actual cognitive state, as opposed to any of the agent's possibly non-actual or hypothetical or derived cognitives states, as basic. It is, as far as I can tell, impossible to define the latter sort of updates in terms of the former, even with regular operations at our disposal.⁷⁸

4.5.3 Challenging the Dynamic Program: The Modal Analysis

It is, moreover, far from clear that the Standard Account is *mistaken* about this. Why think that there *is* any neat and tidy operation, which might have a plausible claim to the mantle of 'speech act,' and by which a specific kind of contingency plan might be induced in an agent? Work in computer science and artificial intelligence attests to the fact that contingency planning (indeed, planning generally) is, as a psychological phenomenon, an enormously complicated endeavor. Even if it is acknowledged (as it clearly should be) that agents who accept CIs typically come to have corresponding contingency plans, it is by no means clear that there *is* any neat and tidy cognitive instruction (as we've been understanding speech acts to be), the following of which by an agent would yield the relevant state of mind for her.

This stands in marked contrast to the sorts of necessitation acts that we argued were expressed by ordinary, unembedded imperatives on their conventional interpretations. Necessitation acts, I suggested, involve proposing that an agent come to recognize a requirement. In standard cases, this can modeled simply as the agent adding a preference for the relevant state of affairs to her extant preferences (while seeing her realization of this state of affairs as a live practical option).⁷⁹ No such account seems to be forthcoming for CIs and the states of mind they are supposed to inculcate. And, coincidentally, any such account, if it did exist, would run afoul of some of the most cherished commitments of contemporary theorizing about the nature of speech acts. This is, by itself, sufficient to warrant a considerable amount of skepticism about whether there can be a viable speech act analysis of the meaning of CIs (and, thereby, about whether there can be a viable speech act analysis of linguistic meaning at all).

Into this explanatory vacuum steps the Modal Analysis of Imperatives.

78. We will see a stronger reason for thinking this is the case when I put forward my own analysis, which opts to define the former sorts of updates in terms of the latter.

79. Even in non-standard cases—cases requiring that the agent revise some of her prior preferences—the nature of the revision can be represented with an intuitive and natural cognitive operation.

The Modal Analysis. We first discussed the Modal Analysis of Imperatives (MAI) in Section 4.3.1. The MAI has a venerable pedigree,⁸⁰ and has been defended, in various forms, by Åqvist (1964); Lewis (1979a); Han (1998); Schwager (2006); Aloni (2007); Kaufmann & Schwager (to appear). The MAI’s account of the meaning of CIs is given by a simple semantic identity:

$$\llbracket(\text{if } \phi)(!\psi)\rrbracket := \llbracket(\text{if } \phi)(\Box\psi)\rrbracket$$

How does the MAI improve on the speech act-theoretic analyses of CIs explored just above? Not at all (or so I’ll eventually argue). As a prerequisite to that argument, it is worth getting clearer about the apparent advantages of the MAI over speech act accounts. We will spend some time doing that now.

Cognitive force as non-semantic side-effect. According to the MAI, the meaning of a CI like (4.97) is to be identified with the meaning of a corresponding modal sentence—in this case, something like (4.98)—rather than in terms of an *action* on a cognitive state that secures that state’s acceptance of the corresponding modal sentences. This is not to say that, on the MAI, updating on a CI does not secure a state’s recognition of the corresponding conditional requirement. The MAI, however, has no need to posit any *speech act* α such that the result of updating a state with α reliably ensures the state’s recognition of the corresponding conditional requirement. That, at first glance, seems to be a welcome fact. The evidence seems, at this point, to indicate that there may be no such speech act at all.

Here is an analogy, to better get the gist of the argument for the MAI I’m suggesting here.⁸¹ Here are two platitudes about indicative conditionals with declarative consequents:

- Updating with an indicative conditional causes an agent to believe a conditional proposition—perhaps the modal proposition that all near antecedent-possibilities are consequent-possibilities, possibly something more prosaic (it doesn’t really matter)
- Adams’ Thesis (Adams 1975): it is reasonable to identify the probability of an indicative conditional with the probability of the consequent conditional on the antecedent

Given these two platitudes, we would think it a natural methodology for theorizing about the meaning of indicative conditionals to attempt to locate a proposition such that its probability was equal to the relevant conditional probability. The problem is that, according to the Triviality Result of Lewis (1976), *there is provably no proposition* such that its probability

80. Green (1997) suggests that it dates to Kant’s *Groundwork!*

81. The arguments I describe here are in fact much more developed than those offered by fans of the MAI for their accounts, and against speech act accounts, of CIs. Schwager (2006), for instance, simply presupposes that there is no speech act that a speech act-theoretic account might assign to a CI except one constructed out of Standard speech acts. In particular, a speech act-theoretic account of CIs must make use of conditionalized speech acts, in the sense of Section 4.5.2. Kaufmann & Schwager (to appear) simply note variability in imperatives’ use, and assume, on these grounds, the unviability of any use-theoretic account of imperatives. As we have seen at great length (Chapter III), this is much too simplistic.

is equal to this conditional probability. So it is a very real possibility that the correct theory of the meaning for indicative conditionals is non-propositional in nature, even while:

- Updating with an indicative conditional does typically lead an agent to believe a conditional proposition
- Theorizing about meaning in terms of propositions is extremely fruitful for basic declaratives.

The fan of the MAI might exploit this analogy by suggesting the relation between CIs and speech acts to be analogous to the relation between indicative conditionals and propositions: useful, perhaps, for theorizing about easy cases, but misleading and, ultimately, problematic when it comes to more complicated cases. Notice further that, so long as Adams's Thesis is maintained, we should not expect to be able to fruitfully theorize about the meaning of an indicative conditional in terms of the conditional proposition that it typically causes an agent to believe. Further, the fact that an agent who comes to accept an indicative conditional comes to believe a conditional proposition will *not* be explained by saying that the indicative conditional conventionally expresses this conditional proposition. Rather, it will be explained as a *side effect* of incorporating the conditional's non-propositional conventional meaning into one's informational state.

Similarly the fact that an agent who comes to accept a CI comes to recognize a corresponding conditional requirement need not be explained by reference to one's favored account of its conventional meaning. This might be a side effect of incorporating the conditional's modal meaning into one's information state (and, in the end, the procedure by which the state's recognition of the conditional requirement is effected may fail to correspond to any suitable denizen of the algebra of speech acts at all).⁸²

Explanatory scope. It is perhaps an under-appreciated virtue of the MAI that, if it can be made to work as an analysis for basic imperatives, it can very likely be made to work for *any* type of imperative. Indeed, the MAI suggests an extremely powerful methodology for giving an account of the meaning for an *arbitrary* natural language imperative (or an arbitrary syntactic environment embedding an imperative)—a methodology which the MAI's account of CIs illustrates nicely. In summary form, the methodology seems to be something like this:

- i. Identify (using the tried-and-true method of data collection in linguistics, i.e., intuition) which modal sentence a natural language imperative, if uttered, would lead an agent to accept (or, more neutrally, which modal sentence an utterance of the imperative would cause to become true)⁸³

82. For our argument against Schwager's account of this procedure, see Section 4.3.1.

83. No version of the MAI of which I'm aware is actually explicit about either how it (a) associates modal LFs with natural language imperatives, (b) evaluates whether there is an intuitive match in meaning between these modal LFs and their analysanda. Insofar as I understand what such theorists are doing (which is not very far, I

- ii. Construct a translation procedure from regimented representations of imperative sentences to corresponding representations of modal sentences.⁸⁴ (Basically: for every ‘!’ in the regimented representation, the translation procedure substitutes the appropriate necessity operator.)
- iii. Associate natural language imperatives with regimented representations (logical forms), according to the modal sentence identified in Step (i) and the translation procedure described in Step (ii)
- iv. Provide an interpretation for the regimented modal language, and stipulate that the value of the interpretation function $\llbracket \cdot \rrbracket$ for any regimented imperative is identical to the value of the interpretation function for its modal translation

For CIs, the methodology is applied like this. CIs are closely related to conditional requirements. In particular, a CI like (4.97) is closely related to the conditional requirement described in (4.98) in the following sense: accepting (4.97) typically leads an agent to accept (4.98). We assign (4.98) a regimented representation (and a semantics) according to our best theory of the meaning of conditionalized modal sentences in natural language; Schwager (2006); Kaufmann & Schwager (to appear) favor a treatment along the lines of Kratzer (1981, 1991b,a). Finally, an imperative is assigned a regimented representation by using the translation procedure from its corresponding modal sentence, and the regimented representation is assigned an interpretation identical to that of its corresponding modal sentence.

The methodology extends easily to any natural language imperative you might think of. While we hemmed and hawed about the representation of a quantified addressee imperatives—on account of the fact that they seemed to require quantifiers to out-scope speech act operators—the MAI’s analysis of quantified addressee imperatives is automatic. Consider (4.103).

(4.103) Everyone take their seat

The methodology is applied as follows. Accepting (4.103) typically leads to the acceptance of an modal sentence like (4.104).

(4.104) Everyone should take their seat

must confess), I must construe them as employing the step described here.

84. This allows the MAI to be more neutral about the syntax of natural language imperatives than other commentators (e.g. von Stechow & Iatridou 2009) have suggested is possible. The MAI may or may not think that natural language imperatives contain “covert modals” (although Kaufmann & Schwager seem to think that they do). It may just happen that the interpretation of whatever syntax is distinctive of imperatives in natural language happens to be modal in nature—as, e.g., the interpretation of certain strictly non-modal expressions (e.g., *it is necessary that*) happens to be modal in nature.

The MAI assigns (4.103) the same regimented LF as we did, namely $\forall x : !x : P(x)$, since this is the LF that the translation procedure associates with its corresponding modal sentence:

$$\forall x \Box x P(x)$$

But, in contrast with the speech act account, the MAI has no difficulties interpreting this imperative LF. Its interpretation is simply identical to that of its corresponding modal sentence. No fancy machinery for building complex speech acts out of simple speech acts is required. (Nor, therefore, is any departure from the Standard Account of speech acts required.) All that is required is the off-the-shelf semantics of the sort supplied by quantified modal logic and Kratzer's work on linguistic modality.⁸⁵

4.5.4 Against the Modal Analysis

I have been building up the attractions of the MAI—somewhat misleadingly, it turns out, since I doubt whether any of these things should be regarded as *attractions* at all. My largest worry is that the MAI's claimed explanatory power, as described above, is an illusion. The MAI ultimately explains very little of interest to a theorist interested in the semantics and/or meaning of imperative clauses in natural language. Here I'll advance a few arguments for this claim.

Conditional practical questions. Let's first consider conditional practical questions (CPQs). CPQs are CQs whose consequents question the truth of some sentence bearing on what an agent ought to do.

(4.105) If it rains, should I shut the window?

It turns out that *all* of our worries for speech act accounts of CIs carry over to CPQs. And so it will turn out that similar considerations as those that motivated the MAI seem to motivate a modal analysis of CPQs. In other words, as we'll see, the meaning of a CPQ should, by the lights of the MAI, be identified with the meaning of the corresponding CI (and also with the meaning of a corresponding modal sentence).

85. Kaufmann & Schwager (to appear) advance a further, indirect argument for the MAI. They argue that CIs should typically be understood as strict conditionals whose consequents are imperatives (and so imperative consequents are thereby to be understood as embedded modalities, since speech act operators presumably cannot embed under modalities). There is much that is confusing about their paper, in particular, the fact their main source of data is intuitions about whether or not some CI or other is "true" or "false" in some described scenario. I hope I do not sound daft when I say that I cannot evaluate such intuitions for plausibility (since I do not understand what the states of affairs whose truth they claim to be intuiting are supposed to be like). My strong suspicion is that all of the data they give can be accommodated by the account I go on to develop below (making use, as I do, of dynamic analogues of the context-shifting that accompanies static semantic interpretation of iterated linguistic modalities). Substantiating that would, however, require a sustained effort at reconstructing their methodology—a burden more fairly placed on them than on me.

Wide-scoping? An agent can, of course, ask herself all of the following questions.

- (4.106) a. If it rains, should I bring only the umbrella?
b. If it doesn't, should I bring only the sunglasses?
c. Should I bring both?

She can, moreover, entertain a *yes* answer to each question. There is nothing incoherent about such an answer. Indeed, accepting a *yes* answer to each question intuitively just means planning to bring the umbrella in the contingency where it rains, to bring the sunglasses in the contingency where it doesn't, and to bring both in one's actual contingency. As we argued above, such a planning state is both coherent, and would be the sort of planning state an agent who answered *yes* to (4.106a) and (4.106b) (on their Ramsey Readings) and (4.106c) would ultimately end up in.

The problem is that, for anyone who treats '?' as a Standard Account-friendly speech act operator (and thus commits to wide-scoping this operator at LF), the following are the only real scopal possibilities for the sentences in (4.106):

- (4.107) a. $?(R \supset \Box(U \wedge \neg S))$
b. $?(¬R \supset \Box(¬U \wedge S))$
c. $?\Box(U \wedge S)$
- (4.108) a. $?\Box(R \supset (U \wedge \neg S))$
b. $?\Box(¬R \supset (¬U \wedge S))$
c. $?\Box(U \wedge S)$

Either scopal possibility predicts the incoherence of entertaining a *yes* answer to each of the questions in (4.106). For (4.107), the following reasoning is available to such an agent.

- Either it will rain or it won't. Suppose it rains. Then $\Box(U \wedge \neg S)$. Suppose it doesn't. Then $\Box(¬U \wedge S)$. So either $\Box(U \wedge \neg S)$ or $\Box(¬U \wedge S)$. But this contradicts $\Box(U \wedge S)$.


And, for (4.108), the following reasoning is available to such an agent.

- I have answered *yes* to each of these questions. So I will plan $(R \supset (U \wedge \neg S))$ and $(¬R \supset (¬U \wedge S))$ and $(U \wedge S)$. But then my plans are inconsistent—the only way they can be jointly realized is if $(R \wedge \neg R)$.

In either case, the agent is able to reason a priori (albeit in different ways) to some sort of inconsistency on her part. But this runs strongly counter to intuition—an agent who accepts *yes* answers to the questions in (4.106) would not be guilty of any sort of inconsistency.

Narrow-scoping? Another possibility is to narrow-scope the question operator—to treat (4.105)'s questioning force as conditioned, somehow, on whether it is raining. We will ignore the treatment on which the force is conditioned on the bare truth of the antecedent,

since it is, for reasons we've already seen, a non-starter. According to the analysis under consideration, then, (4.105) will receive something like the analysis given below.



This, I suggest, simply gets the “subject matter” of (4.105) wrong (in much the same way that the narrow-scoping analysis of CIs gets their subject matter wrong). An agent who entertains the CPQ in (4.105) is entertaining a question about what her contingency plans should look like *now*, i.e., at the time of entertainment. Specifically, she is entertaining whether or not to recognize a conditional requirement: in the contingency where it rains, should I plan to shut the window? She is not adjusting her cognitive state so that the question gets asked *on update with* the information that it is raining. She is asking herself a question *now*, about how to plan if it is raining. She is capable of answering this question now, even if the information available to her fails to decide whether or not it is raining.

Unsurprisingly, I'll suggest instead that the correct account of what's involved in entertaining a CPQ like (4.105) is given by the Ramsey Procedure. As we noted above, the Ramsey Procedure “tells an agent how to evaluate an indicative in the (standard) situation where the agent's actual cognitive state does not satisfy the property expressed by its antecedent.” According to the Ramsey Procedure, an agent evaluating (4.105) first shifts to a hypothetical information state from whose vantage point it is known to be raining. The agent then evaluates the consequent of (4.105) from the vantage point of this information state. Combining the Ramsey account with our analysis of practical questions (Section 4.3.9), then, evaluating (4.105) consists in raising the question whether or not to shut the window from the vantage point of this hypothetical information state. This is simply not what the narrow-scoping analysis of (4.105) has the agent doing.

I won't belabor the point. As was the case with CIs, a speech act-theoretic account of CPQs like (4.105) requires an apparatus suited to representing contingency planning, as well as an analysis of CPQs on which they involve doing stuff to an agent's plans relative to non-actual contingencies. Neither the wide-scoping nor narrow-scoping treatment of (4.105) succeeds here—even comes close—for reasons unsurprisingly similar to those that doomed similar treatments of CIs.

For the modal analysis? Into this explanatory vacuum steps... the Modal Analysis of Practical Questions (MAPQ). According to the MAPQ, the meaning of (4.105) is given by a simple semantic identity.

$$\llbracket (\text{if } \phi)(?\Box\psi) \rrbracket := \llbracket (\text{if } \phi)(\Box\psi) \rrbracket$$

The MAPQ, like the MAI, understands the cognitive force of a conditional practical question—which, it allows, the Ramsey Procedure has correctly identified—as a non-semantic side effect of its conventional interpretation. The MAPQ is also claimed to have wide explanatory scope; it can be made to work for any type of practical question, using something like the following methodology.

- i. Identify which modal sentence a practical question, if uttered, would lead an agent to entertain
- ii. Construct a translation procedure from regimented representations of practical questions to corresponding representations of modal sentences
- iii. Associate practical questions with regimented representations (logical forms), according to the modal sentence identified in Step (i) and the translation procedure described in Step (ii)
- iv. Provide an interpretation for the regimented modal language, and stipulate that the value of the interpretation function $\llbracket \cdot \rrbracket$ for any regimented imperative is identical to the value of the interpretation function for its modal translation

Puzzling practical questions—e.g., quantified subject practical questions like the following—are no match for this simple and elegant methodology.

(4.109) Should everyone take their seat?

There is no fancy machinery required for stating the MAPQ's analysis of such a construction (as would be required for stating a speech act-theoretic analysis). The LF of this sentence is a familiar quantified modal sentence, with an equally familiar interpretation.

Hooray for the MAPQ?

Against the modal analysis. Instead, let me voice an emphatic *boo* for the MAPQ. In assimilating practical questions to modal practical statements (the practical statements that practical questions introduce as issues for an agent), the MAPQ obviously obscures a distinctive and central feature of their meaning: the fact that they introduce practical *issues*. No one (I hope) would dream of entertaining the sort of analysis favored by the MAPQ, not least because, when paired with the MAI, the MAPQ ends up asserting an identity in meaning between imperatives, their corresponding practical statements, and their corresponding practical questions. In other words, the following sentences are predicted to be equivalent in meaning.

(4.110) Everyone take their seat!

(4.111) Should everyone take their seat?

(4.112) Everyone should take their seat.

While these sentences do, in some sense, share a common propositional core—a common “phrastic” (see Section 4.1)—identifying their meanings is a *prima facie* theoretical embarrassment. It is by no means an obvious, damning embarrassment; the theoretical identification of the meanings of a variety of kinds of clause, based on a shared propositional core, has occasionally been defended (see, again, Section 4.1). Nevertheless, I take it to be obvious that, if we can do better, we should. (I will, in Section 4.5.5, argue that we can.)

Is it possible to retain the MAI while avoiding the MAPQ? It might seem so, on the grounds that imperatives, maybe, just have a modal *flavor*, while their corresponding practical questions do not. Here are two reasons for thinking that this is misguided.

- As I have spent some time trying to establish here, the very considerations that tell in favor of the MAPQ also tell in favor of the MAI. If there is an argument for the MAI that cannot be transformed into an argument for the MAPQ, I am unaware of its existence.
- Insofar as imperatives have a modal “flavor,” practical questions do as well. The modal flavor of an imperative consists in its bearing a certain relationship to a corresponding modal sentence, namely, the relationship of being such that when the former is evaluated, the latter is typically *accepted* (the *making-accepted* relationship). But practical questions also bear a relationship to a corresponding modal sentence, namely, the relationship of being such that when the former is evaluated, the latter is typically *entertained* (the *making-entertained* relationship). These are different species of modal flavor, but I see no principled reason for privileging the modal flavor of imperatives over the modal flavor of practical questions on this basis.

Explanatory “power” through explanatory inertness. As the prior argument suggests, the MAI’s only real explanatory leg-up on the competition is its quietism.⁸⁶ Its strategy for avoiding the troubling predictions of wide- and narrow-scope accounts of CIs can be summed up as follows. The analysis has the right to remain silent. What it does not say cannot be used against it in a linguistic court of law. *Wovon man nicht richtig sprechen kann, darüber soll man schweigen.*

The problem is that what the analysis *does* say is not very explanatory. The analysis notes the modal flavor of imperatives, develops an analysis of basic imperatives on that basis (on which they are semantically just modal statements), and extends the template used in developing that analysis to imperatives of arbitrary complexity. This is a theory, to be sure, but a theory which does not tell us very much of interest about the meaning of imperatives in natural language.⁸⁷ What it does tell us about their meaning is this: imperatives have something semantically and structurally in common with their corresponding modal statements. But it turns out that practical questions, by the arguments

86. Again, I am ignoring the improvements in data coverage claimed in Kaufmann & Schwager (to appear). For my reasons, see note 85.

87. I think this is probably what Portner has in mind when he remarks: “A modal which only had a performative use might as well not be called a modal at all. The performative aspect of its meaning... would explain everything that needs to be explained about its meaning” (2007: 366).

used to motivate the MAI, will also have these things in common with their corresponding modal statements. And so the obvious distinctions in conventional meaning between an imperative, its corresponding modal claim, and its corresponding practical question remain theoretically obscure.⁸⁸

In conclusion, we—not just the fan of the MAI—confront a serious theoretical challenge. For the fan of the MAI, the primary challenge is to adduce arguments for the MAI that cannot be adapted into arguments for the MAPQ. (I am dubious that there are any such arguments, to say nothing of any *good* ones; even the arguments that are adaptable into arguments for the MAPQ are unconvincing, as I've argued here.) For the fan of a speech act-theoretic account of imperatives, the challenge is to do better. Doing better will require alleviating the pressures that pushed us to the MAI and away from a speech act-theoretic account. Specifically, the speech act theorist will require an account of CIS that does justice to their obvious meaning. We will put forward such an account in Section 4.5.5.

4.5.5 A New Account

Our methodology. Like the MAI, we, too, have an automatic “account” (albeit one with a few promissory notes) of the meaning of arbitrary imperatives in natural language (and, so, of conditional imperatives). The methodology for constructing such an account, in fact, looks a good deal like the methodology we suggested for the MAI.

- i. Identify which modal sentence a natural language imperative, if uttered, would lead an agent to accept.
- ii. Construct a translation procedure from regimented representations of imperative sentences to corresponding representations of modal sentences. (As before, the basic idea is that, for every ‘!’ in the regimented representation, the translation procedure substitutes the appropriate necessity operator.)
- iii. Associate natural language imperatives with regimented representations (logical forms), according to the modal sentence identified in Step (i) and the translation procedure described in Step (ii)
- iv. Provide an interpretation for the regimented modal language (specifically, conditions on when a state accepts a sentence in the language).
- v. Provide an interpretation for the regimented imperative language, on which an imperative’s interpretation consists in a suitable speech act mapping an input state to an output state that satisfies the corresponding modal sentence.

88. Or, perhaps worse, relegated to the pragmatics. Although, as we noted in Section 4.3.1, Kaufmann & Schwager do try to account for the typical “performative” force of imperatives by appeal to lexicalized presuppositions of imperative operators, their account does not work. Perhaps some alternative account could be devised, which distinguishes imperatives from modal statements from practical questions, not in terms of their at-issue meaning, but in terms of their presuppositions or conditions of use. But I strongly doubt that it will be as theoretically attractive as an account which distinguishes these kinds of sentences as a matter of their at-issue meaning.

The crucial difference with the MAI’s methodology lies in step (v). The meaning of an arbitrary imperative is still intimately related to the meaning of its corresponding modal sentence. But it is not *identified* with the meaning of its corresponding modal sentence. Instead, its meaning consists in the fact that it leads a state that accepts it to accept the corresponding modal sentence (a state of mind, I’ve argued, that should be identified as recognizing a corresponding requirement).

It’s not hard to see that this is the sort of methodology we were implicitly employing in Sections 4.3 and 4.4 (although we weren’t as fussy about the regimented modal language, or about the translation procedure between that and the regimented imperative language, as we might have been). Where the account went off track with CIs was in step (v). In the case of the other constructions we considered—conjoined, disjoined, and quantified imperatives—defining an appropriate sort of cognitive update relation was not difficult to accomplish on the account of cognitive states—as, effectively, a set of beliefs and a set of desires—with which we were working. The sorts of accounts of CIs that seemed most natural within this paradigm—wide- and narrow-scoping, respectively—were suited only to performing updates (whether deferred or immediate) on such cognitive states. And so we were left wondering, legitimately, if there was any sort of account within the paradigm that could yield a satisfying treatment of CIs.

This section will develop, perhaps more briefly than would be ideal,⁸⁹ such an account. Step (i), of course, is trivial. A CI of surface form $(if \phi)(!\psi)$ leads an agent to recognize a conditional requirement—to accept a modal sentence of the form $(if \phi)(must(\psi))$. We will assign the latter kind of a regimented representation of the following form.

$$\underline{\phi}(\blacksquare\psi)$$

The notation is chosen to represent the Ramsey Procedure; evaluating an indicative involves supposing the antecedent and evaluating the consequent from the standpoint of this supposition. The black box is chosen to distinguish it from a normal box; any state that accepts a sentence of the form $\blacksquare\phi$ also accepts $\square\phi$ (i.e., it requires ϕ). Additionally, though, a state that accepts $\blacksquare\phi$ also recognizes this requirement. (The black box thus fills the role that *must* filled in the discussion of Section 4.3.) The imperative is accordingly assigned the following regimented representation.

$$\underline{\phi}(!\psi)$$

All of this is basically automatic. The remainder is somewhat less so.

Consider a CI that causes a state to accept a sentence of the form $\underline{\phi}(\blacksquare\psi)$. The CI is thus assigned an imperative LF of the form $\underline{\phi}(!\psi)$. At this point we can say at least that, *if* we can

89. For further discussion of this account, see Charlow (2010a,b, 2011a) and Appendix B.

find a suitable value for $\underline{\phi}(!\psi)$, it will meet the following condition:

$$\Sigma[\underline{\phi}(!\psi)]\Sigma' \Rightarrow \Sigma' \models \underline{\phi}(\blacksquare\psi)$$

That is to say, the update potential of the CI must express a relation R between an input state Σ and output state Σ' such that $\Sigma R \Sigma'$ only if Σ' accepts the CI's corresponding modal sentence. We cannot, however, say much more than that.

Our tasks. Earlier, we said that CIs update intuitively an agent's contingency plans (while leaving the notion of contingency planning murky). Now we are saying that CIs update an agent's cognitive state so that it comes to accept a corresponding modal sentence. It is, however, *far from clear that these are the same thing*. My first task will be to develop a rough account of contingency planning, and to say how imperatives update contingency plans. After that, my task will be to show that such an update does in fact correspond to a serviceable (indeed, independently motivated) notion of a state coming to accept the CI's corresponding modal sentence. This latter task is not so easy as it might at first seem. The standard treatment of deontic conditionals—due to Angelika Kratzer—has nothing to do with the model I have developed to represent contingency planning. I will argue briefly (drawing primarily on arguments I have made in other work) that this is a fault of her treatment, rather than a fault of our account of contingency planning.

Contingency planning. The notion of contingency planning is given by the following notions.

- A *contingency*, for an agent, is a possible information state for that agent.
- A *plan for C*, for an agent, is a plan relative or indexed to C , where C is a contingency for that agent.

Earlier arguments have, I think, amply motivated the notion that contingencies should be typed as information states, and I won't belabor the point here. We will make the idealizing assumption that information-change is generally monotonic in nature. What this means, formally, is this:

Definition 4.81. Given a basic cognitive state σ , $I_{\sigma'}$ is a **possible future information state** for σ iff (i) $I_{\sigma'}$ partitions $\bigcup I_{\sigma}$, (ii) $\forall p \in I_{\sigma'} : \exists q \in I_{\sigma} : p \subseteq q$ and (iii) $\forall q \in I_{\sigma} : \forall p \in I_{\sigma'} : p \cap q = \emptyset$ or $p \cap q = p$.

That's to say: I is a possible future information state only if (i) I partitions the possibilities that are live at I , (ii) every cell in I contains at least as much information as some cell in I_{σ} (this means no information loss, since $\bigcup I \subseteq \bigcup I_{\sigma}$), (iii) every issue recognized by I_{σ}

is either recognized at I or resolved at I . Information-change, in a slogan, is treated as information-growth.

Planning or preference states can no longer be typed as pairs of an agent's preferences and the practical issues the agent recognizes. Instead, we should think of an agent's *global planning state* as a contingency plan, specifying, for each possible future information state of the agent, a plan (or, to use our old terminology, an action-guiding state) for that contingency.

Definition 4.82. \mathfrak{P}_σ is a **global planning state** for σ iff \mathfrak{P}_σ is a function from information states into action-guiding states (in the sense of Definition 4.40), such that, for each possible future information state for σ $I_{\sigma'}$, $\mathfrak{P}_\sigma(I_{\sigma'})$ is an action-guiding state for σ' .

Definition 4.83. Π is a **local planning state** for σ iff, for some possible future information state for σ $I_{\sigma'}$, $\Pi = \mathfrak{P}_\sigma(I_{\sigma'})$.

Definition 4.84. $\sigma = \langle I_\sigma, \mathfrak{P}_\sigma \rangle$ is a cognitive state iff I_σ is an information state and \mathfrak{P}_σ is a global planning state for σ .

Contingency planning, in a slogan, is planning what to do, given monotonic changes in one's information.

Evaluating indicatives. Indicative conditionals are generally evaluated using the Ramsey Procedure.⁹⁰ That is to say: evaluating a CI, relative to a basic information state, involves supposing its antecedent and evaluating its imperative consequent with respect to this hypothetical information state. In our framework, this can be represented using a three-step procedure:

- Select a possible future information state, namely, an information state from whose vantage point the CI's antecedent is "known". More formally, if the starting state is σ , the state moves temporarily to $\sigma[\triangleright\phi] = \langle I_{\sigma[\triangleright\phi]}, \mathfrak{P}_\sigma \rangle$.
- Evaluate the imperative consequent at the selected information state

90. The treatment developed here obviously owes a great deal to the analyses of indicatives suggested in Karttunen (1973, 1974); Heim (1983); Kaufmann (2000); Veltman (2005); Isaacs & Rawlins (2008); Gillies (2010). I also owe thanks to Alex Silk and Paul Portner for discussion that helped me to realize that the more complicated treatment of CIs and deontic conditionals I defend in Charlow (2010a) can be elegantly and simply rendered using the dynamic or context-shifting analysis of indicative conditionals (indeed, the dynamic apparatus I make use of here was suggested to me by Portner in conversation). Although the formal implementation differs, I wish to stress that the analysis of CIs presented here is in all important respects the same as the analysis presented in my (2010a) (although my understanding of necessitation acts has become more nuanced), and the analysis of deontic conditionals is the same as the analysis presented in my (2011b). The current analysis, in essence, repackages a static treatment of indicative antecedents as context-shifters as a dynamic treatment of indicative antecedents as context-shifters.

For similar, independently developed views about the relationship between indicative antecedents and preferences/plans, see Silk (ms) and Cariani et al. (ms). Although similar in inspiration, there are fundamental differences between these approaches and mine. I regret that I lack the space to explore them here.

- Return to the original information state (retaining the modification performed at the selected informations state)

There are at least two ways of implementing this sort of picture, and not a lot, as far as I can tell, to decide between them.

Method 1: Specifying a target. The first way represents the described process in the formal representation of the speech act expressed by the CI, but does not actually define a staged update procedure that is expressed by the CI. The CI's antecedent simply selects a contingency for update and modifies the state's local planning state for that contingency accordingly. In other words, letting $C = I_{\sigma[\triangleright\phi]}$:

Definition 4.85. $\mathfrak{P}_\sigma[\underline{\phi}(!\psi)] = \mathfrak{P}_\sigma \cup \{\langle C, \mathfrak{P}_\sigma(C)[!\psi] \rangle\} \setminus \{\langle C, \mathfrak{P}_\sigma(C) \rangle\}$

Of course, an agent who goes about executing this sort of update may engage in the diachronic procedure outlined above. But, on this way of thinking, these sorts of psychological kinematics need not be part of the theory of meaning for CIs. The meaning of a CI specifies a target property for the agent's cognitive state, and the agent goes about modifying her cognitive state, so that it comes to realize that property, in the appropriate way.

Method 2: Getting to the target. The second way actually does define a staged update procedure that is expressed by the CI. Interpreting a CI at σ involves (i) moving to a hypothetical state, namely, $\sigma[\triangleright\phi]$, (ii) updating the hypothetical state by evaluating the imperative consequent at $\sigma[\triangleright\phi]$, (iii) returning to the original information state, in such a way that the resulting information state satisfies the property of recognizing a requirement that ψ in the contingency given by $I_{\sigma[\triangleright\phi]}$.

Formalizing the second strategy requires defining a formal apparatus that can distinguish the agent's actual state from hypothetical states, represent temporary shifts to hypothetical states (and updates on those states), and represent how the agent's actual state is modified when the hypothetical state is relinquished. Here I'll adopt the "stack" metaphor (and its associated formalism) favored by [Kaufmann \(2000\)](#); [Isaacs & Rawlins \(2008\)](#). Stacks represent the cognitive state of an agent, construed to include the suppositions that are in force for that agent, as well as whatever cognitive facts hold under those suppositions. Stacks are typed as ordered sequences of states, whose last element of the sequence is the agent's actual state, the penultimate element of the sequence is the hypothetical state resulting from making a supposition p , the next element is the hypothetical state resulting from making a supposition q on top of the original supposition p , and so on. . .

Definition 4.86. \mathcal{S} is a stack iff

- $\mathcal{S} = \langle \Sigma \rangle$ (where Σ is a state), or

ii. $S = \langle \Sigma, S' \rangle$, where Σ is a state and S' is a stack

In general, an update instruction will either:

- Add a state to the top of the stack (introduce a supposition)
- Modify the state at the top of the stack
- Remove a state from the top of the stack (relinquish the supposition), colloquially known as “popping” the stack.

The function of an indicative antecedent is to introduce a supposition, i.e., to add a hypothetical state to the top of the stack.

$$\langle \Sigma, S \rangle[\underline{\phi}] = \langle \Sigma[\triangleright\phi], \langle \Sigma, S \rangle \rangle$$

Update on normal kinds of clauses—declaratives, interrogatives, and imperatives—always involves modification of the state at the top of the stack. Restricting our focus to imperatives, the idea is that:

$$\langle \Sigma, S \rangle[!\phi] = \langle \Sigma[!\phi], S \rangle$$

The function of an utterance occurring within the “scope” (whether syntactic, as in the case of indicative consequents, or discourse-level, as in the case of modal subordination) of a supposition is to modify the hypothetical state introduced by that supposition. Formally, then, the proposal for CIs is this.

$$\langle \Sigma, S \rangle[\underline{\phi}(!\psi)] = \langle \Sigma[\triangleright\phi][!\psi], \langle \Sigma, S \rangle \rangle$$

What happens next? In the case of conditional *questions*, as Isaacs & Rawlins (2008) persuasively argue:

- Updating the top element of a stack with a question makes the stack resistant to popping that element, until the issue introduced by the question is resolved. The supposition introduced by a conditional question cannot be relinquished until the issue introduced at the suppositional state is resolved. (Their “Inquisitiveness Constraint,” p. 294.)
- Any information introduced at a hypothetical context (by way of resolving the issue introduced there by the conditional question’s consequent) will percolate downward, so that, when the hypothetical state is removed from the stack (when the stack is “popped,” the supposition relinquished), any possibility eliminated as live at the hypothetical state is also eliminated as live from all lower states.

These facts do not, however, extend straightforwardly to CIs. Imperatives, we might say (incorrectly, in my view), do not introduce issues. They tell an agent how to plan at the contingency corresponding to the suppositional state. And so there is no reason to think that the supposition introduced by a CI stays around any longer than it takes to update

the hypothetical state with the CI's imperative consequent. Nor is it immediately clear how the update performed by an imperative at a hypothetical state should percolate to the lower states. Imperatives do not introduce information; they update plans. But it is not, for instance, the case that a possibility that is no longer permissible (eliminated from the resolved permissibility sphere) at a hypothetical state should be regarded as impermissible at the agent's actual state; its hypothetical impermissibility may be a function of the circumstances introduced by the relevant supposition, and it may, for all that, be perfectly acceptable with respect to the actual contingency.⁹¹

The *what happens next?* account of CIs that I would favor has the following two parts.

- If an imperative introduces a practical issue at the top of the stack (if, e.g., it is a choice-offering disjunctive imperative), the stack cannot be popped until the agent resolves the practical issue.⁹² Otherwise, the stack is popped immediately.
- Updating the top element of a stack with an imperative modifies all lower states Σ of the stack so that, in the contingency C corresponding to the information available at the top element, the local planning state for C at Σ is the result of updating with the imperative

The first part should probably not be represented in our linguistic analysis of imperatives; it is more plausibly thought of as a pragmatic or rational constraint on how best to deal with practical questions. The second part is expressed with the following equation. Letting $!\psi$ be an imperative and ϕ an arbitrary supposition, the idea is this:

$$\langle\langle I_{\triangleright\phi}, \mathfrak{P} \rangle, \mathcal{S} \rangle[!\psi] = \langle\langle I_{\triangleright\phi}, \mathfrak{P} \rangle[!\psi], \mathcal{S}' \rangle,$$

(Where $\mathcal{S}' = \langle \Sigma'_1, \dots, \Sigma'_n \rangle$ and $\Sigma'_i = \Sigma_i[\underline{\phi}(!\psi)]$, as defined in Def. 4.85)

Deciding between these? Certain linguistic constructions introduce temporary states or contexts that can nevertheless *persist*; temporary does not need to be interpreted as ephemeral. This fact is often cited to explain the fact that context-sensitive expressions uttered after the introduction of a temporary context can sometimes be interpreted with respect to that temporary context, rather than the actual or basic context—a phenomenon known as *modal subordination* (see esp. Roberts 1989).

(4.113) A: A thief might break in.

B: Would she steal the silver? (Isaacs & Rawlins 2008: 296)

91. This is a fact that the Kratzer account of deontic conditionals gets badly wrong, as we'll see a little below.

92. But can't an agent just postpone deciding what to do in a contingency until the contingency is actual? Yes, but, similarly, an agent can postpone answering a question raised at a hypothetical context (or, indeed, any context at all) until she has more information. There is nevertheless *pressure* for the agent to resolve the issue before she pops the stack. If she cannot resolve the issue, it is stored (although this is as a second-best alternative to resolution). For a sketch of how *issues* might percolate and thereby be stored once the top stack is popped (which is readily extendable to our analysis of practical issues), see Isaacs & Rawlins (2008: 301–3).

(4.114) A: If a thief breaks in, she'll steal the silver.

B: She'd take the jewelry, too.

It seems at the very least quite handy to have the stack device at our disposal to represent such discourse-level phenomena. Modal subordination (or discourse subordination, generally) can then be explained as a reflex of the fact that a temporary context is introduced but not popped prior to subsequent utterances.

That is not to say that stacks are necessarily an indispensable part of a theory of conventional linguistic meaning for indicative conditionals (and other temporary context-inducing kinds of clauses), even for a discourse-centric theory of meaning. It might be a pragmatic effect of indicative antecedents that they are capable of introducing temporary contexts, which are subsequently available (depending on various context-dependent facts about whether the pretense of speaking from the point of view of a temporary context is conversationally warranted) for the interpretation of later utterances. The conventional meaning of, e.g., a conditional question might be exhausted by the fact that it selects a local information state for update. The conventional meaning of a CI might be exhausted by the fact that it selects a local planning state for update. Stacks may figure in the representation of the relevant states of mind, but they will not be exploited by the theory of meaning for the various kinds of indicative conditionals; the theory of meaning consists solely in specifying a target for the "basic" state or context to meet. So far, I see little to decide between these ways of doing things, at least for indicatives with declarative or interrogative consequents.

With CIs, though, it might seem that there is a bit more empirical traction. Notice that modal subordination data is *very* different for imperatives. As the above examples show, a speaker can both ask a question (by uttering an interrogative) and make an assertion (by uttering a declarative) relative to a temporary context. But a speaker *cannot* happily issue a request (by uttering an imperative) relative to a temporary context. And the same holds for corresponding strong necessity modals.

(4.115) a. A thief might break in.

b. #Run away from her!

(4.116) a. A thief might break in.

b. #You must run away from her!

(4.117) a. If a thief breaks in, hide the silver!

b. #Run away from her!

This might seem to suggest that it is a mistake to treat an utterance of a CI as somehow equivalent in meaning to the utterance of an imperative within the "scope" of a temporary context. For, were that the case, then the sequence in (4.115) should be roughly equivalent

in meaning to the CI in (4.118).⁹³

(4.118) If a thief breaks in, run away from her!

But clearly they are not equivalent in meaning. More damningly, the contrast here seems, at first glance, to be explained by the resistance of imperative clauses (and corresponding strong necessity modals) to interpretation with respect to non-actual or hypothetical states or contexts. That, on its face, seems like pretty bad news for the stack analysis of CIs, which analyzes CIs as expressing necessitation acts relative to such temporary contexts, and pretty okay news for the target-specifying analysis).

That, though, is a bit too quick. Notice that other kinds of clause display roughly the same sort of contrast, when the relevant modal is stripped of its counterfactual morphology. So while (4.114) is fine, (4.119) sounds somewhat odd.

(4.119) A: If a thief breaks in, she'll steal the silver.

??B: Will she take the jewelry?

??B: She'll take the jewelry.

This, I think, suggests that counterfactual morphology is, in such cases, typically used to signal the intended persistence of a temporary context introduced by a prior utterance. This isn't necessarily to say (contra Isaacs & Rawlins) that the default, even in the case of conditional questions, is for a stack to be popped subsequent to the interpretation of an utterance that introduced a temporary context. Nor is it to say that counterfactual morphology must subsequently be used to recall that context, no longer being treated as actual for the sake of conversation, for further update. We might equally say that certain kinds of clausal kinds—imperative, future-tensed indicative—*induce* popping of a stack *when they do not subordinate* (syntactically or discursively) a temporary context-inducing clause. Certain clausal kinds have what we might call a sobering effect; they typically remove pretense or supposition and refocus an agent and her interlocutors on the actual circumstances.

The terminology here is liable to be confusing. The type of subordination I'm talking about here is an extension of the familiar syntactic notion. Indicative antecedents are syntactically subordinated to their consequents; the antecedent, in the standard terminology, is the *subordinate clause*, while the consequent is the *matrix clause*. On the stack account, however, the *interpretation* of the consequent is subordinated to the temporary context introduced by the antecedent. Modal subordination involves the latter sort of subordination. The kind of subordination I'm talking about here involves the former. It extends beyond an

93. On the idea that indicative antecedents might be analyzed as raising possibilities, in the same manner as modals expressing epistemic possibility, see Gillies (2007, 2010).

exclusively syntactic notion to cover sequences like the following.

- (4.120) a. If a thief breaks in, hide the silver.
b. Then run away.

I call this subordination because of the intuitive equivalence of this sequence with the following CI, in which the imperative (4.120b) is part of a complex imperative that syntactically subordinates the antecedent:

- (4.121) If a thief breaks in, hide the silver and run away.

The idea, then, is that utterances of clausal kinds with a “sobering” effect must somehow signal (perhaps via some device expressly conveying continuation, like *then* or *too*, perhaps via non-verbal cues) that they are to be interpreted with respect to a previously introduced temporary context (cf. Roberts 1989). The continuation, in effect, indicates a kind of cross-sentential subordination relationship—one holding, in the paradigm case, between (i) the result of composing the later utterance with the matrix clause of the prior utterance and (ii) the subordinate clause of the prior utterance. Thus an interpreter interprets the sequence in (4.120) in roughly the same way she interprets (4.121). Similar explanations hold for the felicity of the sequences below.

- (4.122) A: If a thief breaks in, she’ll steal the silver.
B: Will she *also* take the jewelry?
B: She’ll *also* take the jewelry.

This explanation, in fact, actually seems more explanatory to me than the alternative (on which a stack is popped by default, and the counterfactual morphology of a subsequent modal serves to recall it). Not only does it predict the infelicity of the sequences in (4.115) and (4.119). It correctly predicts that the infelicitous part of the sequence can be repaired by adding a device signaling formation of a cross-sentential subordinating structure.

Our official stance between the target-specifying and stack analyses of CIs will, then, be neutral. Stacks are nice for making phenomena like the ones explored here a bit more transparent, but it is not strictly essential for the theory of meaning to make use of them. That said, I will adopt the stack analysis for the remainder of our discussion of CIs, as it has the welcome, if not theoretically essential, property of encoding the evaluation procedure for a CI as part of its associated update instruction.

4.5.6 The Semantics of Conditional Requirements

The real test for either analysis of CIs is whether it meets the criteria of the methodology laid out in Section 4.5.5 (in particular, whether the suggested interpretation for CIs predicts that update on an arbitrary CI causes a state to accept its corresponding modal sentence). It

is hard to tell, at present, since we don't yet have an interpretation for sentences of the form $\underline{\phi}(\blacksquare\psi)$ on the table.

Kratzer's standard analysis. The standard theory of interpretation for conditional requirement sentences is the one exploited by the modal analysis of CIs—a theory due to Kratzer (1981, 1991b,a). It is an elegant and extremely successful treatment of such sentences, and indicative conditionals with modal consequents generally, and for that reason carries a very strong presumption in its favor.

Kratzer famously presents a collection of linguistic evidence (which I won't recap here) for the thesis that “If-clauses are devices for *restricting the domains of [quantificational, e.g., modal] operators*” (Kratzer 1991b: 656).⁹⁴ We will call her treatment of modals a *Generalized Quantifier* treatment (since she analyzes modals as generalized quantifiers, taking both restrictor and scope arguments), and her treatment of indicative antecedents a *Restrictor* treatment (since they fill the restriction argument for a generalized quantifier). Together, they comprise a *Restrictor Semantics* for modalized indicatives.

The Restrictor Semantics for natural language indicatives of the form $(if \phi)(must(\psi))$ interprets the necessity modal as a binary universal quantifier over deontically best indices in the modal base, the domain of which is restricted by the indicative's antecedent. In a slogan: $(if \phi)(must(\psi))$ is true iff the best ϕ -possibilities are ψ -possibilities. More formally:⁹⁵

Definition 4.87. $\llbracket \underline{\phi}(\blacksquare\psi) \rrbracket^\sigma = 1$ iff $min(\cup I_\sigma \cap \llbracket \phi \rrbracket^\sigma, \preceq_{\Pi_\sigma(I_\sigma)}) \subseteq \llbracket \psi \rrbracket^\sigma$

Informally, Definition 4.87 says that $\underline{\phi}(\blacksquare\psi)$ is true, relative to a state or context σ , iff the best (according to σ 's preferences) σ -live ϕ -possibilities are all ψ -possibilities. The strong deontic necessity modal \blacksquare is interpreted as expressing a binary quantifier Q . When $\blacksquare\psi$ occurs as the consequent of an indicative conditional, the antecedent specifies a domain restriction on Q , and the prejacent of \blacksquare specifies the condition asserted for Q -many individuals in the restricted domain. Unembedded (“bare”) modal sentences are treated as vacuously restricted, so that $\llbracket \blacksquare\psi \rrbracket := \llbracket \top(\blacksquare\psi) \rrbracket$. The relationship between bare and embedded modals is analogous to that between *everyone loves their mother* and *everyone under 12 loves their mother*.

Now, we have of course been arguing for construing preferences—the source of the ordering on live possibilities relative to which the domain of quantification is determined—as relativized to contingencies. On the Kratzer analysis, the ordering source is simply supplied by a context or a state (rather than a state and a contingency), but that's no problem. As is standard, we simply allow the state or context or evaluation to fill in the values of all open parameters, hence allow the state or context of evaluation to supply

94. For indicatives whose consequents lack overt modals (e.g., *if he has his umbrella, it's raining*), Kratzer posits a covert quantifier at LF. For critical discussion of this part of Kratzer's proposal, see Gillies (2010: Sect. 9).

95. Kratzer does not recognize practical issues or the notion of requirement-recognition, and since they are not germane to the discussion here, I will ignore them here.

the contingency of evaluation (as we have in fact done in Definition 4.87). The best ϕ -possibilities, according to σ 's preferences, are the best ϕ -possibilities, according to σ 's preferences relative to her actual circumstances.

But there is a problem. Our analysis of $\phi(!\psi)$ has this CI modifying a state σ 's contingency plans at the contingency supplied by $\sigma[\triangleright\phi]$. And, so, there is a mismatch between the update performed by a CI on a state and the satisfaction conditions for its corresponding conditional requirement sentence. A CI modifies an agent's preferences at a contingency, rather than the preferences that instruct her how to plan in her current state. But the [Kratzer](#) semantics for conditional requirements does not look at an agent's non-actual, not-yet-realized preferences for the ordering source. It looks at what the agent actually wants. A conditional requirement of the form $(if\ \phi)(must(\psi))$ expresses that every actually permitted way of realizing ϕ will also realize ψ . And that, if you think a bit about it, seems very, very intuitive. In short, our suggested update procedure for a given CI does nothing to alter an agent's stance toward the corresponding conditional requirement, on the standard [Kratzer](#) semantics for conditional requirements. That is a serious problem.

A problem for the Kratzer semantics. Taking it as given that CIs establish corresponding conditional requirements, it must be the case that either our update for CIs is wrong, or the [Kratzer](#) semantics for conditional requirements is wrong. Here, I'll argue it is the latter.

Here is a case I discuss in much detail elsewhere, but which makes the point against the [Kratzer](#) semantics vivid.⁹⁶ Consider the following case from [Kolodny & MacFarlane \(2010\)](#). Ten miners are trapped in a single shaft—A or B, although the agent does not know which—and threatened by rising waters. The agent can block one shaft or neither, but not both. If she blocks the shaft they are in, all are saved. But, if she guesses wrong, all die. However, if she does nothing, water will distribute between the shafts and exactly one will die. Now consider the following set of conditional requirements.

(4.123) If they're in A, I ought to block A $in_A(\blacksquare block_A)$

(4.124) If they're in B, I ought to block B $in_B(\blacksquare block_B)$

(4.125) I may leave both shafts open $\neg\blacksquare(block_A \vee block_B)$

Given the case, informants reliably hear each of these obligation-descriptions as *true* (so, *a fortiori*, consistent). But, using the [Kratzer's](#) semantics (and supposing that the agent knows that the miners are all in A or all in B) these sentences are *provably inconsistent*.

Proof. Suppose $\llbracket(4.123)\rrbracket^\sigma = \llbracket(4.124)\rrbracket^\sigma = \llbracket(4.125)\rrbracket^\sigma = 1$ and $\cup I_\sigma \subseteq \llbracket in_A \vee in_B \rrbracket^\sigma$.

- i. Choose any $v \in \min(\cup I_\sigma, \preceq_{\Pi_\sigma(I_\sigma)})$.
- ii. Since $\cup I_\sigma \subseteq \llbracket in_A \vee in_B \rrbracket^\sigma$:

96. The discussion here draws some material from my (2011b).

$$v \in \min(\cup I_\sigma \cap \llbracket in_A \rrbracket^c, \preceq_{\Pi_\sigma(I_\sigma)}), \text{ or}$$

$$v \in \min(\cup I_\sigma \cap \llbracket in_B \rrbracket^c, \preceq_{\Pi_\sigma(I_\sigma)})^{97}$$

- iii. By Definition 4.87, since $\llbracket (4.123) \rrbracket^\sigma = \llbracket (4.124) \rrbracket^\sigma = 1$:
 $\min(\cup I_\sigma \cap \llbracket in_A \rrbracket^\sigma, \preceq_{\Pi_\sigma(I_\sigma)}) \subseteq \llbracket block_A \rrbracket^\sigma$, and
 $\min(\cup I_\sigma \cap \llbracket in_B \rrbracket^\sigma, \preceq_{\Pi_\sigma(I_\sigma)}) \subseteq \llbracket block_B \rrbracket^\sigma$
- iv. So $v \in \llbracket block_A \rrbracket^\sigma \cup \llbracket block_B \rrbracket^\sigma$.
- v. So $\min(\cup I_\sigma, \preceq_{\Pi_\sigma(I_\sigma)}) \subseteq \llbracket block_A \rrbracket^\sigma \cup \llbracket block_B \rrbracket^\sigma$.
- vi. So, by Definition 4.87, $\llbracket \blacksquare(block_A \vee block_B) \rrbracket^\sigma = 1$. \perp

□

The problem, moreover, stems from the very property discussed above: evaluating conditional requirements, on the Kratzer treatment, involves evaluating the relationship holding between actually preferred antecedent-possibilities and consequent-possibilities. We see this, in particular, at step (ii) of the proof. The fact that conditional requirements utilize the agent's actual preferences to order possibilities allows us to infer the following fact about the agent's actual preferences (given the truth of the conditional requirements). Since, according to her actual preferences, the best worlds where the miners are in A are worlds she blocks A (and likewise for B), and since the miners must either be in A or B, she actually does prefer a state of affairs in which she blocks either A or B. But that is incorrect: according to her actual preferences, she does not prefer a state of affairs in which she blocks either A or B.

From the point of view of the Ramsey Procedure, this is not exactly surprising. What actions an agent actually prefers, relative to a stage of a given episode of practical reasoning, are often influenced by the agent's available information. The agent does, of course, strictly prefer the state of affairs in which all ten miners are saved, but, since her actual information does not tell her how to achieve this state of affairs, she works with a less demanding preference in her actual circumstances—perhaps, to save nine. Even if the miners turn out to be in A, it is not appropriate for her, given her actual information, to treat the fact that blocking A would save all ten as a consideration in favor of blocking A. But it is appropriate to treat the fact that blocking neither saves nine as consideration in favor of blocking neither, since her available information tells her how to achieve this state of affairs. And so her actual circumstances, together with the preferences that appropriately guide her decision making in such circumstances, leave this course of action open to her.

But notice that *enriching the agent's information* tends, *ceteris paribus*, to make it *easier for the agent to implement more demanding preferences*, hence, easier for it to be appropriate for the agent to treat the fact that an action realizes a relatively more demanding preference

97. This step relies on a kind of monotonicity property of the Kratzer semantics: if $u \in \min(\Phi, \preceq)$, then for any Ψ such that $\Psi \subseteq \Phi$ and $u \in \Psi$, $u \in \min(\Psi, \preceq)$.

as a consideration in favor of selecting that action. Of course, according to the Ramsey Procedure, the first step in evaluating a conditional requirement is to enrich your information with the information contained in the requirement's antecedent. This changes the agent's preferences over actions accordingly. From the vantage point of an information-state bearing the information that the miners are in shaft A, the agent of course knows how to save everyone; she need only block shaft A. From the vantage point of such an enriched information-state, then, it is appropriate for her to treat the proposition that blocking A will realize the basic end of saving all ten as a reason to block A.

Kratzer's Restrictor Semantic analysis of conditional requirement sentences, then, fails for the simple reason that it does not correctly implement the picture of the analysis of indicative conditionals suggested by the Ramsey Procedure. The antecedents of conditional requirement sentences do not restrict the domain of the requirement operator in the consequent. Rather, from a dynamic perspective, it is better to say that they shift the standpoint of interpretation for the requirement sentence in the consequent. Shifting this standpoint alters the circumstances that are salient in the agent's evaluation of the consequent (and, consequently, may alter the preferences she sees fit to deploy in deciding which possibilities are preferred for the sake of evaluating the consequent). We will, in the next section, see what an analysis in this vein might look like. We will then show that it meshes perfectly with the analysis of CIs suggested above.

Implementing the Ramsey Procedure in an analysis of conditional requirements. The Ramsey Procedure suggests an analysis of conditional requirements along the following lines. Evaluating a conditional of the form $\underline{\phi}(\blacksquare\psi)$ involves:

- i. Moving (temporarily) to $\sigma[\triangleright\phi] = \langle I_{\sigma[\triangleright\phi]}, \mathfrak{P}_\sigma \rangle$
- ii. Checking to see whether $\sigma[\triangleright\phi] \models \blacksquare\psi$.
If *yes*, then $\sigma \models \underline{\phi}(\blacksquare\psi)$. If *no*, then $\sigma \not\models \underline{\phi}(\blacksquare\psi)$.

We will say, in general, that a basic state σ' accepts a sentence of the form $\blacksquare\psi$ iff the best possibilities, relative to $\mathfrak{P}_{\sigma'}(I_{\sigma'})$, all satisfy ψ . In other words, a state accepts a sentence of the form $\blacksquare\psi$ just if updating it with $!\psi$ is redundant.⁹⁸ (Again, we are ignoring the recognitional aspect of requirement sentences for the sake of simplicity.)

$$\sigma' \models \blacksquare\psi \text{ iff } \min(\bigcup I_{\sigma'}, \preceq_{\Pi_{\sigma'}(I_{\sigma'})}) \subseteq \llbracket \psi \rrbracket^{\sigma'} \text{ iff } \sigma' [!\psi] = \sigma'$$

Lifting this to sets of basic states, we have:

$$\Sigma' \models \blacksquare\psi \text{ iff } \Sigma' [!\psi] = \Sigma'$$

98. Implicitly, then, updating on a conditional requirement sentence (or, indeed, any requirement sentence whatever) will involve *testing* one's cognitive state, in the sense of Veltman (1996), to see whether it has some action-guiding property or other. For much more on this, see Appendix B.

Now let us return to the stack apparatus we have been using for CIs. The analysis of conditional requirement sentences suggested by our understanding of the Ramsey Procedure will, then, look like this:

$$\langle \Sigma, \mathcal{S} \rangle [\underline{\phi}(\blacksquare\psi)] = \langle \Sigma[\triangleright\phi][\blacksquare\psi], \langle \Sigma, \mathcal{S} \rangle \rangle$$

To see whether this update has the desired characteristics, we require a notion of what it is for a stack to accept a speech act or typed sentence. This is a generalization of the ordinary dynamic understanding of acceptance (as redundant update).

Definition 4.88. $\langle \Sigma, \mathcal{S} \rangle$ accepts a speech act of the form $\underline{\phi}(\alpha)$, notation $\langle \Sigma, \mathcal{S} \rangle \models \underline{\phi}(\alpha)$, iff $\Downarrow \langle \Sigma[\triangleright\phi][\alpha], \langle \Sigma, \mathcal{S} \rangle \rangle = \langle \Sigma, \mathcal{S} \rangle$

The \Downarrow simply expresses the popping operation: the result of removing the top state from the stack after its updates have been allowed to percolate to lower states in the stack. The idea, then, is that a stack accepts, e.g., an indicative conditional iff supposing the antecedent, then updating on the consequent, and then jettisoning the supposition leaves the state unchanged—iff, that’s to say, there is no residue that remains once the episode of supposition induced by evaluating the indicative is ended.

Putting all of this together, is easy to see that updating on a CI, and subsequently discharging the supposition of its antecedent, leads a cognitive state (here understood as a stack) to accept a corresponding conditional requirement.⁹⁹

Proof. Let $\langle \Sigma, \mathcal{S} \rangle$ be a stack. Then $\langle \Sigma, \mathcal{S} \rangle \models \underline{\phi}(\blacksquare\psi)$ iff (by Definition 4.88)

$$\Downarrow \langle \Sigma[\triangleright\phi][\blacksquare\psi], \langle \Sigma, \mathcal{S} \rangle \rangle = \langle \Sigma, \mathcal{S} \rangle \text{ iff (since, when defined, } \Sigma'[\blacksquare\psi] = \Sigma'[\!|\psi], \text{ for any } \Sigma')$$

$$\Downarrow \langle \Sigma[\triangleright\phi][\!|\psi], \langle \Sigma, \mathcal{S} \rangle \rangle = \langle \Sigma, \mathcal{S} \rangle \text{ iff (by Definition 4.88)}$$

$$\langle \Sigma, \mathcal{S} \rangle \models \underline{\phi}(\!|\psi) \quad \square$$

This is perhaps a bit abstract, but the driving intuitions are quite straightforward. A speech act-theoretic account CIs should, I have suggested, have them modifying contingency plans (by updating an agent’s plan at the contingency characterized by the antecedent with the consequent). And conditional requirements check contingency plans to see if modifying with a CI—updating the agent’s plan at the contingency characterized by the antecedent with the consequent—would be redundant. So it is no surprise that this result holds. Imperatives express necessitation acts. CIs express necessitation at a contingency. And conditional requirements check whether something is necessitated at a contingency.

Ad hoc? Indeed, these analyses fit together so well that it might be worried whether the analysis of conditional requirements is proposed ad hoc, so that it generates the welcome prediction with respect to CIs. Consideration of the miner case shows why this is not so.

⁹⁹ I’m only considering imperatives of the form $\!|\psi$ here, but it would be easy to extend this proof to imperatives embedded under regular operations and quantifiers.

As that case shows, it is critical for an analysis of conditional requirements to allow the ordering on possibilities by which an agent determines how to act to shift with available information, and to allow suppositional episodes of the sort induced by entertaining the antecedent of an indicative conditional to be associated with such shifts. Evaluating a conditional requirement $\underline{\phi}(\blacksquare\psi)$, on our account, means supposing ϕ and then checking to see whether the resulting state is decided for ψ . As I have argued elsewhere (see my 2011b), this is precisely the sort of analysis that is required to give a satisfactory account of our intuitions about the miner case.

4.5.7 Conditional Imperatives as Fundamental

In light of all this, it is useful to ask ourselves how the necessitation analysis of normal, unconditional (and otherwise unembedded) imperative clauses is to be understood. And there is a strong case to be made that, like in Kratzer's treatment, it is LFs for conditional constructions, rather than the LFs for unconditional constructions, that are fundamental. Here I'll advance two arguments for thinking of the sorts of speech acts expressed by simple imperatives as special, rather than fundamental, cases, and for thinking of the sorts of speech acts expressed by CIs as fundamental, rather than special, cases.

Stacks of different lengths. One way to see this is to ask what, in fact, simple imperatives do to a state, when a state is understood as a stack \mathcal{S} (i.e., when we insist on representing contingencies and suppositions in our apparatus for representing cognitive states). When the stack is of length 1—i.e., for some Σ , $\mathcal{S} = \Sigma$ —update with a simple imperative of the form $!\phi$ simply leads Σ to treat ϕ as required (in the sense of Definition 4.74). In other words, the state simply comes to recognize a requirement that ϕ . But when the stack is of length ≥ 2 —i.e., for some ϕ , Σ and \mathcal{S}' , $\mathcal{S} = \langle \Sigma[\triangleright\phi], \langle \Sigma, \mathcal{S}' \rangle \rangle$ —things go differently. Recall the notion that updates to contingency plans percolate to lower states in a stack.

$$\langle \langle I_{\triangleright\phi}, \mathfrak{P} \rangle, \mathcal{S}^{\$} \rangle [!\psi] = \langle \langle I_{\triangleright\phi}, \mathfrak{P} \rangle [!\psi], \mathcal{S}^* \rangle,$$

(Where $\mathcal{S}^* = \langle \Sigma_1^*, \dots, \Sigma_n^* \rangle$ and $\Sigma_i^* = \Sigma_i^{\$}[\underline{\phi}(!\psi)]$, as defined in Def. 4.85)

A proposal to update a stack of length 1 with a simple imperative is the closest thing to a bare necessitation act that is recognized by the contingency-centric account we've developed here. But this sort of speech act should, it seems clear, be treated as a special case of this more general procedure—a special case in that there are no lower states in the stack to which the update can percolate; the percolation is *vacuous*. Similarly, one-place quantifiers—quantifiers taking only a scope argument, and no restriction argument—are generally treated as a special case of two-place generalized quantifiers—quantifiers taking both domain-restriction and scope arguments—because when no restriction argument is explicitly supplied, the restriction is treated as *vacuous*.

The fundamental case, then, in terms of which bare necessitation acts are defined, is the case in which update to the top element of a stack lead to update of contingency plans at every state further down the stack. This is the sort of speech act canonically expressed by a CI. So it is conditional necessitation acts, rather than necessitation acts sans phrase, that are fundamental in the theory of speech acts we have developed here. To be maximally explicit about the nature of the speech act expressed, then, all imperatives should receive LFs of the form $\underline{\phi}(!\psi)$. In the case of a simple imperative, the supposition-indicator will be filled with a tautology.

Selection of contingencies. This first argument might seem to rely too heavily on the stack apparatus—the necessity of which we questioned in Section 4.5.5. Here is a similar argument for the thesis that does not.

CIs propose updates to contingency plans. But of course simple imperatives do this too—they simply propose update to whichever contingency the agent is treating as actual, at the time of update, for the sake of practical reasoning. Indeed, in our setup, all planning is contingency- or information-relative. There are no necessitation acts, sans phrase. There is, strictly speaking, only necessitation *with respect to some contingency or other*. The distinction between conditional and simple imperatives amounts to a different method of specifying the relevant contingency plan for update. In the case of a typical CI, the argument through which a contingency is specified is filled non-vacuously. In the case of a simple imperative, the specification is vacuous. Applying Definition 4.85, a simple imperative of the form $!\phi$ performs the following update on a contingency plan.

$$\mathfrak{P}_\sigma[\underline{\mathbb{I}}(!\phi)] = \mathfrak{P}_\sigma \cup \{ \langle I_{\sigma[\triangleright\top]}, \mathfrak{P}_\sigma(I_{\sigma[\triangleright\top]})[!\phi] \rangle \} \setminus \{ \langle I_{\sigma[\triangleright\top]}, \mathfrak{P}_\sigma(I_{\sigma[\triangleright\top]}) \rangle \}$$

But since, for any σ , $I_{\sigma[\triangleright\top]} = I_\sigma$, this simplifies to:

$$\mathfrak{P}_\sigma[\underline{\mathbb{I}}(!\phi)] = \mathfrak{P}_\sigma \cup \{ \langle I_\sigma, \mathfrak{P}_\sigma(I_\sigma)[!\phi] \rangle \} \setminus \{ \langle I_\sigma, \mathfrak{P}_\sigma(I_\sigma) \rangle \}$$

This—necessitation with respect to the contingency treated by the agent as actual—is the closest thing to a bare necessitation act that the account we've proposed recognizes. But it is, for all that, still necessitation with respect to a contingency.

Although this argument does not rely on the stack apparatus, it may be helpful to see what it means in the context of that analysis. The idea would be that a simple imperative of the form $\underline{\mathbb{I}}(!\psi)$ updates a stack $\langle \Sigma, \mathcal{S} \rangle$ to $\langle \Sigma[\triangleright\top][!\psi], \langle \Sigma, \mathcal{S} \rangle \rangle$. But since the supposition is vacuous—i.e., $\Sigma[\triangleright\top] = \Sigma$ —this simply expresses necessitation with respect to Σ . But necessitation with respect to Σ is still, as we have just seen, necessitation with respect to a contingency.

4.5.8 Upshot

This represents, as it were, the final nail in the coffin of the Standard Account of Speech Acts. Our discussions of conjoined, disjoined, and quantified imperatives all undermined the Standard Account's ontological commitment to a domain of speech acts populated only by force-content complexes. Our discussion of conditional imperatives (and indicative conditional in general) undermines the Standard Account's commitment to the *explanatory fundamentality* of force-content complexes in the theory of speech acts.¹⁰⁰ In the case of imperatives, the fundamental speech act is necessitation with respect to a contingency—which corresponds to a complex operation on a cognitive state, typically involving supposition. Although, once a contingency has been fixed—a supposition made—we still make use of good old Standard-friendly necessitation acts, such necessitation acts are not properly regarded as *speech* acts at all. Such necessitation acts lack a canonical mode of expression in natural language. Simple imperatives express necessitation with respect to a vacuously specified contingency, rather than necessitation sans phrase.

Recall that our objective in this Chapter has been not just to argue against the Standard Account, but additionally to support a dynamic or speech act-theoretic account of the meaning of typed clauses in natural language (and, in particular, imperative clauses in natural language). Regarding this objective, an opponent might object—recalling the argument we used to motivate the Modal Analysis of Imperatives in Section 4.5.3—that we have not really argued that the update we have defined as being conventionally expressed by CIs should be regarded as a *speech act* at all. Our account of update on a CI is perhaps just a very crude model of a very complicated kind of cognitive change (the sort of change involved in contingency planning). There is no reason to think a proposal for an agent to undertake such a cognitive change has much in common with the traditional notion of a speech act at all.

I have several replies to this. First, it is a terminological matter whether or not such a proposal gets labeled as a speech act. We have stated a successful theory for a wide class of natural language imperatives that analyzes their meanings in terms of the sorts of updates on contingency plans they propose. It is indisputably a *dynamic* or *expressivist* account—one that is fundamentally interested in the behavior of the relevant language in discourse (use) and the conventionalized connection between language and states of mind. Whether or not it is really a speech act account is ultimately unimportant.

Second, I don't think there is any good reason to withhold the label of *speech act* from the sort of proposal we have defined. It is a platitude that (simple) imperatives can be used to express speech acts. Since simple imperatives, like CIs, are ultimately analyzed as proposals to update contingency plans, there is no reason to withhold the label of *speech act* from the proposal conventionally expressed by a CI.

Finally, as to the charge of crudeness, every account of CIs and conditional requirements

100. For our exposition of the Standard Account's ontological and explanatory commitments, see Section 4.4.2.

must, I argued, make use of some sort of notion of contingency planning (to account, e.g., for the standard intuitions about the miner case). *No account* that I have considered in this chapter gets this right, besides our own. Our model of contingency planning is, no doubt, crude, and ripe for improvement along any number of dimensions. But it offers *real explanations, of real phenomena*, where other accounts offer *none*. As we argued in our discussion of the Modal Analysis of Imperatives (Section 4.5.3), quietism in the face of complexity is a rather thin basis on which to critique an alternative account.

CHAPTER V

Concluding Remarks

5.1 Taking Stock

The prior chapters have done many things, and I will not try to summarize all (or, really, any) of them here. The overarching theme, though, has been something like this. Dynamic accounts of imperatives are beset by an array of problems that have not, I've argued, been given sufficient attention by their friends. The problems, though, can be resolved. And their resolutions, in most cases, actually make clear that dynamic accounts have an explanatory *edge* over their static, particularly modal, competitors.

Solving these problems *does*, I argued, require fairly large departures from the usual ways of theorizing about such things, and a few new theoretical and formal tools. While such departures might appear as costs, *prima facie*, I think they are ultimately, when considering the scope, power, and unity of the resulting view, properly viewed as virtues. In Chapter III, we developed a formal account of linguistic convention, argued that imperatives conventionally express necessitation acts, and used these ideas to explain a wide array of non-standard interpretations of imperatives. In Chapter IV, we developed a model of necessitation acts on which they were understood as proposals to cause an agent to recognize a requirement (and a novel formal apparatus to represent such a state of mind), and used this account to model a host of closely related phenomena having to do with this state of mind (accepting an imperative, making a decision, entertaining a practical question, weak versus strong permission, etc.). We showed how to extend this account to embedded imperatives of all stripes, by making use of regular operations. Nearly all extant accounts have difficulty with nearly all of the problematic phenomena that these tools and accounts were developed to explain. Our account handles all of them.

5.2 Editorial Digression

I think the most probable reason for the inattention to or failure to resolve the bulk of these problems is that they are generally of a methodological or philosophical, rather than properly "linguistic," nature. The tasks of, e.g., developing a suitable account of

linguistic convention, or of situating a dynamic account of linguistic meaning within (and appreciating the theoretical constraints imposed by) the vast field of speech act theory, are not the sort of tasks one typically sees being pursued in leading linguistics journals.

This has rather distorted the linguistic dialectic. The debate between fans of dynamic or speech act-theoretic or use-theoretic accounts of meaning, on the one hand, and fans of static or model-theoretic accounts of meaning, on the other, although rich and extraordinarily fruitful, has, I think it is fair to say, suffered from inattention to such concerns. Theorists do not, in many cases, even agree on the basic rules of the game. Fans of static accounts, as we saw, fault dynamic accounts on the grounds that clause-type cannot “strictly determine the speech act [a clause is] used for” (without offering arguments that strict determination is a commitment of a dynamic account) (Kaufmann & Schwager to appear). Fans of dynamic accounts fault static accounts for failing to predict the characteristic use of a typed clause (usually without being clear about why we would expect a theory of meaning to give such predictions). There is a lot of talking past each other here, at least some of which might be ameliorated by more careful attention to the sorts of foundational concerns with which I have often been concerned in this dissertation.

Since I am primarily a philosopher, I of course have a professional interest in stressing the relevance of philosophical and methodological concerns to linguistic theorizing—an interest I’m going to indulge a bit here. This interest is nevertheless, I think, well-founded. If this dissertation has accomplished anything, I hope it has shown that philosophers, *qua philosophers*, can offer substantive contributions to the scientific study of meaning in natural language. Many philosophers nowadays actively participate in debates among academic linguists. That’s all to the good. But this participation is sometimes characterized (understandably, given how incredibly technical contemporary work in linguistic theory, in particular formal semantics and pragmatics, has become) by excessive deference to linguistic expertise. But, as with all scientific inquiry, inquiry into linguistic meaning involves a host of presuppositions about the nature of the inquiry, how best to understand the explanatory aims of the inquiry, and so on. Different understandings of such presuppositions, as we’ve seen over and over again, lead to vastly different accounts of the relevant explananda (and often to disagreement about what the relevant explananda actually *are*). Philosophers who understand the relevant empirical work can contribute to this inquiry simply by leveraging their own expertise in foundational matters.

Although a large portion of what this dissertation has done lies in contributing to the empirical side of the inquiry (by giving new accounts of the identified phenomena), that work is implicitly structured and guided by our conception of the inquiry’s subject-matter and aim. Few of the pieces of the account that we’ve endorsed would have been possible to either formulate or appreciate in its absence. So, if we have made any empirical progress here, or gleaned further understanding of the meaning and use of practical language, I would say it is due, in large part, to the attention we’ve paid to foundational matters. This

sort of methodological stance animates nearly all of my work on linguistic meaning, and is especially on evidence in the appendices included below.

5.3 Meaning, Force, Semantics, and Logic

By way of brief introduction of the included appendices, these are intended to illustrate, in different ways, applications of the sort of view of practical language defended in the above chapters to questions about the meaning, cognitive profile, and semantic properties of *normative*, as opposed to imperative, practical language.

As I noted all the way back in Section 3.4.5, one enormous virtue of the sort of view I have been defending here is that it—together with a semantic theory on which the semantic value of an utterance is identified with the set of states that accept it—can be seen to *explain* the default connection between an utterance expressing a semantic value of some type or other and its conventional force. It is no mystery that a syntactic object that semantically expresses the cognitive property of recognizing a requirement would be conventionally associated with the speech act of necessitation. Nor is it a mystery that a syntactic object that semantically expresses the cognitive property of believing that *p* would be conventionally associated with the speech act of asserting that *p*. The connection between clause-type and force is extremely tight. No ancillary machinery to explain it is required.¹

It is very important to note that endorsing such a connection does not in any sense require identifying the *semantic value* of an utterance with the speech act it conventionally functions to express; endorsing such a connection does not, in a slogan, force the adoption of Dynamic Semantics, in the loaded sense of Kamp (1981); Groenendijk & Stokhof (1991); Veltman (1996). Utterances might simply have as their semantic values characteristic functions of cognitive states (equivalently, sets of admissible cognitive states, in the sense of Swanson 2008), rather than update functions (functions mapping input states into output states). Characteristic functions of cognitive states are, of course, *derivable from* update functions (the function mapping a state to 1 iff it is in the update function's co-domain), but it is important to notice that they are capable of playing entirely different theoretical roles, in virtue of what they are supposed to be representing. Update functions *represent speech acts* (forces, cognitive profiles, etc.), and it is highly doubtful, for reasons we will see in Appendix A, whether these representanda have the right sorts of properties for doing the sort of work we would normally expect from a semantic theory—explaining inconsistency, defining a suitable consequence relation, and so forth.

But characteristic functions of cognitive states are not necessarily in the business of representing speech acts. They are instead, I will suggest in Appendix A, in the business of accounting for *synchronic properties* of a cognitive state: whether it is admissible relative

1. Contrast Portner (2004a)'s view, on which imperatives semantically denote ordinary (non-cognitive) properties. On this view, predicting the conventional force of an imperative requires positing an independent pragmatic principle, to the effect that utterances that semantically denote properties are to be interpreted as proposals for addition to a salient ordering source.

to a sentence (satisfaction), whether two sentences can be jointly accepted relative to a single cognitive state (consistency), and so on. Our theory of speech acts tells us what such characteristic functions must look like (and, so, is ultimately explanatorily fundamental in the theory of meaning), but our account of such synchronic properties need not be *part of* the theory of speech acts at all. For all we have said, it may be there is both a diachronic, or transitional, dimension of linguistic meaning, as well as a synchronic, or static, dimension of linguistic meaning. It may, further, be the case that facts within diachronic dimension should be regarded as explaining facts about the synchronic dimension (even though explanations of phenomena within the synchronic dimension are not ultimately provided by the theory of the diachronic dimension at all). A dynamic or speech act-theoretic theory of *meaning* does not, in short, necessitate Dynamic Semantics or a speech act-theoretic theory of, e.g., semantic inconsistency.

This is a fact that has, I will show, serious implications for the debate between meta-ethical Expressivists and Cognitivists over the nature of the meaning of normative language. Similarly for the issue discussed in Appendix B. In that appendix, I give some new arguments for abandoning traditional, quantificational understandings of the meaning of indicative conditionals (Stalnaker 1968, 1975; Lewis 1973; Kratzer 1991a), in favor of a dynamic or Expressivistic account of indicatives with deontic and imperative consequents (together with a dynamic analysis of the meaning of those consequents), much along the lines of the account sketched in Section 4.5. This opens the door to a sophisticated, use-/force-/speech act-theoretic account of the meaning of such conditionals—something that has eluded Expressivism and related forms of meta-ethical Non-Cognitivism since their earliest days. As we saw in Section 4.5, we will see that this sort of account has a real explanatory edge on its quantificational competitors. And, in view of the distinction noted here (and explored in Appendix A) between fundamental linguistic meaning and semantic value, it simply has no need to guard its flank against the most trenchant philosophical criticisms of dynamic theorizing about the meaning of such constructions.

To end on a rather cliché note, these appendices serve, I think, to illustrate, how future work on dynamic theories of practical language (and dynamic theories of linguistic meaning in general) should be conceptualized. They also provide express examples of how future work on practical language within the dynamic program in both philosophy and linguistics might actually go.

In sum, the need for philosophical work to guard the flank of the dynamic program, in both its philosophical and linguistic incarnations, is gone (or at least seriously ameliorated). Let the theory-building begin.

APPENDICES

APPENDIX A

Meaning for Expressivists

In this essay, I defend and develop Expressivism as a serious, empirically plausible theory of the meaning of normative language. I do *not*, however, defend any extant form of Expressivism. Rather, I identify, and ultimately reject, a latent presupposition (“[Meaning Reductionism](#)”) about the relationship between a theory of meaning and a semantic theory—one routinely made by both Expressivists and their critics (in particular, advocates of Frege-Geach trouble for Expressivism). According to [Meaning Reductionism](#), if things of kind *K* are fundamental in a theory of meaning for a language, we must explain any fact about the meaning of expressions of that language by appeal to properties of *K*’s. I show how [Meaning Reductionism](#) leads to unpalatable consequences when applied to theorizing about the meaning and semantics of imperative clauses. Namely, it seems to generate specious Frege-Geach trouble for *non-propositional* treatments of *imperative* clauses. So I suggest replacing it with a new view of the meaning/semantics interface—one on which theorizing about meaning is theoretically *prior* to semantic theorizing, but which allows semantic explanations a kind of theoretically *autonomous* status. I use this new orientation to develop a non-propositional account of imperatives which avoids Frege-Geach trouble, but which is recognizably Expressivist in motivation and substance. I show how this sort of account can serve as a blueprint for an account of *normative* language. Finally, I argue that, while this account is importantly different from extant Expressivist treatments, it is, nevertheless, a clear-cut form of Expressivism.

A.1 The Expressivist Core

Expressivism involves a familiar set of methodological and empirical commitments. In this section and the next, I’ll (i) describe them and (ii) show how the kind of Expressivism that dominates philosophical discussion actually is the result of joining these commitments to a further—and, for the expressivist, *optional*—commitment about the relationship between theorizing about meaning and theorizing about semantics: [Meaning Reductionism](#).

This suggests a tantalizing prospect: an expressivist theory of meaning which does not presuppose Meaning Reductionism.

A.1.1 Meaning as Use

Expressivism is, in part, a theory about the nature of linguistic meaning. The central Expressivist thesis about meaning is that use is *explanatorily fundamental* in explaining linguistic meaning. In Gibbard's well-known formulation, we "explain the meaning of a term" by explaining "what states of mind the term is used to express" (Gibbard 2003: 5–6). Putting this in terms of theoretical desiderata:

CONVENTIONAL MEANING AS USE (CMAU)

The subject matter of the correct theory of meaning for a fragment \mathcal{L} is an account of how competent speakers of \mathcal{L} conventionally use sentences of \mathcal{L} in communication.

This formulation is rather more general than Gibbard's. For Gibbard is not neutral about the sense of 'use' appropriate to theorizing about linguistic meaning: Gibbardian use is a matter of the *state of mind* conventionally expressed by literal utterances of an expression. But other understandings of the notion are certainly available. To give one example, the sense of 'use' appropriate to theorizing about linguistic meaning might be a matter of *discourse role*, so that the subject matter of a correct theory of meaning for sentences of \mathcal{L} is an account of what sort of move in the conversational "language game" (cf. Lewis 1979b) utterances of sentences of \mathcal{L} are conventionally used to make.¹

The Expressivist's posited relation between language and use, I am assuming, must be *conventional* in nature: an expression's use must supervene directly on its *lexical properties*—on the properties that, very roughly, would be recorded in a correct and complete dictionary entry for that expression. Or, what basically amounts to the same thing, its use must be determined by its lexical properties together with "law-like" statements connecting such properties to facts about use). It is not extracted as the conclusion of an inference using general, social-linguistic conventions, together with facts about the context of utterance, as premises (e.g., the sort of Gricean reasoning by which conversational implicatures are usually thought to be generated).² A fact about an expression's interpretation, in a context,

1. "Hybrid" versions of Expressivism typically take this latter form. According to Stevenson (1944); Hare (1952)—and to a lesser extent, Copp (2001); Barker (2000); Boisvert (2008)—an account of meaning for a normative fragment is an account of the distinctive, conventionalized illocutionary force of sentences of that fragment. For recent discussion of Hybrid theories, see Schroeder (2009, 2010); Alwood (2010). As the account developed here will illustrate, there is nothing essentially Hybrid about discourse role Expressivism.

2. One important contrast between the reasoning by which conventional meanings and conversational implicatures are inferred is that the latter sort of reasoning is *defeasible* (by, e.g., the acquisition of further information about the context of utterance), while the former sort is not. This is not to say that if a sentence ϕ has conventional meaning M , ϕ must, in every context, be interpreted, by the speaker's audience, as meaning M . The fact that ϕ means M , together with facts about the context c , is often exploited by a speaker's audience to interpret her as intending to communicate a meaning inconsistent with M via her utterance of ϕ in c . See Asher & Lascarides (2001) for a nice discussion of these themes.

must be conventional, in this sense, to have any real claim to be part of the meaning of that expression. Conversely, any fact about an expression's interpretation, in a context, that is conventional, in this sense, has a strong *prima facie* claim to being part of its meaning, even if that fact is not actually recorded in its lexical entry.

A.1.2 Practicality

Expressivism is not, of course, just a view about the subject-matter of a theory of linguistic meaning. It is distinguished as a substantive view about the meaning of *moral* (and, more broadly, *normative*) language. Expressivists claim a basic divergence in use (hence in meaning) separating descriptive and normative language.

NORMATIVE MEANING AS PRACTICAL

The conventional use of normative language is (principally) *practical*, while the conventional use of descriptive language is (principally) *representational*.

The description of this claim is vague—deliberately so, since Expressivists cash out the practical/representational distinction in different, if broadly related, ways. State of mind Expressivists often draw a distinction in terms of *direction-of-fit*: descriptive and normative language serve, respectively, to express mental states with world-to-mind direction-of-fit and mind-to-world direction-of-fit—e.g., states of accepting a norm (Gibbard 1990) or planning to do something in a given contingency (Gibbard 2003).³ Discourse role Expressivists distinguish between conversational moves that portray the world in a certain way, and which aim mainly at altering the context's depiction of the world (i.e., assertions in the sense of Stalnaker 1978) and those involving some public display of an attitude with mind-to-world direction-of-fit—an attempt to influence the behavior of the addressee (a directive speech act; see, e.g., Stevenson 1944; Hare 1952), or perhaps a mere indication or implication of the speaker's possession of such an attitude, constituting no direct attempt, *as such*, at influencing addressee behavior (see, e.g., Barker 2000; Copp 2001; Boisvert 2008).

It helps to be clear about the direction of explanation here. Expressivists needn't (and typically don't) claim there is any *meaning-independent* criterion or procedure for categorizing expressions as normative or descriptive, the results of which subsequently determine what sort of use (practical or representational) it is appropriate for an Expressivist theory of meaning to assign those expressions. This sort of claim would be problematic: normative constructions (unlike, e.g., directive or interrogative constructions) do not come "labelled" as such. Normative clauses, in particular, lack a distinctive syntax, mood, or clause-type (on the notion of a clause-type, see Portner 2004a; Sadock & Zwicky 1985). Such categorization, in fact, can occur only "post"-interpretation. An expression is categorized according to *how*

3. The distinction is not felicitously drawn in terms of belief/desire psychology. There is a deflationary sense in which the judgments expressed by normative language can be called beliefs (see, e.g., Blackburn 1988, 1998). The natural way of defining an inflationary notion of belief is in terms of direction-of-fit.

it is actually used in thought or communication. Unlike the Expressivist's commitment to **CMAU**, which is *methodological* in nature, the nature of her commitment to the practicality of normative language, properly understood, is *empirical* in nature: whether there is any normative language, in the Expressivist's sense, depends on whether there is any language whose conventional use is principally practical.

A.1.3 Non-Propositionalism

Finally, Expressivism is generally recognized as a form of meta-ethical anti-Realism that is distinct from Error Theory. For my purposes, we can crudely characterize meta-ethical Realism as the conjunction of the following two views:

PROPOSITIONALISM ABOUT NORMATIVE CONTENT

The content, or meaning, of a normative sentence or claim at a context can be faithfully rendered or represented with a *proposition*.

ACCURACY

The content of at least some normative sentences at at least some contexts is true.

Error Theory, as usually understood, agrees with Realism that the content of normative sentences is propositional, but denies that any such contents are accurate.⁴

If Expressivism is to be distinguished as a form of anti-Realism distinct from Error Theory, it must deny that the meaning of a normative sentence can be faithfully represented with a proposition. And, indeed, the Expressivist's other theoretical commitments *commit her to this*. Why? Suppose the meaning of a normative sentence ϕ is represented with a proposition. Either the fact that ϕ 's meaning is a proposition encodes information about ϕ 's conventional use, or it doesn't.

- If it does, that information is compatible with ϕ being used representationally (since descriptive sentences plausibly mean propositions, if anything does, and their meaning makes it the case, according to an Expressivist, that their conventional use is representational). But, in view of Expressivism's baseline commitment to **CMAU**, *meaning must determine use* for an Expressivist. So, since ϕ is by stipulation a normative sentence, ϕ 's meaning is incompatible with its being used in a way that is not primarily practical.⁵

4. See, e.g., Mackie (1977). Error Theory is usually described as claiming that all normative sentences are false. However, since, for any normative sentence ϕ , neither ϕ nor $\neg\phi$ is, for the Error Theorist, true, it seems better to say that Error Theory is committed to describing all normative sentences as exhibiting something like *catastrophic presupposition failure* (cf. Yablo 2006).

5. There are, of course, some normative sentences (e.g., those invoking "thick" terms like 'brave') that, by the Expressivist's own lights, have a use that is *both* representational and practical (see Gibbard 1990: 112–7). But in the case of a sentence ϕ with *purely normative* conventional meaning, the Expressivist is committed to saying that ϕ cannot be used in a way that is not primarily practical.

- If it doesn't, ϕ 's meaning does not determine its use. But *meaning must determine use*, for the Expressivist.

So, for the Expressivist, the meaning of a normative sentence cannot be faithfully represented with a proposition. This is *not* to say that the meaning of a normative sentence cannot be *partially* represented with a proposition. Nor is it to say that the meaning of a normative sentence cannot somehow *characterize* a proposition (in, perhaps, the way that propositions plausibly characterize, without reducing to, functions from worlds to truth-values). Nor, finally, is it to say that the Expressivist cannot allow the possibility of *any theoretical purposes* for which it would be useful to view normative sentences as characterizing propositions.⁶

It is, rather, just to say that, for the Expressivist, the meaning of a normative sentence is not a matter of the proposition, if any, with which it is associated, nor of whatever relationship holds between that sentence and the proposition, if any, it happens to characterize.⁷ Even if there were a particular class of propositions up to the job of determining a distinctive, primarily practical use for sentences associated with them,⁸ that use, rather than the propositions themselves, would be explanatorily fundamental, according to CMAU, in any account of their meaning.

A.2 Dividing the Terrain

These distinctions, together with some others I'll draw presently, allow us to draw a rough taxonomy of the theoretical terrain in the debate between Expressivists and their opponents. But our interest in describing this taxonomy is not just classificatory. Drawn properly, it reveals a region of theoretical space that has gone unnoticed and undiscussed in the contemporary philosophical discussion of Expressivism. Views occupying this region of theoretical space are, I want to suggest, interesting, and not just because of the lack of attention paid to them in contemporary philosophical discussion. Such views, I will go on to argue, offer the best prospects for developing a theory of normative sentence meaning along the lines outlined in the prior section.

A.2.1 The Role of Propositions

Maybe the key variable is the role that a theory of meaning accords to the sorts of entities in terms of which meaning is *standardly* elucidated: propositions, truth-conditions, possible

6. In particular, certain Expressivists ("Quasi-Realists") are interested in affirming ACCURACY (see esp. Blackburn 1988, 1998; Gibbard 2003). This is just one theoretical purpose that regarding normative sentences as characterizing propositions would serve. For similar thoughts, see Alwood (2010).

7. Compare the standard (stronger) view in pragmatics that propositions are "essentially forceless", i.e., a proposition is always apt for being put to a variety of discursive uses, and, depending on the force with which it is paired in a specific illocutionary act, can serve the communicative purposes of speakers with a diverse array of communicative intentions (for discussion, see Charlow 2010a; Green 2000).

8. Expressivists typically endorse an account of motivation (the Humean Theory) that they understand to be incompatible with the existence of such propositions. For classic discussions, see Smith (1987, 1994).

worlds, properties, and the like. For expository ease, I will focus largely on propositions. Propositions can, as we will see, play a variety of theoretical roles in an Expressivist meta-ethic. The claim, then, is that a useful way of dividing the theoretical terrain is according to the theoretical role they assign propositions. Specifically, we will see that it useful to divide them according to the following questions. (ϕ refers to an arbitrary normative sentence.)

Q1. Does ϕ have a proposition as its *semantic value*?

Q2. Are propositions *enlisted to do logical or semantic work* (explaining inconsistency, ϕ 's semantic behavior under embedding, etc.)?

Schroeder (2009)—in remarks that are, it's fair to say, representative of the larger dialectic—proposes to distinguish Hybrid and Pure Expressivists according to their answers to Q1 and Q2. Hybrid Expressivists, in contrast to Pure Expressivists, use propositions “to underwrite [their] answer to the Frege-Geach Problem,” i.e., the problem of explaining inconsistency and embedding for a normative fragment (261). Later, he elaborates:

[P]ure expressivism promises a particular way of accounting for the semantics of moral sentences. Rather than accounting for their semantics by saying what they are about or what their truth conditions are, the expressivist program is to account for their semantics by saying what kind of thoughts they express. I think of it this way: for pure expressivists, the right kind of semantics for moral sentences doesn't assign them propositions as values. It assigns them mental states instead (264).

As these remarks make clear, Schroeder *collapses Q1 and Q2* (or, at least, does so for purposes of his discussion of Pure Expressivism). But Q1 and Q2 are, we should note, different, although related, questions. A 'yes' answer to Q1 plausibly entails the 'yes' answer to Q2 (likewise, 'no' to Q2 entails 'no' to Q1), but the status of the reverse entailment is less clear: *prima facie*, it seems possible for one's semantic theory for ϕ to make use of the properties of propositions, without having ϕ 's semantic value *be* a proposition. A Pure Expressivism which answers 'no' to Q1 and 'yes' to Q2 may be a theoretical impossibility, but we have seen no reason, as yet, for thinking this is so.

What is going on here? Schroeder is not, I think, guilty of mischaracterizing the commitments of Pure Expressivists like Blackburn and Gibbard. The most sophisticated developments of Pure Expressivism do, indeed, try to explain the semantic properties of normative sentences in terms of properties of their associated mental states (uses), rather than in terms of the properties of propositions. (For a case study, see Section A.2.2.) But why do they do that? There is, lurking in the background, a substantive commitment about linguistic methodology that philosophical discussion of Expressivism has taken for granted. Call that commitment “Meaning Reductionism.”

MEANING REDUCTIONISM

If things of kind K are fundamental⁹ in a theory of ϕ 's meaning (i.e. the

9. I will not say too much about fundamentality (although subsequent discussions will somewhat flesh out

“basic” theory of meaning is given in terms of K 's), any facts about ϕ 's meaning must be explained *by appeal to features of K 's*.

Notice that, if [Meaning Reductionism](#) and [CMAU](#) are each correct, any *semantic* facts about ϕ must be explained by appeal to properties of ϕ 's *use*. Meaning Reductionism is what makes sense of the tendency—common to both Expressivists and their critics—to collapse [Q1](#) and [Q2](#)—to assume that Expressivists cannot avail themselves of the properties of propositions (rather, may only avail themselves of the properties of uses) to explain the semantic properties of normative language. It has remained unarticulated (perhaps because the parties to the debate take it to be obvious or in need of no defense). But it is, nevertheless, a *substantive methodological assumption*. So while Schroeder may be innocent of mischaracterizing the commitments of Pure Expressivists like Blackburn and Gibbard, it may (and, indeed, I will argue that it *does*) turn out that he mischaracterizes the commitments of Pure Expressivism.

I will expand on these points in [Section A.2.4](#). For now, the role of Meaning Reductionism is offered as a neutral observation. There may, in fact, be no objection to ignoring semantic theories which make use of propositions while declining to assign propositions as semantic values. It will depend on whether the coarsened picture of theoretical space that results from conflating [Q1](#) and [Q2](#) yields a dialectically adequate division of that space—whether the only theoretically interesting division is between meta-ethical theories which answer ‘yes’ to both [Q1](#) and [Q2](#), and those which answer ‘no’ to both.

A.2.2 Hyperplan Semantics

Both Expressivists and their opponents appear to presuppose Meaning Reductionism. As an illustration, consider Gibbard's semantics for normative language. (If you are familiar with Gibbard's view, or do not require an illustration, feel free to skip to the end of [Section A.2.3](#).)

Now, at first glance, it might seem that Gibbard's theory, contra Schroeder, actually fails to assign mental states as the semantic values of normative sentences. For, on Gibbard's theory of normative content, the content of any atomic normative sentence (e.g., ‘murder is wrong’) is a *property of a cognitive state*, and is represented as a set of pairings of “Hyperplans” (roughly, fully specified contingency plans) and worlds. In the case of ‘murder is wrong’, this content is apparently represented, *not* as an attitude of disapproving of murder (or, perhaps, of disapproving of the wrong-making features of murder), instead with a GIBBARD CONTENT. Formally, the Gibbard Content of ‘murder is wrong’, for instance, is just the set of pairs $\langle \pi, w \rangle$ such that the sentence ‘according to π , murder is disallowed’—a purely descriptive sentence about the content of the rules of π —is true at w (see esp. [Gibbard 2003](#):

the general sort of notion I have in mind). I will say that fundamentality rankings are meant to be a function of *explanatory* relationships: if, in a theory T , features of K 's explain features of K^* 's (and not vice versa), then K 's are more fundamental than K^* 's in T . To say that K 's are fundamental in T (*sans phrase*) is to say that K 's are maximally fundamental in T , with respect to the fundamentality ranking for T .

Chapters 3–4). More or less equivalently, the Gibbard Content of ‘murder is wrong’ is the property a Hyperplan has when it disallows murder.¹⁰

The reason Gibbard’s theory does, in fact, count as Pure, in Schroeder’s sense, is that sets of Hyperplan-world pairs function, for Gibbard, as an *abstract mathematical representation of an individual’s practical state of mind*. In the representation of practical states of mind, Gibbardian Hyperplans play a role analogous to that played by possible worlds in the representation of belief. So, as the beliefs of someone who is uncertain about p are often represented with a set of worlds (complete and consistent assignments of truth-values to atomic sentences), some of which satisfy p , some of which satisfy $\neg p$, the practical state of mind of someone who is uncertain about whether to do some action in some contingency can be represented with a set of Hyperplans. Gibbard takes the analogy to worlds seriously, understanding Hyperplans as complete and consistent assignments of to-be-done values to actions (relativized to situations or contingencies). A Gibbardian Hyperplan π is thus stipulated to satisfy both of the following constraints.

COMPLETENESS

For any alternative α and situation S , *you do α in S* is a rule of π , or *you do not do α in S* is a rule of π .

CONSISTENCY

For any alternative α and situation S , it is not the case that both *you do α in S* is a rule of π and *you do not do α in S* is a rule of π .

Representing contents in terms of sets allows the usual Boolean treatment of the connectives: \wedge is associated with intersection, \neg with absolute complement, etc.

Gibbard, however, is explicit that this abstract Boolean representation is just that: a *representation* of a more fundamental psychological reality, chosen for (i) its formal well-behavedness and (ii) its ability to roughly approximate the semantically interesting properties of that more fundamental reality:

One way to think of fact-plan content is to mimic truth functions and quantification... These [recursive mental] operations—combining, ruling out, generalizing—mimic standard logical operations on statements: conjunction, negation, and universal generalization (Gibbard 2003: 54).

The formalism itself is, of course, intended to be explanatory, in some sense, of semantic phenomena concerning normative language. But its explanatory status is wholly *derivative*: the semantic *explananda* are ultimately explained by characteristics of *representandum* (mental states), rather than of *representans* (Gibbard Contents). Gibbard’s expectation, of course, is that the characteristics and behavior of representans adequately represent those of representandum. Indeed, the point of the formalism is to *elucidate* the recursive and

10. My discussion here draws freely from both Gibbard’s (1990) theory, which makes use of systems of norms paired with worlds, and his (2003) theory, which invokes fully specified contingency plans.

combinatory nature of psychological space, by claiming that its structure is mirrored by that of a familiar Boolean algebra.¹¹

A.2.3 The Role of Propositions, Again

Recall the distinction drawn above between views assigning a normative sentence a proposition as its *semantic value* and views on which it merely *characterizes* a proposition. Prima facie, there is nothing essentially non-Expressivist about a semantics which has normative sentences characterizing propositions. Indeed, Gibbard's semantics has them doing exactly that: relative to a Hyperplan π , a Gibbard Content trivially characterizes a possible worlds proposition (namely, the set of worlds w such that $\langle \pi, w \rangle$ is in the Gibbard Content of 'murder is wrong') (cf. Dreier 1999). In the case of *thick* normative predicates—predicates like 'brave' blending practical and representational content (in virtue of attributing a non-normative property, like exhibiting warrior-like qualities, to an action, while also expressing approval of that action, or the class of actions exhibiting warrior-like qualities)—the proposition thus characterized is typically *contingent*.¹² That is because whether an action is required at w by a Hyperplan that enjoins warrior-like qualities depends on what sorts of qualities the warriors at w actually exemplify. But, in the case of *thin* normative predicates—predicates whose content is purely practical, which place no conditions whatever on the situation in which an action occurs—the proposition thus characterized is typically *necessary*. Distinctive of Gibbard's semantics is the *explanatory inertness* of the proposition thus characterized. If Q is a thin normative predicate, $\lceil Q(\alpha) \rceil$ and $\lceil Q(\beta) \rceil$ will, on Gibbard's view, characterize precisely the same possible-worlds proposition (\top), for any rigid action-designators α and β , and with respect to any Hyperplan. Hence, any difference in their semantic characteristics cannot be explained by appeal to the possible-worlds proposition thus characterized.

Hybrid accounts, to contrast, characteristically treat normative language as *contentful along the propositional dimension*. A typical way of doing this assigns as a semantic value to a normative predicate Q a purely descriptive property (e.g., maximizing pleasure) that serves to determine *substantive, non-normative application-conditions* on Q .¹³ A sentence of

11. For similar thoughts, see Schroeder (2008c); Dreier (2006).

12. Strictly speaking, this depends on how coarsely available alternatives are individuated. If they are individuated very finely (so that *taking out the trash in S* and *taking out the trash in S'* count as distinct), worlds cease to play any independently interesting role in the semantics, and no Gibbard Contents will end up characterizing contingent propositions. I am going to assume, as is fairly standard (see, e.g., Belnap, Jr. & Perloff 1988), that available alternatives are typically individuated more coarsely than this, and that their extensions sometimes vary according to the situation/world in which they are performed.

13. This sort of view originated with Stevenson (1944); Hare (1952), and is a feature of the recent Hybrid theories of Copp (2001); Barker (2000); Boisvert (2008) (although these views differ in the mechanism by which the descriptive property is associated with the normative predicate; see Schroeder 2009 for discussion). Another sort of view, hinted at in Alwood (2010), is reminiscent of Hare's well-known treatment of imperatives (see, e.g., Hare 1967). Hare represents imperatives with formulas in which an illocutionary operator (expressing directive force; Hare dubs this the *neustic*) scopes over a content (the state of affairs enjoined, which Hare dubs the *phrastic*), but treats the semantics of the imperative as exhausted by the proposition expressed by the phrastic.

the form $\lceil Q(\alpha) \rceil$ would thus have as its semantic value the proposition that α meets the non-normative application-conditions for Q . This proposition would then be invoked to explain the semantic properties of the sentence. So far, so Realist. Where Hybrid accounts differ from Realist accounts is in attributing to normative language a conventionalized practical use, independent of (not determined by) its semantics.¹⁴ Hybrid theories thus answer ‘no’ to Q4, hence ‘no’ also to Q3, while Pure theories answer ‘yes’ to Q3, hence ‘yes’ also to Q4:

Q3. Does ϕ 's semantic value *exhaust its meaning*? Should ϕ 's meaning be *identified* with its semantic value?

Q4. Does ϕ 's meaning *supervene* on its semantic value? Is ϕ 's meaning determined by its semantic value?

Hybrid theories, unlike either Gibbardian Expressivism or Realism, represent normative sentence meaning as *two-dimensional*. The meaning of a normative sentence ϕ is both propositional and practical, and a complete theory of meaning for ϕ must build in machinery for the representation of both dimensions: the usual semantic machinery (for computing ϕ 's propositional semantic value), as well as some non-semantic machinery (whatever is requisite for computing ϕ 's practical meaning).¹⁵

A.2.4 Expressivism without a Semantics of Use

To summarize: the division of theoretical space with which we're left is this.

	Q1	Q2	Q3	Q4
PURE	N	N	Y	Y
HYBRID	Y	Y	N	N

I want to conclude this section with a bold claim (one I will spend the rest of the paper supporting): this division is inadequate and has systematically distorted the debate between Pure Expressivists and their opponents. This is precisely because this division reflects the presupposition that an Expressivist semantics must be a *semantics of use*, in which attitudes, discourse roles, and the like are ultimately doing the semantic heavy lifting. It rests on a misunderstanding that is effectively common ground between Pure Expressivists

14. Different theorists understand this use in different ways. According to Barker (2000), this use takes the form of a *conventional implicature*, to the effect that the speaker bears some non-representational attitude toward actions meeting the application-conditions for the relevant normative predicate. According to Boisvert (2008), this use is represented as the speaker's performance of a purely expressive speech act, of the sort that one performs when one utters a pejorative.


15. It is open to Hybrid approaches to endorse (a suitably qualified of) CMAU. I am wary of Hybrid approaches that claim *no distinction* in semantic type between sentences claimed to diverge in conventional use (e.g., Alwood 2010, perhaps Stevenson 1944; Hare 1952). For reasons I cannot discuss here (but see Portner 2004a for discussion), I endorse a one-one correspondence between semantic types of sentences and conventional uses: conventional use supervenes on a proper subclass of a sentence's lexical properties (its semantic properties) and each use is associated with a single sentential semantic type. Sentences, then, cannot differ in conventional use without differing in semantic type.

and their opponents: that “for pure Expressivists, the right kind of semantics for moral sentences doesn’t assign them propositions as values. It assigns them mental states instead” (Schroeder 2009: 264). It rests, in other words, on [Meaning Reductionism](#).

My favored form of Expressivism does endorse a distinctive, non-propositional semantics for normative language, but it is emphatically *not* a semantics of use. On the theory developed and defended here:

- Normative sentences characterize propositions, which are enlisted to do logical and semantic work. It is the semantic values for normative sentences that are responsible for characterizing these propositions, but the semantic value of a normative sentence isn’t itself a proposition.
- The relationship holding between a sentence’s semantic value and its meaning (use) is one of *mutual determination*, rather than *reduction*: meaning supervenes on, without reducing to, semantic value (and vice versa). Most importantly for the expressivist, an expression’s use is distinct from, although it supervenes on, its properly lexical (semantic) properties. Use is lexically encoded in the semantics: what an expression can be used to do (its “functional potential”) is a function of its semantic value; having a certain semantic value means being apt to be put to use in certain ways.
- The theory of meaning (use) is theoretically fundamental or prior: the theory of meaning *governs* the semantics. Theorizing about language of a certain type *T* begins with an account of the use of language of type *T*. Because meaning is lexically (namely, semantically) encoded, this *constrains* semantic theorizing about *T*-language: the correct use for *T*-language must be derivable from the semantics for *T*-language (perhaps in conjunction with “law-like” statements connecting semantic properties to facts about use).

A theory of this sort is able to exploit a possibility that other theories, because of their commitment to [Meaning Reductionism](#), have ignored: a theory of meaning can govern semantic theorizing, without all semantic explanations being explanations *within* the theory of meaning. The theory developed in this paper, then, will have something like the explanatory structure plotted here. (*a* → *b* represents the claim that *b* is explained by *a*; oval nodes represent object-language explananda; rectangular nodes represent theoretical accounts of these phenomena.)



On this picture, language has both *dynamic* and *static dimensions*. These dimensions are concerned, respectively, with how language is used in communication, as well as its model-theoretic properties (e.g., the logical profile of sentences of that language). Expressivism

begins from the assumption that the dynamic dimension is somehow theoretically prior to the static: a theory of *meaning* for a language is, above all, an account of its dynamic dimension; what people do with sentences is more theoretically central than, e.g., the logical relations those sentences bear to one another. So, an account of the dynamic dimension is explanatorily *fundamental*. But it is not, on this picture, explanatorily *complete*: properties of the entities of the semantic theory, rather than the entities of the basic theory of meaning (i.e., uses), explain ϕ 's semantic profile, in spite of the fact that the basic theory of meaning governs the shape and content of the semantic theory. (Note: As the dashed arrow suggests, I am neutral on the transitivity of explanation. The basic theory of meaning for ϕ may, in a derivative sense, "explain" ϕ 's semantic profile, in virtue of governing the theory in which *the explanation* for ϕ 's semantic profile is actually given. I am simply claiming that the explanation of ϕ 's semantic profile is not *part of* the basic theory of meaning for ϕ .)

So far, this is not an account; it is just a preliminary attempt to carve out a region of theoretical space in which one might be situated. Before stating a concrete theory (which we will not actually do until Section A.4.2), we should get clear about the difficulties that are usually thought to beset Expressivist theories of normative discourse. This will help us see what sorts of challenges Expressivism is up against, as well as to be precise about the ways in which our version of Expressivism strictly improves on the standard version.

A.3 Embedding Problems

This section presents two versions of the major challenge to an Expressivist theory¹⁶ of the meaning of normative discourse: the Embedding Problem. Drawing on linguistic work on clause-type and the semantics of imperatives, I'll argue that, although theories like Gibbard's *do* run into at least one version of the Embedding Problem, it is reasonably clear that there *is* a version of Expressivism which does not. This is, so far, not to say what form such a theory should actually take (although I will go on to suggest such a theory in the next section). It is a case for optimism—one whose success does not depend on the success of the theory I'll go on to propose.

As a theory of meaning, Expressivism is subject to the usual methodological and theoretical constraints on such theories. For instance:

COMPOSITIONALITY

A theory of meaning for a fragment \mathcal{L} should yield a method of computing the meaning of any well-formed expression α of \mathcal{L} , on which that meaning is a function of (i) the meanings of α 's constituents and (ii) the mode in which those meanings are combined.

LOGICALITY

16. From now on, I'll drop the "Pure" qualifier, unless it is needed.

A theory of meaning for \mathcal{L} should give rise to satisfactory notions of *logical consequence* and/or *validity*.

Compositionality, I trust, speaks for itself. Logicality, perhaps, does not. Minimally, of course, a theory of meaning for \mathcal{L} should predict, for any sentence ϕ of \mathcal{L} , that ϕ stands in the consequence relations in which we know it to stand (WEAK LOGICALITY). More strongly, this prediction must be generated *in a theoretically satisfactory way* (STRONG LOGICALITY). Expressivism's critics charge that it has difficulty respecting compositionality and logicality (in its weak and strong guises) in its account of the meaning of constructions embedding normative constructions in "unasserted" environments (under negation, in conditional antecedents, etc.). It generally goes undisputed that Expressivism is subject to such demands. But, as we'll see, what, exactly, these demands require of a theory can be controversial.

A.3.1 Frege-Geach

As initially formulated,¹⁷ the Frege-Geach Problem was just to explain why, if a sentence ϕ 's meaning was identified with its use, certain ways of embedding ϕ seemed to systematically vitiate or cancel that use. Whatever the sentence 'doing meta-ethics is wrong' is typically used by a speaker to do (display disapproval of meta-ethics, say), someone who utters either (A.1) or (A.2) is not using it to do *that*.

(A.1) Doing meta-ethics isn't wrong, not even a little bit.

(A.2) If doing meta-ethics is wrong, I don't want to be right.

The problem originates in the Expressivist's commitment to CMAU, which presupposes that meaningful sentences have conventionalized (in the sense of Sect. A.1.1) uses, and claims that such uses constitute, in some sense or other, their meaning. Sentences like (A.1) and (A.2) seem to be worrying for both of these notions. Indefeasibility is a hallmark of linguistic conventionality (and especially of linguistic meaning): 'the house is red,' as a matter of linguistic convention and meaning, expresses the proposition that the house is red, and this is not defeated by embedding it under negation. If 'doing meta-ethics is wrong,' as a matter of linguistic convention and meaning, expresses or displays disapproval of meta-ethics, we should not expect this to be cancelled by embedding the sentence under negation, as in (A.1). But it *is*.

This worry is not too troublesome. Though it's often suggested to threaten the compositionality of an Expressivist theory of meaning, that is clearly not right.¹⁸ Although compositionality does plausibly require an embedded normative sentence to contribute its conventional meaning to the computation of the meaning of the construction embedding

17. See Searle (1962); Geach (1965). For a critical discussion, see Schroeder (2008c,d).

18. Here we basically follow Schroeder (2008c,d).

it,¹⁹ it does not require that a computation of the meaning of some maximal syntactic projection (e.g., a complete sentence) “pass up” the meanings of its parts, or that those meanings “project” to the whole. It requires only that there be *some function from the meanings of those parts to the meaning of the whole*. Gibbard’s account—which has negation expressing the mental operation of *ruling an attitude out*, and material implication the mental operation of ruling out accepting the attitude expressed by its “left” argument while ruling out the attitude expressed by its “right” argument—is designed, in part, to meet exactly this constraint. It predicts, *prima facie* correctly, that a sentence like (A.1) conventionally expresses the attitude of ruling out disapproving of meta-ethics, rather than disapproval of meta-ethics.²⁰

Linguistic research, moreover, lends cautious support to the Expressivist’s claim that there is a conventionalized relationship between sentences and their uses, specifically between a sentence’s clause-type (declarative, interrogative, imperative, etc.) and its canonical discourse role (assertion, interrogation, instruction) (see esp. [Sadock & Zwicky 1985](#); [Portner 2004a](#)). So long as there is a conventionalized relationship between sentences and their uses, there will *be* (even if we cannot presently identify) some adequate explanation of how a sentence ϕ can conventionally express, for instance, an assertion, without certain constructions that embed ϕ (as, for instance, the antecedent of a conditional) expressing an assertion that ϕ as well.²¹

A.3.2 Inconsistency and Disagreement

More troubling is Logicality. Weak Logicality was long thought to be a problem for Expressivism, in light of the fact that notions like logical consequence and validity were standardly defined in terms of a relationship of truth-preservation holding between propositions, inconsistency in terms of the impossibility of simultaneous truth of two propositions. But Weak Logicality turns out to be no sort of problem for Expressivism at all: Gibbard Contents yield *formally satisfactory* notions of both inconsistency and validity. If $\llbracket \phi \rrbracket$ is the Gibbard Content for ϕ :

GIBBARD INCONSISTENCY

ϕ and ψ are Gibbard-inconsistent iff $\llbracket \phi \rrbracket \cap \llbracket \psi \rrbracket = \emptyset$.

19. [Geach \(1965\)](#) presents the requirement that an embedded occurrence mean the same thing as an unembedded occurrence as due to Logicality (rather than Compositionality): if ϕ has a different meaning when embedded in a conditional of the form $\lceil \phi \rightarrow \psi \rceil$, it’s not clear why modus ponens should be valid.

20. It is overlooked that pioneers of so-called DYNAMIC SEMANTICS developed precisely this sort of account for computing the force of Boolean compounds in the 1970s (and have since added accounts of quantification, modality, anaphora, etc., to the basic framework). On Dynamic accounts, sentences determine *update potentials*—functions from input states to updated output states—and Boolean operators are operations on such functions (\wedge , for instance, corresponds to function composition). Some classics in this tradition are [Kamp \(1981\)](#); [Heim \(1982\)](#); [Groenendijk & Stokhof \(1991\)](#); [Veltman \(1996\)](#). An interesting piece of trivia: the anti-Expressivist arguments we consider here can all be adapted into arguments against the Dynamic program. I take that as an embarrassment for those arguments, rather than Dynamic Semantics.

21. [Alwood \(2010\)](#) makes the same point. For an account of the clause-type/use relationship, see my “Imperatives and Necessitation” (Chapter III).

GIBBARD VALIDITY

ψ is a Gibbard-valid inference from ϕ_1, \dots, ϕ_n (notation: $\phi_1, \dots, \phi_n \models \psi$) iff $[[\phi_1]] \cap \dots \cap [[\phi_n]] \subseteq [[\psi]]$.

The reader familiar with Possible Worlds semantics will recognize that there is nothing special about these definitions at all: Gibbard-inconsistency is defined *exactly* as the property of inconsistency is standardly defined in Possible Worlds semantics.

Strong Logicality is the snag: a by-now-familiar point from the meta-ethical discussion is that the account of inconsistency and validity thus generated, although formally adequate, seems doomed to be theoretically or explanatorily deficient (Dreier 2006, 2009; Schroeder 2008a,c,d; Unwin 1999, 2001). Schroeder's presentation of this problem is especially cutting. He invites us to consider the pairs in (A.3), (A.4).

- (A.3) a. One should commit murder $\approx O\phi$
b. It's not the case that one should commit murder $\approx \neg O\phi$
- (A.4) a. One shouldn't commit murder $\approx O\neg\phi$
b. It's not the case that one shouldn't commit murder $\approx \neg O\neg\phi$

What attitudes do sentences (A.3a) and (A.3b) express? Schroeder suggests two possible Expressivist responses:²²

SAME ATTITUDE ANALYSIS (SAA)

A normative sentence and its negation express the *same propositional attitude* (perhaps toward inconsistent contents).

DIFFERENT ATTITUDE ANALYSIS (DAA)

A normative sentence and its negation *distinct propositional attitudes* (perhaps toward the same content).

On the SAA, it's natural to say that (A.3b) and (A.3a) express the same attitude, toward inconsistent contents; likewise for (A.4a) and (A.4b). This is turned into an account of inconsistency by supposing the attitude expressed is *inconsistency-transmitting*.

INCONSISTENCY-TRANSMITTING ATTITUDES

A propositional attitude *A* is inconsistency-transmitting iff bearing *A* toward inconsistent propositions is itself inconsistent.

Attitudes like belief and intention are paradigmatically inconsistency-transmitting (whereas attitudes like desire are paradigmatically not). On the SAA, the inconsistency of (A.3a) and (A.3b) is thus explained by the fact that (A.3a) and (A.3b) express some inconsistency-transmitting attitude *A* toward inconsistent contents.

22. Schroeder presents this argument in many venues. The presentation here follows closely his (2008c).

But there is a problem. Supposing, as seems natural, that the propositional objects of the attitudes expressed by (A.3a) and (A.4a) are jointly exhaustive (the former, perhaps, expresses an intention to commit murder, the latter, perhaps, an intention *not* to commit murder), it follows that (A.3b) and (A.4b) inconsistent.²³ That, of course, is the wrong prediction: someone indifferent about murder—hence, who accepts (A.3b) and (A.4b)—is morally repellent, but not thereby inconsistent.

DAA fares no better. According to DAA, (A.3b) and (A.3a) express *distinct attitudes*—while (A.3a) expresses, perhaps, disapproval of failure to murder, (A.3b) expresses the attitude of tolerating failure to murder. But what, exactly, is inconsistent about both disapproving and being tolerant of someone’s failure to murder? The account must effectively stipulate rational norms on which it is inconsistent to both disapprove and tolerate ϕ . But, says Schroeder, the Expressivist is not entitled to assume the existence of such norms. There are “few good examples” of rational norms requiring such relations between “logically unrelated” attitudes (2008c: 581).

Gibbard’s account is a version of the DAA (and, as such, encounters difficulties with Strong Logicality). On Gibbard’s account, (A.3a) expresses, roughly, the state of mind of planning not to murder (for any contingency whatever). This state of mind is represented with the set of Hyperplan-world pairs $\langle \pi, w \rangle$ such that ‘according to π , murder is disallowed’ is true at w . This set’s complement represents the state of mind of *disagreeing with planning not to murder*. But, unless this state of mind is just the state of *planning to murder* (and, evidently, it is not²⁴), we cannot explain the inconsistency of (A.3a) and (A.3b) in terms of the inconsistency-transmitting-ness of planning. So, although the state of mind [planning not to murder while disagreeing with planning not to murder] cannot be represented with a Gibbard Content (since their associated Gibbard Contents are, by stipulation, disjoint), Gibbard’s account leaves unexplained why disjoint Gibbard Contents are appropriate ways of representing their representanda. Saying they *disagree* is question-begging. We want an account of *why* the states of mind expressed by (A.3a) and (A.3b) are in disagreement with one another; Gibbard’s account seems simply to stipulate this (cf. Dreier 2006, 2009; Schroeder 2008c).

23. Proof (cf. Schroeder 2008c: 579): By SAA, (A.3a), (A.3b) must express A toward inconsistent contents (respectively: ϕ and some ψ s.t. $\phi, \psi \models \perp$), and (A.4a), (A.4b) must too (respectively: $\neg\phi$ and some χ s.t. $\neg\phi, \chi \models \perp$). Since $\phi, \psi \models \perp$ and $\neg\phi, \chi \models \perp$, it follows that $\psi, \chi \models \perp$. So, since A is inconsistency-transmitting, bearing A toward ψ and ϕ is inconsistent. So, then, are (A.3b) and (A.4b). The way to block this argument is to suppose that the contents of the attitudes expressed by (A.3a) and (A.4a) are not jointly exhaustive (indeed, Schroeder 2008a suggests an Expressivist semantics that does exactly that). I will not be interested in this response here, as I reject the pressures that would lead an Expressivist to appeal to it.

24. A major problem for Gibbard’s account, stressed by Dreier (2006, 2009) and Schroeder (2008c), is that, it actually conflates these states of mind. By the Completeness constraint on Hyperplans, the complement of $\{\langle \pi, w \rangle : \text{‘according to } \pi, \text{ murder is disallowed’ is true at } w\}$ is $\{\langle \pi', w \rangle : \text{‘according to } \pi', \text{ murder is required’ is true at } w\}$. We discuss this further in Section A.5.2.

A.4 A Sketch of a Theory

In this section, I draw on (and expand upon) the sorts of considerations developed in Section A.2 to argue that, while meeting the Strong Logicality constraint is problematic for Gibbard's Expressivism, doing so is by no means problematic for Expressivism *per se* (in particular, an Expressivism that jettisons [Meaning Reductionism](#)). There are two parts to my argument here. In the first, I show how [Meaning Reductionism](#), together with an independently plausible set of claims about the meaning and semantics of directive language (e.g., imperatives), generates a specious Embedding Problem for directives. In the second, I show how to build theories of directive semantics and directive meaning that jettison [Meaning Reductionism](#), while retaining the set of independently plausible claims with which we began.

A.4.1 Part I: The Negation Problem as Pseudo-Problem

In this section, I show that, if one understands Strong Logicality as leading to a problem for Expressivism about normative language, one is committed to understanding Strong Logicality as a problem for Expressivism about imperative (and, more broadly, directive) language. But there are good reasons—*independent, linguistic reasons*, having nothing to do with standard motivations for Expressivism in meta-ethics—for being (something that is reasonably referred to as) Expressivist about directive language. This presents the opponent of Expressivism with a dilemma: give up Expressivism for directive language, or admit there is something defective about an argument against Expressivism built on Strong Logicality.

A.4.1.1 Expressivism for Directive Language.

Building on Lewis' (1979a) classic account of direction, Paul Portner (2004a; 2007) has constructed an influential account of imperative meaning around the following three claims:

CONVENTIONALIZATION OF USE

Sentences typically have a conventionalized use (discourse role, illocutionary force, etc.), and they have that use in virtue of their *clause-type* (declarative, interrogative, or imperative).

IMPERATIVES ARE USED NON-REPRESENTATIONALLY

The conventional use of imperatives should be represented, at the level of discourse, as a tendency to *introduce obligations* on an addressee, via addition to a contextual parameter that determines, in part, what that addressee ought to do (her "To-Do List").

IMPERATIVE MEANING AS USE

A theory of imperative meaning is a theory of their conventional use. An account of an imperative's conventional use explains "everything that needs to be explained about its meaning" (2007: 366).

Portner develops this account in opposition to Modal Accounts of Imperative Meaning. A Modal Account of the meaning of an imperative $!\phi$ (read: *see to it that ϕ*) is defined as any account which assigns it the same semantic content as a modalized obligation-statement (something like: *you must see to it that ϕ*).²⁵ More generally, Portner's account stands in opposition to any propositional account of imperative meaning: any account which proceeds by assigning a proposition as the content of the imperative.²⁶ All such accounts face a major difficulty: the problem of explaining why, if the meaning of an imperative is represented with a proposition, the conventional use of any imperative is *performative* (obligation-creating), rather than *representational* (obligation-describing). Why, if $!\phi$ means that you must see to it that ϕ , can $!\phi$ not be used to *assert that* you must see to it that ϕ ?²⁷

The going theory of imperative meaning—Portner's—is, in short, an Expressivist theory. Portner endorses a CMAU claim for imperatives, claims the conventional use of imperative use is principally performative (rather than representational), and rejects Propositionalism for imperative sentences.

A.4.1.2 The Negation "Problem" for Directive Language.

By loading *all* of an imperative's conventional meaning into its performative use, Portner would seem to face a challenge similar to that faced by the Expressivist, namely, meeting Strong Logicity. Indeed, there's a particularly vivid parallel with the negation problem for Expressivism: an imperative $!\phi$ (e.g., *do not jaywalk*) is obviously *inconsistent* with a contrary grant of permission $!\neg\phi$ (e.g., *you may jaywalk*). It seems clear that these sentences are inconsistent; that anyone who issues them (e.g., a judge or police officer) can rightly be charged with inconsistency; and that anyone who accepts them as binding (by which I do *not* mean that they come to *believe that* they are binding, rather, that they try to adjust their desires and plans accordingly) can rightly be charged with inconsistency.

Assuming Meaning Reductionism, then, the onus would seem to be on the imperative Expressivist to explain the inconsistency of $!\phi$ and $!\neg\phi$ in terms of properties of their associated uses (that is to say, their associated discourse roles, or speech acts). As before, there are two possibilities: $!\phi$ and $!\neg\phi$ express a single inconsistency-transmitting speech act (toward inconsistent contents), or they express distinct, logically unrelated speech acts (commanding

25. For versions of the modal view, see Aloni (2007); Åqvist (1964); Han (1998); Schwager (2006).

26. There are many accounts that fit this bill. For instance, some accounts analyze imperatives with explicit performatives ($!\phi \approx I \text{ command you to see to it that } \phi$) and assign the latter satisfaction conditions (Lewis 1970). Others analyze them in terms of future-tense indicatives ($!\phi \approx \text{you'll do } x$) (Geach 1958). And others still claim that the meaning of $!\phi$ is exhausted by its fulfillment-conditions (i.e., the satisfaction-conditions for ϕ) (Jørgensen 1937-8; Hare 1952, 1967; Bennett 1970).

27. Certain maneuvers are, of course, available. None, so far as I know, is satisfactory. See my 2010a.

and permitting, perhaps). Each way, the imperative Expressivist seems to face precisely the problem faced by the Expressivist about normative discourse.

There is a rather important dialectical difference. For independent reasons, Expressivism has a strong claim to being the correct view about the meaning of imperatives. Most of us are loath to think imperatives have propositions as their contents, or that the meaning of an imperative can be faithfully represented with a proposition (with good reason, if conventional use supervenes on semantic value; cf. footnote 15). Moreover, the conventional use of an imperative is clearly performative, hence non-representational. Finally, given the centrality of the characteristic use and force of imperatives in our understanding of their meaning, there is a *prima facie* case for thinking a theory of their meaning should be an account of their characteristic use and force—that the purpose of such a theory is to illuminate the distinctive role of directive language in discourse and communication.

These sorts of intuitions are, I think, widely shared, and together they define an Expressivist view of imperative meaning. Still, I don't want to rest too much weight on them. While it would be surprising if, for instance, we were required to state a [Propositionalist](#) account of imperative meaning, expectations are, perhaps, not the most reliable guide to philosophical and linguistic truth. So, in the next section, I'll offer a vindication of these expectations, by stating an Expressivist theory of imperative meaning that (i) violates [Meaning Reductionism](#), (ii) meets the requirement of Strong Logicality. The theory is easily adapted to an Expressivist theory of normative meaning.

A.4.2 Part II: Expressivism for Directive and Normative Language

The project of stating an account of the meaning of directive and normative language is an enormous one, and I can only begin to gesture at it here. We'll make things easier by restricting our attention to:

- An imperative fragment consisting of a propositional language \mathcal{L}_P and formulas of the form $\ulcorner !\phi \urcorner$ and $\ulcorner \text{i}\phi \urcorner$ (for any $\phi \in \mathcal{L}_P$).
- A normative fragment \mathcal{L}_N consisting of a propositional language \mathcal{L}_P and formulas of the form $\ulcorner O\phi \urcorner$ (for any $\phi \in \mathcal{L}_N$)
- A modal fragment \mathcal{L}_M consisting of a propositional language \mathcal{L}_P and formulas of the form $\ulcorner \Box\phi \urcorner$ (for any $\phi \in \mathcal{L}_M$)

Formulas of the form $\ulcorner !\phi \urcorner$ and $\ulcorner \text{i}\phi \urcorner$ ('directives') are regimented representations of natural language imperatives and permissives; formulas of the form $\ulcorner O\phi \urcorner$ of *normative claims* like 'one ought to give to charity'; and formulas of the form $\ulcorner \Box\phi \urcorner$ of *non-normative descriptions* of what is required by an independently fixed set of considerations (claims like 'according to the rules, ϕ is required', which in no way commit a speaker to endorsement of those rules). Natural language is not nearly so well behaved as this, of course, but the strategy I outline for these fragments should generalize without much difficulty.

A.4.2.1 Propositionalism for Directives

Propositional accounts of the meaning of directives meet the Strong Logicality requirement. Indeed, they meet it easily: on one common account—the Modal Account (introduced in Section A.4.1.1)—an account of the inconsistency of $!\phi$ and $;\neg\phi$ seems immediate. The Modal Account analyzes $!$ and $;$ in terms of the alethic modality \Box , so that they express requirement and permission, respectively, with respect to an independently fixed body of rules (the To-Do List T). (As before, $\llbracket \cdot \rrbracket$ is an interpretation function mapping an expression to its semantic content.)

THE MODAL ACCOUNT OF DIRECTIVES

$$\llbracket !\phi \rrbracket := \llbracket \Box\phi \rrbracket$$

$$\llbracket ;\phi \rrbracket := \llbracket \neg\Box\neg\phi \rrbracket$$

Inconsistency is immediate: since $\Box\phi \wedge \neg\Box\neg\phi$ is a classical contradiction, $\llbracket !\phi \rrbracket \cap \llbracket ;\neg\phi \rrbracket = \emptyset$. Put differently: suppose $!\phi$ and $;\neg\phi$ were consistent—that $\llbracket !\phi \rrbracket \cap \llbracket ;\neg\phi \rrbracket \neq \emptyset$. Then, on the **Modal Account**, it would follow that $\llbracket \Box\phi \rrbracket \cap \llbracket \neg\Box\neg\phi \rrbracket \neq \emptyset$, thus that $\llbracket \Box\phi \wedge \neg\Box\neg\phi \rrbracket \neq \emptyset$. And that is a *contradiction*! Note that this result is *not* a matter of stipulation (unlike the similar result generated by Gibbard’s account; cf. Section A.2.2). The explanation of inconsistency follows from a substantive claim about the proper semantics for $!$ and $;$. The adequacy of this explanation depends primarily on the empirical adequacy of that semantics. There’s no question of the **Modal Account** yielding a satisfactory account of inconsistency, supposing it’s empirically adequate. (I’ll elaborate on this in Section A.5.1.)

All that’s needed to complete the **Modal Account** is an interpretation of \Box . (Although we’ll reject the **Modal Account**, it will, we’ll see, be useful to have such an interpretation around.) We start by defining a modal accessibility relation on worlds.

THE ACCESSIBILITY RELATION R_T

$$wR_T v \text{ iff } \forall p \in T(w) : v \in p$$

A world v is accessible from world w iff every proposition the To-Do List at w requires allows v . Think of the set of worlds accessible from w as the Sphere of Permissibility at w —the set of worlds that can be permissibly actualized, by T ’s lights, at w (cf. Lewis 1979a). The modality \Box is treated as a universal quantifier over that sphere: $\Box\phi$ is true at w just in case ϕ holds throughout the Sphere of Permissibility at w .²⁸

SEMANTICS FOR \Box

$$\llbracket \Box\phi \rrbracket^w = \text{true} \text{ iff } \forall v : wR_T v \Rightarrow \llbracket \phi \rrbracket^v = \text{true}$$

To sum up: the **Modal Account** secures Strong Logicality, but at the price of an account of the performative dimension of directive meaning. On the **Modal Account**, the content

28. For a more sophisticated treatment of the semantics for \Box , see Kratzer (1981).

of a directive is identified with a representational content, to the effect that certain sorts of requirements issue (or fail to issue) *from* a given To-Do List. To contrast, on Portner’s Expressivist account, the meaning of a directive is understood in terms of its distinctively performative use—in terms of the sort of operation it *performs on* the To-Do List. The Expressivist account illuminates the characteristic use of directive language, but seemingly at the price of a satisfactory account of inconsistency. The two perspectives seem irreconcilable: a theory of imperatives can give a satisfactory account of their semantic properties, or it can give a satisfactory account of their use. But it does not seem that it can do both.

A.4.2.2 Non-Propositionalism for Directives

This picture, I think, rests on a misunderstanding. There is nothing proprietary about the [Modal Account](#)’s explanation of inconsistency. The Expressivist about directives would, of course, face difficulties if she tried to adopt that explanation wholesale. But, so long as she designs her account of directive content in the right way, I’ll show that she can exploit its central insight.

Where the [Modal Account](#) errs is in banishing action-guiding, or practical, content from its depiction of directive content. Directives, unlike descriptive sentences, obviously have such content: they tell people *how to plan*, rather than *what to believe*. It is natural (although not obligatory) to think that this difference (described, so far, as a difference in how descriptive and directive sentences are typically *used*) demands a difference in the semantic analysis; the fact that certain types of sentences are apt for being used in certain ways is something that should be semantically encoded. One natural way of cashing that out: the semantic content of a purely descriptive sentence—one apt for being used to express a belief—is a *proposition*, while that of a directive is... something else.

How can use be semantically encoded? The semantic value of a sentence can be treated as characterizing a *property of a cognitive state*: it partitions cognitive states into those which meet some condition specified by the sentence, and those which fail to meet that same condition.²⁹ A proposition p , for instance, characterizes the property of accepting or believing that p . More precisely, suppose for concreteness that an agent’s doxastic state is represented as a set of possible worlds (the agent’s doxastic alternatives), and that a proposition p is just a set of possible worlds. Then p will partition the space of possible doxastic states into:

- i. The set of doxastic states that accept p : $\{D : D \subseteq p\}$
- ii. The set of doxastic states that do not accept p : $\{D : D \not\subseteq p\}$

29. There are many important precedents for the account I state here. The idea that a semantic value should be thought of as determining a conventional use is a central claim of [Portner \(2004a\)](#). The idea to extract cognitive properties from semantic values comes from [Yalcin \(2007, forthcoming\)](#). Finally, the idea that such properties can be thought of as proposing *constraints* on cognitive states is from [Swanson \(2006\)](#).

An utterance whose semantic value is a proposition p conventionally constitutes a proposal that one's interlocutors should come to accept p . Accepting an utterance whose semantic value is a proposition p involves adjusting one's doxastic state so that it comes to accept p . Rejecting such an utterance involves remaining in a state in which one's doxastic state does not accept p . If, per [Stalnaker \(1978\)](#), "The purpose of expressing propositions is to [distinguish among alternative possible ways that things may be]," accepting an utterance whose semantic value is a proposition involves assenting to that proposition's representation of the way that things are.

The general idea is that the envisioned, semantically encoded partitions also encode conventional uses, in a straightforward way:

THE SUPERVENIENCE OF USE ON SEMANTIC VALUE

An utterance by S of a sentence whose semantic value characterizes a cognitive property F conventionally constitutes a proposal that the cognitive state(s) of S 's addressee(s) come to satisfy F .

On this sort of account, Propositionalism about directives errs by assigning them the *wrong sorts* of semantic values—semantic values, namely, that are incompatible with the way that directives are conventionally used in conversation. For directives have practical, rather than representational, uses: they tell us how to plan and act, rather than what to believe. The semantic value of a directive, then, should not characterize a property of a doxastic state. Instead, it should characterize a property of a *planning state*. It should, in other words, partition the space of possible *To-Do Lists*, rather than the space of possible doxastic states.

A.4.2.3 The Ladom Account of Directives

How should it partition them? A natural idea is that the semantic value of a command directive $!\phi$ partitions the space of possible To-Do Lists into:

- i. Those that, in any relevant situation, require ϕ
- ii. Those that, for some relevant situation, don't require ϕ

A permission directive $;\phi$, on the other hand, partitions the space of To-Do Lists into:

- i. Those that, in any relevant situation, permit ϕ
- ii. Those that, for some relevant situation, don't permit ϕ

A plausible constraint of adequacy on a semantics for directives is that the semantic value it assigns to a directive should determine such a partition. The following semantics for command and permission directives—the *Ladom Account*—is a natural way of doing just that.³⁰

30. Notice that a genuine account for conditional directives—directives applying to a restricted type of

$$\begin{aligned} \llbracket !\phi \rrbracket &:= \{T : \text{when evaluated relative to } T, \forall w : \llbracket \Box\phi \rrbracket^w = \text{true}\} \\ \llbracket ;\phi \rrbracket &:= \{T : \text{when evaluated relative to } T, \forall w : \llbracket \neg\Box\neg\phi \rrbracket^w = \text{true}\} \end{aligned}$$

The desired partitions can, in fact, be read directly off the proposed semantic values: $\llbracket !\phi \rrbracket$ is the set of To-Do Lists that, in any relevant situation, require ϕ ; its complement is the set of To-Do Lists that, for some relevant situation, do not.³¹ Likewise, $\llbracket ;\phi \rrbracket$ is the set of To-Do Lists that, in any relevant situation, permit ϕ , while its complement is the set of To-Do Lists that, for some relevant situation, do not.

The **Ladom Account** is *designed to capture the practical content of a directive*. And it does. An utterance whose semantic value is a set of To-Do Lists, all of which require ϕ in w , conventionally constitutes a proposal that the plans of one's addressee should require ϕ in w . Accepting such a proposal involves adjusting one's plans accordingly. Rejecting such an utterance means declining to adjust one's plans accordingly. Similar conventions are in place for permission directives. (For further detail, see [Charlow 2010a](#) and my "A Dynamic Theory of Imperative Meaning, Chapter IV.)

But it accomplishes this while managing to hold onto the **Modal Account's** strategy for securing Strong Logicality. Suppose, for reductio, that $!\phi$ and $;\neg\phi$ were consistent—that $\llbracket !\phi \rrbracket \cap \llbracket ;\neg\phi \rrbracket \neq \emptyset$. On the **Ladom Account**, it would then follow that there is some T such that, when evaluated with respect to T , $\llbracket \Box\phi \rrbracket^w = \text{true}$ and $\llbracket \neg\Box\neg\phi \rrbracket^w = \text{true}$ (for any w). Supposing there is some such T and at least one world, we arrive at exactly the same contradiction as before: $\llbracket \Box\phi \wedge \neg\Box\neg\phi \rrbracket \neq \emptyset$.

There is, of course, a lot to fill in about this account. How does it manage to be just as explanatory as the **Modal Account**, while nevertheless retaining an Expressivist core? In Section A.5, I'll argue that the key to this question lies in our answers to Q1–Q4 and the rejection of **Meaning Reductionism**. Before doing that, I want to gesture at how this sort of strategy might be extended to properly normative language, by stating a Ladom Account for the normative fragment \mathcal{L}_N .

A.4.2.4 Normative Language

The operator O is intended to represent normative verbal auxiliaries like *should*, *must*, and *ought* (although there are, of course, important differences between these auxiliaries, which we'll ignore here). Modal languages (and their familiar semantics), it's well-known, are useful tools for understanding the semantics of such auxiliaries (cf. [van Fraassen 1972](#); [Lewis 1974](#); [Kratzer 1981](#); [Jackson 1985](#)). From here it might seem a short step to

situation, rather than every situation *sans phrase*—is not too far off. For discussion of conditional directives in this vein, see [Mastop \(2005\)](#); [Charlow \(2010a\)](#).

31. Suppose, for instance, that $T \in \llbracket !\phi \rrbracket$. According to the **Ladom Account**, this holds iff, when evaluated with respect to T , for any world w , $\Box\phi$ is true at w . By the semantics for \Box , this holds iff $\forall v : wR_T v \Rightarrow \llbracket \phi \rrbracket^v$. Iff, that is to say, ϕ holds throughout the Sphere of Permissibility T specifies for w . So $T \in \llbracket !\phi \rrbracket$ iff T requires ϕ , in any relevant situation.

analyzing these auxiliaries *as* modal operators, so that *O* just expresses the modal concept of requirement with respect to an independently fixed body of rules (perhaps a body of rules with a special authority or status, depending on the sort of normativity involved).

THE MODAL ACCOUNT OF *O*³²

$$\llbracket O\phi \rrbracket := \llbracket \Box\phi \rrbracket$$

The **Modal Account**'s account of inconsistency and validity is, of course, immediate. As, however, we saw from our look at directive language, whether a Modal Account of *O* is suitable might well depend on whether:

- i. The property of a cognitive state characterized by a sentence of the form *Oφ* is indeed the property of accepting the proposition that $\Box\phi$.
- ii. An utterance whose semantic value is $\llbracket O\phi \rrbracket$ is conventionally a proposal that one's interlocutors should come to accept some proposition as true.

These, questions, of course, raise substantive philosophical and empirical issues, about which Expressivists and their opponents disagree. But, significantly, it's not the Expressivist's take on *these* issues to which advocates of the various versions of the Embedding Problem are objecting. Like the arguments against Expressivism about directive language, the arguments from Embedding against Expressivism about normative language trade on alleged methodological and theoretical inadequacies of a *theory of linguistic meaning* which tries to cash out Expressivist views about (i) and (ii), rather than Expressivist views about (i) and (ii) *per se*. Supposing Expressivists are right about (i) and (ii), and that normative content is practical, rather than representational, there would be reason to trade in the **Modal Account** for some version of the Ladom Account. (I elaborate on this in Section A.5.1.)

Just to give a flavor how such an account might go, someone who thought the use of normative language has a directive flavor (Stevenson 1944; Hare 1952) might wish to view obligation-sentences as expressing something like *Gibbard Contents*: sets of Hyperplans (with Hyperplans understood as a special kind of To-Do List—those satisfying Completeness and Consistency). Letting Π be the space of Hyperplans:

THE LADOM ACCOUNT OF *O*

$$\llbracket O\phi \rrbracket := \{ \pi \in \Pi : \text{when evaluated relative to } \pi, \forall w : \llbracket \Box\phi \rrbracket^w = \text{true} \}$$

The **Ladom Account of *O*** allows the Expressivist to co-opt the **Modal Account**'s explanation of inconsistency. Suppose that *Oφ* and $\neg O\phi$ were consistent—that $\llbracket O\phi \rrbracket \cap \llbracket \neg O\phi \rrbracket \neq \emptyset$. On

32. The Modal Account is much too simple, for, perhaps, a non-obvious reason. *Oφ* expresses *unconditional* or *categorical obligation*, i.e., obligation holding in every relevant contingency. On the semantics stated here, *Oφ* expresses the weaker claim that ϕ is required in the contingency of evaluation. (Similar remarks hold for the Modal Account of Directives.) In the interests of keeping things simple, I'll decline to fix this. The Ladom Accounts do not have this defect.

the [Ladom Account of \$O\$](#) , it would then follow that there is some π such that, when evaluated with respect to π , $\llbracket \Box \phi \rrbracket^w = \text{true}$ and $\llbracket \neg \Box \phi \rrbracket^w = \text{true}$ (for any w). The contradiction is immediate.³³

A.5 Situating This Theory

In this final section, I want to mount a preliminary defense of the semantic adequacy of Ladom Accounts of directive and normative language. I'll first argue that the [Ladom Account of directives](#) yields an adequate account of the semantics of ! and ;—at least as satisfactory, in any case, as the [Modal Account of directives](#). If the [Ladom Account of directives](#) is theoretically adequate, it will be theoretically adequate for the normative case too. I then do a comparison of the [Ladom Account of \$O\$](#) with Gibbard's theory. Gibbard's theory can, in fact, be assimilated to the Ladom Account, provided that he jettisons his commitment to a semantics of use (hence also his commitment to [Meaning Reductionism](#)). I close with some remarks on why this theory really does constitute a genuinely Expressivist theory.

A.5.1 On Semantic Adequacy

We've been happy to take for granted that the [Modal Account of directives](#) is theoretically adequate. But here is a case for thinking it isn't. The case turns out to be weak. But seeing how it goes wrong will help us articulate more plausible conditions of theoretical adequacy on an account of meaning, which we will then use to assess the adequacy of the [Ladom Account of directives](#).

Hare (1967) notes that embedding an imperative under negation is either ungrammatical (A.5a), or expresses a "a meta-linguistic statement; it reports the second-order fact that somebody has" failed to issue the relevant command (A.5b).³⁴

- (A.5) a. *Not: do your homework
 b. You aren't to do your homework

This turns out to be interesting, for several reasons. Most interesting for us is that this shows that the directive operators we use to represent directive constructions in English are not syntactically inter-definable: ; must be introduced as a syntactic primitive, rather than an abbreviated way of writing $\neg ! \neg$. If that is right, it seems the [Modal Account of directives](#) may have a problem. For the semantic values of ! ϕ and ; $\neg\phi$, the [Modal Account](#) assigns

33. There are, of course, important differences in the ways directive and normative language are respectively used. One difference is that normative language is not generally used to introduce obligations; utterances of normative sentences do not generally display *performative* force. It is not hard for the expressivist to accommodate this, but I will not attempt it here.

34. Whether languages like Dutch or Korean embed imperatives under negation is neither here nor there for us, as we are interested in the semantics of a fragment of English.

inconsistent modal propositions ($\llbracket \Box \phi \rrbracket$ and $\llbracket \neg \Box \neg \neg \phi \rrbracket$). But, as we saw with Gibbard's use of disjoint Hyperplans to model inconsistency, such inconsistency-via-stipulation is easily won. The **Modal Account** offers no motivation for the stipulation—it simply *assumes* the appropriateness of assigning j a semantic value identical to that of $\neg \Box \neg$.

This argument demands too much of semantic theory. We start theorizing equipped with an understanding of the meaning of the object language. The job of semantic theorizing is to provide a formal presentation of that understanding, not to *argue in favor of it* or somehow try to *ground* it (whatever that would entail). In the most straightforward cases of semantic theorizing, the theoretical procedure is something like this (cf. [Larson & Segal 1995](#): Chapter 1).

STAGE I: OBSERVATION

Identify and describe object-language phenomenon P .

STAGE II: METALINGUISTIC ASCENT

State a canonical representation R of P in a theoretical metalanguage.

STAGE III: THEORY-INTERNAL REASONING

Identify R 's consequences in the context of a larger linguistic theory.

STAGE IV: METALINGUISTIC DESCENT

Associate these consequences with predictions about object-language phenomena, and evaluate these predictions for correctness.

Although claims made at the Observation stage are, of course, subject to scrutiny, such scrutiny is best viewed as prior or external to semantic theorizing. The semantic theorist, qua semantic theorist, is concerned with meta-linguistic representation of object-language phenomena, rather than a vindication, defense, or account of judgments about these phenomena.³⁵

In the case of clear-cut object-language inconsistency, job one for the theorist is to state a canonical representation of that inconsistency in her favored metalanguage. Whether that representation is theoretically adequate depends on whether:

METHODOLOGICAL ADEQUACY

It meets independent desiderata for canonical representations of inconsistency in a theoretical metalanguage.

EMPIRICAL ADEQUACY

It generates correct predictions about other object-language phenomena, in conjunction with one's other theoretical commitments.

35. Of course, this is not to rule out that in "the process of constructing a rigorous and explicit theory, we must be prepared for elements in the pretheoretical domain to be reanalyzed and redescribed in various ways" ([Larson & Segal 1995](#): 8).

Regarding methodological adequacy, the central desideratum for canonical representations of inconsistency—perhaps there are others—is that the metalinguistic representation *itself display some sort of contradiction or inconsistency*:

METALINGUISTIC INCONSISTENCY

A canonical representation R of an object-language inconsistency is methodologically adequate only if R itself displays inconsistency.³⁶

The Modal Account, of course, is engineered to satisfy this desideratum. The theoretical adequacy the Modal Account thus turns on its empirical adequacy. This is what I had in mind when I remarked (Section A.4.2.1) that there was no question of the Modal Account yielding a satisfactory account of inconsistency, supposing it was empirically adequate.

This is how a proponent of the Modal Account should reply to the objection that her theory is merely stipulative. The theoretical adequacy of the Modal Account is an empirical, rather than methodological, issue. But note that if the Modal Account meets methodological muster on these grounds, the [Ladom Account of directives](#), which generates the very same metalinguistic inconsistency as the Modal Account, *must as well*. The dispute between these accounts thus reduces to the question of whether each account generates reasonable predictions about other object-language phenomena, in conjunction with our other theoretical commitments. The [Modal Account of directives](#), I've argued, does not. The [Ladom Account of directives](#), engineered to capture the distinctively practical content of directives, does.

What goes for directives goes equally for normative language. Concerns of methodological adequacy do not decide between the [Modal](#) and [Ladom](#) Accounts of O : these accounts share the property responsible for the [Modal Account of directives](#) meeting methodological muster, namely, the fact that their proposed canonical representations of the object-language inconsistency of $O\phi$ and $\neg O\phi$ display an inconsistency (indeed, the very same inconsistency). The dispute between these accounts thus reduces to the question of whether each account generates reasonable predictions about other object-language phenomena, in conjunction with our other theoretical commitments. To be clear, I *haven't* argued that the [Ladom Account of \$O\$](#) does better on this score than its competitor. That will depend on the truth of substantive philosophical and empirical questions about the role of normative language in communication and thought. Expressivists defend a characteristic set of answers to these questions. Contra those who have used Expressivists' difficulties with embedding to express worries about whether Expressivist theories are even semantically tractable, it is according to these answers that an Expressivist account of the meaning of normative sentences should be judged.

36. I remain neutral on (i) what, exactly, displaying inconsistency requires, (ii) whether the [Same Attitude Analysis's](#) proposed representation (Section A.3.2) would display inconsistency in the requisite sense. More important for my purposes is that [Different Attitude Accounts](#) clearly do not do so.

A.5.2 Comparison with Gibbard

There's a sense in which Gibbard's theory of normative content (see Section A.2.2) is a special case of the Ladom Account. The Gibbard Content plausibly expressed by $O\phi$ is the set of Hyperplan-world pairs $\langle \pi, w \rangle$ such that according to π , ϕ is unconditionally (i.e., in every relevant contingency) required at w . The world parameter does no work in this formulation; ignoring it, we're left with the set of Hyperplans π such that according to π , ϕ is unconditionally required. Assuming a modal treatment of 'according to π , ϕ is unconditionally required' (so that it is true iff, when evaluated with respect to π , $\forall w : \llbracket \Box\phi \rrbracket^w = \text{true}$), then, the resulting Gibbard Content is *identical* to the semantic value furnished by the [Ladom Account of O](#) (Section A.4.2.4).

Is there any difference between the accounts? Yes, a foundational one. Recall from Section A.2.2 that Gibbard Contents have wholly derivative explanatory status in Gibbard's theory: their theoretical function is to represent practical states of mind (a consequence of Gibbard's implicit commitment to [Meaning Reductionism](#)). The semantic facts are ultimately explained by characteristics of the *representanda*, rather than of Gibbard Contents. To be sure, Gibbard's account, understood as a Ladom Account, is able to generate the same contradiction as the [Ladom Account](#): supposing the Gibbard Contents expressed by $O\phi$ and $\neg O\phi$ have a non-empty intersection entails that there is some world at which both $\Box\phi$ and $\neg\Box\phi$ are true. But, strictly speaking, this contradiction is not properly viewed as a *deliverance* of Gibbard's theory. The reason is that, for Gibbard, the canonical metalinguistic representation of an object-language inconsistency is not in terms of sets of Hyperplans. It is, rather, in terms of the states of mind that sets of Hyperplans serve to represent. Gibbard's account fails to meet the requirement of [Metalinguistic Inconsistency](#) because the *canonical representations he suggests* for $O\phi$ and $\neg O\phi$ —the states of mind of [planning to see to it that ϕ] and [disagreeing with planning to see to it that ϕ]—*do not themselves display any clear inconsistency*.

To put the point a bit differently, we can see that Gibbard is actually making use of two metalanguages, two ways of presenting canonical representations of object-language phenomena: the metalanguage of states of mind, and the metalanguage of Hyperplans. He assumes, crucially, that it is possible to translate freely between these languages. But, we've seen, this is not obvious: whether the states of mind Gibbard designates as [planning to see to it that ϕ] and [disagreeing with planning to see to it that ϕ] can be faithfully represented with disjoint sets of Hyperplans surely depends on whether those states of mind are themselves inconsistent. But Gibbard offers no obvious account of their inconsistency.

Can Gibbard's account be repaired simply by jettisoning the metalanguage of states of mind? Not quite. As [Dreier \(2006, 2009\)](#) and [Schroeder \(2008c\)](#) note, the Gibbard Content of $O\phi$ is the set of Hyperplans that unconditionally require ϕ . This set's complement (the Gibbard Content of $\neg O\phi$) is the set of Hyperplans that do not unconditionally require ϕ . But, by the [Completeness](#) constraint on Hyperplans, the latter is just the set of Hyperplans

that require ϕ in some situation(s) S . So, it's predicted that $\neg O\phi$ (which says *only* that ϕ is not unconditionally required) and *if S , ϕ is required* express the same Gibbard Content—not a happy result! To avoid this, Dreier proposes that Gibbard re-work the Completeness constraint on Hyperplans:

COMPLETENESS*

For any alternative α and situation S , one of the following is a rule of π : *you do α in S , you do not do α in S , or you may choose either to do α or not to do α in S .*

The original formulation of [Completeness](#) misses the fact that one way to be fully decided about what to do in a situation is to decide that it does not matter what you do in a situation. [Completeness*](#) avoids this mistake, and thereby allows $\neg O\phi$ and $O\neg\phi$ to express different Gibbard Contents.

This tack conflicts with the rationale we've used to motivate the [Ladom Account](#). Consider some Hyperplan π_i that, for any situation, is indifferent between ϕ and $\neg\phi$. Obviously, π_i should be in the Gibbard Content of $\neg O\phi$, and not in the Gibbard Content of $O\phi$. But what motivates a semantics that assigns sets of Hyperplans to sentences in this way? The only available answer, as best I can tell, is that a planning state of being indifferent between ϕ and $\neg\phi$ is simply *inconsistent* with one which unconditionally prefers ϕ —that these planning states disagree. But this is exactly the sort of claim which a proponent of the [Ladom Account](#) would prefer to avoid making—such planning states do not, after all, display their inconsistency in the right way (cf. [Schroeder 2008c](#): 586–7).

The real way around this is to abandon Completeness, in any of its forms. The [Ladom Account](#) is entitled to do this because, within the [Ladom Account](#) and the Gibbard account, Hyperplans are *playing completely different theoretical roles*. In the [Ladom Account](#), it is not the job of a Hyperplan to play a role in the representation of practical states of mind analogous to that played by possible worlds in the representation of belief. Rather, the semantic role of the parameter filled by a Hyperplan is simply to characterize an *accessibility relation* (or, in a fancier version of the semantics, a partial ranking) on worlds. Characterizing such a relation in no way demands that a plan be complete, in any sense of that notion; an ordinary To-Do List can do the job just as well as a Hyperplan.³⁷

A.5.3 Its Place in the Metaethical Landscape

To close, some remarks about why the view proposed above, for both directive and normative language, is a clear-cut, full-throated Expressivism, rather than a form of watered-down (“Hybrid”) Expressivism. Recall the standard division of theoretical space sketched in Section [A.2.3](#).

37. Indeed, it's well-known that, if we move to an ordering-semantics, on the model of [Kratzer \(1981\)](#); [Lewis \(1981\)](#), so that the domain of quantification is restricted to the worlds that are *best* according to some set of rules π (roughly, worlds satisfying as many rules in π as any other world), rather than worlds that are ideal, the set of rules or imperatives that determines the accessibility relation needn't even be consistent.

	Q1	Q2	Q3	Q4
PURE	N	N	Y	Y
HYBRID	Y	Y	N	N

The accounts of directives and normative language defended here are, I'll suggest, best plotted as follows.

	Q1	Q2	Q3	Q4
LADOM	N	Y	N	Y

Although I'll focus on the [Ladom Account of directives](#), what I say goes equally for the [Ladom Account of O](#).

The Role of Propositions. The [Ladom Account](#) assigns directives sets of To-Do Lists (partial contingency plans) as their semantic values. It does not assign propositions as their semantic values (hence answers “no” to Q1). But the [Ladom Account](#) does enlist propositions to do logical and semantic work (hence answers “yes” to Q2): the explanation for the inconsistency of $!\phi$ and $;\neg\phi$ is that (i) the contents of these sentences can be faithfully represented as sets of To-Do Lists that make corresponding modal sentences ($\Box\phi$ and $\neg\Box\neg\neg\phi$, respectively) true (a substantive empirical claim about the proper canonical metalinguistic representation for directives), (ii) the supposition that any To-Do List could make both $\Box\phi$ and $\neg\Box\neg\neg\phi$ true leads to contradiction. There is, then, a clear sense in which the semantic properties of *modal propositions explain* the semantic properties of directives. It's just this fact that the [Ladom Account of directives](#) exploits in its account of directive inconsistency.


Semantics, Use, and Meaning. The [Ladom Account of directives](#) builds in a tight link between the semantic value of a directive and its conventional use: the fact that directives conventionally have a non-representational use *motivates a semantics on which they are practically contentful*: a semantics on which they induce a non-trivial partition on sets of possible plans. That they lack a representational use motivates a semantics on which they are not propositionally contentful: they fail to introduce a non-trivial partition on ways the world might be. That descriptive sentences, on the other hand, typically have a conventionalized representational (but lack a conventionalized practical) use motivates a semantics on which they are propositionally contentful (but not practically contentful).³⁸

38. To explain syntactic coordination, subordination, and various other kinds of embedding, it is important (though, pace [Schroeder 2008b](#): 94–5, not essential) that the Expressivist assign normative and descriptive declaratives semantic values of the *same semantic type*. The natural thought is to have all declaratives express a pair of a set of To-Do Lists and a proposition. Normative declaratives will, according to the Expressivist, be distinguished from descriptive declaratives, not by semantic type, instead according to whether their semantic values induce non-trivial partitions on plans or possibilities.

What is the best way to categorize the relationship between semantic facts and facts about conventional use, on this account? The relationship is, I think, best understood in terms of mere *supervenience* or *determination*, rather than *reduction*. On the account we've suggested, the fact that a directive with a given semantic value has a conventionalized non-representational use is a *further fact* about it, albeit one that follows from facts about its semantic value (that it is practically, rather than propositionally, contentful), together with "law-like" linguistic conventions associating ways of being contentful with communicative uses.³⁹

How best to understand the relationship between *conventional meaning* and *conventional use* on this account? Whether a directive's use can claim to be part of its conventional meaning is a largely terminological issue. As Tyler Burge notes, "[T]he term 'meaning' has always been vague, multi-purposed, and to some extent adaptive to the viewpoint of different theories" (1979: 398). One may, of course, define theoretical terms as one likes. Someone for whom facts about conventional meaning are exhausted by *semantic* facts will resist the notion that "dynamic" facts about a sentence's conventional use, although distinct from "static (i.e., properly semantic) facts about that sentence, nevertheless constitute part of its conventional meaning. But, insofar as the [Ladom Account](#) is an account of directive *meaning*, this understanding of "meaning" seems not to accurately reflect its subject matter.

In fact, the [Ladom Account](#) seems to satisfy [CMAU](#). It is, after all, a desire *to get the conventional use of directives right* (together with the claim that a propositional theory of directive content will fail to do this) that motivates the [Ladom Account](#) of directive content. The explanatory structure of the account is, to re-use the diagram from Section [A.2.4](#), something like the following (for an arbitrary directive ϕ):



The [Ladom Account](#) begins from the assumption that an account of directive meaning is fundamentally an account of directive use—that the central constraint of adequacy on canonical representations of directive meaning is that they predict their canonical uses. That the conventional uses of directives are not themselves suitable for certain theoretical jobs—e.g., the job of accounting for facts within the "static" dimension of directive meaning—is not in tension with the account's overarching theoretical aim: to build a theory of meaning

39. I want to stay neutral on the further question of how to classify such conventionalized use-facts. I can say they are probably not classifiable as conventional implicatures (Barker 2000) or purely expressive content (Potts 2007; Boisvert 2008). Such facts, although lexicalized, tend to be *independent of semantic value* (e.g., the fact that 'but' denotes conjunction has no bearing on its conventionally implicating a contrast between its conjuncts), in a way that I'm supposing conventionalized use-facts are not.

that explains the role of directive language in communication and the expression of thought.

In addition to satisfying CMAU, the [Ladom Account of directives](#) also clearly satisfies [Practicality](#) and [Non-Propositionalism](#). So it is a clear-cut form of non-Hybrid Expressivism for directives. But, in answering “no” to Q3 (while still endorsing the supervenience of use on semantic content), it diverges sharply from the sort of theory favored by Expressivism’s major proponents. If our account is on target, then, it suggests a new strategy for Expressivists who are interested in developing a serious, empirically estimable theory of the meaning of normative language. It is, of course, only a start. But it is a promising one. It is explicit about how Expressivism should discharge its linguistic commitments. And it actually, in large part, discharges them.

APPENDIX B

Conditional Utility, Preference, Obligations, and Imperatives

B.1 Introduction

The degree to which a rational agent wants some state of affairs *A* can depend on what is the case about *A*, in at least the following ways:

- **Means-end:** whether *A* causally conduces to some further thing she wants.
- **Rule-case:** whether *A* is constitutive of some thing she wants.

To illustrate means-end dependence, the degree to which an agent wants to quit smoking can depend on whether quitting would make her live longer. And to illustrate rule-case dependence, the degree to which an agent wants to be virtuous can depend on whether virtue is constitutive of her own wellbeing. Or, more mundanely, the degree to which an agent reflectively prefers *A* can depend on whether she *A* is something she wants at all. The same goes, plausibly enough, for *non-degreed* preferences. Whether *A* is preferred, with respect to an arbitrary kind of preference-ranking, can depend, in both of the above ways, on what is the case about *A*.

Paul Weirich (1980) has argued that these sorts of dependencies demand introduction of a notion of *conditional utility* into decision-theoretic frameworks.¹ This essay argues that the notion of conditional utility is philosophically important, not just for decision theorists interested in the correct formulation of the Law of Expected Utility—something on which I will not spend any time here—but also for philosophers interested in the meaning of natural language constructions that function to *express conditional preferences*, including deontic conditionals (*if you want A, you should see to it that A*) and conditional imperatives (*if you want A, see to it that A*). Weirich argues that the conditional utility of a state of affairs *B* on *A* is to be identified as the degree to which it is desired under indicative supposition

1. For related discussions of the notion, see Jeffrey (1983); Joyce (2000).

that *A*. Similarly, I will argue that the proper analysis of indicative deontic and imperative conditionals is in terms of what is planned, desired, or preferred, given monotonic changes to an agent's information. Though my main focus in this paper is on questions of linguistic meaning, also interesting, and susceptible of the sort of methodology that I employ, is the related question of how to represent the conventional *cognitive force* of constructions that express conditional preference. Here, too, I argue that Weirich's idea bears fruit.

Interestingly, implementing Weirich's conception of conditional preference in a semantic analysis of indicative deontic conditionals turns out to be incompatible with a host of popular philosophical accounts of their meaning. These include Wide-Scope accounts (a popular choice in the ethical literature on normative requirements), the philosophically standard "Variably Strict" analysis of Stalnaker and Lewis, and even many sophisticated versions of Context-Shifting accounts (on which indicative antecedents function, in part, to shift the context of interpretation for material in their consequents). In fact, I will ultimately argue that it is incompatible with any *quantificational* approach to the semantics of indicative conditionals—any approach which treats the indicative conditional as expressing non-vacuous universal quantification over some domain of relevant antecedent-possibilities. The only analysis that does a satisfactory job with simple deontic conditionals like *if you want A, you should see to it that A* is an account that blends (i) a Context-Shifty account of indicative antecedents with (ii) an Expressivistic treatment of their deontic consequents. To put the point somewhat differently, if we treat *if you want A, you should see to it that A* as expressing that seeing to it that *A* is best, given that you want *A*—i.e., that the best possibilities, conditional on your wanting *A*, are *A*-possibilities—quantificational analyses of indicative conditionals encode a fundamental misunderstanding of what it is to be *best, given some condition*.

Section B.2 introduces the notion of conditional utility, and briefly recapitulates Weirich's arguments for treating conditional utilities as suppositional degrees of desire. It goes on to develop one of these arguments into an argument against a Wide-Scope analysis of the logical forms of deontic conditionals—an analysis on which the logical form of a deontic conditional like *if you want A, you should see to it that A* is represented with a deontic operator scoping over a conditional. Section B.3 goes on to introduce three different ways of Narrow-Scoping—the Variably Strict approach, Context-Shifty Cognitivism, and Context-Shifty Expressivism—and argues that all of these manage to avoid the sort of Weirich-inspired concerns that troubled Wide-Scope accounts. However, Section B.4 argues that Variable Strictness and Context-Shifty Cognitivism—both of which treat deontic conditionals as expressing, roughly, *known* preferability, on update with the relevant condition—make unpalatable semantic predictions. Specifically, both accounts counterintuitively license bootstrapping—inferring from *if you want A, you should see to it that A* that you should see to it that *A*, when you simply happen to want *A*—as well as antecedent-strengthening—inferring from *if you want A, you should see to it that A* that if you want both *A* and *B*, you

should see to it that A. Context-Shifty Expressivism—which treats deontic conditionals as expressing preferability, on update with the relevant condition—licenses neither. Various damage-control strategies (including denials of the adverted semantic data) are considered for Variable Strictness and Context-Shifty Cognitivism. None, I argue, happens to be very convincing.

In Section B.5, I turn to meta-semantic matters—in particular, to clarifying the extent to which the envisaged account is really a full-fledged form of Expressivism. Contrary, perhaps, to appearances, the account is compatible with a variety of different views in meta-ethics. Hence, its empirical success is not to be construed as any sort of positive argument for Expressivism in meta-ethics. (It is, nevertheless, *inspired by* an Expressivistic treatment of normative judgment, and, as such, Expressivists may—indeed, should be eager to—make free use of it. So, although meta-ethically neutral, the account is nevertheless of substantial meta-ethical interest.)

I go on to argue that, whether or not one accepts meta-ethical Expressivism, one should accept the Expressivist's account of the *cognitive force* of these deontic conditionals. (Such acceptance is, I suggest, meta-ethically neutral.) Entertaining a deontic conditional, I argue, is a matter of *consulting one's preferences*, on update with its antecedent. This, importantly, is distinct from consulting one's beliefs about one's preferences, on update with the antecedent. Correspondingly, deciding that a deontic conditional is the case is a matter of ending up in a cognitive state characterized by a specific, suppositional preference. This, I argue, is distinct both from (i) deciding *that* there is such a suppositional preference (hence distinct from a Cognitivist account of endorsement of deontic conditionals), and (ii) *coming to have* such a suppositional preference. This is a good thing. Any account that conflates a decision to accept a deontic conditional with either (i) or (ii) is incorrect as an account of their cognitive force; the former because it misconstrues the subject matter of such a decision, the latter because it makes it mysterious how a decision to accept a deontic conditional can *resolve a question* for the agent (namely, the issue of *whether* the deontic conditional is the case). On our favored account of the cognitive force of deontic conditionals, it is possible to entertain whether a deontic conditional is the case (hence, accepting a deontic conditional is potentially issue-resolving for an agent). Nevertheless, deciding that a deontic conditional is the case is not a matter of deciding that some body of preferences has some property or other; it is, rather, a matter of responding to a successful query or test on one's suppositional preferences.

B.2 Conditional Utility and Deontic Conditionals

Weirich (1980) considers three ways of conceptualizing the notion of conditional utility. In this section, I describe them briefly, and lay out Weirich's arguments in favor of treating conditional utility as degree of desire under indicative supposition. Since conditional utility, per se, is not our focus, the discussion shall be spare (but still, I hope, suggestive).

I will then show how similar phenomena arise in the case of deontic conditionals; such conditionals intuitively express the preferability of some state of affairs, conditional on a relevant situation obtaining. Unsurprisingly, then, Weirich's argument against understanding the utility of A conditional on B as reducing to the utility of the conditional ($B \rightarrow A$) extends straightforwardly to an argument against treating deontic conditionals expressing conditional preferences in terms of preferences that some conditional be made the case.

B.2.1 Representing Conditional Utility

Proposal 1: Conditional Utility Reduces to Unconditional Utility. Specifically, on this proposal, conditional utility is identified with the utility of the corresponding conditional: $u(A|B) = u(B \rightarrow A)$. (Here and throughout, u is a relative utility function.)

Weirich objects to reduction on roughly the following grounds (1980: 704). Suppose John wants p and also wants q conditional on p . So the corresponding relative utilities for these things— $u(p)$ and $u(q|p)$ —are both positive. But John has decided to make it the case that $\neg q$, since he doesn't know whether or not p , and q has *negative* relative utility conditional on $\neg p$. Then $u(p \rightarrow q)$ is negative, since given what John knows, the only way that $p \rightarrow q$ is if $\neg p$ (recall that p has positive relative utility). So $u(q|p) \neq u(p \rightarrow q)$.

Proposal 2: Conditional Utility is the Conditionally Predicted Utility. In particular, if u_r gives the agent's utility function, *were she to learn r* , then $u(A|B) = u_B(A)$.

Weirich objects on these grounds:

In general, conditional utilities may differ from the corresponding conditionally predicted utilities because the relative utility of an action given a condition may depend on one's being ignorant of the condition, whereas the conditionally predicted relative utility does not (1980: 705).

Here is a case that illustrates the point:

- John wants to be surprised, and he would be surprised to observe p .
- If q , he would observe p . So $u(q|p)$ is positive.
- If John were simply to learn p , he would, of course, *not* be surprised to observe p . And suppose that q is kind of a hassle to John, so he has no independent reason to do it. Then $u_p(q)$ is negative.
- So $u(q|p) \neq u_p(q)$.

A quick dialectical aside to preliminarily situate Weirich's discussion in the philosophical dialectic on conditionals. Something like this point is, I think, familiar from philosophical discussions of the Ramsey Test. The naive formulation of the Ramsey Test for an indicative conditional [*if* p][q] says pretend you know p , then assess whether q . That version runs into

well-known problems (*if my wife is cheating, I won't find out*), which are structurally *just* like the problem for Proposal 2 described here.

A more sophisticated version of the Ramsey Test—the version recently popularized by Irene Heim and a family of dynamic semantic treatments of conditionals (see, e.g., Karttunen 1973; Heim 1983; Yalcin 2007; Gillies 2010)—says that an indicative antecedent p characterizes a property of an information state: the property of accepting p . To evaluate the indicative is a matter of moving to the nearest information state that satisfies that property (which we can gloss roughly as *supposing that p*), and then evaluating whether q , from the vantage point of that enriched information state. That is crucially different from evaluating whether q , under the supposition that *I have accepted p* ; supposing that p leaves it open whether or not one *also accepts p* , under this supposition. Similarly for conditional utilities: Proposal 2 errs by suggesting that a relative utility conditional on p is a matter of evaluating relative utility on the supposition that p is known/learned.

All this would seem to suggest an account of conditional utility in terms of properties of an agent's utility function, under the supposition that the relevant condition holds (where supposition that p is understood in terms of having an information state that accepts p). And, indeed, that's exactly what Weirich goes on to suggest.

Proposal 3: Conditional Utility is the Degree of Desire, Under (Indicative) Supposition that the Condition Holds. In particular, if $u_{|s|}$ gives the agent's "hypothetical" utility function, when she supposes s , then $u(A|B) = u_{|B|}(A)$. Notice that the unviability of Proposal 1 means that the notion *degree of desire under supposition that the condition holds* cannot be understood as the degree of desire that the corresponding material implication hold. Rather, hypothetical utility functions are understood, irreducibly, as degrees of desire that are read off from the agent's preferences when she is making the relevant supposition. Pretty clearly Proposal 3 gets the right results about the cases that troubled Proposals 1-2.²

Interestingly, similar sorts of issues arise for constructions in natural language that are used to *express conditional preferences* (of which descriptions of conditional requirements and obligations—hereafter "deontic conditionals—are a special case). In the remainder of this section, I'll outline a central case and argue that considerations much like those that Weirich uses to argue against reductionism about conditional utility argue against Wide-Scope treatments of deontic conditionals (and corresponding reductionism about conditional deontic selection functions).

2. Precising this proposal gives rise to a host of complications, but addressing them would bog us down in decision-theoretic issues that are not strictly relevant to my project here. The proposal that we eventually articulate for deontic conditionals will be sufficiently precise and well-developed to constitute a constructive account of their meaning and cognitive force.

B.2.2 Against Wide-Scoping

Here is a case much like the case that causes trouble for the proposal to treat conditional utility as a special kind of unconditional utility. You're something of a loner, and this is the source of considerable pain for you. You know that it would be best (for you) if you wanted a family: you should (in view of what's good for you) want a family. And, given that you want a family, you should (also in view of what's good for you) see to it that you have a family. But you don't know whether or not you want a family: you're the kind of loner who wants to want to a domestic life, but who is unsure whether such a life will actually make you happy. Indeed, if you don't actually want a family, then having one will end up making you fantastically unhappy.

- (B.1) a. $O(\text{want}(A))$
You should want A
- b. $O(A|\text{want}(A))$
 \approx Conditional on wanting A , you should see to it that A
 \approx If you want A , you should see to it that A
- c. $O(\neg A)$
All things considered, you shouldn't see to it that A
(Since A has high negative relative goodness, conditional on *not* wanting A)

There is a question about how to represent the logical form of the "hypothetical imperative" expressed by (B.1b). This question parallels the question about how to represent the conditional utility of A conditional on wanting A : what $O(A|\text{want}(A))$ says is that, given that you want A , the relative utility of seeing to it that A is high. There are broadly two kinds of option here:³

- **Wide-scoping.** The logical form of (B.1b) is $O(\text{want}(A) \rightarrow A)$.
- **Narrow-scoping.** The logical form of (B.1b) is $[\text{if } \text{want}(A)][O(A)]$. ($[\text{if } \phi][\psi]$ is a regimented indicative conditional whose antecedent is ϕ and consequent is ψ .)

Wide-scoping in this case runs into roughly the same sort of problem that afflicted Proposal 1 about conditional utility. (B.1b) sounds true. But $O(\text{want}(A) \rightarrow A)$ must, given our assumptions, be false. Why is that? Suppose that $O(\text{want}(A) \rightarrow A)$ were true. If O expresses a kind of modal necessity, as I will assume here, then (B.2) is a substitution

3. Actually, there are three: as is standard in the semantics literature (cf. Kratzer 1981), we could treat O as a dyadic quantifier, with both restriction and scope arguments. But (i) this runs into problems that are beyond the scope of our discussion here (and remedying those problems requires making the dyadic approach effectively equivalent to the narrow-scoping approach; cf. von Stechow & Iatridou 2005), (ii) it has no direct parallel in the conditional utility case.

instance of Modal Axiom **K**.⁴

$$(B.2) \quad O(\text{want}(A) \rightarrow A) \rightarrow (O(\text{want}(A)) \rightarrow O(A))$$

Given that $O(\text{want}(A))$ is true, $O(A)$ follows, on the assumption that $O(\text{want}(A) \rightarrow A)$. This represents a strong argument against Wide-Scoping the operator O in rendering the logical form of (B.1b).

B.2.3 Conditional Selection Functions

It's instructive to recreate this result at the level of the semantics—specifically, at the level of *selection functions* (which, in the semantics of natural language deontic modals, play a role rather analogous to the role of utility functions in decision theory).

The deontic modal O is standardly interpreted with respect to a *modal base* and some sort of *selection function*, both furnished by some sort of *state* (which can be identified, depending on one's metasemantic inclinations, with a context, an abstract representation of a state of mind, etc.) (see, e.g., Kratzer 1981). The modal base provides the set of possibilities that are *relevant* at a state, while the selection function selects the best state-relevant possibilities, yielding the modal O 's domain of quantification.

Definition 1. A *state* S determines a pair $\langle f_S, \sigma_S \rangle$. $f_S = \{i : i \text{ is a relevant possibility at } S\}$ is a *modal base*, supplying the set of S -relevant possibilities (the S -relevant information). And σ_S is a *selection function* selecting the S -best possibilities from a set.

We will assume, as is fairly standard:

REALISM

$$\sigma_S(f_S) \subseteq f_S \quad (\text{the preferred possibilities are always relevant})$$

DEFINEDNESS

$$\sigma_S(f_S) \text{ is defined and non-empty}^5$$

Deontic necessity modals are assumed, as is also standard, to be universal quantifiers over the domain defined by the modal base and selection function: $O(\phi)$ is true at S just if the S -preferred possibilities are ϕ -possibilities.

Definition 2. A modal formula of the form $\ulcorner O(\phi) \urcorner$ is true at S (notation: $S \models O(\phi)$ or $\llbracket O(\phi) \rrbracket^S = 1$) iff $\sigma_S(f_S) \subseteq \llbracket \phi \rrbracket$. (Note: $\llbracket \phi \rrbracket$ gives the set of ϕ -possibilities.)

4. This is a plausible assumption, but the argument does not require it. As with the case that made trouble for Proposal 1, we might suppose that you take notice of the risk of acquiring a family, and decides against it. So, given what you know, the only way to make it the case that $(\text{want}(A) \rightarrow A)$ is not to want A . But that conflicts with our supposition that you should want A . So, $O(\text{want}(A) \rightarrow A)$ is actually, given what you know, false. But, intuitively, $O(A|\text{want}(A))$ is true. So the logical form of the latter cannot be given by the former.

5. Definedness is the Limit Assumption (Lewis 1973). Rejecting it affects none of my arguments, and would complicate the semantics. See Kratzer (1981) for discussion.

On this semantics, if we suppose that the sentences in (B.1) are true at S , and help ourselves to some notation of the form $\sigma_S(f_S|B)$ —designating, roughly, the relevant σ_S -preferred possibilities, given B —we have the following properties of σ_S .

- (B.3) a. $\sigma_S(f_S) \subseteq \llbracket \text{want}(A) \rrbracket$ (*want*(A)-possibilities are preferred)
 b. $\sigma_S(f_S|\text{want}(A)) \subseteq \llbracket A \rrbracket$ (Given your wanting A , A -possibilities are preferred)
 c. $\sigma_S(f_S) \subseteq \llbracket \neg A \rrbracket$ ($\neg A$ -possibilities are preferred)

The conditions on σ_S found in (B.3) are, I claim, plausible representations of the satisfaction-conditions of the obligation-sentences in (B.1). However, the as-yet-uninterpreted notion of a **conditional selection function**, introduced to represent the meaning of the conditional obligation-sentence (B.1b), cries out for clarification.

The Wide-Scoper—or, rather, the Wide-Scoper’s semantic ally—has a proposal at the ready: a preference for A possibilities, given B , is defined equivalent to an unconditional preference for $(B \rightarrow A)$ -possibilities.

REDUCTION FOR CONDITIONAL SELECTION FUNCTIONS

$$\sigma_S(f_S|B) \subseteq \llbracket A \rrbracket \text{ iff } \sigma_S(f_S) \subseteq \llbracket B \rightarrow A \rrbracket$$

The reductive proposal for the conditional selection function cannot, though, be correct. Given (B.3b) and the reductive proposal for conditional selection functions, we have:

- (B.4) $\sigma_S(f_S) \subseteq \llbracket \text{want}(A) \rightarrow A \rrbracket$ (*want*(A) $\rightarrow A$ -possibilities are preferred)

The difficulty here, of course, is that (B.4) cannot hold if (B.3a) and (B.3c) hold. It is a consequence of (B.3a) and (B.3c) that $(\text{want}(A) \wedge \neg A)$ -possibilities are preferred. Whence it follows that $(\text{want}(A) \rightarrow A)$ -possibilities must be strictly *dispreferred*.

The parallel with Weirich’s trouble case for Proposal 1 is now fully explicit. In both cases, a common reductive maneuver is suggested: with conditional utility, it is suggested that conditional utility is the utility of the corresponding conditional, while with conditional selection functions, it is suggested that conditional preference for possibilities is equivalent to a preference for possibilities satisfying the corresponding conditional. But it is not difficult to contrive circumstances in which the corresponding conditional has negative relative utility or is strictly dispreferred, even when the relevant conditional utility is intuitively positive and the relevant conditional obligation intuitively is the case. So conditional utility does not reduce to the unconditional utility of a conditional, and the verdict of a conditional deontic selection function for some state of affairs cannot be understood as the verdict of an unconditional deontic selection function for some conditional state of affairs.

B.3 Ways of Narrow-Scoping

Having rejected Wide-Scoping, we have apparently committed ourselves to assigning (B.1b) a logical form in which the modal O is interpreted *in situ*, something like: $[if\ want(A)][O(A)]$. But this is well short of an analysis. A complete semantics for (B.1b) would provide at least:

- An interpretation for the conditional construction⁶ $[if \cdot][\cdot]$
- An interpretation for the modal O
- An explanation of the intuitive equivalence of (B.1b) and the metalinguistic gloss in terms of conditional preference given in (B.3b): *given your wanting A , A -possibilities are preferred*. (Hence, a complete semantics would provide an interpretation of the notion of a conditional selection function.)

In this section, I want to present (and, in the subsequent section, critically evaluate) two distinct approaches to the semantics of (B.1b), all of which strictly improve on the Wide-Scope approach:

- The *Variably Strict* approach (due originally to Stalnaker 1968; Lewis 1973)
- The *Context-Shifting* approach (which has various⁶ sources, but see Karttunen 1973; Heim 1983; Yalcin 2007; Gillies 2010; Kolodny & MacFarlane 2010)

We will ultimately develop two versions of the Context-Shifting approach: one on which indicatives express universal quantification over a domain of possibilities, but shift the context of interpretation for proposition-expressing consequents (*Context-Shifty Cognitivism*), and one on which indicatives are treated as non-quantificational devices for shifting the context of interpretation for non-proposition-expressing consequents (*Context-Shifty Expressivism*).

B.3.1 Variably Strict

The Variably Strict analysis of Stalnaker (1968); Lewis (1973) for the indicative conditional can be informally glossed thus: $[if \phi][\psi]$ is true just if the closest ϕ -possibilities are ψ -possibilities. More carefully, the Variably Strict analysis interprets indicatives relative to a similarity ordering \leq_i that is *weakly centered* on a point of evaluation $i \in W$:

WEAK CENTERING

$$\forall p \subseteq W : i \text{ is } \leq_i \text{ minimal in } p \text{ (i.e., } \neg \exists j \in p : j \leq_i i \wedge i \not\leq_i j)$$

6. I don't want to suggest that a complete semantics for (B.1b) requires an analysis of all indicative conditionals in natural language (or even of all indicative conditionals with modalized consequents). (B.1b) might (indeed, many have claimed it *does*) have special properties that make, e.g., a Wide Scope analysis seem plausible, whereas such an analysis might not be generally correct for indicative conditionals with modalized consequents.

Weak centering codifies the intuition that, if i is a ϕ -world, there are no ϕ -worlds closer to i than i itself. Let \rightsquigarrow represent a variably strict conditional operator at logical form. On the Stalnaker-Lewis analysis, then, an indicative of the form $\lceil \phi \rightsquigarrow \psi \rceil$ will be true at i iff the \leq_i -minimal ϕ -worlds are ψ worlds.⁷

Definition 3. $\llbracket \phi \rightsquigarrow \psi \rrbracket^{S,i} = 1$ iff $\forall j \in f_{S_i} \cap \llbracket \phi \rrbracket : j$ is \leq_i -minimal in $\llbracket \phi \rrbracket \Rightarrow \llbracket \psi \rrbracket^{S,j} = 1$.

As before, $\llbracket \phi \rrbracket$ gives the set of ϕ -possibilities (whenever no state is made explicit, ϕ 's interpretation is assumed to be state-invariant). Note that truth is now (explicitly) relativized to points of evaluation—possibilities—in addition to states. Truth, of course, is always (implicitly) relative to points of evaluation, but they often play no real semantic role. Not so with the Variably Strict analysis, given the role of Weak Centering in the semantics. (Throughout this paper, the relativity of truth to points of evaluation will be suppressed, unless semantically interesting.)

The Variably Strict analysis of (B.1b) comes in two parts:

- A rendering of (B.1b)'s logical form as $want(A) \rightsquigarrow O(A)$
- A statement of truth-conditions for the embedded modal sentence $O(A)$. (Here, the fan of Variable Strict-ness is free to simply appropriate the semantics for deontic modals given in Definition 2.)

How does this fare? Pretty well. Informally, $want(A) \rightsquigarrow O(A)$ is true at just if, in the closest possibilities where you want A , you ought to see to it that A . And this holds just if, in all the closest possibilities where you want A , the preferred possibilities there are all A -possibilities.⁸ But, unlike the truth-conditions assigned by the Wide-Scoper, these conditions are completely compatible with the *actually preferred* possibilities being $\neg A$ possibilities, even given an actual preference for possibilities where you want A . $want(A) \rightsquigarrow O(A)$ says that in (close) possibilities where you want A , A is preferred. A conclusion about what is *actually preferred* follows when you do want A , but not, in general, otherwise. Recall that uncertainty about whether you want A is built into (indeed, integral to the plausibility of) the case. So it is indeterminate, given the description of the case, whether A is actually preferred. And that, according to the fan of Variable Strictness, is as it ought to be. (There is some *serious* fudging going on in this explanation, but I will hold off on that until the bootstrapping discussion in Section B.4.1.)

The Variably Strict analysis thus invalidates the move from (B.1a) and (B.1b) to $O(A)$. We can see this as due to its proprietary understanding of conditional selection functions: what is preferred, given B , is identified with what is preferred throughout the relevant B -possibilities. (For simplicity, we here ignore the role of the similarity ordering \leq .)

7. Again, we make the Limit Assumption. Nothing important turns on this.

8. More carefully, by Definition 3, $\llbracket want(A) \rightsquigarrow O(A) \rrbracket^{S,i} = 1$ just if $\forall j \in f_{S_i} \cap \llbracket want(A) \rrbracket : j$ is \leq_i -minimal in $\llbracket want(A) \rrbracket \Rightarrow \llbracket O(A) \rrbracket^{S,j} = 1$. By Definition 2, this just if $\forall j \in f_{S_i} \cap \llbracket want(A) \rrbracket : j$ is \leq_i -minimal in $\llbracket want(A) \rrbracket \Rightarrow \sigma_{S_i}(f_{S_i}) \subseteq \llbracket \phi \rrbracket$.

$$\sigma_{S_i}(f_{S_i}|B) \subseteq \llbracket A \rrbracket \text{ iff } \forall j \in f_{S_i} \cap \llbracket B \rrbracket : \sigma_{S_j}(f_{S_j}) \subseteq \llbracket A \rrbracket$$

The Wide-Scoper's difficulties traced to her reduction of conditional preference to unconditional preference for a corresponding conditional state of affairs. The Variably Strict analysis avoids such a reduction: for possibilities to be preferred, given B , is for them to be preferred throughout the modal base, when it has been incremented with the information that B . And, so, the Variably Strict analysis has a prima facie plausible claim to being a semantic implementation of Weirich's thesis that being preferred, conditional on A , is just being preferred on the supposition that A .

So far, so good.

B.3.2 Context-Shifting

What I will call the Context-Shifting approach draws on the intuition that checking for the truth of $[if \phi][\psi]$ at a state S involves checking whether ψ holds at the result of updating S with ϕ .⁹ A variety of approaches in this vein are possible, given various possibilities for understanding the intuitive notions of:

- A sentence *holding at a state*
- The *result of updating* a state with a sentence

Indeed, as we will see, the Context-Shifting approach admits both *truth-conditional* and *non-truth-conditional* (or *expressivist*, or *dynamic*) understandings of the relevant deontic conditionals.¹⁰

It is natural (and, in dynamic frameworks, standard) to understand the notions of holding at and updating a state as interrelated. First, updating S with ϕ typically yields a state $S|\phi|$ at which ϕ holds. We will write this:

$$(B.5) \quad S|\phi| \Vdash \phi$$

Second, a sentence ϕ holds at a state S just if updating S with ϕ does nothing to S :

$$(B.6) \quad S \Vdash \phi \text{ iff } S|\phi| = S$$

9. For riffs on this theme, see Karttunen (1973); Heim (1983); Yalcin (2007); Gillies (2010); Kolodny & MacFarlane (2010). I will be developing the Context-Shifting approach within a dynamic (update-theoretic) framework, but that is not strictly essential to the Context-Shifting story. Indeed, Yalcin (2007); Gillies (2010); Kolodny & MacFarlane (2010) show how to develop the same sort of account within a non-dynamic, truth-conditional framework (although it is unclear how such an account might be extended to indicative conditionals with non-declarative consequents—e.g., conditional imperatives).

10. This is fortunate, for it will allow us to generalize our treatment of these conditionals to corresponding conditional imperatives, which exhibit phenomena in all respects similar to the phenomena for the corresponding deontic conditionals, but seem not to admit of truth-conditional analysis. It should also be of interest to the meta-ethical Expressivist.

On the Context-Shifting approach, then, checking for the truth of $[if\ \phi][\psi]$ involves checking whether ψ holds at a state at which ϕ also holds.¹¹

$$(B.7) \quad S \Vdash [if\ \phi][\psi] \text{ iff } S|\phi| \Vdash \psi \text{ iff } S|\phi||\psi| = S|\phi|$$

In dynamic systems, a state S usually is limited to what we have characterized as its “informational” parameter f_S . In these systems, the function mapping sentences into update functions can be given the following recursive characterization.

(B.8) THE DYNAMIC INTERPRETATION FUNCTION $|\cdot|$

$$\begin{aligned} S|p| &= S \cap \llbracket p \rrbracket \\ S|\phi \wedge \psi| &= S|\phi||\psi| \\ S|\neg\phi| &= S - S|\phi| \end{aligned}$$

It is easy to see that, on this definition, when the object language is propositional or first-order (i.e., all sentences are Boolean combinations of atoms), updating with ϕ amounts just to incrementing the information with ϕ :

$$(B.9) \quad S|\phi| = S \cap \llbracket \phi \rrbracket$$

This system must be revised, if we understand the object language as containing sentences of the form $\ulcorner O(\phi) \urcorner$ (hereafter, “deontic sentences”): (B.8) must be augmented with a clause for such formulas. And this requires broadening our conception of the parameter of the state on which discourse can “bear” beyond the usual dynamic understanding. We have, recall, assumed that interpretation of such sentences at a state S is sensitive, not just to the S -relevant information, but also the selection function σ_S . We cannot accommodate deontic sentences unless we allow that the content of such sentences relates, somehow, to properties of the selection function.

B.3.2.1 Cognitivism

On one conception of that relation (which I will somewhat tendentiously label “Cognitivism”), the “content” of a deontic sentence is *informational* in nature: updating a state with a deontic sentence yields a state where the information entails that sentence.

$$(B.10) \quad f_{S|O(\phi)|} = f_S \cap \llbracket O(\phi) \rrbracket^S$$

Cognitivism, then, holds that updating a state with a deontic sentence yields a new state at which it is “known,” roughly speaking, that that deontic sentence is true. Supposing that updating with a deontic sentence yields a state at which it holds, we have that a deontic

11. In non-dynamic versions of this approach, truth-conditions for indicatives are given roughly as follows: $S \models [if\ \phi][\psi]$ iff $S|\phi| \models \psi$ (cf. Gillies 2010).

sentence holds at a state just if the information already entails it.

$$(B.11) \quad S \Vdash O(\phi) \text{ iff } f_S \subseteq \llbracket O(\phi) \rrbracket^S$$

For a complete account, the Cognitivist can appropriate the standard semantics for O :

$$(B.12) \quad \llbracket O(\phi) \rrbracket^S = \{i : \sigma_{S_i}(f_{S_i}) \subseteq \llbracket \phi \rrbracket\}$$

The Cognitivist thus models update on a deontic sentence $O(\phi)$ as *learning that the selection function has some property*, namely, the property of preferring ϕ .

We may combine Cognitivism with the Context-Shifting account of indicatives to obtain an analysis of the problematic sentence (B.1b). The Cognitivist suggests rendering the logical form of the conditional *ought* as the regimented indicative conditional $[if\ want(A)][O(A)]$. This yields the following analysis of (B.1b): $[if\ want(A)][O(A)]$ holds at S just if $O(A)$ holds at $S|want(A)$.

$$(B.13) \quad S \Vdash [if\ want(A)][O(A)] \text{ iff } S|want(A) \Vdash O(A)$$

This avoids the Wide-Scoper's problem, in much the same way as the Variably Strict account: $[if\ want(A)][O(A)]$ holds at S just if, should S learn $want(A)$, S thereby learns $O(A)$. This is the case just if, should S learn $want(A)$, at all the possibilities compatible with S 's information, A -possibilities will be preferred. That is completely compatible with the actually preferred (i.e., S -preferred) possibilities being $\neg A$ possibilities, even given an actual preference for possibilities where A is wanted. Suppose, for instance, that some possibilities are $\neg want(A)$ possibilities: $f_S \cap \llbracket \neg want(A) \rrbracket \neq \emptyset$. In *these* possibilities, recall, $\neg A$ is preferred. So since both $O(A)$ and $O(\neg A)$ are true at some relevant possibilities, $f_S \not\subseteq \llbracket O(A) \rrbracket^S$. Hence, according to Cognitivism—specifically, claim (B.11)— $S \not\Vdash O(A)$.

The Cognitivist account thus avoids validating the move from (B.1a) and (B.1b) to $O(A)$. As before, we can see this as due to its proprietary understanding of conditional selection functions: what is preferred, given B , is identified with what is *known to be preferred* on update with B (i.e., the hypothetical elimination of $\neg B$ possibilities from consideration).

PREFERABILITY-GIVEN- B AS KNOWN PREFERABILITY ON UPDATE WITH B

$$\sigma_{S_i}(f_{S_i}|B) \subseteq \llbracket A \rrbracket \text{ iff } \forall j \in f_{S|B_i} : \sigma_{S|B_j}(f_{S|B_j}) \subseteq \llbracket A \rrbracket$$

Note that the Cognitivist's conception of conditional preferability is *not* interchangeable with the Variably Strict conception (repeated here).

PREFERABILITY-GIVEN- B AS PREFERABILITY THROUGHOUT B -WORLDS

$$\sigma_{S_i}(f_{S_i}|B) \subseteq \llbracket A \rrbracket \text{ iff } \forall j \in f_{S_i} \cap \llbracket B \rrbracket : \sigma_{S_j}(f_{S_j}) \subseteq \llbracket A \rrbracket$$

Parts of these definitions can be interchanged, since, for the Cognitivist, update on any kind

of sentence is informational in character.

- $f_{S|B|_i} = f_{S_i} \cap \llbracket B \rrbracket^S$
(Update on a sentence adds its content to the relevant information)
- $\sigma_{S|B|_j} = \sigma_{S_j}$
(By definition, update on a sentence never alters the selection function)

But parts are not interchangeable: for an arbitrary B -possibility j , it may be the case at j that the best B -possibilities are A -possibilities, without it being the case at j that the best possibilities *simpliciter* are A -possibilities (although not vice versa, since $f_{S|B|_j}$ may be properly included in f_{S_j} , but not the other way around).

$$(B.14) \quad \sigma_{S|B|_j}(f_{S|B|_j}) \subseteq \llbracket A \rrbracket \not\Rightarrow \sigma_{S_j}(f_{S_j}) \subseteq \llbracket A \rrbracket$$

$$(B.15) \quad \sigma_{S_j}(f_{S_j}) \subseteq \llbracket A \rrbracket \Rightarrow \sigma_{S|B|_j}(f_{S|B|_j}) \subseteq \llbracket A \rrbracket$$

The Variably Strict notion of conditional preferability is thus *stronger* than the Cognitivist notion. On Variable Strictness, if A is preferable-given- B , then A is *preferred at all relevant B -possibilities*. On Cognitivism, if A is preferable-given- B , then A does not have to be preferred at all relevant B -possibilities; at all relevant B -possibilities, A -possibilities only have to be preferred *among the B -possibilities*, not preferred *simpliciter*. This may seem like a small difference, but in Section B.4.1, we'll see that it is actually significant.

B.3.2.2 Expressivism

Another version of the Context-Shifting analysis—owing to a different notion of the relationship between the content of deontic sentences and properties of the selection function—is possible. According to Cognitivists, deontic sentences have informational content: they *proffer* or *assert* information about the selection function. But deontic sentences might have content that bears on properties of the selection function, without that content being informational in character. Deontic sentences might *express properties of a selection function*, without *asserting that* the selection function has those properties. This sort of account has evident affinities with the Expressivist meta-ethic defended in, e.g., Blackburn (1984, 1998); Gibbard (1990, 2003). So it's natural to refer to this version of the Context-Shifting analysis as Expressivist.¹²

What, then, is it for a deontic sentence $O(\phi)$ to hold at a state S ? For the Cognitivist, it is for the S -relevant information to entail that the deontic selection has a certain property: the property of requiring ϕ . For the Expressivist, it is simply for the selection function *to have*

12. Gibbard, especially, makes systematic use of the distinction between expressing acceptance of a norm and asserting that one accepts a norm. The analysis sketched in this section (and in Section B.5) can be seen as an attempt to implement the key tenets of Gibbard's expressivism in a formal semantics for natural language. For more on this, see my "Meaning for Expressivists" (Appendix A).

a certain property, namely, requiring ϕ , as in (B.16). (There is, by the way, no real need to relativize the parameters in the Expressivistic definitions to individual points of evaluation. So, to save ink, we'll decline to do so.)

$$(B.16) \quad S \Vdash O(\phi) \text{ iff } S \models O(\phi) \text{ iff, by Definition 2, } \sigma_S(f_S) \subseteq \llbracket \phi \rrbracket$$

Thinking, for the moment, of states as representations of states of mind, the Cognitivist might understand acceptance of a deontic sentence $O(\phi)$ as analogous to holding a *higher-order* attitude: if σ_S represents the content of an agent's plans (or goals, or desires), the attitude might be glossed as thinking ϕ is a requirement of one's plans (or goals, or desires). The Expressivist, to contrast, might analogize it to a fact about one's *first-order* attitudes: S 's acceptance of $O(\phi)$ is analogous to ϕ just *being* required, in view of the agent's plans (goals, desires). To think $O(\phi)$ is for ϕ to be required, relative to a state representing the agent's information and plans (cf. Yalcin 2007).

There is some difficulty with relating this conception of acceptance of a deontic sentence to a definition of the update performed by such a sentence on a state—i.e., to a definition of $|O(\phi)|$, and we will return to this in Section B.5. But that is not strictly relevant to saying how Expressivist implementation of the Context-Shifting account would go. Indeed, since the Expressivist analysis pertains to the content of deontic sentences, rather than to the content of indicative conditionals *per se*, the Expressivist may simply take on the Cognitivist's account of what it is for a deontic conditional to hold at a state. So, for (B.1b), both Context-Shifty Cognitivists and Context-Shifty Expressivists claim that $[if\ want(A)][O(A)]$ holds at S just if $O(A)$ holds at $S|want(A)|$.

$$(B.17) \quad S \Vdash [if\ want(A)][O(A)] \text{ iff } S|want(A)| \Vdash O(A)$$

That is not, of course, to say that Cognitivism and Expressivism are equivalent for (B.1b). Cognitivists hold that: $S|want(A)| \Vdash O(A)$ just if A is preferred throughout $f_{S|want(A)|}$.

$$(B.18) \quad S|want(A)| \Vdash O(A) \text{ iff } \forall j \in f_{S|want(A)|_j} : \sigma_{S|want(A)|_j}(f_{S|want(A)|_j}) \subseteq \llbracket A \rrbracket$$

Expressivists hold that: $S|want(A)| \Vdash O(A)$ just if A is preferred at $S|want(A)|$.

$$(B.19) \quad S|want(A)| \Vdash O(A) \text{ iff } \sigma_{S|want(A)|}(f_{S|want(A)|}) \subseteq \llbracket A \rrbracket$$

According to both Cognitivists and Expressivists, evaluating (B.1b) is a matter of checking whether the hypothetical state $S|want(A)|$ has a certain property. For Cognitivists, that property is *informational*; for Expressivists, it is not.

Expressivism thus embraces a distinctive understanding of conditional selection functions: what is preferred, given B , is identified simply—and, it must be said, very intuitively—with what is preferred on update with B .

$$\sigma_{S_i}(f_{S_i}|B) \subseteq \llbracket A \rrbracket \text{ iff } \sigma_{S|B}(f_{S|B}) \subseteq \llbracket A \rrbracket$$

Compare the Cognitivist conception, which understands preferability-given- B as *known* preferability, on update with B . Expressivism, meanwhile, understands conditional selection functions as exactly analogous to conditional utility functions (on Weirich’s suggested understanding): something is conditionally preferred just if it is preferred *under the supposition* that the condition holds—just if it is, in fact, preferred from the vantage of the hypothetical state that results from updating with the relevant condition. Pretty clearly, this proposal also does the job for (B.1b).

B.4 For Expressivism

We have seen that both Variably Strict and Context-Shifting accounts turn out to yield a *prima facie* attractive analysis of (B.1b). That is not surprising. For both arguably implement, albeit in different ways, Weirich’s suggestion that being preferred, conditional on A , is just being preferred on supposition that A . It is not, though, that there is nothing to decide between these two approaches. Indeed, as I will argue in this section, two pieces of data seem to tell independently against the Variably Strict analysis (and, ultimately, against the Cognitivist version of the Context-Shifting account).

- Variably Strict accounts, as well as Context-Shifty Cognitivist accounts, are seen to license—contrary to intuition (and a body of literature on the logic of rational requirements)—**bootstrapping**: when (i) (B.1b) is true, and (ii) you want A , the Variably Strict analysis predicts—regardless of whether you know that you want A —that you ought to see to it that A . The Variably Strict analysis relinquishes a key attraction of the Wide-Scope account of rational requirements: the ability to block detachment of a normative conclusion when the antecedent is true.
- As is well-known, the Variably Strict analysis counterintuitively validates **antecedent-strengthening**: $[if \phi][\psi] \models [if (\phi \wedge \chi)][\psi]$. Context-Shifty Cognitivism does too. The apparent invalidity of antecedent-strengthening is often given a pragmatic or presuppositional explanation, so that antecedent-strengthening is semantically *valid* (given a fixed domain of quantification), but the corresponding inferences are not generally pragmatically *licensed*. Although I endorse these explanations in the cases they are designed to cover (so-called “Sobel Sequences”), I will argue that they cannot be carried over to deontic conditionals whose antecedents articulate *goals* or *desires* of an agent. Such conditionals are semantically *non-monotonic*.

Context-Shifty Expressivism is argued to resolve both of these difficulties. That is not surprising. Properly interpreted, the cases that distinguish Variably Strict and Context-Shifting approaches (and Context-Shifty Cognitivism from Context-Shifty Expressivism) highlight the superiority of Expressivism as an implementation of Weirich’s notion that conditional preferability is suppositional preferability (and ultimately as an implementation

of Ramsey’s idea that interpretation of $[\textit{if } \phi][\psi]$ proceeds by evaluating whether ψ from the standpoint of a state that accepts ϕ).

B.4.1 Bootstrapping

B.4.1.1 What Is It?

Both the Variably Strict and Context-Shifty Cognitivist analyses turn out to license some form of bootstrapping:

BOOTSTRAPPING

Whenever $[\textit{if want}(A)][O(A)]$ holds at a state, and the agent *actually does want* A , the agent *actually ought* to see to it that A .

In the case of the Variably Strict semantics, this is a corollary of the fact that variably strict conditionals (hence, indicative conditionals whose logical forms are rendered using variably strict conditionals) validate modus ponens.

$$(B.20) \quad \text{If } \llbracket \phi \rightsquigarrow \psi \rrbracket^{S,i} = \llbracket \phi \rrbracket^{S,i} = 1, \text{ then } \llbracket \psi \rrbracket^{S,i} = 1$$

Proof. Suppose $\llbracket \phi \rightsquigarrow \psi \rrbracket^{S,i} = \llbracket \phi \rrbracket^{S,i} = 1$. Then, by Def. 3, for any $j \in \llbracket \phi \rrbracket$ that is \leq_i -minimal in $\llbracket \phi \rrbracket$, $\llbracket \psi \rrbracket^{S,j} = 1$. By Weak Centering, i is \leq_i -minimal in $\llbracket \phi \rrbracket$. By supposition, $\llbracket \phi \rrbracket^{S,i} = 1$. So $\llbracket \psi \rrbracket^{S,i} = 1$. \square

So, supposing a Variably Strict analysis of $[\textit{if want}(A)][O(A)]$, it follows immediately that $\llbracket [\textit{if want}(A)][O(A)] \rrbracket^{S,i} = \llbracket \textit{want}(A) \rrbracket^{S,i} = 1$ implies $\llbracket O(A) \rrbracket^{S,i} = 1$. So the Variably Strict analysis licenses bootstrapping.

A limited version of this result also holds for Context-Shifty Cognitivism, if we make the (standard) assumption that the actual world is always practically relevant: that the accessibility relation corresponding to relevance at a possibility is reflexive.¹³

REFLEXIVITY OF RELEVANCE

$$\forall i, S : i \in f_{S_i}$$

The limited result is that, if i_S is the “actual” world (i.e., the world locating the agent whom S represents), $S \Vdash [\textit{if } \phi][O(\psi)]$, and $\llbracket \phi \rrbracket^{S,i_S} = 1$, then $\llbracket O(\psi) \rrbracket^{S,i_S} = 1$.

$$(B.21) \quad \text{If } S \Vdash [\textit{if } \phi][O(\psi)] \text{ and } \llbracket \phi \rrbracket^{S,i_S} = 1, \text{ then } \llbracket O(\psi) \rrbracket^{S,i_S} = 1$$

13. The assumption can be motivated by noting that the possibilities relevant for an agent’s decision are typically the agent’s *epistemic* alternatives—the possibilities compatible with what the agent knows. Since knowledge is factive, the actual world can never be eliminated from an agent’s epistemic alternatives; this amounts to saying the agent knows a falsehood.

Proof. Suppose $S \Vdash [if \phi][O(\psi)]$, and $\llbracket \phi \rrbracket^{S, i_S} = 1$. Then, by the Cognitivist account, for any $j \in f_{S|\phi|_{i_S}} : \sigma_{S|\phi|_j}(f_{S|\phi|_j}) \subseteq \llbracket \psi \rrbracket$. By Reflexivity, $i_S \in f_{S, i_S}$. By supposition, $\llbracket \phi \rrbracket^{S, i_S} = 1$. So, $i_S \in f_{S, i_S} \cap \llbracket \phi \rrbracket$. Whence it follows that $i_S \in f_{S|\phi|_{i_S}}$. But then $\sigma_{S|\phi|_{i_S}}(f_{S|\phi|_{i_S}}) \subseteq \llbracket \psi \rrbracket$. So, by the semantics for O , $\llbracket O\psi \rrbracket^{S, i_S} = 1$. \square

What this means, informally, is that whenever ϕ holds at the world locating the agent (or context, as the case may be), and $[if \phi][O(\psi)]$ holds at the agent's state-representation, the agent *really ought* to see to it that ψ . And that, too, is bootstrapping.

B.4.1.2 Against Bootstrapping

Contrary to the Variably Strict and Context-Shifty Cognitivist analyses, bootstrapping does *not* seem to be licensed in case (B.1). Given that A has high negative relative goodness, conditional on the agent not wanting A , it seems to me that the agent—who, while aware of all of these things is also, recall, basically unsure whether she does in fact want A —actually ought *not* to see to it that A . Even if we grant that there is a reading of conditional deontic claims like (B.1b) on which they license bootstrapping, there is also a reading—a clear, salient reading—on which they do not. That, I wish to suggest, is the actual *conditional preference reading*, as Weirich would suggest we understand it: the reading according to which A is preferred, on the supposition that *want*(A).

If there is such a reading available for (B.1b) (as, I maintain, there clearly is), we would not actually expect that it would license bootstrapping. When, at some possibility i , something (say, ψ) is preferred on *supposition* that ϕ , this does not generally imply anything about what is preferred at i , even if ϕ is true at i . For the i -relevant information may be compatible with both ϕ and $\neg\phi$. To say that, at i , ψ is preferred on supposition that ϕ , is only to say that ψ is preferred at i when the hypothetical context that is induced by supposing ϕ and the actual, i -relevant information *agree on* ϕ . Supposing, then, that...

- The deontic conditional $[if \text{want}(A)][O(A)]$ expresses conditional preferability
- Conditional preferability is to be understood in this suppositional sense

...we will, then, expect such a conditional to license a merely ersatz kind of bootstrapping.

ERSATZ BOOTSTRAPPING

Whenever $[if \text{want}(A)][O(A)]$ holds at a state, and that state *accepts* or *comes to accept* *want*(A), the agent ought to see to it that A .

The semantic analogues of ersatz bootstrapping would be the following:

(B.22) If $S \Vdash [if \phi][O(\psi)]$ and $S \Vdash \phi$, then $\llbracket O(\psi) \rrbracket^{S, i_S} = 1$ and $S \Vdash O(\psi)$.

(B.23) If $S \Vdash [if \phi][O(\psi)]$, then $\llbracket O(\psi) \rrbracket^{S|\phi|, i_S} = 1$ and $S|\phi| \Vdash O(\psi)$

More generally, if we are convinced that indicatives whose consequents contain supposition-sensitive syntactic items do not generally allow the inference of their consequents when the relevant supposition is not in force, we will want such indicatives to validate *only* an ersatz form of modus ponens.¹⁴

GENUINE MODUS PONENS (NO!)

$$\llbracket [if \phi][\psi] \rrbracket^{S,i} = \llbracket \phi \rrbracket^{S,i} = 1 \text{ implies } \llbracket \psi \rrbracket^{S,i} = 1$$

ERSATZ MODUS PONENS (YES!)

- i. If $S \Vdash [if \phi][\psi]$ and $S \Vdash \phi$, then $S \Vdash \psi$
- ii. $\llbracket [if \phi][\psi] \rrbracket^{S,i} = \llbracket \phi \rrbracket^{S|\phi,i} = 1 \text{ implies } \llbracket \psi \rrbracket^{S|\phi,i} = 1$

The point that bears emphasis is this. A state S 's having some property P when a supposition is in force (when whether S has P is, in part, a function of available information) ought to imply that S has P when the content of that supposition is accepted at S , *not* when the content of that supposition happens just to be the case. Given their embrace of bootstrapping, this point, together with a suppositional analysis of (B.1b), makes trouble for Variable Strictness and Context-Shifty Cognitivism. Fans of these accounts may react to this observation in one of two ways:

- Allow that conditional preferability is to be understood as Weirich suggests (i.e., suppositionally), but deny there is a reading of $[if \text{ want}(A)][O(A)]$ on which it expresses preferability of A conditional on wanting A .
- Allow that there is a reading of $[if \text{ want}(A)][O(A)]$ on which it expresses preferability of A conditional on wanting A , but deny conditional preferability is to be understood as preferability under a supposition.

Both reactions are, I'll take it, *prima facie* unpalatable. Variable Strictness and Context-Shifty Cognitivism were pursued as attempts to implement the notion of conditional preference suggested by the discussion of conditional utility in Section B.2.1, given the intuitive equivalence of a conditional like $[if \text{ want}(A)][O(A)]$ with a statement of conditional preferability like (B.1b). Taking either of the bulleted routes amounts to an admission that any such attempt in either vein is doomed to failure. That may, in the end, be tolerable. We may discover that appearances are misleading. For present purposes, however, I take it as given that things are more or less as they appear. We want an analysis of deontic conditionals that does justice to (i) the intuitive equivalence of deontic conditionals and corresponding

14. It might seem unacceptably radical to deny modus ponens for even a *limited* class of natural language indicatives. See Kolodny & MacFarlane (2010); Charlow (2011b) for some arguments that it's not. Notice that the Wide-Scoper must also deny modus ponens for the same class of indicatives. The Wide-Scoper renders $[if \text{ want}(A)][O(A)]$ as $O(\text{want}(A) \supset A)$; but $O(\text{want}(A) \supset A)$ and $\text{want}(A)$ do not validly entail $O(A)$. Indeed, in the ethical literature on normative or rational requirements, this is generally regarded as a *virtue* of the Wide-Scope representation of indicatives describing normative requirements (e.g. Broome 1999, 2001).

statements of conditional preferability and (ii) a suppositional understanding of conditional preferability. Neither Variable Strictness nor Context-Shifted Cognitivism is up to the task.

Here, briefly, is a related way of making the point. Both the Variably Strict and Context-Shifted Cognitivist treatments of (B.1b) seem to predict the following claims *inconsistent*, when you want *A*, even when you have no idea whether you actually want *A*. (Small caps are used to indicate intended focal stress.)

- (B.24) a. IF you want *A*, THEN you should see to it that *A*.
 b. BUT, all things considered, you shouldn't see to it that *A*.

But there is, I insist, a suppositional reading of (B.24a) on which it is at least *possibly* true, even when (B.24b) is true and you happen to want *A*. What you should do, given that you want *A*, concerns what is preferable from the standpoint of the supposition that you do want *A*. What you should actually do concerns what is preferable from the standpoint of your actual information.

Again, this is not to deny that there is a reading of (B.24a) on which it is incompatible with (B.24b), when it happens that you do want *A*. It is only to insist that *there is* what we might call a “deliberative” *should*: one that tracks preferability from the standpoint of whatever information is relevant to a given episode of practical deliberation, and on which (B.24a) and (B.24b) are compatible. The same phenomenon is what plausibly licenses the following sequences of deontic conditionals and bare deontics.

- (B.25) a. IF the drug prevents infection, THEN you should take it.
 b. BUT, since no clinical studies have shown this, it's better that you don't.
 (B.26) a. IF it's sunny, THEN, of course, you should leave the umbrella.
 b. BUT, since we won't know until we're outside, you should bring it.

Deontic conditionals employing the deliberative *should*, when their antecedents are entertained as relevant to the consequent, have a Weirich-style, suppositional semantics. Neither the Variably Strict nor Context-Shifted Cognitivist analyses can accommodate this. So they cannot be correct accounts of this class of deontic conditionals.

B.4.1.3 Expressivism to the Rescue

Expressivism easily (indeed, trivially) accommodates suppositional readings of deontic conditionals. According to Expressivists, evaluating $[if\ want(A)][O(A)]$ and its intuitive conditional-preference equivalent (B.1b) at a state *S* involves checking whether *A* is preferred at $S|want(A)$: $S \Vdash [if\ want(A)][O(A)]$ just if $\sigma_{S|want(A)}(f_{S|want(A)}) \subseteq \llbracket A \rrbracket$, just if the possibilities that are preferable, under the supposition that *A* is wanted, are *A*-possibilities.

It's apparent that Expressivism doesn't validate genuine bootstrapping: given that $S \Vdash [if\ want(A)][O(A)]$, whether *A* actually happens to be wanted, at the possibility

locating S , is, in a sense, *irrelevant* to whether A -possibilities are preferable. There is, of course, a sense in which whether A actually happens to be wanted is relevant to whether A -possibilities are preferable. To clarify these senses, let us distinguish *factual* and *informational* relevance.

- ϕ is **factually relevant** to the evaluation of $[if \phi][\psi]$ at S just if $\llbracket \phi \rrbracket^{S, i_s} = 1$ implies that $\llbracket \psi \rrbracket^{S, i_s} = 1$ (when $\llbracket [if \phi][\psi] \rrbracket^{S, i_s} = 1$ or $S \Vdash [if \phi][\psi]$)
- ϕ is **informationally relevant** to the evaluation of $[if \phi][\psi]$ at S just if ψ holds at S , given that S accepts ϕ

Both Variable Strictness and Context-Shifty Cognitivism construe the antecedents of conditionals like $[if \textit{want}(A)][O(A)]$ as factually relevant; hence they tolerate genuine bootstrapping.

Context-Shifty Expressivism construes such antecedents as informationally relevant; hence it (correctly) tolerates merely ersatz bootstrapping. To illustrate this, we'll prove (B.22) and (B.23).

Proof. Suppose $S \Vdash [if \phi][O(\psi)]$ and $S \Vdash \phi$. So, $S|\phi \Vdash O(\psi)$. Then, by Expressivism, $\sigma_{S|\phi}(f_{S|\phi}) \subseteq \llbracket \psi \rrbracket$. But since $S \Vdash \phi$, $S|\phi = S$, by item (B.6). So, $\sigma_S(f_S) \subseteq \llbracket \psi \rrbracket$. Hence, by Expressivism, $S \Vdash O(\psi)$. \square

This extends to a proof that, on these assumptions, $\llbracket O(\psi) \rrbracket^{S, i_s} = 1$, provided the Expressivist endorses a corresponding “minimalist” or “acceptance-conditional” theory of truth for deontic sentences, on which a deontic sentence is “true” relative to a state just if it’s accepted by that state (cf. esp. Blackburn 1984; Gibbard 2003):

$$(B.27) \quad \llbracket O(\phi) \rrbracket^{S, i_s} = 1 \text{ iff } S \Vdash O(\phi) \text{ iff } \sigma_S(f_S) \subseteq \llbracket \phi \rrbracket$$

The proof of (B.23) is trivial:

Proof. Suppose $S \Vdash [if \phi][O(\psi)]$. Then, by definition, $S|\phi \Vdash O(\psi)$. By the acceptance-conditional truth-condition, $\llbracket O(\psi) \rrbracket^{S|\phi, i_s} = 1$. \square

B.4.1.4 Expressivism, Amended

Expressivism has a minor deficiency which the Variably Strict and Context-Shifty Cognitivist accounts, rather fortuitously, lack. The latter accounts reliably predict the *truth* of $[if \textit{want}(A)][O(A)]$ and its conditional-preference equivalent (B.1b).

- On the Variably Strict account, $[if \textit{want}(A)][O(A)]$ is true just if, if i is any sufficiently close possibility where A is wanted, the best (in view of what’s wanted) possibilities at i are A -possibilities. But, of course, if A is wanted at i , then the best (in view of what’s wanted) possibilities at i will typically be A possibilities. (There are exceptions to this, hence the “typically” qualification, which we will discuss in

Section B.4.2. Typically, though, $[if\ want(A)][O(A)]$ is true, and the Variably Strict account *predicts* this.)

- On Context-Shifty Cognitivism, $[if\ want(A)][O(A)]$ holds at S just if, if i is any S -relevant possibility where A is wanted, then the best (in view of what's wanted) possibilities at i where A is also wanted are A -possibilities. Again, if A is wanted at i , then the best (in view of what's wanted) possibilities at i in which A is wanted will typically be A -possibilities.

Expressivism cannot mimic these predictions directly. Why should updating S on $want(A)$ make it the case that the favored possibilities are A -possibilities—why should we expect that $S|want(A)|$ favors A ? Recall that, for the Expressivist, $O(A)$ expresses a property of the selection function, rather than information. Supposing that the selection function at S favors $\neg A$, why should the selection function at $S|want(A)|$ favor A , given that updating with $want(A)$ simply temporarily eliminates $\neg want(A)$ -possibilities from consideration?

There are both short and long answers to this question. I will sketch the short one here. The short answer is that updating on $want(A)$ induces a change in the selection function: it makes it the case that A is wanted.¹⁵ We can model this as a kind of general *coherence constraint* on the relationship between one's information about the selection function and the properties of the selection function. Let $\sigma_{F,S}$ be a selection function characterized by a set of propositions $\llbracket F \rrbracket^S$ —the set of propositions satisfying F at S (equivalently, a function mapping a proposition to true just if that proposition is F at S). So, for instance, $\sigma_{want,S}$ is the selection function characterized by $\llbracket want \rrbracket^S = \{p : p \text{ is wanted at } S\}$. The constraint I have in mind is this:

$$(B.28) \quad \llbracket A \rrbracket \in \llbracket F \rrbracket^{S|F(A)|}$$

That is to say, if there is a selection function $\sigma_{F,S}$ (whose job is to select the best possibilities, in view of what is F at S), and if S is updated with the information that A is F , then the selection function $\sigma_{F,S|F(A)|}$ comes to regard A as F . A state is coherent only if, when it thinks A desirable, the selection function regards A as desirable. So updating with $want(A)$ makes it the case that the selection function regards A as wanted. So updating S on $want(A)$ will typically make it the case that the favored possibilities are A -possibilities.

That is our short answer. The longer answer—which will amount to an extension of Expressivism to judgments or descriptions of what is wanted (so that thinking A desired is just to desire A , and determining whether you think A desired is a matter of determining whether you desire A)—is given in Appendix B.

15. The intended status of this claim is as an idealization, on a par with the idealization that a state accepting each member of a set Γ accepts every classical consequence of any $\phi_1, \dots, \phi_n \in \Gamma$.

B.4.2 Defeasibility and Antecedent-Strengthening

B.4.2.1 Sobel Sequences

A familiar argument against accounts of the indicative conditional that validate antecedent-strengthening appeals to so-called Sobel Sequences, like (B.29).

- (B.29) a. If Pedro goes to the parade, he'll see Nomar.
b. But, if Pedro goes to the parade and gets stuck behind David, he won't see Nomar.

Any analysis of the indicative conditional as a generalized universal quantifier—that is to say, any analysis which gives $[if \phi][\psi]$ truth- or acceptability-conditions along the lines of *all ϕ -possibilities meeting some condition C meet some condition D*—will validate antecedent-strengthening (since all ϕ -possibilities meeting C meet D only when all $(\phi \wedge \chi)$ -possibilities meeting C meet D). Any such account will counterintuitively predict the Sobel Sequence (B.29) inconsistent. And both the Variably Strict and Context-Shifty Cognitivist accounts give indicatives truth- or acceptability-conditions along these lines.

A popular explanation of this phenomenon, from fans of generalized universal quantifier analyses of $[if \cdot][\cdot]$, is that indicative antecedents may *raise*, *make salient*, or *make relevant* possibilities that were not raised/salient/relevant at the original context of utterance (see von Fintel 2001; Gillies 2007; Moss to appear). So the sequence in (B.29) is fine because the possibilities where Pedro goes to the parade and is stuck behind David are not initially relevant; an utterance of (B.29b) makes them relevant. The initially relevant possibilities where Pedro goes to the parade (and which meet some condition C) may all meet condition D, since none of these are possibilities where Pedro goes to the parade and is stuck behind David. But the possibilities that are relevant *after* an utterance of (B.29b) will include some where Pedro goes to the parade and is stuck behind David. And it is easy to see how some (indeed, all) of *these* possibilities might fail to meet D.

This account's plausibility is enhanced by the badness of so-called Reverse Sobel Sequences, like (B.30) (see esp. von Fintel 2001; Gillies 2007).

- (B.30) a. If Pedro goes to the parade and gets stuck behind David, he won't see Nomar.
b. #But, if Pedro goes to the parade, he'll see Nomar.

On this account, (B.30a) raises the possibility that Pedro goes to the parade and is stuck behind David, and asserts that, in any relevant such possibility, Pedro does not see Nomar. So (B.30a) makes relevant possibilities in which Pedro goes to the parade and is stuck behind David (and asserts that in none of these does Pedro see Nomar). Since (B.30b) asserts that all relevant possibilities where Pedro goes to the parade are possibilities where he sees Nomar, (B.30b) asserts something that cannot be true, given a prior utterance of

(B.30a). Similar reasoning will account for the badness of the sequence in (B.31).

- (B.31) a. Pedro may go to the parade and get stuck behind David.
b. #But, if Pedro goes to the parade, he'll see Nomar.

B.4.2.2 Anankastics

Not surprisingly, we notice many apparent failures of antecedent-strengthening with indicative conditionals whose consequents are deontic.

- (B.32) a. If you want to put sugar in your soup, you should see a doctor!
b. But, of course, if you want to put sugar in your soup because you want extra fuel for the marathon, you shouldn't see a doctor.
- (B.33) a. If you want to go to Harlem, you should take the A-train.
b. But, of course, if you want to go to Harlem without using the MTA, you should take a cab.
- (B.34) a. If you want to eat out, you should go to Fette Sau.
b. But, of course, if you want to eat out and dislike BBQ, you should go to Egg.

Generalized universal quantifier accounts are, of course, capable of explaining these failures, by appeal to the same resources they use to explain Sobel Sequences like (B.29). But, although this works well for examples like (B.32)—which says, roughly, that any possibility where you want sugared soup is one where the best thing for you to do is see a doctor—it leaves much to be desired for so-called *anankastic* conditionals like (B.33-B.34).¹⁶

A conditional is anankastic, in my sense, if its antecedent specifies a goal for some agent, and its consequent expresses, as a matter of its *semantics*, a view about which courses of action best realize this goal. Not surprisingly, as other authors have noticed, such conditionals, are readily paraphrased as (and intuitively equivalent to) “teleological modal[s] with a purpose *to*-clause,” as in (B.35) and (B.36) (von Fintel & Iatridou 2005; Sæbø 2001). Non-anankastic interpretations can be identified by their resistance to such paraphrase, as in (B.37).

- (B.35) a. To go to Harlem, you should take the A-train.
b. But, of course, to get there without using the MTA, you should take a cab.
- (B.36) a. To eat out, you should go to Fette Sau.
b. But, of course, to eat out without eating BBQ, you should go to Egg.

16. The term, as well as example (B.32), is due to Hare (1971). For an accessible recent discussion of anankastics, see von Fintel & Iatridou (2005). One quibble about their treatment: von Fintel and Iatridou make the mistake of conflating genuine anankastics (which, as a matter of their semantics, *do not* validate antecedent-strengthening) with descriptions of what a given body of imperatives (e.g. the law) requires, given certain suppositions about its properties (which, as a matter of their semantics, *do* validate antecedent-strengthening).

- (B.37) a.??To put sugar in your soup, you should see a doctor.
 b.??But, of course, to put sugar in your soup so as to have extra fuel for the marathon, you shouldn't see a doctor.

Clearly, the salient reading for the conditionals in which we're interested—conditionals of the form $[if\ want(A)][O(A)]$ —is anankastic: *if you want A, you should see to it that A* is paraphrasable as *to realize A, you should see to it that A*, and vice versa.

I do not want to take a stand here on the question of whether Sobel Sequences are to be given a semantic account (as in Gillies 2007) or a pragmatic account (as in Moss to appear). What I want to argue is that the failure of antecedent-strengthening for anankastic conditionals:

- Is an evidently semantic phenomenon, deserving a semantic explanation
- Cannot admit of the same explanation as Sobel Sequences (regardless of whether that explanation is semantic or pragmatic in character)

These points are made in service of a larger point: Variable Strictness and Context-Shifty Cognitivism validate antecedent-strengthening, but the available explanation of the failure of antecedent-strengthening for anankastics (the explanation of Sobel Sequences) is inadequate; Sobel Sequences and the failure of left-downward-monotonicity for anankastics are different phenomena, with different explanations. A semantics for anankastics which accounts for their peculiar failure to license antecedent-strengthening is to be preferred. As we'll see, the Expressivist account is just such an account.

I am going to provisionally assume that anankastic conditionals and their teleological modal paraphrases are semantically equivalent. So we can approach the semantics of anankastics indirectly, by exploring what the teleological modal paraphrases *say*. What they say is that the best way of realizing the goal in the subordinated purpose-clause is to perform the action described in the matrix clause. Of course, a course of action A may be the best way to realize a desire D , without it being the case that A is the best way to realize a *stronger* desire $(D \wedge D')$. Desire-implementation thus has a decidedly *non-monotonic* flavor: what is best for a given end may fail to be best for a set of ends properly including that end. Since teleological modals clearly express best-ness, relative to a given set of ends, we would expect the following:

$$(B.38) \quad [to\ \phi][O(\psi)] \not\equiv [to\ \phi \wedge \chi][O(\psi)]$$

This failure is not to be explained by the contextual salience of the relevant goals (as a function of the content of the *to*-clause). The *to*-clause does not affect interpretation by making possibilities *relevant*. Rather, it supplies the *argument* for the selection function with

respect to which the modal O in the matrix clause is interpreted.¹⁷

$$(B.39) \quad \llbracket [to \phi][O(\psi)] \rrbracket^{S,i} = 1 \text{ iff } \sigma_{\{\phi\},S}(f_S) \subseteq \llbracket \psi \rrbracket$$

That's to say, $[to \phi][O(\psi)]$ is true at S just if the selection function characterized by the goal ϕ (and perhaps some body of secondary considerations; cf. von Fintel & Iatridou 2008) recognizes only ψ possibilities as sufficiently good. This sort of semantics readily predicts the fact in (B.38): the preferred possibilities relative to a set of goals $\{\phi\}$ may cease to be preferred relative to a stronger set of goals $\{\phi \wedge \chi\}$. Possibilities that are best relative to $\{\phi\}$ may fail to be best relative to $\{\phi \wedge \chi\}$; hence, $[to \phi][O(\psi)]$ does not imply $[to \phi \wedge \chi][O(\psi)]$.¹⁸ This, I submit, is what explains the acceptability of the sequences (B.35) and (B.36). Likewise, given the assumed equivalence between anankastics and their teleological modal paraphrases, for anankastics.

This argument has considerable intuitive plausibility. Here is some linguistic data that bolsters it. Were failures of *to*-clause-strengthening for teleological modals explained as a reflex of the contextual salience of relevant goals, we would predict that reversing the order of the modals should wreck the sequence's acceptability (as reversing Sobel Sequences wrecks their acceptability). This prediction is not borne out by the data: both (B.40) and (B.41) sound fine. Similarly for the anankastics which they paraphrase, (B.42) and (B.43). Making a strengthened goal salient, as in (B.44) and (B.45) does not affect this.

- (B.40) a. To go to Harlem without using the MTA, you should take a cab.
b. To go to Harlem, you should take the A-train.

- (B.41) a. To eat out without eating BBQ, you should go to Egg.
b. To eat out, you should go to Fette Sau.

- (B.42) a. If you want to go to Harlem without using the MTA, you should take a cab.
b. If you want to go to Harlem, you should take the A-train.

- (B.43) a. If you want to eat out without eating BBQ, you should go to Egg.
b. If you want to eat out, you should go to Fette Sau.

- (B.44) a. I don't know whether you wish to go to Harlem without using the MTA.
b. Still, to go to Harlem, you should take the A-train.

17. Cf. von Fintel & Iatridou (2008). For simplicity, I obscure the semantic distinction between weak deontic modalities (*should*) and strong deontic modalities (*must*). These truth-conditions are actually more appropriate for strong deontic modals; weak deontic modals have a more complicated semantics, making use of ordered sequences of selection functions (as argued in von Fintel & Iatridou 2008). Introducing this complexity would serve no purpose for my discussion here.

18. On various pictures about the relationship between sets of goals and selection functions, this is easily proven. For instance, when G is a set of goals, Kratzer (1981) holds that a possibility i is at least G -good as j just if everything in G satisfied by j is satisfied by i . The selection function σ_G selects the possibilities that are minimal relative to this ordering; these are the possibilities that are good enough relative to G . Clearly a possibility i that is good enough relative to G can fail to be good enough relative to a set G^+ such that $G \subset G^+$ (since some other j equally G -good to i may satisfy some $p \notin G \in G^+$ that i fails to satisfy, so that i ceases to be minimal relative to the ordering characterized by G^+).

- (B.45) a. I don't know whether you wish to go to Harlem without using the MTA.
 b. Still, if you want to go to Harlem, you should take the A-train.

The acceptability of sequences (B.40) and (B.41) is explained by the semantics' failure to validate *to*-clause-strengthening. The acceptability of sequences (B.42) and (B.43) is explained by the fact that (B.40) and (B.41) are their intuitive semantic equivalents. And, finally, the acceptability of (B.44) and (B.45) is explained, in part,¹⁹ by the fact that anankastics and their paraphrases are not sensitive to salient epistemic possibilities in the way that non-anankastic conditionals are typically sensitive.

B.4.2.3 Expressivism to the Rescue

Where does this leave us? We have seen that Variable Strictness and Context-Shifty Cognitivism each (i) validate antecedent-strengthening for anankastics, (ii) have no obvious resources for explaining intuitive failures of antecedent-strengthening (due to differences in the behavior of sequences of anankastics and classic Sobel Sequences). Does Context-Shifty Expressivism do better? It does. Whether a particular conditional licenses antecedent-strengthening depends on the acceptance-conditions of its consequent.

Indicative conditionals with *information-bearing* consequents—consequents that are accepted just if the information entails them, that express a property of the information-state, rather than the selection function—will license antecedent-strengthening. If ψ is information-bearing, $S \Vdash [\textit{if } \phi][\psi]$ iff $S|\phi| \Vdash \psi$ iff $f_{S|\phi|} \subseteq \llbracket \psi \rrbracket$ iff, informally, updating with ϕ yields a state in which the information expressed by ψ is borne by the information-tracking parameter of the state. But if every possibility in $f_{S|\phi|}$ is a ψ -possibility, then every possibility in $f_{S|\phi \wedge \chi|} = f_{S|\phi||\chi|}$ is too.²⁰ So $S \Vdash [\textit{if } \phi][\psi]$ implies $S \Vdash [\textit{if } \phi \wedge \chi][\psi]$. With respect to indicatives with information-bearing consequents, this leaves Expressivism in the same position as its competitors.

However, indicative conditionals with *deontic* consequents—which, according to Expressivists, are accepted just if the selection function *has* a certain property—do not validate antecedent-strengthening. According to Expressivists, $S \Vdash [\textit{if } \phi][\psi]$ iff $S|\phi| \Vdash \psi$ iff $\sigma_{S|\phi|}(f_{S|\phi|}) \subseteq \llbracket \psi \rrbracket$. The fact that $\sigma_{S|\phi|}(f_{S|\phi|}) \subseteq \llbracket \psi \rrbracket$ (i.e., that the best possibilities, on update with ϕ , are ψ -possibilities) does *not*, however, imply that $\sigma_{S|\phi||\chi|}(f_{S|\phi||\chi|}) \subseteq \llbracket \psi \rrbracket$ (i.e., that the best possibilities, on update with ϕ and ψ , are ψ -possibilities). By the coherence constraint in (B.28), update with χ may *introduce new goals* into the set of imperatives relative to which the selection function is characterized. If, for instance, $\chi = \textit{want}(A)$, then the selection function $\sigma_{\textit{want}, S|\textit{want}(A)|}$ (i.e., the selection function for S updated with χ) comes to regard A

19. A more thorough diagnosis of these phenomena would be interesting, but what is important for my purposes is just that they have a different source than Sobel Sequence phenomena.

20. This follows from a more general fact about dynamic frameworks, which we have not bothered to prove (but whose proof would be trivial, given an explicit characterization of the object language and dynamic interpretation function $|\cdot|$). Namely, information-growth is *monotonic*: $\forall S, \phi : f_{S|\phi|} \subseteq f_S$.

as wanted, in view of the following:

$$(B.46) \quad \llbracket A \rrbracket \in \llbracket want \rrbracket^{S|want(A)}$$

So ψ -possibilities may be sufficiently good relative to $\sigma_{want,S}$ without being sufficiently good relative to $\sigma_{want,S|want(A)}$ (if, for instance, ψ is incompatible with the goal that updating with $want(A)$ introduces).²¹

The long and short of it is this: update may—and, in cases involving interpretation of anankastic conditionals, typically does—*enlarge* the set of goals relative to which the relevant selection function is characterized. Courses of action that are good enough relative to a comparatively *modest* goal-set may cease to be good enough relative to a comparatively *ambitious* goal-set. Since anankastic antecedents function to *strengthen* the relevant goal-set, we would not generally expect an anankastic with a comparatively strong antecedent to follow from an anankastic with a comparatively weak antecedent.²²

B.5 Clarifying Expressivism: Imperative versus Deontic Conditionals

Having argued for Context-Shifty Expressivism’s superiority over its competitors, I want to spend some time interpreting Context-Shifty Expressivism. In particular, I want to clarify the extent to which the envisaged account is really a full-fledged form of Expressivism. Contrary, perhaps, to appearances, I will suggest that the account is compatible with a variety of different views in meta-ethics. Its success, then, is not to be construed as any sort of positive argument for Expressivism in meta-ethics. (It is, nevertheless, *inspired by* an Expressivistic treatment of normative judgment, and, as such, Expressivists may—indeed, should be eager to—make free use of it. So, although meta-ethically neutral, the account is nevertheless of meta-ethical interest.)

I will go on to argue that, regardless of one’s meta-ethical commitments, the discussion of the prior section encourages us to embrace the Expressivist’s account of the *cognitive force* of these deontic conditionals. (Such acceptance is, I suggest, meta-ethically neutral.) Entertaining a deontic conditional, I argue, is a matter of *consulting one’s preferences*, on update with its antecedent. This, importantly, is distinct from consulting one’s beliefs about one’s preferences, on update with the antecedent. Correspondingly, deciding that a deontic conditional is the case is a matter of ending up in a cognitive state characterized by

21. The Context-Shifty Cognitivist can actually mimic this, if she agrees that (B.28) is a plausible constraint on the relationship between information and goals. There is, moreover, nothing in her view that would prevent her from saying this. (This will not help with her Bootstrapping problem, though.)

22. What, then, is the difference between anankastics and non-anankastic conditionals with goal-introducing antecedents like (B.32)? It is not, I think, a difference in the way in which their antecedents are entertained. It is, rather, a difference in the interpretation of the modal in their consequents. Anankastic consequents contain a goal-oriented modal whose goal is *supplied by* their antecedents. The consequents of non-anankastic conditionals with goal-introducing antecedents contain a modal interpreted with respect to a *different goal-set* than that on which the antecedent bears. There is some work required to turn this into a compositional account, but it would take us too far afield to pursue that here.

a specific, suppositional preference. This, I argue, is distinct both from (i) deciding *that* there is such a suppositional preference (hence distinct from a Cognitivist account of endorsement of deontic conditionals), and (ii) *coming to have* such a suppositional preference.

B.5.1 Cognitive and Non-Cognitive Semantic Methodology

Bootstrapping and the peculiar non-monotonicity of anankastic conditionals motivate a semantics for the class of deontic conditionals in which we're interested that is outside the standard philosophical paradigm. The proposed semantics comes in two, in principle independent, parts.

- A Context-Shifty account of their antecedents
- An Expressivist account of their deontic consequents

The second part is labelled “Expressivist” because it *originates in* an Expressivistic understanding of what it is for an agent to accept a statement of the form $O(\phi)$: to accept $O(\phi)$ is for the relevant selection function (plan, system of norms, ...) to have some property P , rather than for the relevant information to entail that the relevant selection function has P . I have also argued that a semantics for these conditionals characterized by Context-Shiftness for indicative antecedents and Expressivism for deontic consequents *strictly improves* on quantificational approaches to their semantics (of which we have considered two—Variably Strictness and Context-Shifty Cognitivism).

This may, however, portray a misleading picture of the relationship between semantic and meta-normative questions, which I will try to correct here. In fact, it is difficult either to read off a specific meta-normative view from the features of a specific semantic proposal, or to cash out a specific meta-normative view in a semantic proposal for the sorts of constructions we've been considering. Although inspired by meta-normative Expressivism, the semantic apparatus developed here does *not* presuppose (or otherwise require) meta-normative Expressivism. To support this, it suffices to outline a non-Expressivist meta-ethic that makes appropriate use of the proposed semantic apparatus.

A meta-normative Expressivist will give the semantic apparatus a *cognitive gloss*:

- A state $S = \langle f_S, \sigma_S \rangle$ is representative of an agent's cognitive state. The relation \Vdash represents *acceptance* by the agent.
- World and will are kept apart. f_S is semantic kin to the agent's *representational* states of mind (her picture of the world), while σ_S is semantic kin to her non-cognitive *motivational* states. An agent's plan or norm-system Π straightforwardly characterizes a selection function, namely, a function which selects from a set of possibilities those possibilities that are best, with respect to Π .²³

23. Gibbard (1990) makes extensive use of Norm-Systems, while Gibbard (2003) makes extensive use of Hyperplans; in both works, Gibbard thinks of a deontic sentence $O(\phi)$ as holding or failing to hold relative to a Norm-System or Hyperplan just if ϕ is *required* by the relevant Norm-System or Hyperplan. My “Meaning for

But this kind of cognitive gloss is strictly optional. For instance, a selection function σ_N appropriate to a kind of normativity N might be characterized relative to a body of considerations that are given by the *facts* about what is N (or, better, what is N , *prima facie* or *pro tanto*). Outlining a distinctive semantic role for such bodies of considerations in determining the N -facts is meta-ethically inert. Descriptions of N -facts are like descriptions of modal facts, in the sense that whether they hold at a possibility i is not (semantically, anyway) a matter of the intrinsic features of i . It is, rather, a matter of what holds at possibilities that σ_N recognizes as deontically *accessible*, given the facts about what is N , *prima facie*. Similarly, descriptions of modal facts hold when a set of worlds, determined by an independent semantic parameter (typically an accessibility relation), meets a specified condition. Nevertheless, for σ_N to have a property P is for reality to be a certain way, just as for a modal fact to be true is for reality to be a certain way. Accepting (or expressing) that something is N thus amounts to adjusting (or expressing) one's beliefs about reality—one's information. The semantic relation \Vdash , then, will represent, not acceptance, but something more akin to the familiar notion of satisfaction in a model.²⁴

In developing our analysis of deontic conditionals, we have, of course, been working within a semantic tradition—the Dynamic (Update) Semantic tradition—that makes liberal use of the sort of cognitive talk that the meta-normative Expressivist will find congenial. But it is important to realize that this talk represents a philosophically loaded gloss on semantic apparatus that is itself philosophically neutral.

B.5.2 The Cognitive Profile of Deontic Conditionals

All that said, I take it that it is worth exploring how to pursue meta-normative Expressivism as a theory of the kinds of deontic conditionals we've been considering. Why is that? The project has intrinsic interest, of course, for meta-normative Expressivists who are serious about building a working theory of meaning in natural language. But that is not all. Expressivism offers, I'll argue, an appealing treatment of the *cognitive profile* and *force* of deontic conditionals like $[if\ want(A)][O(B)]$ (and self-directed anankastics more generally)—a treatment we would be liable to miss if not for Expressivism's insights.

Specifically, Expressivism offers an appealing treatment of the following questions:

- How an agent decides whether $[if\ want(A)][O(B)]$
- What happens to someone's mind when they decide that $[if\ want(A)][O(B)]$ —i.e., when they come to accept $[if\ want(A)][O(B)]$. How does an agent *update* on such conditionals?

I do not, however, mean to imply that these answers *require* meta-normative Expressivism.

Expressivists" (Appendix A) discusses one Expressivist-friendly way of conceiving the relationship between selection functions, deontic accessibility relations, and plans.

24. Dreier (1999, 2009) makes similar points.

Here is a helpful analogy. One can be (as I am) a fan of plain-vanilla, truth-conditional semantics for belief-reports, while still thinking:

- One typically decides whether one believes p by examining one's information *about* p (rather than one's information about one's beliefs)
- Deciding that one does believe p is typically a matter of ending up in an information state that entails p (rather than an information state that entails that one believes p)

A standard, truth-conditional semantics for belief-reports—on which *I believe p* expresses the proposition whose truth-conditions are that I believe p —would lead us to anticipate the parenthetical answers. Expressivism for belief-reports notices that avowing a belief that p typically expresses belief that p , rather than (or rather than merely) belief that one believes that p , and constructs a theory of meaning accordingly.²⁵ Truth-conditional semantics, however, can also utilize this data about use: perhaps it is a conventional (but non-semantic) fact about first-personal belief-reports that they express the belief that they semantically report (as, perhaps, *but* conventionally implicates something more than *and*, while semantically expressing just \wedge). Only if use-facts are somehow identified with (or required to supervene on) semantic facts, as the Expressivist suggests,²⁶ does truth-conditional semantics for belief-reports run into trouble in utilizing this data about use. How use bears on semantic theorizing is an interesting issue, but one that is strictly independent of one's understanding of the cognitive profile of belief-reports.

It would be exceedingly natural, given a truth-conditional semantics for deontic language, and a quantificational semantics for conditionals, to think that:

- Deciding whether $[if\ want(A)][O(B)]$ is deciding whether $O(B)$ holds—that B is preferred—throughout a range of relevant $want(A)$ -possibilities.
- Deciding that $[if\ want(A)][O(B)]$ is a matter of *adjusting your information* so that $O(B)$ holds throughout the range of relevant $want(A)$ -possibilities.

But, importantly, the discussion of Section B.4 would seem to suggest that this idea, however natural, is *unworkable*. For reasons we saw in Section B.4.1, from the fact that the actual world is a $want(A)$ -possibility, it does not follow that, from the standpoint of the cognitive state of an agent who accepts $[if\ want(A)][O(B)]$, she should see to it that B . Nor, for reasons we saw in Section B.4.2, does it follow that, from the standpoint of the same agent, if she wants A and C , she should see to it that B . But treating the cognitive force of $[if\ want(A)][O(B)]$ as learning that B is preferred throughout a range of relevant $want(A)$ -possibilities will wrongly predict both of these things to be the case.

In the remainder of this section, I want to further explore a different intuition, due in its essentials (but not, therefore, necessarily proprietary) to the meta-normative Expressivist.

25. Something like this view is endorsed by Gillies (2001).

26. For discussion of the Expressivistic understanding of the relationship between meaning and use, and how best to understand it, see my "Meaning for Expressivists" (Appendix A).

- An agent decides whether $[if\ want(A)][O(B)]$ by supposing $want(A)$ and seeing, under this supposition, whether B is best
- Deciding that $[if\ want(A)][O(B)]$ is a matter of ending up in a cognitive state that prefers B , on supposition that $want(A)$

B.5.3 The Cognitive Profile of Imperative Conditionals

Implementing these intuitions—the second, in particular—turns out to present a challenge for Expressivism. For it is a natural idea—one, indeed, that I have defended in other work (Charlow 2010a) (and earlier discussion in this thesis)—that conditional *imperatives* like (B.47), which I will represent schematically as sentences of the form $[if\ want(A)][!B]$, are interpreted as proposals to establish suppositional or hypothetical preferences—preferences holding hypothetically, under possible future incrementations of one’s information.

(B.47) If you want an apple, take one.

On this sort of picture, deciding to accept a conditional imperative is a matter of adopting a suppositional or hypothetical preference. Accepting an unconditional imperative $!B$ is a matter of adopting an actual preference (i.e., a preference given one’s actual information).²⁷ Accepting a conditional imperative is, in other words, a matter of accepting an unconditional imperative in a suppositional cognitive state.

The reason this is problematic is that it would seem to suggest that the cognitive profiles of a deontic conditional and a corresponding conditional imperative are to be identified (at least partially so). But there is good reason to *reject* any such identification. True, both deontic and imperative conditionals can be accepted by or integrated into a cognitive state. But the phenomenon of accepting a deontic conditional is fundamentally distinct from the phenomenon of accepting an imperative conditional. On the picture that I favor, accepting a deontic conditional $[if\ want(A)][O(B)]$ is a natural upshot of:

- Entertaining whether $[if\ want(A)][O(B)]$
 \approx supposing $want(A)$ and entertaining whether $O(B)$
- Deciding that $[if\ want(A)][O(B)]$
 \approx deciding that $O(B)$, on the supposition that $want(A)$

But accepting a corresponding imperative conditional cannot result from this sort of process. One cannot entertain whether $[if\ want(A)][!B]$, since one can simply never entertain whether $!B$, regardless of one’s cognitive state. Nor can one decide that $[if\ want(A)][!B]$, since one can simply never decide that $!B$.²⁸

27. The primary proponent of this account of unconditional imperatives is Portner (2004a, 2007).

28. Generally, entertaining whether ϕ and deciding that ϕ are possible only when ϕ is declarative.

According to the Expressivist account of the cognitive profile of deontic conditionals that I have been entertaining, deciding that $[if\ want(A)][O(B)]$ is a matter of ending up in a cognitive state that prefers B , on supposition that $want(A)$. But there is good reason to think updating on the corresponding imperative conditional is also a matter of ending up in a cognitive state that prefers B , on supposition that $want(A)$. So, it doesn't seem that this can be the right story about accepting a deontic conditional. Acceptance of a deontic conditional is generally *issue-resolving*, in a way that acceptance of an imperative conditional is not—indeed, cannot be. Expressivism as a theory of deontic cognitive force seems not to admit a sufficient distinction between deontic and imperative cognitive force.

This is fair enough, at least as a description of the account as developed to this point. In the next section, we will show how to correct this fault.

B.5.4 Updates and Queries

It is not difficult, if we are willing to be impressionistic for a moment, to identify a difference between acceptance of a deontic conditional and acceptance of an imperative conditional. That difference, I will suggest, is just this: accepting a deontic conditional is a matter of *verifying* (as a result of successfully *querying* or *testing*) current properties of a cognitive state, while accepting an imperative conditional is a matter of *altering* current properties of a cognitive state.

Nevertheless, those cognitive states in which $[if\ want(A)][O(B)]$ is accepted—those that pass the query or test—are just those in which $[if\ want(A)][!B]$ is accepted: the deontic conditional checks the state for just the property that would lead update with the imperative conditional to idle. Executing a query on a state's hypothetical preferences, however, is a fundamentally different operation than adjusting the state's hypothetical preferences. Executing a query on a state's preferences can *resolve issues* concerning its properties (since it is inherently sensitive to features of the state), while altering the state's preferences—acting by way of changing the state's preferences—cannot (since it is insensitive to its features²⁹).

Here is a critical distinction. Querying a state for a property P , in the intended sense, is different from *asking whether* the state has P . Likewise, successfully querying the state for P —verifying that it has P —is distinct from deciding *that* it has P . The latter induces a change to one's state (by updating the state's information about itself), while the former simply leaves the state unchanged. The issue of whether $[if\ want(A)][O(B)]$ is resolved by querying a state, without it being the case that properties of the state are what it is at issue. One decides whether $[if\ want(A)][O(B)]$ by testing one's hypothetical preferences, but the question of whether $[if\ want(A)][O(B)]$ is not a question *about* one's hypothetical preferences. It is, rather, a non-psychological question—one about what is best, on the

29. Changes to states often incorporate preconditions or presuppositions, so that update crashes when those preconditions or presuppositions fail to be met by the state. So there is a sense in which preference-changes might be state-sensitive. But these preconditions are best understood as queries of a state: the state-sensitivity of preference-change is due to their incorporation of such queries.

supposition that *want*(A). It is a question that is decided *from* the standpoint of hypothetical preference, rather than a question *about* that standpoint.

As it happens, an account very much in this vein has been developed for unembedded epistemically modalized clauses like (B.48) (whose logical forms I will represent as $\diamond\phi$) by semanticists in the dynamic tradition (see esp. [Veltman 1996](#); [von Fintel & Gillies 2007](#)).

(B.48) It might be raining outside.

I will try briefly to motivate that account here, and then show how to extend it to deontic conditionals in the subsequent section.³⁰ On the account, the cognitive force of accepting an epistemically modalized clause $\diamond\phi$ is to verify ϕ 's compatibility with the relevant information. To utter $\diamond\phi$ is to *proffer* a query / test on the relevant information. To decide that $\diamond\phi$ is to respond to the successful performance of this test, by remaining in a state compatible with ϕ ; we will say that such a state accepts $\diamond\phi$.

(B.49) $S \Vdash \diamond\phi$ iff $f_S \cap \llbracket \phi \rrbracket \neq \emptyset$

On the other hand, to decide that $\neg\diamond\phi$ (or to reject $\diamond\phi$) is to respond to the unsuccessful performance of this test, by remaining in a state which fails to accept $\diamond\phi$ (a state incompatible with ϕ). The account's motivating intuitions are well-stated by [Veltman \(1996: 229\)](#).

[A]ll you can do when told that it might be the case that ϕ is to agree or to disagree. If ϕ is acceptable in your information state S , you must accept $\diamond\phi$. And if ϕ is not acceptable in S , neither is $\diamond\phi$. Clearly, then, sentences of the form $\diamond\phi$ provide an invitation to perform a test on S rather than to incorporate some new information in it.

Formally, tests are traditionally specified as vacuous (or else catastrophic) updates on states: integrating $\diamond\phi$ involves checking to see whether the input-state accepts $\diamond\phi$, returning that state in the event that it does, returning an absurd state otherwise.

(B.50) $S|\diamond\phi| = \begin{cases} S, & \text{if } S \Vdash \diamond\phi \\ \emptyset, & \text{otherwise} \end{cases}$

Note that Definition B.50 is *not* to be interpreted as suggesting that $\diamond\phi$ proffers an update-instruction. Given a state as input, $|\diamond\phi|$ returns information about that state—namely, about its compatibility with ϕ —as output. Executing the test yields this information without attempting to induce the state to either (i) *represent* that information as true (contra a descriptivist pragmatics for epistemic modals), or (ii) induce the state to become compatible with ϕ . If S passes the test, S accepts $\diamond\phi$, and remains unchanged. If S fails to pass the

³⁰ I should note that whether or not this is right story for epistemic modals is neither here nor there for my purposes. Our concern is deontic conditionals; the dynamic account of epistemic modals is functioning only as inspiration.

test, S rejects $\diamond\phi$ (since accepting it would lead to an absurd state) and, likewise, remains unchanged.³¹ Although a state's interpretation of $\diamond\phi$ may *set the stage* for attempts to coordinate on the property of being informationally compatible with ϕ (if, for instance, interpretation brings to light a failure in coordination), there is no proposal that aims, as such, at such coordination *and* which is also part of the conventional cognitive force of $\diamond\phi$.

B.5.5 Deontic Conditionals, Redux

This is the sort of template I want to pursue for deontic conditionals. More precisely, I want to pursue it for bare or unembedded deontic sentences, while treating deontic conditionals as a special case. The gist is this: a deontic sentence $O(\phi)$, like an epistemically modalized sentence $\diamond\phi$, will test a state for a property: the property of accepting $O(\phi)$ or $\diamond\phi$, as the case may be.

$$(B.51) \quad S|O(\phi)| = \begin{cases} S, & \text{if } S \Vdash O(\phi) \\ \emptyset, & \text{otherwise} \end{cases}$$

Acceptance of $O(\phi)$ treated Expressivistically, as in Definition (B.16), repeated here.

$$(B.52) \quad S \Vdash O(\phi) \text{ iff } S \models O(\phi) \text{ (iff, by Definition 2, } \sigma_S(f_S) \subseteq \llbracket \phi \rrbracket)$$

The cognitive force of accepting $O(\phi)$, then, is *not* to enforce coordination with respect to the property of accepting $O(\phi)$; this, rather, is the cognitive force of accepting the corresponding imperative sentence $!\phi$.³² The cognitive force of accepting $O(\phi)$ is, rather, to verify that ϕ is preferred: $O(\phi)$ proffers a query/test on a cognitive state (more specifically, a test on its action-guiding parameter). To entertain whether $O(\phi)$ is to execute this test on one's own cognitive state. To decide that $O(\phi)$ is to respond to the successful performance of this test, by remaining in a state that prefers ϕ ; it is *not* to decide that one's cognitive state has the property of preferring ϕ . On the other hand, to decide that $\neg O\phi$ is to respond to the unsuccessful performance of this test, by remaining in a state which fails to prefer ϕ ; it is *not* to decide that one's cognitive state lacks the property of preferring ϕ .

There is much more to say about this proposal, but, however it gets said, notice that, contra the suggestion in Section B.5.3, embracing Expressivism as a theory of deontic cognitive force still allows a wide separation between deontic and imperative cognitive force. Indeed, the proposal developed here has both of the following components.

- Deciding that $O(\phi)$, like accepting the corresponding imperative $!\phi$, is a matter of ending up in a cognitive state that prefers ϕ (rather than a matter of ending up in a

31. This is contra Yalcin (2007)'s proposal for epistemic modals, on which the conventional cognitive force of $\diamond\phi$ is to *induce coordination* on the property of treating ϕ as compatible with the relevant information. The conflation of testing a state S for a property P with proposing that S come to satisfy P is, I suggest, responsible for the running together of imperative and deontic force discussed in Section B.5.3.

32. See Charlow 2010a and the discussion in my "A Dynamic Theory of Imperative Meaning" (Chapter IV).

cognitive state in which ϕ is known to be preferred).

- Deciding that $O(\phi)$ can be issue-resolving, when it occurs in the context of entertaining whether $O(\phi)$. Accepting a corresponding imperative cannot be issue-resolving,³³ although both updates terminate in a state that prefers ϕ .

I emphasize that entertaining whether $O(\phi)$ is distinct from entertaining whether one's cognitive state prefers ϕ . Entertaining whether $O(\phi)$, cognitively, is a matter of executing a test on one's preferences. An agent assesses whether ϕ is best by *consulting* her preferences. Such an assessment does yield information about the agent's preferences, but the agent's inquiry is not necessarily about *this* information. It is typically, rather, about the deontic "worldview" that those preferences encode or represent.

Extending this story to deontic conditionals proceeds much as you'd expect. A deontic conditional $[if\ want(A)][O(B)]$ tests S 's hypothetical or suppositional preferences: its preferences under the supposition that $want(A)$.

$$(B.53) \quad S|[if\ want(A)][O(B)]| = \begin{cases} S, & \text{if } S \Vdash [if\ want(A)][O(B)] \\ \emptyset, & \text{otherwise} \end{cases}$$

As the Expressivist said in Section B.3.2.2, S accepts a deontic conditional $[if\ want(A)][O(B)]$ just if S prefers B on supposition that $want(A)$.

$$(B.54) \quad S \Vdash [if\ want(A)][O(B)] \text{ iff } S|want(A)| \Vdash O(B) \text{ iff } \sigma_{S|want(A)|}(f_{S|want(A)|}) \subseteq \llbracket B \rrbracket$$

This is, I note, a "compositional" story about the cognitive force deontic conditionals. It is a direct consequence of the following two components, the first familiar from our earlier discussion, the second developed in this section.

- A Context-Shifty account of indicative antecedents
- A Test-Expressing Expressivist account of their deontic consequents

It also has the following desirable features.

- Deciding that $[if\ want(A)][O(B)]$, like accepting the corresponding imperative $[if\ want(A)][!B]$, is a matter of ending up in a cognitive state that prefers B , on the supposition that $want(A)$

33. This is a bit sloppy, but I think harmlessly so. In Section 4.3.9, I did suggest that imperatives are the answer to what I there termed "practical questions." So accepting an imperative can, in this sense, be issue-resolving. But not in the following sense: the issue that is salient for an agent who wonders whether she should realize ϕ is one answered by deciding, e.g., that she should realize ϕ . Imperatives cannot resolve issues about what one *should do*, at least in the sense under consideration here. They resolve, rather, issues about what to plan, prefer, intend, and so on. The questions answered by imperatives are, on my way of thinking, *about the agent's psychology* in a way that questions answered by *should*-claims are decidedly not. Although this is my official view, I will not worry about the sloppiness here (or in Section 4.3.9), since these discussions are so far removed from one another.

- Deciding that $[if\ want(A)][O(B)]$ can be issue-resolving, when it occurs in the context of entertaining whether $[if\ want(A)][O(B)]$. Accepting a corresponding imperative cannot be issue-resolving.

Entertaining whether $[if\ want(A)][O(B)]$, cognitively, is a matter of executing a test on one's non-actual, suppositional preferences. An agent assesses whether $[if\ want(A)][O(B)]$ is best by *consulting* her suppositional preferences. Such an assessment does yield information about the agent's suppositional preferences, but, again, the agent's inquiry is not necessarily about *this* information.

BIBLIOGRAPHY

BIBLIOGRAPHY

- Adams, Ernest W. 1975. *The Logic of Conditionals: An application of probability to deductive logic*. Dordrecht.
- Aloni, Maria. 2007. Free choice, modals, and imperatives. *Natural Language Semantics* 15: 65–94. doi:10.1007/s11050-007-9010-2.
- Alonso-Ovalle, Luis. 2006. Disjunction in alternative semantics. <http://semanticsarchive.net/Archive/TVkY2ZIM/alonso-ovalle2006.pdf>. Ph.D. Dissertation, University of Massachusetts, Amherst.
- Alwood, Andrew. 2010. Imperative clauses and the Frege-Geach problem. *Analysis* 70: 105–17. doi:10.1093/analys/anp146.
- Anderson, A. R. 1966. The formal analysis of normative systems. In N. Rescher (ed.) *The Logic of Decision and Action*. University of Pittsburgh Press.
- Åqvist, Lennart. 1964. Interpretations of deontic logic. *Mind* 73: 246–53. doi:10.1093/mind/LXXIII.290.246.
- Åqvist, Lennart. 1965. Choice-offering and alternative-presenting disjunctive commands. *Analysis* 25: 182–4. doi:10.1093/analys/25.5.182.
- Asher, Nicholas & Alex Lascarides. 2001. Indirect speech acts. *Synthese* 128: 183–228. doi:10.1023/A:1010340508140.
- Asher, Nicholas & Alex Lascarides. 2003. *Logics of Conversation*. Cambridge: Cambridge University Press.
- Austin, J. L. 1975. *How to Do Things with Words*. Cambridge: Harvard University Press.
- Barker, Stephen. 2000. Is value content a component of conventional implicature? *Analysis* 60: 268–79. doi:10.1093/analys/60.3.268.
- Belnap, Jr., Nuel D. 1991. Backwards and forwards in the modal logic of agency. *Philosophy and Phenomenological Research* 51: 777–807. doi:10.2307/2108182.
- Belnap, Jr., Nuel D. & Michael Perloff. 1988. Seeing to it that: a canonical form for agentives. *Theoria* 54: 175–99.
- Belnap, Jr., Nuel D. & Michael Perloff. 1992. The way of the agent. *Studia Logica* 51: 463–84.

- Bennett, Jonathan. 1970. [untitled omnibus review]. *Journal of Symbolic Logic* 35: 314–18.
- Bhatt, Rajesh & Roumyana Pancheva. 2006. Conditionals. In M. Everaert & H. van Riemsdijk (eds.) *The Syntax Companion: Volume I*. Blackwell. http://www-rcf.usc.edu/~pancheva/bhatt-pancheva_syncom.pdf.
- Blackburn, Simon. 1984. *Spreading the Word*. Oxford: Blackwell.
- Blackburn, Simon. 1988. Attitudes and contents. *Ethics* 98: 501–17. <http://www.jstor.org/stable/2380964>.
- Blackburn, Simon. 1998. *Ruling Passions*. Oxford: Oxford University Press.
- Boisvert, Daniel R. 2008. Expressive-assertivism. *Pacific Philosophical Quarterly* 89: 169–203. doi:10.1111/j.1468-0114.2008.00315.x.
- Bratman, Michael. 2009. Intention, belief, and instrumental rationality. In D. Sobel & S. Wall (eds.) *Reasons for Action*. Cambridge University Press.
- Broome, John. 1999. Normative requirements. *Ratio* 12: 398–419. doi:10.1111/1467-9329.00101.
- Broome, John. 2000. Incommensurable values. In R. Crisp & B. Hooker (eds.) *Well-Being and Morality: Essays in Honour of James Griffin*. Clarendon Press.
- Broome, John. 2001. Normative practical reasoning. *Aristotelian Society Supplementary Volume* 75: 175–193. doi:10.1111/1467-8349.00085.
- Burge, Tyler. 1979. Sinning against Frege. *The Philosophical Review* 88: 398–432. <http://www.jstor.org/stable/2184957>.
- Cariani, Fabrizio, Stefan Kaufmann & Magdalena Schwager. ms. Deliberative modality under epistemic uncertainty.
- Castañeda, Hector Neri. 1958. Imperatives and deontic logic. *Analysis* 19: 42–8. doi:10.2307/3326749.
- Charlow, Nate. 2009. Imperative statics and dynamics. http://www-personal.umich.edu/~ncharlo/work/imperative_statics_and_dynamics.pdf. Ms., University of Michigan.
- Charlow, Nate. 2010a. Restricting and embedding imperatives. In M. Aloni, H. Bastiaanse, T. de Jager & K. Schulz (eds.) *Proceedings of the 17th Amsterdam Colloquium*. ILLC. doi:10.1007/978-3-642-14287-1_23.
- Charlow, Nate. 2010b. Steps toward a theory of imperative meaning. http://www.natecharlow.com/work/NYU_IMPERATIVES_2010.pdf. Talk delivered to NYU Semantics Group.
- Charlow, Nate. 2011a. The language of direction. <http://www.natecharlow.com/work/toronto.pdf>. Talk delivered at University of Toronto.
- Charlow, Nate. 2011b. What we know and what to do. Ms.

- Chellas, Brian F. 1971. Imperatives. *Theoria* 37: 114–29. doi:10.1111/j.1755-2567.1971.tb00064.x.
- Copp, David. 2001. Realist-expressivism: A neglected option for moral realism. *Social Philosophy and Policy* 18: 1–43. doi:10.1017/S0265052500002880.
- Culicover, Peter W. & Ray Jackendoff. 1997. Semantic subordination despite syntactic coordination. *Linguistic Inquiry* 28: 195–217. <http://www.jstor.org/stable/4178974>.
- Davidson, Donald. 1980. *Essays on Actions and Events*. New York: Oxford University Press.
- Dreier, James. 1999. Transforming expressivism. *Notus* 33: 558–72. doi:10.1111/0029-4624.00193.
- Dreier, James. 2006. Negation for expressivists: a collection of problems with a suggestion for their solution. In R. Shafer-Landau (ed.) *Oxford Studies in Metaethics: Volume 1*. Oxford University Press.
- Dreier, James. 2009. Relativism (and expressivism) and the problem of disagreement. *Philosophical Perspectives* 23: 79–110. doi:10.1111/j.1520-8583.2009.00162.x.
- Dummett, Michael. 1973. *Frege: Philosophy of Language*. London: Duckworth.
- von Fintel, Kai. 1998. The presupposition of subjunctive conditionals. In U. Sauerland & O. Percus (eds.) *The Interpretive Tract*, Vol. 25 of *MIT Working Papers in Linguistics*, 29–44. Cambridge: MITWPL.
- von Fintel, Kai. 2001. Counterfactuals in a dynamic context. In M. Kenstowicz (ed.) *Ken Hale: A Life in Language*. MIT Press. <http://mit.edu/fintel/fintel-2001-counterfactuals.pdf>.
- von Fintel, Kai. 2004. Would you believe it? The King of France is back! Presuppositions and truth-value intuitions. In M. Reimer & A. Bezuidenhout (eds.) *Descriptions and Beyond*, 269–96. New York: Oxford UP.
- von Fintel, Kai & Anthony S. Gillies. 2007. An opinionated guide to epistemic modality. In T. Gendler & J. Hawthorne (eds.) *Oxford Studies in Epistemology: Volume 2*. Oxford University Press. <http://mit.edu/fintel/fintel-gillies-2007-ose2.pdf>.
- von Fintel, Kai & Irene Heim. 2007. *Intensional Semantics*. Unpublished ms.: MIT.
- von Fintel, Kai & Sabine Iatridou. 2005. What to do if you want to go to Harlem: Anankastic conditionals and related matters. <http://mit.edu/fintel/www/harlem-rutgers.pdf>. Ms., MIT.
- von Fintel, Kai & Sabine Iatridou. 2008. How to say *ought* in foreign: The composition of weak necessity modals. In J. Guéron & J. Lecarme (eds.) *Time and Modality*. Springer Netherlands. doi:10.1007/978-1-4020-8354-9. <http://mit.edu/fintel/www/ought.pdf>.
- von Fintel, Kai & Sabine Iatridou. 2009. Covert modals? One particular case. <http://web.mit.edu/fintel/lisa220-class-6-handout.pdf>. Lecture delivered at LSA 2009, Berkeley.

- Fox, Danny. 2006. Free choice and the theory of scalar implicatures. http://web.mit.edu/linguistics/people/faculty/fox/free_choice.pdf. Ms., MIT.
- van Fraassen, Bas C. 1972. The logic of conditional obligation. *Journal of Philosophical Logic* 1: 417–38. doi:10.1007/BF00255570.
- Geach, P. T. 1958. Imperative and deontic logic. *Analysis* 18: 49–56. doi:10.2307/3326785.
- Geach, P. T. 1965. Assertion. *The Philosophical Review* 74: 449–65. <http://www.jstor.org/pss/2183123>.
- Geurts, Bart. 2005. Entertaining alternatives: Disjunctions as modals. *Natural Language Semantics* 14: 383–410. doi:10.1007/s11050-005-2052-4.
- Gibbard, Allan. 1990. *Wise Choices, Apt Feelings*. Cambridge: Harvard University Press.
- Gibbard, Allan. 2003. *Thinking How to Live*. Cambridge: Harvard University Press.
- Gillies, Anthony S. 2001. A new solution to Moore's Paradox. *Philosophical Studies* 105: 237–50. doi:10.1023/A:1010361708803.
- Gillies, Anthony S. 2007. Counterfactual scorekeeping. *Linguistics and Philosophy* 30: 329–60. doi:10.1007/s10988-007-9018-6.
- Gillies, Anthony S. 2010. Iffiness. *Semantics and Pragmatics* 3: 1–42. doi:10.3765/sp.3.4.
- Ginzburg, Jonathan & Ivan A. Sag. 2000. *Interrogative Investigations: the form, meaning, and use of English interrogatives*. Stanford: CSLI Publications.
- Gordon, David & George Lakoff. 1975. Conversational postulates. In P. Cole & J. Morgan (eds.) *Syntax and Semantics Volume 3: Speech Acts*. Academic Press.
- Green, Mitchell S. 1997. The logic of imperatives. In E. Craig (ed.) *The Routledge Encyclopedia of Philosophy*, 717–21. New York: Routledge. <http://people.virginia.edu/~msg6m/implogic.html>.
- Green, Mitchell S. 2000. Illocutionary force and semantic content. *Linguistics and Philosophy* 23: 435–73. doi:10.1023/A:1005642421177.
- Grice, H. P. 1989. *Studies in the Way of Words*. Cambridge: Harvard University Press.
- Groenendijk, Jeroen. 1999. The logic of interrogation. In T. Matthews & D. L. Strolovitch (eds.) *Proceedings of SALT IX*. CLC Publications.
- Groenendijk, Jeroen & Floris Roelofsen. 2009. Inquisitive semantics and pragmatics. <http://www.illc.uva.nl/inquisitive-semantics>. Ms., ILLC.
- Groenendijk, Jeroen & Martin Stokhof. 1984. Studies on the semantics of questions and the pragmatics of answers. <http://dare.uva.nl/record/123669>. Ph.D. Dissertation, University of Amsterdam.
- Groenendijk, Jeroen & Martin Stokhof. 1989. Dynamic montague grammar. In *Papers from the Second Symposium on Logic and Language*, 3–48. Akademiai Kiadoo.

- Groenendijk, Jeroen & Martin Stokhof. 1991. Dynamic predicate logic. *Linguistics and Philosophy* 14: 39–100. doi:10.1007/BF00628304.
- Grosz, Barbara J., Martha Pollack & Candace Sidner. 1989. Computational models of discourse. In M. Posner (ed.) *Foundations of Cognitive Science*. MIT Press.
- Hamblin, Charles. 1973. Questions in Montague English. *Foundations of Language* 10: 41–53. <http://www.jstor.org/stable/25000703>.
- Hamblin, Charles. 1987. *Imperatives*. Oxford: Basil Blackwell.
- Han, Chung-hye. 1998. The structure and interpretation of imperatives: Mood and force in universal grammar. <http://www.sfu.ca/~chunghye/papers/dissertation.pdf>. Ph.D. Dissertation, University of Pennsylvania.
- Hansson, Sven Ove. 1990. Preference-based deontic logic. *Journal of Philosophical Logic* 19: 75–93. doi:10.1007/BF00211186.
- Hare, R. M. 1949. Imperative sentences. *Mind* 58: 21–39. doi:10.1093/mind/LVIII.229.21.
- Hare, R. M. 1952. *The Language of Morals*. Oxford: Oxford University Press.
- Hare, R. M. 1967. Some alleged differences between imperatives and indicatives. *Mind* 76: 309–26. doi:10.1093/mind/LXXVI.303.309.
- Hare, R. M. 1971. Wanting: Some pitfalls. In *Practical Inferences*. Macmillan.
- Hare, R. M. 1989. Some sub-atomic particles of logic. *Mind* 98: 23–37. doi:10.1093/mind/XCVIII.389.23.
- Harel, David, Dexter Kozen & Jerzy Tiuryn. 2000. *Dynamic Logic*. Cambridge: MIT Press.
- Harman, Gilbert. 1976. Practical reasoning. *Review of Metaphysics* 29: 431–63.
- Heim, Irene. 1982. The semantics of definite and indefinite noun phrases. <http://semanticsarchive.net/Archive/Tk0ZmYyY/>. Ph.D. Dissertation, University of Massachusetts, Amherst.
- Heim, Irene. 1983. On the projection problem for presuppositions. In M. Barlow, D. Flickinger & M. Wescoat (eds.) *WCCFL 2: Second Annual West Coast Conference on Formal Linguistics*. Stanford: CSLI Publications.
- Heim, Irene & Angelika Kratzer. 1998. *Semantics in Generative Grammar*. Oxford: Blackwell.
- Isaacs, James & Christopher Potts. 2003. Hidden imperatives. <http://people.umass.edu/potts/talks/isaacs-potts-nels34-handout-1up.pdf>. Talk delivered at NELS 34, SUNY Stony Brook.
- Isaacs, James & Kyle Rawlins. 2008. Conditional questions. *Journal of Semantics* 25: 269–319. doi:10.1093/jos/ffn003.
- Jackson, Frank. 1985. On the semantics and logic of obligation. *Mind* 94: 177–95. doi:10.1093/mind/XCIV.374.177.

- Jacobson, Pauline. 1999. Towards a variable-free semantics. *Linguistics and Philosophy* 22: 117–85. doi:10.1023/A:1005464228727.
- Jeffrey, Richard. 1983. *The Logic of Decision*. Chicago: University of Chicago Press, 2nd edn.
- Jørgensen, Jørgen. 1937-8. Imperatives and logic. *Erkenntnis* 7: 288–96.
- Joyce, James. 2000. Why we still need the logic of decision. *Philosophy of Science* 67: S1–13. <http://www.jstor.org/pss/188653>.
- Kadmon, Nirit & Fred Landman. 1993. Any. *Linguistics and Philosophy* 16: 353–422. doi:10.1007/BF00985272.
- Kamp, Hans. 1981. A theory of truth and semantic representation. In J. Groenendijk, T. Janssen & M. Stokhof (eds.) *Truth, Interpretation, and Information: Selected Papers from the Third Amsterdam Colloquium*. Amsterdam: Mathematisch Centrum.
- Karttunen, Lauri. 1973. Presuppositions of compound sentences. *Linguistic Inquiry* 4: 169–93. <http://www.jstor.org/stable/4177763>.
- Karttunen, Lauri. 1974. Presuppositions and linguistic context. *Theoretical Linguistics* 1: 181–94. doi:10.1515/thli.1974.1.1-3.181.
- Kaufmann, Stefan. 2000. Dynamic context management. In M. Faller, S. Kaufmann & M. Pauly (eds.) *Formalizing the Dynamics of Information*. Stanford: CSLI Publications. <http://faculty.wcas.northwestern.edu/kaufmann/Papers/dcm.pdf>.
- Kaufmann, Stefan & Magdalena Schwager. to appear. A unified analysis of conditional imperatives. In *Proceedings of SALT 19*. CLC Publications. semanticsarchive.net/Archive/TYzZWYyM/.
- Kenny, A. J. 1966. Practical inference. *Analysis* 26: 65–75. doi:10.2307/3326284.
- Kolodny, Niko & John MacFarlane. 2010. Ifs and oughts. *Journal of Philosophy* 107: 115–43. <http://johnmacfarlane.net/ifs-and-oughts.pdf>.
- Kratzer, Angelika. 1977. What ‘must’ and ‘can’ must and can mean. *Linguistics and Philosophy* 1: 337–55. doi:10.1007/BF00353453.
- Kratzer, Angelika. 1981. The notional category of modality. In H. Eikmeyer & H. Rieser (eds.) *Words, Worlds, and Contexts*, 38–74. Berlin: De Gruyter.
- Kratzer, Angelika. 1991a. Conditionals. In A. von Stechow & D. Wunderlich (eds.) *Semantics: An International Handbook of Contemporary Research*, 651–56. Berlin: De Gruyter.
- Kratzer, Angelika. 1991b. Modality. In A. von Stechow & D. Wunderlich (eds.) *Semantics: An International Handbook of Contemporary Research*, 639–51. Berlin: De Gruyter.
- Kratzer, Angelika & Junko Shimoyama. 2002. Indeterminate pronouns: the view from Japanese. In Y. Otsu (ed.) *Proceedings of the Third Tokyo Conference on Psycholinguistics (TCP 2002)*, 1–25. Tokyo: Hituzi Syobo. <http://semanticsarchive.net/Archive/WEwNjc4Z/Indeterminate%20Pronouns.pdf>.

- Krifka, Manfred. 2001a. Quantifying into question acts. *Natural Language Semantics* 9: 1–40. doi:10.1023/A:1017903702063.
- Krifka, Manfred. 2001b. Speech acts and truth-conditional meaning. <http://amor.rz.hu-berlin.de/~h2816i3x/Talks/QuestActsIATL.pdf>. Talk delivered at IATL 17.
- Krifka, Manfred. 2004. Semantics below and above speech acts. <http://amor.rz.hu-berlin.de/~h2816i3x/Talks/StanfordLecture2004.pdf>. Talk delivered at Stanford University.
- Larson, Richard & Gabriel Segal. 1995. *Knowledge of Meaning*. Cambridge: MIT Press.
- Larson, Richard K. 1990. Double objects revisited: Reply to Jackendoff. *Linguistic Inquiry* 21: 589–632. <http://www.jstor.org/pss/4178697>.
- Lascarides, Alex & Nicholas Asher. 2004. Imperatives in dialogue. In P. Kuehnlein, H. Rieser & H. Zeevat (eds.) *The Semantics and Pragmatics of Dialogue for the New Millenium*. Benjamins. <http://homepages.inf.ed.ac.uk/alex/pubs/imps.html>.
- Lewis, David. 1969. *Convention*. Cambridge: Harvard University Press.
- Lewis, David. 1970. General semantics. *Synthese* 22: 18–67. doi:10.1007/BF00413598.
- Lewis, David. 1973. *Counterfactuals*. Malden: Basil Blackwell.
- Lewis, David. 1974. Semantic analyses for dyadic deontic logic. In S. Stenlund (ed.) *Logical Theory and Semantic Analysis*. Dordrecht: D. Reidel.
- Lewis, David. 1976. Probabilities of conditionals and conditional probabilities. *Philosophical Review* 85: 297–315. <http://www.jstor.org/pss/2184045>.
- Lewis, David. 1979a. A problem about permission. In E. Saarinen, R. Hilpinen, I. Niiniluoto & M. B. Hintikka (eds.) *Essays in Honour of Jaakko Hintikka*. Dordrecht: D. Reidel.
- Lewis, David. 1979b. Scorekeeping in a language game. *Journal of Philosophical Logic* 8: 339–59. doi:10.1007/BF00258436.
- Lewis, David. 1981. Ordering semantics and premise semantics for counterfactuals. *Journal of Philosophical Logic* 10: 217–34. doi:10.1007/BF00248850.
- Mackie, J. L. 1977. *Ethics: Inventing Right and Wrong*. New York: Penguin.
- Mastop, Rosja. 2005. What can you do? Ph.D. Dissertation, ILLC.
- Merin, Arthur. 1991. Imperatives: linguistics vs. philosophy. *Linguistics* 29: 669–701.
- Montague, Richard. 1973. The proper treatment of quantification in ordinary English. In Jaakko Hintikka, Julius M.E. Moravcsik & Patrick Suppes (eds.) *Approaches to Natural Language: Proceedings of the 1970 Stanford Workshop on Grammar and Semantics*, 221–242. Dordrecht, Holland: D. Reidel Publishing Co. Reprinted in *Formal Philosophy*, by Richard Montague, Yale University Press, New Haven, CT, 1974, pp. 247–270.
- Morgan, J. L. 1975. Some interactions of syntax and pragmatics. In P. Cole & J. Morgan (eds.) *Syntax and Semantics Volume 3: Speech Acts*. Academic Press.

- Morgan, J. L. 1978. Two types of convention in indirect speech acts. In P. Cole & J. Morgan (eds.) *Syntax and Semantics Volume 9: Pragmatics*. Academic Press.
- Moss, Sarah. to appear. On the pragmatics of counterfactuals. *Noûs* doi:10.1111/j.1468-0068.2010.00798.x.
- Ninan, Dilip. 2005. Two puzzles about deontic necessity. In J. Gajewski, V. Hacquard, B. Nickel & S. Yalcin (eds.) *New Work on Modality*, Vol. 51 of *MIT Working Papers in Linguistics*. Cambridge: MITWPL. <http://semanticsarchive.net/Archive/WZINTg0Y/>.
- Portner, Paul. 2004a. The semantics of imperatives within a theory of clause types. In K. Watanabe & R. Young (eds.) *Proceedings of SALT 14*. CLC Publications. <http://semanticsarchive.net/Archive/mJlZGQ4N/>.
- Portner, Paul. 2004b. Vocatives, topics, and imperatives. https://www9.georgetown.edu/faculty/portnerp/my_papers/Stuttgart_handout.pdf. Talk delivered at IMS Workshop on Information Structure, Bad Teinach.
- Portner, Paul. 2007. Imperatives and modals. *Natural Language Semantics* 15: 351–83. doi:10.1007/s11050-007-9022-y.
- Portner, Paul. 2010. Permission and choice. <http://semanticsarchive.net/Archive/jl1YjMyY/>. Talk delivered at Semantics and Philosophy in Europe III.
- Potts, Christopher. 2003. Keeping world and will apart: A discourse-based semantics for imperatives. <http://people.umass.edu/potts/talks/potts-nyu-handout.pdf>. Talk delivered at NYU Syntax/Semantics Lecture Series.
- Potts, Christopher. 2005. *The Logic of Conventional Implicatures*. Oxford: Oxford University Press.
- Potts, Christopher. 2007. The expressive dimension. *Theoretical Linguistics* 33: 165–98. doi:10.1515/TL.2007.011.
- Price, Huw. 1994. Semantic minimalism and the Frege point. In L. Tsohatzidis (ed.) *Foundations of Speech Act Theory: Philosophical and Linguistic Perspectives*. London: Routledge. <http://www.usyd.edu.au/time/price/preprints/sem-min.pdf>.
- Roberts, Craige. 1989. Modal subordination and pronominal anaphora in discourse. *Linguistics and Philosophy* 12: 683–721. doi:10.1007/BF00632602.
- Roberts, Craige. 1996. Information structure in discourse: Towards an integrated formal theory of pragmatics. In *OSU Working Papers in Linguistics, Vol 49: Papers in Semantics*. Ohio State University. <http://semanticsarchive.net/Archive/WYzOTRkO/>.
- van Rooy, Robert. 2000. Permission to change. *Journal of Semantics* 17: 119–45. doi:10.1093/jos/17.2.119.
- Ross, Alf. 1941. Imperatives and logic. *Theoria* 7: 53–71. doi:10.1111/j.1755-2567.1941.tb00034.x.
- Ross, W. D. 1930. *The Right and the Good*. Oxford: Oxford University Press.
- Russell, Benjamin. 2007. Imperatives in conditional conjunction. *Natural Language Semantics* 15: 131–66. doi:10.1007/s11050-007-9012-0.

- Sadock, Jerrold. 1978. On testing for conversational implicature. In Peter Cole (ed.) *Syntax and Semantics 9: Pragmatics*. New York: Academic Press.
- Sadock, Jerrold M. & Arnold M. Zwicky. 1985. Speech act distinctions in syntax. In T. Shopen (ed.) *Language typology and syntactic description: Clause structure*. Cambridge University Press.
- Sæbø, Kjell Johan. 2001. Necessary conditions in a natural language. In C. Féry & W. Sternefeld (eds.) *Audiatur Vox Sapientiae: A Festschrift for Arnim von Stechow*. Akademie Verlag. <http://vivaldi.sfs.nphil.uni-tuebingen.de/~arnim10/Festschrift/Saeboe-8-komplettfertig.pdf>.
- Schroeder, Mark. 2008a. *Being For*. Oxford: Oxford University Press.
- Schroeder, Mark. 2008b. Expression for expressivists. *Philosophy and Phenomenological Research* 76: 86–116. doi:10.1111/j.1933-1592.2007.00116.x.
- Schroeder, Mark. 2008c. How expressivists can and should solve their problems with negation. *Noûs* 42: 573–99. doi:10.1111/j.1468-0068.2008.00693.x.
- Schroeder, Mark. 2008d. What is the Frege-Geach problem? *Philosophy Compass* 3/4: 703–20. doi:10.1111/j.1747-9991.2008.00155.x.
- Schroeder, Mark. 2009. Hybrid expressivism: Virtues and vices. *Ethics* 119: 257–309. doi:10.1086/597019.
- Schroeder, Mark. 2010. Getting noncognitivism out of the woods. *Analysis* 70: 117–29. doi:10.1093/analys/anp156.
- Schulz, Katrin. 2003. You may read it now or later: A case study on the paradox of free choice permission. <http://www.ilc.uva.nl/Publications/ResearchReports/MoL-2004-01.text.pdf>. M.A. Thesis, University of Amsterdam.
- Schulz, Katrin. 2005. A pragmatic solution for the paradox of free choice permission. *Synthese* 147: 343–77. doi:10.1007/s11229-005-1353-y.
- Schwager, Magdalena. 2005. Permitting permissions. In P. Egré & L. Alonso i Alemany (eds.) *Proceedings of the Ninth ESSLLI Student Session*. EALLI. <http://user.uni-frankfurt.de/~scheiner/papers/permission.ps>.
- Schwager, Magdalena. 2006. Conditionalized imperatives. In M. Gibson & J. Howell (eds.) *Proceedings of SALT 16*. CLC Publications. <http://user.uni-frankfurt.de/~scheiner/papers/schwagerFEB07.pdf>.
- Searle, John R. 1962. Meaning and speech acts. *The Philosophical Review* 71: 423–32. <http://www.jstor.org/pss/2183455>.
- Searle, John R. 1969. *Speech Acts: An Essay in the Philosophy of Language*. Cambridge: Cambridge University Press.
- Searle, John R. 1975a. Indirect speech acts. In P. Cole & J. Morgan (eds.) *Syntax and Semantics Volume 3: Speech Acts*. Academic Press.
- Searle, John R. 1975b. A taxonomy of illocutionary acts. In K. Gunderson (ed.) *Minnesota studies in the philosophy of language*, 334–69. Minneapolis: University of Minnesota Press.

- Searle, John R. & Daniel Vanderveken. 1985. *Foundations of illocutionary logic*. Cambridge: Cambridge University Press.
- Segerberg, Krister. 1990. Validity and satisfaction in imperative logic. *Notre Dame Journal of Formal Logic* 31: 203–21. doi:10.1305/ndjfl/1093635415.
- Segerberg, Krister. 1992. Getting started: Beginnings in the logic of action. *Studia Logica* 51: 347–78. doi:10.1007/BF01028968.
- Silk, Alex. ms. Evidence-sensitivity in deontic modals.
- Smith, Michael. 1987. The humean theory of motivation. *Mind* XCVI: 36–61. doi:10.1093/mind/XCVI.381.36.
- Smith, Michael. 1994. *The Moral Problem*. Oxford: Blackwell.
- Stalnaker, Robert. 1968. A theory of conditionals. In N. Rescher (ed.) *Studies in logical theory (American Philosophical Quarterly Monograph Series 2)*, 98–112. Blackwell.
- Stalnaker, Robert. 1975. Indicative conditionals. *Philosophia* 5: 269–86. doi:10.1007/BF02379021.
- Stalnaker, Robert. 1978. Assertion. In Peter Cole (ed.) *Syntax and Semantics 9: Pragmatics*. New York: Academic Press.
- Starr, William B. 2011. A preference semantics for imperatives. Ms.
- Stenius, Erik. 1967. Mood and language game. *Synthese* 17: 254–74. doi:10.1007/BF00485030.
- Stevenson, C. L. 1944. *Ethics and Language*. New Haven: Yale University Press.
- Swanson, Eric. 2006. Interactions with context. <http://www-personal.umich.edu/~ericsw/Swanson,%20Interactions%20with%20Context.pdf>. Ph.D. Dissertation, MIT.
- Swanson, Eric. 2008. Constraint semantics and its application to conditionals. <http://www-personal.umich.edu/~ericsw/research/Swanson,%20Constraint%20Semantics%20handout.pdf>. Talk delivered at First Formal Epistemology Festival, Konstanz.
- Thomason, Richmond H. 1981a. Deontic logic and the role of freedom in moral deliberation. In R. Hilpinen (ed.) *New Studies in Deontic Logic*. D. Reidel.
- Thomason, Richmond H. 1981b. Deontic logic as founded on tense logic. In R. Hilpinen (ed.) *New Studies in Deontic Logic*. D. Reidel.
- Unwin, Nicholas. 1999. Quasi-realism, negation and the Frege-Geach problem. *The Philosophical Quarterly* 49: 337–52. doi:10.1111/1467-9213.00146.
- Unwin, Nicholas. 2001. Norms and negation: A problem for Gibbard's logic. *The Philosophical Quarterly* 51: 60–75. doi:10.1111/1467-9213.00214.
- Velleman, J. David. 1989. *Practical Reflection*. Princeton: Princeton University Press.
- Velleman, J. David. 2000. *The Possibility of Practical Reason*. New York: Oxford University Press.

- Veltman, Frank. 1996. Defaults in update semantics. *Journal of Philosophical Logic* 25: 221–61. doi:10.1007/BF00248150.
- Veltman, Frank. 2005. Making counterfactual assumptions. *Journal of Semantics* 22: 159–80. doi:10.1093/jos/ffh022.
- Veltman, Frank. 2008. Imperatives at the semantics/pragmatics interface. Ms., ILLC.
- von Wright, G. H. 1963. *Norm and action: a logical inquiry*. London: Routledge.
- Vranas, Peter B. M. 2010. In defense of imperative inference. *Journal of Philosophical Logic* 39: 59–71. doi:10.1007/s10992-009-9108-8.
- Weirich, Paul. 1980. Conditional utility and its place in decision theory. *The Journal of Philosophy* 77: 702–15.
- Williams, B. A. O. 1963. Imperative inference. *Analysis* 23, Supplement 1: 30–6. doi:10.2307/3326619.
- Wilson, Deirdre & Dan Sperber. 1988. Mood and the analysis of non-declarative sentences. In J. Dancy, J. Moravcsik & C. Taylor (eds.) *Human agency: Language, duty and value*. Stanford University Press.
- Wittgenstein, Ludwig. 1958. *Philosophical Investigations*. Oxford: Blackwell.
- Yablo, Stephen. 2006. Non-catastrophic presupposition failure. In J. Thomson & A. Byrne (eds.) *Content and Modality: Themes From the Philosophy of Robert Stalnaker*. Oxford University Press.
- Yalcin, Seth. 2007. Epistemic modals. *Mind* 116: 983–1026. doi:10.1093/mind/fzm983.
- Yalcin, Seth. forthcoming. Nonfactualism about epistemic modality. In B. Weatherson & A. Egan (eds.) *Epistemic Modality*. Oxford: Oxford University Press.
- Zarnić, Berislav. 2003. Imperative logic, moods and sentence radicals. <http://www.ffst.hr/~logika/ac14zarnic.pdf>. Ms., University of Split.
- Zimmerman, David. 1980. Force and sense. *Mind* 89: 214–33. <http://www.jstor.org/pss/2253170>.
- Zimmermann, Thomas Ede. 2000. Free choice disjunction and epistemic possibility. *Natural Language Semantics* 8: 255–90. doi:10.1023/A:1011255819284.