

Subjunctive Credences and Semantic Humility

SARAH MOSS

University of Michigan

Suppose that I am holding a delicate china plate. For any plate that drops from my hands, there is an extremely small objective chance that a sudden gust of hot air will cushion its fall, preventing it from breaking. Consider the sequence of conditionals:

(1a) If I were to drop this plate, it might be saved by a sudden gust of hot air.

(1b) If I were to drop this plate, it would break.

In light of the chance that the plate will be saved by a sudden gust of hot air, what should we say about the truth conditions of the ordinary conditional (1b)?

Some theories predict that (1b) is true. For instance, Lewis 1973b says that the plate being saved by friendly air currents is a remarkable low-probability event, and that worlds where such events happen are farther away than worlds where I drop the plate and nothing special happens. Hence the plate breaks in all the closest worlds where I drop it, and so (1b) comes out true. Williams 2008 develops and defends a neo-Lewisian account which similarly excludes wacky events from worlds closest to ours. And Bennett 2003 develops a “near-miss” proposal, according to which the plate may be saved by friendly air currents in some of the closest worlds where I drop it, but (1b)

I am grateful to John Hawthorne, David Manley, Bob Stalnaker, Jason Stanley, Eric Swanson, Steve Yablo, and the 2010 Midwest Epistemology Workshop for helpful comments and discussion. I am especially grateful to Robbie Williams for helpful comments on earlier drafts and for ongoing fruitful conversations about this material.

nevertheless counts as true because the plate breaks at a very high proportion of the closest plate-dropping worlds.

Some theories predict that (1b) is not true. For instance, Hájek 2007 argues that (1a) is true and incompatible with (1b), and so he concludes that (1b) is false. Gillies 2007 argues that conditionals like (1a) expand the domain that subjunctives quantify over to include some possibilities where the plate is saved by friendly air currents, so that (1b) fails to be true in the resulting context.

I argue that both predictions are misguided. In particular, the above theories are incompatible with justified intuitions about what credences we ought to have in subjunctive conditionals.¹ In §§1–2, I argue that the theories fail to accommodate ordinary judgments about embedded subjunctives. In §§3–4, I examine our ordinary judgments in more detail. Rather than endorsing armchair verdicts about the truth values of subjunctives, ordinary speakers endorse constraints linking credences in subjunctives with objective chances. I demonstrate that we can often derive such constraints from justified premises. In §§5–6, I consider an objection, namely that my argument supports a subjunctive analog of Adams' Thesis that is undermined by an analog of a familiar triviality result for indicatives. In response, I outline the limits of the premises that ground constraints on subjunctive credences. Belief in these premises is justified in many ordinary contexts, but some premises fail to hold in the context of the threatening triviality proof. Hence my derivation of constraints on subjunctive credences shows both why we are often justified in accepting those constraints, and why the constraints are not threatened by a subjunctive triviality result modeled after Lewis 1976.

1. Embedding Data for a Theory of Subjunctive Conditionals

In order to highlight the correct response to (1b), it helps to consider our ordinary language judgments about another case involving wacky events. Suppose I have a fair coin which I am not going to flip one million times. Consider the following sequence of conditionals:

(2a) If I were to flip this coin one million times, it might land tails at least once.

(2b) If I were to flip this coin one million times, it would land heads each time.

¹ In what follows, 'subjunctive conditional' denotes those conditionals commonly associated with updating by imaging rather than by conditionalization, i.e. "ontic conditionals" in the terminology of Lindström & Rabinowicz 1995 and Rott 1999. I use 'subjunctive' rather than 'counterfactual' since many of the relevant conditionals are future subjunctives with true antecedents.

The theories surveyed above assign the same truth value to (2b); they agree that (2b) is false.² In particular, each of the theories says that we can know the truth value of (2b) without doing any coin flipping, since its truth value is determined by facts that we can know without doing any coin flipping: whether the coin landing heads each time is a remarkable and improbable event, whether the coin lands heads each time in a very small proportion of the closest worlds where it is flipped one million times, whether (2a) is true and incompatible with (2b), or whether (2a) expands the domain that conditionals quantify over to include some tails worlds. All of these facts are commonly taken to be plainly determined in the case described above.

Here is the challenge for the theories in question: even though these facts are plainly determined in the coin case, our ordinary judgments contradict the verdict that (2b) is plainly false. Saying that (2b) is false fails to predict the behavior of (2b) in embedded contexts. For instance, (3b) sounds distinctly less felicitous than hedged assertions such as (3a):

(3a) I can't say for sure what would happen if I were to flip this coin one million times.

(3b) #It's not the case that this coin would land heads each time.

And (2b) does not embed in attitude ascriptions as if it were plainly false:

(4) I wonder if this coin would land heads each time if I were to flip it one million times.

(5) I think it's possible—though of course extremely unlikely—that this coin would land heads each time if I were to flip it one million times.

(6) I doubt that this coin would land heads each time if I were to flip it one million times. But of course, you never know for sure until you try.

Each of these judgments suggests that without doing any coin flipping, the semanticist can know that (2b) deserves very low credence, but not that (2b) is actually false. In more detail: as we test semantic theories

² Strictly speaking, Gillies 2007 argues only that (2b) is not true, cf. his official account on p. 351. The distinction between predicting that (2b) is false and predicting that it is not true is not relevant for my argument, as neither prediction exhibits appropriate semantic humility.

against our intuitions, we take it for granted that we can tell which events are improbable and atypical and quasi-miraculous. Suppose our semantics says the truth of (2b) depends just on facts about how probable or typical or quasi-miraculous it would be for a coin to land heads one million times. Then as soon as we knew the semantics for (2b), we would be able to tell whether (2b) was true. But intuitively, it should not be so easy to tell whether (2b) is true. Quite the contrary: if we do not actually flip the coin one million times, then intuitively it should be impossible to say how it would have landed if we had done so. It is overreaching for our theory of conditionals to make a pronouncement about how a fair coin would have landed in any number of merely possible flips, whether one or one million.

The case of the fragile plate is similar. Even if we are certain that the plate being saved by friendly air currents would be a remarkable and improbable event, we still cannot be certain that the plate would break if dropped. Hence our theory of conditionals should not predict that (1b) is true, but should instead accommodate the intuition that we should simply give it very high credence:

(1b) If I were to drop this plate, it would break.

In other words, our theory of subjunctives should display *semantic humility*, i.e., our semantic theory should deliver the truth conditions of sentences without pronouncing on whether those conditions actually obtain.

The case for semantic humility is even stronger in cases where the objective chances relevant to our judgments are more conspicuous. For instance, consider a coin biased 3:1 in favor of heads. (7a) sounds much better than (7b):

(7a) Probably, the coin would land heads if I were to flip it.

(7b) #It's not the case that the coin would land heads if I were to flip it.

Our acceptance of (7a) is at odds with the overreaching verdict delivered by some theories of subjunctives, namely that the prejacent of 'probably' is plainly false.³ Intuitively, we should think that it is more likely than not that the biased coin would land heads if I were to flip it, even in light of information that is alleged to plainly determine that this conditional is false.

³ For instance, see p. 173 of Joyce 1999 and p. 331 of Lewis 1981.

Judgments supporting semantic humility can help us decide between otherwise similar theories of subjunctives. For example: in Moss 2010, I defend a variably strict conditional semantics for subjunctives against recent objections from advocates of dynamic semantic theories.⁴ In particular, I argue that an independently motivated pragmatic theory predicts the infelicity of reverse Sobel sequences, which was alleged to be a challenge for the variably strict semantics. In addition, I offer an olive branch to the dynamic semanticists at the end of that paper, acknowledging that they may predict a variety of intuitions about reverse Sobel sequences if they supplement their semantic proposal with a variant of the pragmatic theory I develop. The present paper whittles away at that olive branch. A semantic theory of the infelicity of (1b) is distinguished from a pragmatic theory in virtue of predicting that (1b) fails to be true or that it generates some semantic inconsistency after (1a). And these sorts of semantic verdicts are challenged by the case for semantic humility that I present here.

In diagnosing theories that contradict our judgments about (1)–(7), it is helpful to understand one common motivation for delivering semantically immodest verdicts. For instance, consider the counterintuitive prediction that (8) is plainly true:

(8) If I were to flip this coin one million times, it would land heads at least once.

Predicting that (8) is true can seem like a forced alternative to giving a widespread error theory. Here is a characteristic statement of this reasoning from Williams 2010b:⁵

Quite generally it is plausible that any theory that makes [(8)] false, will have a hard time avoiding making ordinary counterfactual judgments false. The relevant facts seem so similar—the antecedents counterfactually imply that the consequents are overwhelmingly probable, but not that it is absolutely certain that they will occur. ... I shall assume, therefore, that if we are to avoid an error theory of ordinary counterfactuals, counterfactuals such as [(8)] must also be true. (12)

As we have seen, we should reject this argument for the truth of (8). In particular, this argument presupposes that semanticists should decide between theories that either “make (8) false” or “make (8) true.” The embedding data we have considered suggest that neither is appropriate.

⁴ These objections are developed at length in von Stechow 2001 and Gillies 2007.

⁵ For further examples of similar reasoning, see §98 of Bennett 2003 and p. 396 of Hawthorne 2005.

On at least one reading, the truth of (8) intuitively depends on facts that are not plainly accessible to the semanticist.

However, the reasoning in Williams 2010b does suggest a more serious worry that we must address. If our theory of conditionals does not itself determine that (8) is true, then it may be inappropriate to say that one *knows* that the coin would land heads at least once if it were flipped one million times. One might worry that for the same reason, it will be inappropriate to self-ascribe knowledge of any ordinary subjunctive conditional. In other words, it may appear as if semantic humility demands a unified treatment of knowledge ascriptions embedding the following:

(1b) If I were to drop this plate, it would break.

(8) If I were to flip this coin one million times, it would land heads at least once.

After all, we cannot rule out that a dropped plate will be saved by friendly air currents, just as we cannot rule out that a fair coin flipped one million times will land tails each time. So even if avoiding making (8) true does not force us to give an error theory of ordinary subjunctives, it may appear to force us to give an error theory of the knowledge ascriptions that embed them. For similar reasons, it may appear as if semantic humility demands that ordinary subjunctives should be unassertable, for the same reason that some assertions of (8) are infelicitous, namely that they express an inappropriate confidence that an unlikely event simply would not obtain.

But in fact, semantic humility does not demand a uniform treatment of (1b) and (8). Furthermore, what explains the contrast in the knowability or assertability of these two sentences has nothing to do with their conditional nature. In aiming to predict that (1b) and (8) are plainly true, the opponent of semantic humility mistakes an epistemic problem for a semantic one. Suppose that you receive testimony that a certain fragile plate was dropped yesterday. On that basis, you have a justified high credence that the plate broke. (9) will generally be felicitous for you to utter in most ordinary conversations about the plate:

(9) That plate broke.

By contrast, suppose that you receive testimony that someone flipped a certain fair coin one million times yesterday. On that basis, you have a justified high credence that the coin landed heads at least once. (10) will

not generally be felicitous for you to utter in most ordinary conversations about the coin:

(10) That coin landed heads at least once.

(10) expresses confidence that a certain unlikely event did not obtain. And this expression of confidence seems inappropriate. Given that the contrast between (1b) and (8) is closely paralleled by the contrast between (9) and (10), the former contrast begs for an explanation that is not limited to conditional statements.

In fact, comparing (9) with (10) suggests a simple and natural explanation of both contrasts, namely that (10) sets up a “lottery” context, while (9) does not. In other words, (9) is like (11), while (10) is like (12):

(11) I will not be able to afford to go to Paris next summer.

(12) This lottery ticket is going to lose.

Furthermore, some general theories of the unassertability of lottery statements may be easily extended to account for the contrast between (9) and (10).⁶ For example: in typical contexts of utterance of (8), (10), and (12), certain wacky possibilities are salient. These possibilities are incompatible with the truth of the uttered sentence. According to general principles defended in Moss 2010, this explains why utterances of (8), (10), and (12) are typically infelicitous. By contrast, utterances of (1b), (9), and (11) typically do not raise wacky possibilities to salience. Given this theory of lottery statements, we may account for the unassertability of (8) without endorsing the claim that ordinary subjunctives are unassertable. And similarly, we may appeal to general theories of knowledge ascriptions in ordinary contexts and lottery contexts to explain the contrast between knowledge ascriptions embedding (1b) and (8).

2. Alternative Accounts of Our Ordinary Language Judgments

Some opponents of semantic humility simply reject the foregoing summary of our ordinary language judgments. For example, Hájek 2007 sometimes suggests that embedding data are no problem for his theory, because it is perfectly fine to assert conditionals such as:

⁶ Some readers have reported that they hear (8), (10), and (12) as felicitous even in lottery contexts. It is not important for my discussion that these sentences are heard as infelicitous by all speakers, but rather that many speakers find them less acceptable than (1b), (9), and (11).

(13) It's not the case that if I had flipped the coin, it **WOULD** have landed heads.

So Hájek is not motivated to give a theory that predicts that negations of ordinary subjunctive conditionals are unassertable.⁷

However, intuitions about (13) should not direct our general theory of subjunctive conditionals. By capitalizing 'would' when giving examples of unassertable conditionals, Hájek prompts interpretations of the conditionals where 'would' is focused. But focusing constituents may change the semantic values of conditionals, or suppress readings that would be available in the absence of focus. By concentrating on readings of conditionals made salient by focal stress, Hájek neglects other available readings. If a speaker merely utters (13') without any focal stress, then there is at least one reading on which the sentence is unassertable.

(13') It's not the case that if I had flipped the coin, it would have landed heads.

In order to isolate this reading, it helps to think of (13') as answering a question:

(14a) Would this coin have landed heads if it had been flipped?

(14b) I don't know what would have happened. / #No.

If it is common ground that the coin in question is fair, (13') sounds bad in response to (14a). Insofar as (13') sounds okay on this reading, it sounds as if the speaker is implicitly communicating that she knows that neither side of the coin is marked heads. On this reading, far from asserting or assuming (13'), it is felicitous to ask someone to guess whether the embedded subjunctive holds. (14b) highlights a shortcoming of error theories of subjunctives, namely that they fail to predict that on at least one reading, (13') is unassertable.

Another common response to arguments for semantic humility about subjunctives is that sentences such as (5') and (7a) merely appear to ascribe credence in a conditional:

(5') I think it's extremely unlikely that the coin would land heads each time if I were to flip it one million times.

(7a) Probably, the coin would land heads if I were to flip it.

⁷ See p. 84 of Hall & Hájek 1994 for further sympathetic comments.

Here is a characteristic statement of this response from Bennett 2003:

Admittedly, we often find it natural to say things like ‘There’s only a small chance that if he had entered the lottery he would have won’, and ‘It’s 50% likely that if he had tossed the coin it would have come down heads’. In remarks like these, the speaker means something of the form $A > (P(C) = n)$ —if the antecedent were true, the consequent would have a certain probability; yet the sentence he utters means something of the form $P(A > C) = n$ When we use one to mean the other, we employ a usage that is idiomatic but not strictly correct. (251)

This clever reinterpretation strategy accommodates our judgments about (5') and (7a), but still fails to explain a range of ordinary judgments about subjunctive conditionals. For instance, the unembedded conditional (8) is plainly true on the near-miss proposal developed by Bennett 2003:

(8) If I were to flip this coin one million times, it would land heads at least once.

The reinterpretation strategy under consideration does not explain why (8) nevertheless sounds unacceptably hubristic in a variety of contexts. Furthermore, it remains difficult to explain why one cannot simply answer (14a) in the negative:

(14a) Would this coin have landed heads if it had been flipped?

(14b) I don't know what would have happened. / #No.

The most natural extension of the reinterpretation strategy to (14b) would be to claim that the negation in the infelicitous answer only apparently takes wide scope over the conditional, so that the logical form of that answer is as follows:

(15) #If it had been flipped, this coin would have not landed heads (i.e., it would have landed tails).

But it is implausible that the negation in the answer ‘No’ in (14b) scopes under covert material made explicit in (15). Similar readings are absent in structurally parallel cases. For instance, it is not possible to read (16b) as expressing the proposition that exactly three people did not read the newspaper:

(16a) Did exactly three people read the newspaper?

(16b) No.

Hence it is not clear how Bennett's reinterpretation strategy accounts for our judgments about subjunctives embedded in questions.

In light of these difficulties for the reinterpretation strategy, one might instead respond to my defense of semantic humility by saying that accommodating ordinary language judgments about subjunctives is supererogatory. Hájek 2007 endorses an error theory of subjunctive conditionals, suggesting that it is not a significant cost for his theory to contradict systematic ordinary language judgments.⁸ Hájek does concede one crucial point, though: even the error theorist must explain the success of our practice of using subjunctive conditionals. In response, Hájek claims that our utterances of conditionals such as (1b) are legitimized by the existence of nearby conditionals such as (17):

(1b) If I were to drop this plate, it would break.

(17) If I were to drop this plate, it would very probably break.

Hájek says that this is his best explanation of what vindicates our practice of uttering ordinary subjunctive conditionals.

But our ordinary judgments about (1b) do not in fact pattern with our judgments about (17). For instance, if we are certain that the conditional chance of the plate breaking if dropped is merely very high, then intuitively we ought to be certain of (17) and less than certain of (1b). In fact, as I argue in the following section, our expectation of the chance of the plate breaking if dropped should match our credence in (1b), while intuitively it should not match our credence in surrogate conditionals like (17). Hence our judgments about 'would' conditionals are not explained by the claim that speakers treat 'would' conditionals as if they expressed the propositions literally expressed by their 'would probably' surrogates. The error theorist lacks an account of our highly systematic ordinary language judgments about subjunctive conditionals.

⁸ Of course, this strategy incurs costs common to all error theories. In discussing his rejection of the conditional law of excluded middle, Lewis concedes that his semantics does not respect the "offhand opinion of any ordinary language speaker," and that *ceteris paribus*, one should aim to respect such judgments (Lewis 1973a, 80).

3. An Argument Supporting Intuitive Constraints on Subjunctive Credences

The ordinary language judgments that conflict with extant theories of subjunctives are not only pervasive, but highly systematic. Recall that if a coin is biased 3:1 in favor of heads, (7a) sounds much better than (7b):

(7a) Probably, the coin would land heads if I were to flip it.

(7b) #It's not the case that the coin would land heads if I were to flip it.

(7a) suggests that our credence that the coin would land heads should be determined by our opinions about the bias of the coin. If we are certain that the objective chance of the coin landing heads if flipped is .75, then intuitively we should have .75 credence that the coin would land heads if it were flipped. The same reasoning yields the intuitively correct credence in the other subjunctives mentioned in §1:

(1b) If I were to drop this plate, it would break.

(2b) If I were to flip this coin one million times, it would land heads each time.

(8) If I were to flip this coin one million times, it would land heads at least once.

Several facts about objective chances govern our assessment of these conditionals. The objective chance of a fragile plate breaking if dropped is extremely high. The objective chance of a fair coin landing heads one million times if tossed one million times is $\frac{1}{2^{1,000,000}}$. And the objective chance of a fair coin landing heads at least once if tossed one million times is $1 - \frac{1}{2^{1,000,000}}$. On reflection, it is intuitive that our acceptance of these objective chance facts intuitively constrains our credence in (1b), (2b), and (8):

(1b') Our credence in (1b) should be extremely high.

(2b') Our credence in (2b) should be $\frac{1}{2^{1,000,000}}$.

(8') Our credence in (8) should be $1 - \frac{1}{2^{1,000,000}}$.

(1b'), (2b'), and (8') are each instances of the following schematic principle:

(BRIDGE) A subject's credence in 'if it were going to be that *A*, it would be that *C*' should match her subjective expectation of the current conditional objective chance that *C* given that *A*.

(BRIDGE) delivers the credence that we should have in each conditional stated above, even in light of information that is alleged to plainly determine its truth value.⁹

Furthermore, each of the above instances of (BRIDGE) is not only intuitively correct, but also supported by a simple and compelling argument. Let ‘ $A \Box \rightarrow C$ ’ be one of the above subjunctive conditionals, and suppose you are a rational subject with a well-defined credence in this conditional. Let \mathcal{C} be your current credence distribution, and let $\varepsilon(x)$ be the expected value of $x \in [0,1]$ given these credences.¹⁰ Let ch be the current objective chance function. The argument for each of the above instances of (BRIDGE) proceeds as follows:

$$\begin{aligned}
 \mathcal{C}(A \Box \rightarrow C) &= \varepsilon(ch(A \Box \rightarrow C)) && \text{Superintendent Principle} \\
 &= \varepsilon(ch(A \Box \rightarrow C|A)) && \text{expectation of (INDEPENDENCE)} \\
 &= \varepsilon(ch(A \wedge A \Box \rightarrow C|A)) && \text{probability calculus} \\
 &= \varepsilon(ch(A \wedge C|A)) && \text{expectation of Centering} \\
 &= \varepsilon(ch(C|A)) && \text{probability calculus}
 \end{aligned}$$

Two of the three substantive premises in this argument are familiar. The Principal Principle says that your credences should reflect certainties about objective chances: if you are certain of the objective chance of P , your credence in P should equal your subjective expectation of the objective chance of P . The first premise of the above derivation is a natural generalization of the Principal Principle defended by Lewis 1980 and van Fraassen 1980. Call it the *Superintendent Principle*: your credence that P should always equal your subjective expectation of the objective chance of P . Just as with the Principal Principle, the Superintendent Principle must be restricted to cases where subjects lack inadmissible evidence. But intuitively you do not have any inadmissible evidence about how an ordinary coin would land if it were flipped, or whether an ordinary plate would break if it were dropped. Hence your credence that the coin would land heads if flipped should match your expectation of the chance of it landing heads if flipped. And your credence that the plate would break if dropped should match your expectation of the chance of it breaking if dropped. These instances of

⁹ Skyrms 1980 defends a similar principle connecting estimates of chance and the “basic assertability value” of a subjunctive conditional, and his principle entails (BRIDGE) for subjects uncertain about which of several deterministic states obtains. See Skyrms 1994, 1998 for further discussion. I use expressions such as “credences in subjunctives” in order to avoid foreclosing on the possibility that subjunctive conditionals do not express propositions; if subjunctives do express propositions, “credences in subjunctives” simply refers to credences in propositions expressed by subjunctives.

¹⁰ In other words: $\varepsilon(x) = \int_{r=0}^1 r \cdot \mathcal{C}(x = r)$.

the Superintendent Principle are intuitively just as good as similar instances of the Principal Principle, e.g., the claim that your credence that the coin will be flipped and land heads should equal your expectation of the chance that the coin will be flipped and land heads.

The other familiar premise in the above derivation of (BRIDGE) is roughly that you accept axioms for subjunctive conditionals that are valid under standard centering assumptions. In the axiomatization for VC given in Lewis 1973a, the relevant axioms are as follows:

$$(vc6) (\phi \Box \rightarrow \psi) \supset (\phi \supset \psi)$$

$$(vc7) (\phi \wedge \psi) \supset (\phi \Box \rightarrow \psi)$$

The axiom (vc6) suffices to establish that if $A \wedge A \Box \rightarrow C$, then $A \wedge C$. The axiom (vc7) establishes the opposite direction.¹¹ The centering axioms are commonly accepted principles governing the logic of subjunctive conditionals. For example, it is commonly accepted that ‘if you were to flip the coin, it would land heads’ is true in worlds where you flip the coin and it lands heads. And this is in accord with the ordinary intuition that flipping a coin and seeing how it lands is a good way to see whether that coin would land heads if you flipped it. Given the centering axioms, necessarily, $A \wedge A \Box \rightarrow C$ holds if and only if $A \wedge C$ holds. Hence conditional chances concerning whether $A \wedge A \Box \rightarrow C$ must equal conditional chances concerning whether $A \wedge C$, and so your rational expectation of such conditional chances will be equal.

The less familiar premise in the above argument concerns your expectations about the nature of the objective chance function: $\varepsilon(ch(A \Box \rightarrow C)) = \varepsilon(ch(A \Box \rightarrow C | A))$. This premise holds whenever you are certain that relative to the objective chance function, a subjunctive conditional is probabilistically independent of its antecedent, i.e. that the following principle holds:¹²

$$(INDEPENDENCE) ch(A \Box \rightarrow C) = ch(A \Box \rightarrow C | A)$$

¹¹ In the context of a variably strict semantics for subjunctives, (vc6) is valid if the actual world is at least as close to itself as any other world, and (vc7) is valid if the actual world is closer than any other. For further discussion of centering assumptions, see Lewis 1973a, p. 29ff.

¹² In more detail: when you are certain that $ch(A \Box \rightarrow C) = ch(A \Box \rightarrow C | A)$, we can conclude:

$$\begin{aligned} \varepsilon(ch(A \Box \rightarrow C)) &= \int_{r=0}^1 r \cdot \mathcal{C}(ch(A \Box \rightarrow C) = r) = \int_{r=0}^1 r \cdot \mathcal{C}(ch(A \Box \rightarrow C | A) = r) \\ &= \varepsilon(ch(A \Box \rightarrow C | A)) \end{aligned}$$

For instance, dropping a plate normally makes no difference to the objective chance that the plate would break if dropped. Similarly, flipping a coin normally makes no difference to the objective chance that the coin would land heads if flipped. It is difficult even to construct a realistic imaginative picture of a coin that would land heads if flipped just in case I flip it, and would land tails if flipped otherwise. In fact, in most normal cases, a subjunctive and its antecedent are not only objectively independent, i.e., independent relative to the current objective chances, but evidentially independent, i.e. independent relative to your current credences. For instance, when I flip an ordinary coin, the information that I flip the coin intuitively does not give you any evidence about how likely it is that the coin would land heads if I were to flip it.

Though it is not directly relevant to the present argument, it is worth noting that even in a large number of less normal cases where a subjunctive is not evidentially independent of its antecedent, (INDEPENDENCE) still holds. For example, consider the following simplified version of a case from Kaufmann 2004.¹³ Urn A contains 49 red balls with spots and one black ball. Urn B contains one red ball without spots and 49 black balls. One of the two urns has been placed in front of you. This urn was chosen according to the outcome of a fair coin toss: Urn A is in front of you just in case the coin landed heads. Consider the following conditional:

(18) If I were to draw a red ball, I would draw a ball with spots.

Intuitively, (18) is equivalent with the statement that Urn A rather than Urn B is in front of you, which is equivalent with the statement that the fair coin landed heads. So you should have .5 credence in (18). If you were to learn just that you will draw a red ball, you could rationally take this as evidence that Urn A is in front of you, and raise your credence in (18) accordingly. Hence (18) is not evidentially independent of its antecedent.¹⁴ However, (INDEPENDENCE) still governs the objective

¹³ Kaufmann says that his case is a variation of a case suggested by Dorothy Edgington in conversation.

¹⁴ Some have argued that cases similar to Kaufmann's challenge Adams' Thesis for indicative conditionals, e.g. see McGee 2000 and Morton 2004. Rothschild 2010 notes that an indicative conditional will satisfy Adams' Thesis if it is evidentially independent from its antecedent, e.g. $C(R \square \rightarrow S) = C(R \square \rightarrow S \mid R)$. Rothschild claims that in cases where evidential independence fails, it is intuitive to say that Adams' Thesis fails also. If this claim is correct, then one could respond to the first triviality result in Lewis 1976 by restricting Adams' Thesis to cases where evidential independence holds. The claim is contentious, though; emulating Gibbard 1981, one might object that our intuitive rejection of Adams' Thesis in such cases is based on fallacious reasoning.

chance of (18). Let R be the proposition that you draw a red ball, and let S be the proposition that you draw a ball with spots. In the Kaufmann case, there are two ways the world could be: either Urn A is in front of you, or Urn B is. In the former case, we have $ch(R \square \rightarrow S) = ch(R \square \rightarrow S \mid R) = 1$. In the latter case, $ch(R \square \rightarrow S) = ch(R \square \rightarrow S \mid R) = 0$. In either case, (INDEPENDENCE) holds.

Accordingly, (BRIDGE) also holds for your credences in the Kaufmann case. The conditional chance of S given R is 1 just in case $R \square \rightarrow S$. The conditional chance of S given R is 0 just in case it is not the case that $R \square \rightarrow S$. Hence your expectation of the conditional chance of S given R should match your expectation of the truth value of $R \square \rightarrow S$, just as (BRIDGE) dictates. In a similar fashion, (BRIDGE) holds in a number of cases where the antecedent of a conditional simultaneously gives you information about what conditionals are true and what the objective conditional chances are.¹⁵

4. A Complication Concerning (BRIDGE)

(INDEPENDENCE) intuitively holds for many ordinary chance functions, and the principle is useful in deriving instances of (BRIDGE). But there is a caveat that complicates the present discussion: an exception to (INDEPENDENCE) might arise when the antecedent of a subjunctive conditional has no chance of occurring, as the conditional chance of its consequent given its antecedent may then be undefined. This exception creates a problem for (BRIDGE): if a subject has some credence that the antecedent of a future subjunctive $A \square \rightarrow C$ has chance 0, her expectation of the conditional chances $ch(A \square \rightarrow C \mid A)$ and $ch(C \mid A)$ would also then be undefined. Williams 2010a notes that one could respond to such potential counterexamples by positing primitive conditional chances or by limiting the domain of (BRIDGE):

It's perfectly consistent to take conditional chances to be well-defined even if the conditioned proposition is chance zero. [...] And even if the relevant conditional chances in some hard cases were in fact undefined, so that the conditional chance norm was inapplicable, that wouldn't mean that it was inaccurate within its domain of applicability. (8)

Even if one limits the domain of applicability of (BRIDGE), though, the possibility of chance 0 antecedents still poses an underappreciated challenge for the principle. The relevant challenge is that virtually no

¹⁵ For a more familiar case, consider the standard Newcomb problem (cf. Nozick 1969), where the fact that you will take one box is evidence that you would get rich if you did, and also evidence that the conditional chance of your getting rich if you take one box is 1 rather than 0.

subject is *certain* that the antecedent of a subjunctive has some chance of occurring, and that therefore the limited domain of application of (BRIDGE) would be virtually empty.

A more complete response is needed. The advocate of (BRIDGE) has several respectable options. First, as Williams notes, she may posit primitive conditional chances. Second, she may say that (BRIDGE) governs your credences insofar as you approximate a subject who is certain that whether A is true has not yet been settled. For instance, suppose that for practical purposes, you should act as if you are certain that A has some chance of occurring. Then for practical purposes, you should act as if (BRIDGE) constrains your credences. Third, she may endorse a variant of (BRIDGE) that merely *bounds* the credences of ordinary subjects. For example, even if your expectation of the conditional chance of C given A is not well-defined, your credence in ' $A \square \rightarrow C$ ' may still be bounded by the credences of subjects who think the conditional chance of C given A is some extreme value, i.e. either 0 or 1.

To spell out this last suggestion: suppose that you have credence α that A has no chance of occurring, and that your remaining credence is divided among various chance hypotheses in \mathcal{H} . Then on reflection, your credence may intuitively be constrained to satisfy the following principle, which we may call (BOUND):

$$\sum_{H_i \in \mathcal{H}} \mathcal{C}(H_i) \cdot \varepsilon(\text{ch}(C|A)|H_i) \leq \mathcal{C}(A \square \rightarrow C),$$

$$\mathcal{C}(A \square \rightarrow C) \leq \sum_{H_i \in \mathcal{H}} \mathcal{C}(H_i) \cdot \varepsilon(\text{ch}(C|A)|H_i) + \alpha,$$

where $\varepsilon(x | P)$ is just the expected value of $x \in [0,1]$ given your credences conditional on P .¹⁶ For example, suppose you have .05 credence that it is already determined that you will not flip a certain coin, and .95 credence that there is a chance you will flip the coin and that it is .5 likely to land heads if you do. Conditional on the latter chance hypothesis, the expected conditional chance that the coin will land heads if flipped is a well-defined value, namely .5. (BOUND) entails that your credence that the coin would land heads if flipped should be near this value:

$$.95(.5) + .05(0) = .475 \leq \mathcal{C}(F \square \rightarrow H) \leq .525 = .95(.5) + .05(1)$$

In a similar manner, we may clarify our earlier discussion of Centering. I said that the second familiar premise in the derivation of (BRIDGE) was “roughly” that you accept axioms (vc6) and (vc7). Strictly speaking,

¹⁶ In other words: $\varepsilon(x | P) = \int_{r=0}^1 r \cdot \mathcal{C}(x = r | P)$.

deriving (BRIDGE) requires that you accept these axioms with certainty. In response, the meticulous theorist should endorse a qualified version of (BRIDGE). She may say that (BRIDGE) governs your credences insofar as you approximate a subject who is certain that Centering holds. Or she may say that your credence in Centering determines the degree to which your subjunctive credences are bounded by your expectations of conditional chances.

This discussion highlights why principles like (BRIDGE) do not directly yield substantive constraints on your credence in past subjunctives, if corresponding primitive conditional chances are indeed not well-defined. The antecedent of a past subjunctive usually concerns some past event that did or did not occur, and you are usually well aware of this fact when you evaluate a past subjunctive. In fact, in many cases, you may be confident that the antecedent of a past subjunctive did not obtain. In this respect, you will not at all approximate a subject who is certain that the antecedent of a subjunctive has some chance of occurring. And (BOUND) will not impose any substantial constraint on your credences. As you evaluate the present subjunctive conditionals in §1, by contrast, (BOUND) does impose a substantial constraint on your credences.

As a side note, the fact that (BOUND) does not directly constrain credences in past subjunctives does not mean that the principle has no implications for your credence in such conditionals.¹⁷ In considering a past subjunctive, it is important to remember that claims like (INDEPENDENCE) and (BRIDGE) may have both held at an earlier time, constraining your earlier credence in what you later express using the past subjunctive. As long as you have not gotten information relevant to the truth of the subjunctive in the meantime, your credence in that conditional should match your earlier credence, and therefore match your earlier expectation of the conditional chance of its consequent given its antecedent. For example: suppose that just a moment ago, you decided not to make a bet on the outcome of a fair coin toss. You have tossed the coin but have not yet looked to see how it landed. Let us suppose that your earlier utterance of (19) has the same truth conditions as your later utterance of (20):

(19) If I were to bet on heads, I would win my bet.

(20) If I had bet on heads, I would have won my bet.

¹⁷ I am grateful to an anonymous referee for prompting me to clarify the import of my discussion for credences in past subjunctives.

(BRIDGE) says that before you failed to bet on heads, you should have had .5 credence in (19). In the meantime, you have gotten no information relevant to its truth. Hence you should later have .5 credence in (20). On this proposal, (BRIDGE) yields constraints on credences about past events in much the same way as the Principal Principle does. For example: the Principal Principle says that before you flip a fair coin, you should have .5 credence that it will land heads. If you flip the coin and do not look at it, you get no information about how the coin landed. Hence you should later have .5 credence in the claim that the coin landed heads.

Of course, intuitively, you can usually get information relevant to the truth of (20) after a coin is flipped. Simply looking at the coin and seeing that it landed heads may give you relevant information; indeed, it may demonstrate that the past subjunctive is true. Slote 1978 and Edgington 2003 claim that in this sort of case—originally due to Sydney Morgenbesser—your later rational credence in a past subjunctive comes apart from your earlier rational credence in the corresponding future subjunctive. The present suggestion is that one can account for this claim about the Morgenbesser case by saying that your later credence in the past subjunctive results from conditionalizing your earlier credence in that subjunctive on the observed result of the actual coin toss. By contrast, in a case where you bet on tails and a different coin would have been flipped if you had bet on heads, learning the result of the actual coin toss intuitively does not give you relevant information about the truth of (20), and so (BRIDGE) may still inform your credence in (20).

Let us survey where we stand. In evaluating the subjunctives mentioned in §1, it is reasonable to assume that the centering axioms are valid and that the objective chance function satisfies (INDEPENDENCE). But it may not be reasonable to endorse either of these assumptions with certainty. Hence strictly speaking, our credences are constrained by bounds determined by our credence in Centering and (INDEPENDENCE), or they are constrained to the degree to which we approximate subjects that are certain of these assumptions. In what follows, for ease of exposition, I will adopt two assumptions: that subjects are in fact certain that Centering holds, and that they are certain that antecedents of subjunctives are not determinately false. I rely on the reader to bear in mind that one may ultimately correct for these assumptions in the manner just described. Having established this caveat, we may conclude: the Superintendent Principle, an expectation of (INDEPENDENCE), and an expectation of Centering are all justified premises concerning our subjunctive credences in (1b), (2b), and (8). As we have demonstrated, it follows that instances of (BRIDGE) constrain

our subjunctive credences in each case. The derivation in §3 thereby confirms the ordinary language judgments that challenge many extant theories of subjunctives. Conditionals such as (8) do not call for an error theory of subjunctive conditionals:

- (8) If I were to flip this coin one million times, it would land heads at least once.

They call for a theory that accommodates our reluctance to say whether conditionals like (8) are true or false. In particular, we should prefer theories consistent with our systematic and justified acceptance of (BRIDGE) as a constraint on our credences in the subjunctive conditionals considered so far.

5. Subjunctive Triviality and the Limits of (BRIDGE)

The triviality results given by Lewis 1976 demonstrate that indicative conditionals do not express propositions such that your credence in a conditional proposition matches your conditional credence in the consequent given the antecedent.¹⁸ One might worry that (BRIDGE) implicitly commits us to the dubious claim that subjunctive conditionals cannot express propositions.¹⁹ Suppose that I am rational and that I am certain of the objective chances. Let A and C be arbitrary propositions. By the Principal Principle, the objective chance that $A \Box \rightarrow C$ matches my credence that $A \Box \rightarrow C$. By (BRIDGE), that credence matches the conditional chance of C given A . So the objective chance that $A \Box \rightarrow C$ matches the conditional chance of C given A . But for non-trivial chance distributions, subjunctive conditionals cannot express propositions such that the chance of each proposition matches the conditional chance of the consequent given the antecedent. The triviality proof is just as for indicatives, but with chance distributions rather than credence distributions as the relevant probability functions.²⁰

With Edgington and Skyrms as notable exceptions, it is generally accepted that subjunctives do express propositions. The potential objection to my argument is as follows: if that claim is correct and this argument is a *reductio* of (BRIDGE), then theories of conditionals should not

¹⁸ See Hall & Hájek 1994 for a helpful catalog of a variety of triviality results.

¹⁹ Edgington 1995 expresses the related worry that if we must take the same attitude toward subjunctive conditionals as toward previously uttered indicatives, triviality results about the latter demonstrate that the former attitude is not an attitude toward a proposition.

²⁰ This objection was first brought to my attention by Bob Stalnaker in conversation. Here I discuss the objection as developed by Robbie Williams in Williams 2010a.

be faulted for yielding verdicts about ordinary subjunctives that are inconsistent with (BRIDGE). For instance, in his discussion of a closely related principle, Williams 2010a suggests that it is not reasonable to expect theories to yield the verdicts entailed by claims like (BRIDGE), since “it will be *very hard* to satisfy the Ramsey Bounds for any conditional over a wide range of antecedents and consequents—so hard that consensus opinion in the indicative debate is that the enterprise is quixotic” (16).

However, the argument in §3 shows why we should be responsible for accommodating the instances of (BRIDGE) highlighted by our judgments in §1. (BRIDGE) is a general schema with many substitution instances, and some of these instances are indeed justified. The §3 derivation of (BRIDGE) depends on three substantive premises: the Superintendent Principle, an expectation of (INDEPENDENCE), and an expectation of Centering. As we evaluate the subjunctives discussed in §1, we are justified in accepting each of these premises, and the instances of (BRIDGE) that they entail.²¹ Conversely, as we will see, our rejection of each of these three premises patterns with our rejection of (BRIDGE). And the context of the triviality proof is just one of those contexts in which (BRIDGE) and its grounding premises are not justified.

For starters, cases in which we reject the Superintendent Principle are cases in which we should intuitively reject (BRIDGE) as well. For example, as mentioned in §3, the Superintendent Principle fails in cases where a subject has *inadmissible* information, in the sense of Lewis 1980. Intuitively, (BRIDGE) also fails in such cases. As Williams 2010a acknowledges, if a trusted oracle tells you that a certain fair coin would land heads if it were flipped, you may raise your credence in that conditional without changing your expectations of conditional objective chances involving the coin. For instance, you may remain confident that the coin is fair, so that your expectation of the conditional chance of it landing heads if flipped remains .5. In response to such examples, Williams endorses a qualified version of (BRIDGE) that ranges only over cases where subjects lack inadmissible information.

The Superintendent Principle may also fail for subjects who are simply skeptical about whether there are any such things as objective chances. Such subjects may not have any well-defined expectations for objective chances. Since the Superintendent Principle constrains credences according to expectations for objective chances, it seems inappropriate to apply the principle in cases where the latter values are plainly undefined. This restriction on the principle arises naturally if we take the objective chance function to play the role of an expert

²¹ I am grateful to David Manley for helping me get clear on the dialectic here.

in constraining your credences, as suggested in Gaifman 1988, van Fraassen 1989, and Hall 1994. If someone does not think that there are any experts on a certain subject, her credences should not be constrained by her estimates of current expert opinion on that subject. Here again, (BRIDGE) intuitively fails along with the Superintendent Principle. The skeptic about objective chance may have well-defined credences in subjunctives, though her expectations of conditional objective chances are not well-defined. Or more carefully, skepticism about chances simply introduces another respect in which both the Superintendent Principle and (BRIDGE) must be qualified: both constrain your credences only insofar as you give credence to the claim that objective chances exist.²²

Second, standard centering assumptions may be counterintuitive in some cases, and such cases may count as additional exceptions to (BRIDGE). For example, Lewis 1973a imagines his interlocutor protesting that Centering intuitively fails for subjunctives with unrelated antecedents and consequents, saying “what would we make of someone who saw fit ... to assert that if the sky were blue then grass would be green?” (28). Suppose that the interlocutor convinces you that ‘if the sky were blue, then grass would be green’ is false. Suppose you are certain that it is currently the case that the sky is blue and grass is green. Then your expectation of the conditional objective chance that grass is green if the sky is blue is 1. But you should not be certain of the corresponding subjunctive. By rejecting Centering for this case, you also thereby reject (BRIDGE). As in cases where subjects have inadmissible evidence or have doubts about the existence of objective chances, such failures of (BRIDGE) are easily understood and circumscribed.

Finally, it may be argued that (INDEPENDENCE) has exceptions. Here is an example due to John Hawthorne in conversation: suppose that it is unlikely that you perform a certain physical movement *M* tomorrow, though in the unlikely event that you contract a rare disease *D*, the chance of your performing *M* is high. Suppose also that the combination of contracting *D* and performing *M* causes death. Then many judge that the objective chance of ‘if you were to perform *M* tomorrow, you would die’ is low, but the conditional

²² This restriction need not severely limit the scope of (BRIDGE). For one thing, it seems unlikely that many ordinary subjects would avow skepticism about the existence of objective chances. And if they did, they might still engage in enough ordinary talk about chances so that on balance, charity would demand that we deny their self-ascriptions of skepticism. I am grateful to an anonymous referee for prompting me to discuss skepticism about chances and to clarify this point.

objective chance of this subjunctive given that you perform M is high. Insofar as (INDEPENDENCE) intuitively fails in this case, (BRIDGE) fails as well. Since you are fairly certain that you do not have the rare disease, your credence in ‘if you were to perform M , you would die’ should be low. But since you are certain that your death is objectively likely given that you perform M , your expectation of the conditional objective chance of death given that you perform M should be high.

In addition to this exception, one can imagine fanciful chance functions that could in principle provide exceptions to (INDEPENDENCE) and (BRIDGE). For instance, suppose that you are given a coin and told that the following conditional has objective chance $\frac{2}{3}$: that if you were to flip it, the coin would land tails. But the coin is extremely strange: its disposition to land tails is extreme if you flip it, and eliminated if you do not. In other words, the objective chances are as follows, where F is the proposition that you flip the coin and H is the proposition that it lands heads:

$$\begin{aligned} ch(F \wedge F \square \rightarrow H) &= 0 \\ ch(F \wedge \overline{F \square \rightarrow H}) &= ch(\overline{F} \wedge F \square \rightarrow H) = ch(\overline{F} \wedge \overline{F \square \rightarrow H}) = \frac{1}{3} \end{aligned}$$

If you are certain of these facts about the objective chances of F , H , and $F \square \rightarrow H$, then intuitively, (BRIDGE) should not constrain your credence that $F \square \rightarrow H$. It is perfectly compatible with your beliefs that in fact it will turn out to be the case that $F \square \rightarrow H$. Since you are certain that the objective chance of that conditional is $\frac{1}{3}$, you should have $\frac{1}{3}$ credence in the conditional. And since you are certain about what the conditional objective chances are, the expected value of those chances is also straightforwardly defined. Since you are certain that it is *objectively certain* that the coin will not heads if flipped, the expected value of $ch(H | F)$ should be 0 rather than $\frac{1}{3}$.

6. Answering Subjunctive Triviality

Having discussed the limits of (BRIDGE), let us return to the subjunctive triviality result outlined earlier. For simplicity, define conditional objective chance functions $ch'(x) = ch(x | C)$ and $ch''(x) = ch(x | \overline{C})$. Let \mathcal{C}' and \mathcal{C}'' be rational credence distributions of hypothetical subjects that are certain that ch' and ch'' give the actual objective chances, respectively. Here is a slight adaptation of the subjunctive triviality proof in Williams 2010a:

$$\begin{aligned}
& ch(A \square \rightarrow C) \\
&= ch(A \square \rightarrow C|C)ch(C) + ch(A \square \rightarrow C|\overline{C})ch(\overline{C}) && \text{probability theory} \\
&= ch'(A \square \rightarrow C)ch(C) + ch''(A \square \rightarrow C)ch(\overline{C}) && \text{probability theory} \\
&= \mathcal{C}'(A \square \rightarrow C)ch(C) + \mathcal{C}''(A \square \rightarrow C)ch(\overline{C}) && \text{(Principal Principle)} \\
&= ch'(C|A)ch(C) + ch''(C|A)ch(\overline{C}) && \text{(BRIDGE)} \\
&= 1 \cdot ch(C) + 0 \cdot ch(\overline{C}) && \text{probability theory} \\
&= ch(C) && \text{calculation}
\end{aligned}$$

(BRIDGE) is applied twice in the proof: once to replace $\mathcal{C}'(A \square \rightarrow C)$ with $ch'(C | A)$ and once to replace $\mathcal{C}''(A \square \rightarrow C)$ with $ch''(C | A)$. But given our more complete understanding of (BRIDGE), it is clear that these instances of (BRIDGE) are not justified. In particular, (INDEPENDENCE) does not always hold for the constructed chance functions ch' and ch'' , and so expectations about these chances may not always constrain hypothetical credences in ' $A \square \rightarrow C$ '.

For example, let F be the proposition that you flip an ordinary fair coin, and let H be the proposition that the coin lands heads. Suppose that there is a .5 objective chance that you will flip the coin. The intuitive objective chances in this case are not hard to calculate:

$$\begin{aligned}
ch(F \wedge F \square \rightarrow H) &= \frac{1}{4} \\
ch(F \wedge \overline{F \square \rightarrow H}) &= \frac{1}{4} \\
ch(\overline{F} \wedge F \square \rightarrow H) &= \frac{1}{4} \\
ch(\overline{F} \wedge \overline{F \square \rightarrow H}) &= \frac{1}{4}
\end{aligned}$$

(INDEPENDENCE) holds for this chance function. But consider the chance function that results from conditionalizing the chances on the proposition that it is not the case that the coin will land heads. The resulting objective chances are exactly those described in the final exception to (INDEPENDENCE) in the previous section:

$$\begin{aligned}
ch(F \wedge F \square \rightarrow H) &= 0 \\
ch(F \wedge \overline{F \square \rightarrow H}) &= ch(\overline{F} \wedge F \square \rightarrow H) = ch(\overline{F} \wedge \overline{F \square \rightarrow H}) = \frac{1}{3}
\end{aligned}$$

In other words, according to the resulting objective chance function, the coin's disposition to land tails is extreme if you flip the coin, and eliminated if you do not. Whether the coin would land heads if flipped is not objectively independent of whether you flip the coin. As we have

seen, (BRIDGE) intuitively fails in such cases, and so it fails here, in the context of the proof of the triviality result.

As we have seen, the class of chance functions for which (INDEPENDENCE) holds is not closed under conditionalization. Hence (BRIDGE) may constrain the credences of a subject certain that ch describes the objective chances, while not constraining the credences of a subject certain that ch'' describes the objective chances. The existence of such exceptions to (BRIDGE) does not undermine the fact that instances of (BRIDGE) hold in a wide range of ordinary cases. I should clarify the dialectic here: ultimately it does not matter for my argument that any of the cases considered in §5 are genuine exceptions to (BRIDGE). Rather, what matters is that the premises that ground our acceptance of (BRIDGE) are reasonable in many ordinary cases. Hence theories of subjunctive conditionals remain responsible for accommodating our acceptance of justified instances of (BRIDGE). Semantically immodest theories yield verdicts that are incompatible with justified instances of (BRIDGE), and such verdicts are a significant cost of such theories.

For all I have said so far, it might be that any viable theory of subjunctives necessarily yields verdicts incompatible even with justified instances of (BRIDGE). It is beyond the scope of the present paper to defend any particular theory of subjunctives consistent with justified instances of (BRIDGE). But the §3 derivation may aid in the search for such theories, namely by demonstrating that as long as a theory accommodates the premises that ground our acceptance of (BRIDGE), the theory will accommodate our intuitions as well. And several recent theories of subjunctives indeed meet these conditions. For instance, the simplest strategy for accommodating justified instances of (BRIDGE) is briefly advocated in Hawthorne 2005: simply accepting the limit and uniqueness assumptions and giving a closest-worlds semantics for subjunctive conditionals. In light of these assumptions, our opinions or ignorance about a subjunctive conditional are simply opinions or ignorance about the character of the closest world where the antecedent is true.

Other theories of subjunctives aim to account for (BRIDGE) without positing metaphysically brute facts about similarity. For instance, Schulz 2010 argues that ' $A \Box \rightarrow C$ ' is true if and only if C is true at an arbitrarily selected relevant A world. Drawing on the radical semantics of arbitrary reference in Magidor & Breckenridge 2011, Schulz advances the theory that a particular arbitrarily selected world is uniquely relevant to the truth value of a subjunctive, and that our principled ignorance of which world is arbitrarily selected is responsible for our uncertainty about the truth value of that subjunctive. A second metaphysically conservative but semantically profligate family of

strategies makes use of the sort of constructions found in van Fraassen 1976. In particular, van Fraassen says that we can construct extended chance functions over possible worlds supplemented with similarity orderings, which we can then use to define truth values for subjunctives at ordinary possible worlds. Williams 2010c gives a helpful catalog of non-standard semantic theories that make use of such constructions. For example: one might say that ‘if I were to flip this coin, it would land heads’ is .5 true, provided that the coin lands heads in the closest world where I flip it according to exactly half of the similarity orderings supplementing the actual world.²³ Theories in this spirit avoid strong metaphysical assumptions, but they do give up on the project of giving a standard truth-conditional semantics for subjunctive conditionals.

In light of these alternatives, it seems that the best argument for rejecting justified instances of (BRIDGE) consists of the following pair of claims: first, that any theory that accommodates instances of (BRIDGE) will necessarily be either metaphysically or semantically profligate, and second, that metaphysical and semantic conservatism are theoretically more valuable than accommodating systematic and justified ordinary language judgments. Neither claim may be taken for granted. For instance, Swanson 2010 outlines a metaphysically and semantically conservative theory that constitutes a challenge to the first claim. Williamson 1994 challenges conservative theories about what ingredients may participate in determining the truth conditions of ordinary sentences; his arguments for epistemicism pave the way for semantically profligate theories that say that the context of utterance may contribute a unique similarity ordering to the truth conditions of a subjunctive. And van Fraassen 1974 suggests that one might justifiably accept metaphysically profligate theories of subjunctives on the grounds that they “add explanatory power *vis-à-vis* the relevant language game” (189). The upshot of this discussion is that a complete theory of subjunctives must be not only a semantic theory, but a metaphilosophical theory as well. The debate over the correct theory of subjunctive conditionals cannot be settled by dismissing (BRIDGE) as unjustified or by explaining away its instances with reinterpretation strategies. On the contrary, the debate ultimately concerns the relative value of various virtues of first-order theories of subjunctive conditionals.

References

Bennett, Jonathan. 2003. *A Philosophical Guide to Conditionals*. Oxford University Press, Oxford.

²³ Briggs 2010 outlines a theory in this spirit, drawing on proposals in Gärdenfors 1982 and Stalnaker & Jeffrey 1994.

- Briggs, Rachael. 2010. "Two Interpretations of the Ramsey Test." Ms., University of Sydney.
- Edgington, Dorothy. 1995. "On Conditionals." *Mind*, vol. 104: 235–329.
- . 2003. "Counterfactuals and the Benefit of Hindsight." In *Causation and Counterfactuals*, P. Dowe & P. Noordhof, editors, 12–27. Routledge, London.
- von Fintel, Kai. 2001. "Counterfactuals in a Dynamic Context." In *Ken Hale: A Life in Language*, Michael Kenstowicz, editor. MIT Press, Cambridge.
- van Fraassen, Bas C. 1974. "Hidden Variables in Conditional Logic." *Theoria*, vol. 40: 176–190.
- 1976. "Probabilities of Conditionals." In *Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science*, W. Harper & C.A. Hooker, editors. D. Reidel, Dordrecht.
- 1980. "A Temporal Framework for Conditionals and Chance." *Philosophical Review*, vol. 89: 91–108.
- 1989. *Laws and Symmetry*. Oxford University Press, Oxford.
- Gaifman, Haim. 1988. "A Theory of Higher Order Probabilities." In *Causation, Chance, and Credence*, Bryan Skyrms & William Harper, editors. Kluwer, Dordrecht.
- Gärdenfors, Peter. 1982. "Imaging and Conditionalization." *Journal of Philosophy*, vol. 79 (12): 747–760.
- Gibbard, Allan. 1981. "Indicative Conditionals and Conditional Probability: Reply to Pollock." In Harper et al. (1981), 253–255.
- Gillies, Thony. 2007. "Counterfactual Scorekeeping." *Linguistics and Philosophy*, vol. 30: 329–360.
- Hájek, Alan. 2007. "Most Counterfactuals are False." Ms., Research School of the Social Sciences, Australian National University.
- Hall, Ned. 1994. "Correcting the Guide to Objective Chance." *Mind*, vol. 103 (412): 505–518.
- Hall, Ned. & Alan Hájek, 1994. "The Hypothesis of the Conditional Construal of Conditional Probability." In Skyrms & Eells (1994), 75–110.
- Harper, William L., Robert Stalnaker & Glenn Pearce, editors. 1981. *Ifs: Conditionals, Belief, Decision, Chance, and Time*. D. Reidel Publishing Company, Dordrecht.
- Hawthorne, John. 2005. "Chance and Counterfactuals." *Philosophy and Phenomenological Research*, 70 (2): 396–405.
- Joyce, James M. 1999. *The Foundations of Causal Decision Theory*. Oxford University Press, Oxford.
- Kaufmann, Stefan. 2004. "Conditioning Against the Grain." *Journal of Philosophical Logic*, vol. 33 (6): 583–606.
- Lewis, David K. 1973a. *Counterfactuals*. Basil Blackwell Ltd., Malden, MA.

- 1973b. “Counterfactuals and Comparative Possibility.” In Harper et al. (1981), 57–85.
- 1976. “Probabilities of Conditionals and Conditional Probabilities.” In Lewis (1986), 305–339. With postscript.
- 1980. “A Subjectivist’s Guide to Objective Chance.” In Lewis (1986), 83–132. With postscript.
- 1981. “Causal Decision Theory.” In Lewis (1986).
- 1986. *Philosophical Papers*, vol. 2. Oxford University Press, Oxford.
- Lindström, Sten & Włodzimierz Rabinowicz. 1995. “The Ramsey Test Revisited.” In *Conditionals: From Philosophy to Computer Science*, Gabriella Crocco, Luis Farinas & Andreas Herzig, editors, 147–191. Oxford University Press, Oxford.
- Magidor, Ofra & Wylie Breckenridge. 2011. “Arbitrary Reference.” Ms., Oxford University. Forthcoming in *Philosophical Studies*.
- McGee, Vann. 2000. “To Tell the Truth About Conditionals.” *Analysis*, vol. 60 (1): 107–111.
- Morton, Adam. 2004. “Against the Ramsey Test.” *Analysis*, vol. 64: 294–299.
- Moss, Sarah. 2010. “On the Pragmatics of Counterfactuals.” Forthcoming in *Noûs*.
- Nozick, Robert. 1969. “Newcomb’s Problem and Two Principles of Choice.” In *Essays in Honor of Carl G. Hempel*, Nicholas Rescher, editor, 114–146. D. Reidel Publishing Company, Dordrecht.
- Rothschild, Daniel. 2010. “Do Indicative Conditionals Express Propositions?” Ms., Department of Philosophy, Oxford University. Forthcoming in *Noûs*.
- Rott, Hans. 1999. “Moody Conditionals: Hamburgers, Switches, and the Tragic Death of an American President.” In *JFAK. Essays Dedicated to Johan van Benthem on the Occasion of his 50th Birthday*, Jelle Gerbrandy, Maarten Marx, Maarten de Rijke & Yde Venema, editors. Amsterdam University Press, Amsterdam.
- Schulz, Moritz. 2010. “Counterfactuals in a Chancy World.” Ms., Department of Philosophy, Oxford University.
- Skyrms, Bryan. 1980. “The Prior Propensity Account of Subjunctive Conditionals.” In Harper et al. (1981), 259–265.
- 1994. “Adams’ Conditionals.” In Skyrms & Eells (1994), 13–26.
- . 1998. “Subjunctive Conditionals and Revealed Preference.” *Philosophy of Science*, vol. 65 (4): 545–574.
- Skyrms, Bryan. & Ellery Eells, editors. 1994. *Probabilities and Conditionals: Belief Revision and Rational Decision*. Cambridge University Press, Cambridge.

- Slote, M. 1978. "Time in Counterfactuals." *Philosophical Review*, vol. 87: 3–27.
- Stalnaker, Robert C. & Richard C. Jeffrey. 1994. "Conditionals as Random Variables." In Skyrms & Eells (1994), 31–46.
- Swanson, Eric. 2010. "Conditional Excluded Middle without the Limit Assumption." Forthcoming in *Philosophy and Phenomenological Research*.
- Williams, J. Robert G. 2008. "Chances, Counterfactuals, and Similarity." *Philosophy and Phenomenological Research*, vol. 77 (2).
- 2010a. "Counterfactual Triviality: A Lewis-Impossibility Argument for Counterfactuals." Ms., Department of Philosophy, University of Leeds. Forthcoming in *Philosophy and Phenomenological Research*.
- 2010b. "Defending Conditional Excluded Middle." *Noûs*, vol. 44 (4): 650–668.
- 2010c. "Tenable Conditionals." Ms., Department of Philosophy, University of Leeds.
- Williamson, Timothy. 1994. *Vagueness*. Routledge, London.