

Locating Practical Normativity

by

David J. Plunkett

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Philosophy)
in The University of Michigan
2010

Doctoral Committee:

Professor Allan F. Gibbard, Co-Chair
Professor Peter A. Railton, Co-Chair
Emeritus Professor Stephen L. Darwall
Assistant Professor Scott A. Hershovitz
Professor Frank C. Jackson, Princeton University
Professor Scott J. Shapiro, Yale University

© David J. Plunkett

2010

Acknowledgments

My work in this dissertation has benefited enormously from a wide range of conversations and comments. I would particularly like to thank Allan Gibbard, Peter Railton, Stephen Darwall, Frank Jackson, Scott Shapiro, Scott Hershovitz, Andy Egan, Nadeem Hussain, Sarah Buss, Tristram McPherson, David Braddon-Mitchell, Michael Smith, Sharon Street, Daniel Nolan, Tim Sundell, Kenny Easwaran, Howard Nye, John Ku, Tom Dougherty, Elisa Mai, Katia Vavoa, Daniel Stoljar, Tamar Schapiro, Alex Silk, Richard Holton, Peter Galison, Nishi Shah, Dave Chalmers, Eduardo-Garcia-Ramirez, Fritz Warfield, Ben McKean, Melissa Barry, Michael Bratman, Matt Kotzen, Chandra Sripada, Agustin Rayo, Adam Plunkett, Neil Mehta, Mark Schroeder, Sam Liao, Alejandro Perez Carballo, Christine Korsgaard, Peter Gordon, Arnold Davidson, Boris Kment, Louisa Thomas, Michelle Kosch, David Velleman, Julia Markovits, Kristie Miller, Paolo Santorio, Mike Titelbaum, Rachel Briggs, Adina Roskies, Robbie Williams, Damien Rochford, Lina Jansson, Aaron Bronfman, Sally Haslangar, Moises Vaca Paniagua, Derek Baker, Josh Foer, Alex Plakias, Ralph Wedgwood, Nic Southwood, Marie Jayasekera, Steve Campbell, David Mahfouda, Justin Reynolds, Alex Gourevitch, Brett Lockspeiser, Richard Tuck, Raymond Geuss, Taylor Carman, David Dick, Alexa Forrester, Michael Friedman, Gabriel Zamos-Regueros, Jeff Sebo, Jay Wallace, Patrick Tomlin, Laura Valentini, Jamie Tappenden, and Gordon Belot. I presented some of the material in this dissertation at talks given at the Australian National University RSSH philosophy department in July 2009, the 2009 Australasian Association of Philosophy Conference, and the University of Minnesota, Duluth in November 2010. Thanks to everyone who asked questions and gave comments at those talks.

I had the great fortune of writing this dissertation in many different wonderful locations. This was in large part made possible by family and friends. I would particularly like to thank Marcy Plunkett, James Plunkett, Marty Peretz, Anne Peretz, Dick Sheahan, Dottie Sheahan, Caroline Meyer, Judy Sabot, Kenny Easwaran, and Ziva Cooper.

Table of Contents

Acknowledgments	ii
Abstract	iv
Chapter One: Can Metaethical Non-naturalism be Made Science-Friendly?	1
Chapter Two: The Metaethical Role of Evaluative Attitudes	49
Chapter Three: Ethical and Metaethical Agent-Attitude-Dependence	94

Abstract

A central feature of ethical thought is that it appears to involve not only descriptive belief, belief about what *is* the case, but also normative belief about what *should* be done. Suppose we take this at face value and understand normative thought in ethics to consist of attitudes that, at the most basic explanatory level, are genuine beliefs. What then should we say about the basic nature of the normative properties that such beliefs are about? I argue that normative properties are complex naturalistic properties of psychology. In the first chapter, I consider the *non-naturalistic realist* position, according to which our world contains the instantiation of irreducibly normative, metaphysically *sui generis* properties. I argue that proponents of non-naturalistic realism have not successfully shown that this view is compatible with confidence in the claims and methodologies of the natural sciences. This gives us powerful (if ultimately defeasible) reason to reject this view. In the second chapter, I consider metaethical *ideal attitude theory*, exemplified in the work of Michael Smith, according to which normative properties about what an agent A should do concern what an ideal version of A would desire that non-ideal A do. In order a) to maintain a naturalistic account of normative properties, b) to avoid radical skepticism about ethical knowledge, and c) to explain the motivational force of normative judgment, I argue that ideal attitude theorists should hold that what it is for an agent A to be ideal is derived from A's own evaluative attitudes. I call this a fully *agent-attitude-dependent* version of ideal attitude theory. In the third chapter, I consider Sharon Street's recent arguments in favor of *metaethical constructivism*, according to which normative properties concern what is entailed by an agent's practical standpoint. I argue that Street's metaethical constructivism is best developed as a version of agent-attitude-dependent ideal attitude theory.

Chapter One:
Can Metaethical Non-naturalism be Made Science-Friendly?

Introduction.

Many of my ethical beliefs strike me as obviously correct. For instance, I find it obviously correct that I should not poke puppies in the eyes for fun and that I should help to save another's life if it involves only a minimal cost to myself. Many of these beliefs are the same ones that other people find to be obviously correct as well. Yet, many people throughout history have also had ethical beliefs – i.e. beliefs about how agents *should* live or act – that they have taken to be obviously correct that I think are badly mistaken. These include, for instance, beliefs that I think are sexist, racist, homophobic, and/or chauvinistic. I also know that, had I been born and raised in a different social-historical context – for instance, in a religious community in 15th-century England rather than in a secular community in late 20th-century America – I would likely have had many of those same strongly held ethical beliefs that I condemn today. Moreover, I probably would have thought that some of those beliefs were obviously correct. Because of this, the sheer fact that an ethical belief now strikes me as obviously correct won't be able to separate those cases where I think people got things wrong from my own current case – it seems reasonable to hold that, in both cases, the basic phenomenology involved will be the same. What, then, should give me confidence that I am getting things right now?

One thing that we want from ethical theory is an answer to this question. Or, to put it somewhat more precisely, one thing that we want is some justification for having confidence in the subset of our ethical beliefs that we think we *should* be confident in – for instance, the belief that I should not poke puppies in the eyes for fun.

Perhaps our hopes here will be dashed. For instance, perhaps we will discover through careful ethical reasoning that most of our current strongly held ethical beliefs are incorrect. Or, more drastically, perhaps we will discover that there are *no* ethical judgments that could be correct in our world – perhaps because all ethical judgments presuppose the

instantiation of irreducibly normative properties of a sort that simply are not instantiated in our world. Yet, insofar as we are creatures who actively engage in ethical thinking about how we should act and live – something that is in practice inescapable for creatures like us – we can't help but to think that we are correct in *some* of our fundamental ethical beliefs that we inevitably must reason from in making ethical judgments, or at least to act under the hypothesis that this is so. Moreover, insofar as we try to form new ethical beliefs, we can't help but to think that we are not *hopeless* in coming to form new ethical beliefs that could also be correct. Vindicating such theses is part of what will make a theory in ethics compatible with the perspective we inevitably adopt in practical reasoning. It is thus, we might say, part of what a theory in ethics must do in order for it to be *practically adequate* for creatures such as ourselves.

At the same time, we want our overall understanding of ethics to be not only practically adequate but also *theoretically adequate*. We want it to be one that – in broad terms – is acceptable not only from the perspective we adopt in practical reasoning about what to do and how to live, but also from a theoretical perspective from which we consider how our ethical thought and talk fit into our broader understanding of reality. For many of us – and likely for most people reading this paper – part of this understanding involves a basic confidence in the natural sciences. At the core of this confidence is both the acceptance of certain theses about the way things are – theses that we take to be supported by the contemporary natural sciences – as well as the more basic idea that, whatever problems the natural sciences have, they nonetheless do a generally good job of studying those aspects of reality with which they are concerned. At the same time, this basic confidence in the natural sciences involves some level of confidence in the *import* of the central claims of the natural sciences – a confidence, we might say, that even if scientific claims could never exhaust all there is to be known about reality, they nonetheless represent a privileged place in our overall account of how things are. If a theory of ethics is going to be theoretically adequate, it must be one that respects this basic confidence in the central claims and methodologies of the natural sciences. Call this the *naturalistic constraint* on ethical theories. Call those ethical theories that meet the naturalistic constraint *science-friendly* ethical theories.

The naturalistic constraint will not always have direct import for theories in ethics. Most importantly, take theories that concern *substantive* first-order questions about how agents should live and act. Such theories might address specific questions about how agents

should live and act in specific circumstances (e.g. the theory that Susie should not steal Bugsy's motorcycle) or general questions about how agents should live and act in general (e.g. the theory that all agents should always act to maximize well-being). Following one standard convention, call such theories ones in *normative ethics*. Such theories will not necessarily aim to answer questions about why we should be confident in ethical thought in general – questions, for instance, that we might be worried about in light of wondering how ethical thought and talk is compatible with a basic confidence in the central claims and methodologies of the natural sciences. Hence, the naturalistic constraint will not necessarily impact (or at least not *directly* impact) these theories in normative ethics.

In contrast, now consider *metaethical* theories that, in basic terms, take ethical thought and talk as an *object of inquiry* that needs to be explained. In broad terms, we can say that *metaethical* theories concern the question of how ethical thought and talk – and the ontological commitments involved in that thought and talk – should be integrated into a broader conception of reality.¹ Metaethical theories need to be both practically and theoretically adequate. Or, more precisely, metaethical theories need to be both practically and theoretically adequate *if* we want to avoid radical skepticism about our ability to form correct ethical judgments.²

¹ For a similar gloss on the task of metaethics see (McPherson Forthcoming, 1). Insofar as the task I have here identified concerns *all parts* of ethical thought and talk, I take this task to be what we can designate as the *fundamental* question of metaethics – as opposed to what we can designate as *non-fundamental* questions about how different parts of ethical thought and talk are related to each other. I thank Howard Nye for first suggesting to me this basic way of distinguishing between the fundamental and non-fundamental tasks of metaethics. For Nye's own way of drawing this basic distinction, see (Nye 2009). In some cases, the term "metaethics" is taken simply to *mean* what I am here designating as the "fundamental" task of metaethics. For instance, see (Hussain and Shah 2006). In this paper, I will use the term "metaethics" to refer to what I am here calling the "fundamental" task of metaethics even though, in contrast to Hussain and Shah, I take both fundamental and non-fundamental questions to be genuinely *metaethical* ones insofar as both types of question concern, in the first instance, how to make sense of ethical thought and talk as an object of inquiry rather than the issue of which ethical thoughts are correct.

² It should be stressed that, as I am using the term here, *ethics* involves the broad practical questions of how agents should live and act *full stop*. To put it another way, ethics concerns how one should answer these questions *simpliciter* – and not, for instance, how one should answer these questions from a distinctively *moral* point of view (which one might or might not take to be definitive for answering the former questions depending on how one understands the meaning of the term "moral"). In turn, when I use the term "metaethics," I intend to keep this use of the term "ethics" fixed. Thus, when I use the term "metaethics," I intend to signal that I am concerned with the project of taking *ethical* thought and talk as an object of inquiry rather than that of taking specifically *moral* thought and talk as an object of inquiry.

In this paper, I am going to argue that one prominent type of theory in contemporary metaethics – namely, a familiar type of metaethical view that I will call *non-naturalistic realism* – does not successfully accomplish this task. In short, I will argue that we do not have sufficient evidence that non-naturalistic realism can be developed into a metaethical view that is both practically adequate and that respects the naturalistic constraint. Insofar as we want a metaethical theory that can do so – something that most contemporary non-naturalists agree we should want – I conclude that my argument gives us a strong (if ultimately defeasible) reason to reject metaethical non-naturalistic realism.

I will break up my work in this paper into four main sections. In §1, I clarify which version of metaethical non-naturalism is my target in this paper. In broad terms, my target is the metaethical view that holds a) that there are irreducibly normative properties that are metaphysically *sui generis*, b) that such properties are instantiated in our world, and c) that normative judgments express straightforward truth-apt beliefs about the instantiation of these properties. Following one standard convention in metaethics, I call the conjunction of these three claims *non-naturalistic realism*. In §2, I introduce a central question that all theories in metaethics face: namely, are normative properties themselves part of the causal-explanatory order? This question gives us a basic way of distinguishing between two different types of non-naturalistic realist metaethical theories: *non-causal* versions, which answer “no” to this question, and *causal* versions, which answer “yes” to it.

In §3 and §4, I argue that proponents of both non-causal and causal versions of non-naturalistic realism face significant challenges in showing how non-naturalistic realism can be made into a science-friendly version of the view. Moreover, I argue that none of the current proposals on offer by non-naturalistic realists successfully meet these challenges in full. On the one hand, in §3, I argue that non-causal non-naturalists have yet to deliver a naturalistically viable account of how creatures like us could ever have epistemic access to these non-causally efficacious properties. Moreover, I argue that non-causal non-naturalists face a significant challenge in explaining what science-friendly grounds there are for this sort of ontological commitment. On the other hand, in §4, I argue that causal non-naturalists lack a good account of what makes their view a genuine version of *non-naturalism* about normative properties in the first place. Furthermore, I argue that causal non-naturalism runs into trouble with what I argue is one of the core substantive ideas many have about the natural

sciences; namely, the idea that the natural world is *causally closed*.³ Thus, I argue that either way metaethical non-naturalism is developed, it runs into trouble with delivering a metaethical theory that is both practically adequate and respectful of the *naturalistic constraint*. Such problems might be capable of being solved. But I argue that we have yet to see from non-naturalists how they can be solved. Furthermore, I argue that solving these problems would involve significant work beyond what contemporary non-naturalistic realists have done so far. This provides powerful (if ultimately defeasible) reason to hold that non-naturalistic realism fails to meet one of the central aspirations we should have for metaethical theories: namely, to provide a metaethical theory that is both theoretically and practically adequate.

1. *Metaethical Non-naturalism.*

When contemporary philosophers use the term “non-naturalism” to describe a view in metaethics, they usually intend to signal that the view in question has the following, broad feature: it, in some basic way, asserts a fundamental un-bridgeable divide between ethics and the natural sciences. Not all so-called “non-naturalistic” views agree, however, on what this “un-bridgeable divide” is. In this paper, I am going to be concerned exclusively with one of the standard – one might say the *canonical* – ways of articulating what this divide amounts to.

In order to clarify *which* basic type of non-naturalist view I am concerned with in this paper, we can consider three different important so-called “non-naturalist” views in metaethics: those of G.E. Moore, Allan Gibbard, and T.M. Scanlon. These views are representative of three broad versions of non-naturalism in contemporary metaethics and help to bring out some of the important distinctions that will matter for my argument in this paper.

Start with Moore’s view in *Principia Ethica* from the turn of the 20th-century.⁴ When Moore introduced the term “non-naturalism” to describe his metaethical view, he understood the term to capture something important about how he understood the

³ Much of what I have to say in this fourth section of the paper draws extensively on both personal communication with Hussain as well as a 2008 presentation he gave in Jerusalem entitled “Normativity and the Causal Order” (Hussain 2008). See also (Hussain Forthcoming).

⁴ (Moore 1993).

metaphysics of normativity. According to Moore, ethical beliefs are straightforward truth-apt beliefs that aim to correctly represent a genuine feature of reality. Yet, he argued, unlike scientific beliefs, which aim to correctly represent features of the natural world, ethical beliefs aim to correctly represent a completely different part of reality: a distinct realm of what we can call normative *properties*. Such properties, claimed Moore, are *irreducibly normative* properties different in kind than any properties studied by the natural or social sciences such as *being an elephant* or *weighing five pounds*. Furthermore, he claimed that they are different in kind from any *supernatural* properties such as *being commanded by God* or *being disapproved of by Zeus*. They are, he insisted, *non-naturalistic* properties that are fundamentally different in kind from properties of any other sort.⁵ Thus, we can say that, for Moore, normative properties are metaphysically of a kind: they are metaphysically *sui generis*.

In contrast to Moore – who aimed to explain the fundamental gulf between ethics and the natural sciences in terms of a difference in the *metaphysics* of their respective subject matters – Allan Gibbard has recently attempted to do so in terms of a difference in the *concepts* one uses in the sciences and those one uses in ethics. According to Gibbard in his recent *Thinking How to Live*, there are no metaphysically *sui generis*, irreducibly normative properties of the sort that Moore envisioned. However, Gibbard claims that there are irreducibly normative *concepts* that we use in ethical thought.⁶ These concepts, he claims, cannot be fully analyzed in non-normative terms. Furthermore, these concepts are fundamentally different *in kind* from what he calls the *naturalistic* concepts that we use in the natural sciences, concepts such as WHALE or QUARK.⁷ They are, we might say, *sui generis* concepts. When this core thesis is developed into his larger expressivist account of normative thought and talk, Gibbard maintains that he can explain the unbridgeable gulf that Moore envisioned between ethics and the natural sciences – and, moreover, that he can do so without taking on any of Moore’s more ambitious metaphysical claims.

⁵ I use the term “non-naturalistic” here rather than Moore’s original “non-natural” in order to avoid confusion with the “natural” vs. “non-natural” properties distinction that is used in metaphysics (for instance, that is used when discussing so-called “grue-some” properties). See (Lewis 1983). This distinction between “non-naturalistic” and “non-natural” will become especially important later in this dissertation.

⁶ (Gibbard 2003).

⁷ Following one standard convention, I will henceforth use smallcaps to designate concepts. On this convention, see (Margolis and Laurence 1999).

Like Moore and Gibbard, Scanlon argues that there are – in some sense – irreducibly normative truths.⁸ Both Moore and Gibbard hold that non-naturalists face a significant challenge in *explaining* the possibility of the divide between ethics and the natural sciences, and moreover, that this challenge must be met by taking on robust commitments about the special nature of normative properties or concepts. However, in contrast to Moore and Gibbard, it is not entirely clear that Scanlon takes his belief in irreducibly normative truths to involve taking on any major commitments of this sort. Some of what he says would suggest that he does. Yet, in many of his explicit statements about the nature of metaethics, Scanlon seems to want to *undercut* or to *dissolve* the perceived need for such explanations in order to maintain a belief in irreducibly normative truths. This suggests a reading of Scanlon on which he rejects the thesis endorsed by both Moore and Gibbard that metaethical non-naturalism involves taking on robust theses about the special nature of normative properties or concepts. Instead, it suggests that Scanlon accepts what we might call a *quietist* position with respect to the fundamental questions in metaethics.⁹

In this paper, I am going to be concerned exclusively with non-quietist versions of metaethical non-naturalism. Furthermore, I am going to be concerned exclusively with the non-expressivist versions of the view according to which the purportedly unbridgeable divide between ethics and the sciences is explained in terms of the instantiation of properties that are irreducibly normative and that can only be studied by normative inquiry. Moreover, in order to differentiate the *non-naturalistic* position here from a form of non-reductive naturalism, on which the claim that normative properties are irreducible might be no more radical than holding that the chemical or biological properties are irreducible, I will be concerned exclusively with metaethical views that hold that this irreducibility is entailed by the deeper fact that normative properties are metaphysically *sui generis*. This claim that normative properties are metaphysically *sui generis* amounts to the claim that the normative properties are different from any other sort of property (including, the non-naturalist asserts, *all* of the naturalistic properties that science is concerned with) and is, as I see it, the heart of Moore's original style of non-naturalism.

⁸ (Scanlon 2003) and (Scanlon Forthcoming).

⁹ For a reading of Scanlon along these lines, see (McPherson Forthcoming). See (McDowell 1998) for another prominent quietist position of this sort.

For Moore, this claim about *sui generis* properties was intended as a claim about *what there is* in ultimate metaphysical reality – and not just, for instance, what there is *from our perspective* given certain necessary features of cognition shared by all finite creatures. In other words, this was intended as a metaphysical thesis in fundamental ontology *full stop* and not just one in something such as “phenomenological ontology” or “Kantian critical ontology.” In this paper, I am going to be concerned exclusively with versions of non-naturalism that share this commitment and, more generally, that share the same basic structure as Moore’s original position.¹⁰

With this in mind, I can now state more precisely the form of metaethical non-naturalism with which I will be concerned in the rest of this paper. I will be concerned with views that endorse the following three theses:

Sui Generis Properties. The fundamental normative properties are *irreducibly normative* properties that are metaphysically *sui generis*.¹¹

Descriptivism. To make a normative judgment is to express an attitude that, at the most basic explanatory level, consists in straightforward truth-apt belief.

Realism. Some normative properties are instantiated in the actual world. This makes it the case that some possible normative judgments that concern the actual world are correct.¹²

¹⁰ Tristram McPherson argues in (McPherson Forthcoming) that all non-naturalists who believe in irreducible normative truths but who deny expressivism face a basic metaphysical burden (even if they want to call it something else) that can’t be avoided by a “quietest” route of the sort he attributes to Scanlon. I agree. If this is right, then my arguments in this paper – or at least my arguments in some slightly modified form – also apply to Scanlon’s sort of position as well. However, it is beyond the scope of this paper to argue for the claim that Scanlon’s quietism does not succeed as a way to sidestep robust commitments in ontology.

¹¹ By a “fundamental” normative property, I mean a normative property that is metaphysically basic in terms of its relationship to other normative properties. More specifically, it is a normative property that can be used to explain the nature of other, further non-fundamental normative properties but that itself cannot be explained by the nature of any other normative properties. For our purposes here, it does not matter what normative properties one takes to be basic in this sense (a sense, it should be noted, that Moore was after when he called goodness a “simple” property). For instance, it won’t matter whether or not one takes the class of fundamental normative properties to include the properties of *being good* and *being a normative reason* and *being a duty*, or whether one thinks only one of these properties (or another property altogether) is needed to explain the basic nature of these other properties.

¹² It should be noted that I use the phrase “concern the actual world” rather than “about the actual world” to leave room for Gibbard’s expressivist view according to which normative judgments express plans about what to do in the actual world rather than judgments that are strictly speaking *about* the world. It should also be noted that if realism is true, then correct normative judgments that

Call the conjunction of these three theses the metaethical view of *non-naturalistic realism*. Within the last decade, Derek Parfit, Jonathan Dancy, Ralph Wedgwood, David Enoch, and Russ Shafer-Landau have all defended metaethical views that – although they are often presented under different labels – are versions of the view that I am characterizing here as “non-naturalistic realism.”¹³

As stated, one could in principle adopt a *non-naturalistic realist* view to account for many different parts of normative thought and talk. In this paper, given that I am concerned with *metaethics* and not *metanormative* theory as a whole, I will restrict myself to the parts of normative thought and talk that concern the basic ethical questions of how agents should live and act. Because of this, when I use the term “normative,” I will usually implicitly take this to be restricted to the domain of ethics. For instance, when I use the phrase “normative properties,” I will implicitly take this to be restricted to the normative properties about how one should live and act – excluding, for instance, normative properties about how one should form beliefs and normative properties about how one should assign meanings to speakers’ utterances.

In order to underscore what is distinctive about non-naturalistic realism as a basic type of metaethical view, it is worth pausing here to clarify what I take each of these three theses to amount to. Let’s start with the *sui generis* properties thesis. In order to clarify what this thesis amounts to – in particular, how it is different from Gibbard’s thesis that there are *sui generis concepts* – let me say a bit about the basic way that I will use the term “property” and “fact” in this paper (as well as in the other papers in this dissertation). For the purposes of differentiating different metaethical views, we do *not* need – nor should we want – a fully developed substantive theory of what properties or facts are. Rather, what we need is a basic schematic understanding of these things that allows us to fruitfully differentiate between different sorts of metaethical views regardless of the initial labels that are used by their proponents to present those views.

concern the actual world will be *non-trivially* correct in virtue of the fact that there are instantiated normative properties in the actual world. In contrast, imagine that there are no normative properties instantiated in the actual world. If so, then the statement “Andy does not have a normative reason to eat the pizza” is trivially correct in the actual world.

¹³ See (Parfit 2006), (Parfit Forthcoming), (Dancy 2006), (Wedgwood 2009), (Enoch 2007), and (Shafer-Landau 2005).

Let's start with the term "property." Following one convention in contemporary philosophy, one exemplified in Frank Jackson's *From Metaphysics to Ethics*, I will use the term "property" to refer to the following: the features of an object in virtue of which that object does or does not fall under a given concept.¹⁴ To see what I have in mind here, consider the concept CAT. When an object falls under this concept – something reflected in the fact that the predicate "is a cat" correctly applies to it – it does so because of features of that object above and beyond the simple fact that it falls under the concept. Call those features – whatever those are – the property of *being a cat*. To generalize from this, we can say that the A-type properties (or "A-properties" for short) concern the features of an object in a possible world in virtue of which that object falls under A-type concepts.

With this basic schema on the table we can, in turn, distinguish different types of properties – for instance, "normative properties," "scientific properties," "chemical properties," or "physical properties" – based on the parts of thought, discourse, and/or practice where those properties are referenced and/or play a genuine explanatory role. After investigation, we might find that certain properties are in fact reducible or identical to other properties – for instance, that the chemical properties studied by chemistry can be reductively explained in terms of the physical properties studied by physics. Nothing in my use of the term "normative properties" is meant to rule out that such a possibility might be correct in terms of the relationship between normative properties and non-normative ones. Instead, as I see it, much of the debate in metaethics concerns this sort of metaphysical issue – and, moreover, it is precisely this issue that the *sui generis* properties thesis addresses. According to the *sui generis* properties thesis, the fundamental normative properties are irreducible *in virtue of* the fact that normative properties are metaphysically *sui generis*.

As Jackson emphasizes in putting forward this basic type of schema, there are lively debates in metaphysics about how to understand the nature of properties themselves. Such debates include, for example, the question of whether or not one should be a so-called

¹⁴ For more on the use of the term "property" in this way, see (Jackson 1998, 15-16) and (Chalmers and Jackson 2001). As Jackson underscores, it should be noted that a similar minimal schema is also needed for the term "relation" – roughly, a connection between different objects. Since the issues involved are essentially the same – and since the framework I lay out in what follows could be modified to cover features of relations as well as features of objects – I will stick with Jackson's lead and focus on the question of properties (understood, as he does, as features of objects rather than relations). See (Jackson 1998, 15-16). In so doing, I do not mean to privilege the idea that the fundamental account in normative metaphysics – nor in metaphysics more generally – will privilege properties rather than relations.

“realist” or a so-called “nominalist” about properties. Like Jackson, as I see it, such debates are best understood as concerning questions about the metaphysics of properties over which non-naturalistic realists in metaethics can reasonably differ. For, as I see it, the heart of the *sui generis* properties thesis is the claim that there are certain *sui generis*, irreducibly normative features of a world in virtue of which our normative concepts in ethics ultimately apply or not. Debates about how to best understand what it is for something to be a genuine feature of the world are thus a further matter. Thus, in principle, even though it would be an uncommon position, one could even adopt a Quinean skepticism about the language of “properties” and still accept the *sui generis* properties thesis given my broad definition of the term “property” here.

One might worry that my use of the term “property” in this way stacks the deck in favor of naturalism about normative properties. To see why one might have this worry, consider that most contemporary philosophers accept some version of a *supervenience* thesis about the relationship between the non-normative, naturalistic properties and the normative properties (of the sort that we invoke in ethics). There is a lively debate about which supervenience thesis should be accepted by all sides – and, indeed, whether there is in fact a non-controversial one that can be taken as a starting point at all.¹⁵ For our purposes here, we can put the thesis as a *global* supervenience thesis about the relationship between different possible worlds: no two possible worlds are identical with respect to their non-normative properties while differing in their normative properties. Call this the *ethical supervenience thesis*.¹⁶ One of the key challenges that non-naturalistic realists face is to explain why such a tight metaphysical connection holds between two radically different types of properties. For our purposes here, however, there is a more basic issue at hand. This issue is this: given that most non-naturalistic realists accept the ethical supervenience thesis, one might worry that it makes it hard to see how it is even conceptually coherent for them to buy the *sui generis* properties thesis given my account of properties. If this sort of global supervenience really holds between normative and non-normative properties, then, as Jackson argues in *From Metaphysics to Ethics*, there is no way to learn about the existence of a purportedly additional normative property by pulling it apart modally from its supervenience base – even

¹⁵ For some of the important different views on this debate, see (Jackson 1998), (Smith 1994), (Dancy 2006), (Sturgeon 2009), and (Wedgwood 2009).

¹⁶ For a similar way of putting forward the thesis that the normative properties in ethics supervene on the non-normative, see (McPherson Manuscript).

using a two-dimensional framework where one considers possible worlds (roughly, a complete description of the way things are) both as counterfactual possibilities and as alternatives to actuality (e.g. as alternative worlds from which we evaluate the counterfactual possibilities from). Since this is so, Jackson argues that there is good reason to doubt that there really is an additional property here at all – the cost of saying otherwise, he thinks, might well be an unprincipled proliferation of extra properties.¹⁷ More generally, in rough terms, Jackson argues that if two purportedly distinct features of the world can *never* be modally pulled apart on a two-dimensional framework, we lack good reason to posit that there really are two distinct properties here in the first place. Thus, one might think that, when combined with the *ethical supervenience thesis*, my Jackson-style definition of properties stacks the deck in favor of precisely the sort of *naturalism* about normative properties that the non-naturalistic realist denies.

However, this thought is in fact mistaken. Instead of stacking the deck in favor of the naturalist, my way of framing things in fact helps bring out an underlying metaphysical issue that is crucial for the debate between naturalists and non-naturalists about normative properties. In order to block Jackson’s conclusion, the non-naturalist about normative properties does not need to reject either my proposed definition of properties or the ethical supervenience thesis. Instead, the non-naturalist could offer a different account of *property individuation* than Jackson’s – namely, one in which global supervenience of the B-properties on the A-properties on a two-dimensionalist framework does *not* guarantee property identity.

In turn, with such a view of property individuation in hand, the non-naturalist would have a good way of defending a thesis that Jackson rightly identifies as a central component of Moore’s original non-naturalism: namely, the thesis that “what is left of language after we cull the ethical terms is in principle inadequate to the task of ascribing the properties we ascribe using the ethical terms.”¹⁸ Call this the *inadequacy thesis*. The reason that the inadequacy thesis would hold, is because, in short, on such a proposed view of property individuation, the normative properties could be genuine features of reality above and beyond their supervenience base – and, in turn, the non-naturalistic realist could potentially hold that it is only the ethical terms in our language that refer to these properties instead of the non-normative terms that refer only to the supervenience base of the normative

¹⁷ (Jackson 1998).

¹⁸ (Jackson 1998, 121).

properties. Such a view would be importantly different from Gibbard's version of metaethical non-naturalism – according to which there are no genuine features of reality that only ethical terms refer to, but instead *sui generis* normative concepts. Thus, rather than stacking the deck in favor of the naturalist about normative properties, my Jackson-inspired way of talking about “properties” helps bring out what the debate *is* between non-naturalists and non-naturalists about normative properties – and, furthermore, how this debate is different from the one that Gibbard is engaged in about the nature of normative concepts.¹⁹

Whether the view of property individuation that the non-naturalistic realist needs here is defensible is part of a much larger debate in metaphysics. Thus, even if there are strong reasons to think that it is plausible – for instance, concerns about preserving the intuitive difference between the properties of *being trilateral* and *being triangular* – it might be that, in the end, Jackson's style of property individuation will win out. However, the point here is simply that these are part of the further debates that matter for the non-naturalistic realist in metaethics – and, thus, non-naturalistic realism is not ruled out as in any way *conceptually incoherent* given the minimal definition of properties I have given so far.

Let's now turn to the term “facts.” In order to understand how I will use this term in the context of differentiating different types of metaethical theory, let me start by saying something about how I will differentiate different *types* of facts. Following the lead of David Chalmers, I will hold that “The *A-facts* and *B-facts* about the world are the facts concerning the instantiation and distribution of A-properties and B-properties.”²⁰ Combined with the

¹⁹ It should be stressed that my way of putting the issues involving properties is perfectly compatible with Gibbard's version of non-naturalism. Since Gibbard marks the bifurcation between normative and non-normative thought in terms of *concepts* rather than *properties*, he can freely grant that there is a sense of the term “property” according to which a given property is a feature of the world in virtue of which a given concept applies or not. In turn, he just needs to deny – as he does – that the nature of normative concepts is exhausted by the property that the concept picks out. In fact, this is precisely what Gibbard does when he denies that normative concepts are the sorts of things that should be understood, as Jackson proposes, in terms of the function from possible worlds to extensions that the concept effects. As with the view of property individuation that non-naturalistic realist needs, it is a further question whether or not Gibbard's views of the nature of normative concepts and concept-individuation are defensible or not. Indeed, it is a further question whether or not his views are even conceptually coherent given our concept CONCEPT. I discuss these issues at length in (Plunkett Forthcoming). My point here is simply that nothing in how I am using the term “properties” settles this issue.

²⁰ (Chalmers 1996, 33). It should also be noted that the relevant notion of “world” here in Chalmers' characterization of facts is a *possible world* – which can be understood, roughly, as a *complete description of the way things are*. Since such a description could include descriptions of both concrete and abstract objects, nothing in his use of the term “world” here is meant to signal a privileging of something

basic Jackson-style definition of properties I have given, this means that, as I am using the term “facts” in this paper, A-facts are the substantive, non-conceptual considerations in virtue of which A-judgments are correct or incorrect. This type of definition of facts – in which they are understood as the “correctness-makers” for correct judgments – is different from other ways in which one might want to use the term (for instance, to refer to the correct judgments themselves, to refer to things that play a causal-explanatory role, etc.). For our purposes in this paper, I do not intend to take any stand on which other uses of the term “facts” are also appropriate. Rather, all that matters is that using the term “facts” in the way that I propose is helpful for clarifying the core issues at hand. Given this way of using the term “facts,” we can say that, according to the non-naturalistic realist, the *normative facts* (facts concerning the instantiation of normative properties) are irreducible to the *non-normative facts* (facts concerning the instantiation of non-normative properties) because normative properties are *sui generis*.

With a better understanding of the *sui generis* properties thesis in hand, let’s now turn to the second thesis that I am taking to be part of the constitutive definition of non-naturalistic realism. To repeat, this is:

Descriptivism: To make a normative judgment is to express an attitude that, at the most basic explanatory level, consists in straightforward truth-apt belief.

Expressivists such as Gibbard and Blackburn deny this claim.²¹ Instead, they hold that the attitudes that we express in normative judgment should be understood, at the most basic explanatory level, as desire-like attitudes with a world-to-mind direction of fit rather than a mind-to-world direction of fit. Gibbard argues that, through a careful explanation of the constraints involved in expressing such attitudes, expressivists can “earn the right” to meaningfully talk, at least in a so-called minimalist way, of normative thoughts as consisting in *beliefs* with representational content. In turn, according to a minimalist account of truth-aptness, they might also “earn the right” to think of these beliefs as genuinely truth-apt.²² Non-naturalistic realists of the sort I am interested in here do not use an indirect pattern of explanation such as Blackburn’s and Gibbard’s to “earn the right” to talk of normative

such as the *natural world*.

²¹ (Gibbard 2003) and (Blackburn 1998).

²² (Gibbard 2003).

judgments expressing truth-apt beliefs. Rather, non-naturalistic realists hold that normative judgments express attitudes that, *at the most basic explanatory level*, are truth-apt beliefs.

Now consider the last thesis of *realism*. To repeat, this thesis is as follows:

Realism. Some normative properties are instantiated in the actual world. This makes it the case that some possible normative judgments that concern the actual world are correct.

There is a wide range of ways the term “realism” is used in metaethics – not to mention, of course, the range of ways the term is used in philosophy more broadly. As I use the term here, “realism” is intended to pick out a fairly minimal thesis.²³ Given that I am understanding normative properties simply as the features of objects that make normative judgments correct, we can say that the *realism* thesis amounts to the following: there are some features of how things are in the world that make it the case that some possible normative judgments that concern the actual world are correct. The crucial point is that, whether one likes this paraphrase in terms of properties or not, my minimal use of the term “realism” is one that makes the *realist thesis* compatible with a wide range of metaethical theories. For instance, as I am putting it forward, the realist thesis is compatible with both naturalism or non-naturalism about normative properties, and it is compatible with both expressivism and descriptivism about normative judgment. In effect, the thesis is simply the opposite thesis of that endorsed by *error theories* that, in one way or another, deny that there are any possible normative judgments that concern the actual world that are non-trivially correct.²⁴ Since this is the main contrast that matters for us here, as I see it, accepting this minimal thesis of realism does not in any way dictate what one believes about the epistemological issue of how we come to learn which normative judgments are correct – let alone the substantive issue of how many of our current normative judgments *made in* the actual world are correct.

As I see it, what unites normative realists is the affirmation that normative properties are instantiated in the actual world and, thus, the denial of an eliminativism about normative properties.

²³ My understanding of realism here is close to – but not identical to – Michael Smith’s way of putting the thesis of realism forward in (Smith 2004).

²⁴ I say “non-trivially” correct for the following reason that I noted earlier in a footnote (n.13) when I first introduced the thesis of realism. Imagine that the error theorist is right that there are no normative properties instantiated in the actual world. If so, then the statement “Andy does not have a normative reason to eat the pizza” is trivially correct.

Therefore, it should be stressed that accepting this minimal thesis of realism does not in any way dictate whether or not possible normative judgments that concern the actual world are made correct or not in virtue of *irreducibly normative* properties. This, as I see it, is simply a further debate between different theories of normative properties – just as, for instance, I believe that debates about whether instantiated chemical properties are irreducible or not is best understood as an intramural debate between realists about chemical properties rather than “realists” about chemical properties and their foes.²⁵

With this clarification of the realist thesis of non-naturalistic realism in mind, it is also worth noting here how I will take the realist thesis to line up with my use of the terms “naturalistic realism” and “non-naturalistic realism.” As I use the term “naturalistic realism,” the view consists in the conjunction of three theses: the descriptivist thesis, the realist thesis, and the thesis that normative properties are naturalistic properties. In other words, naturalistic realism is the same as non-naturalistic realism *except* that it denies the crucial *sui generis properties* thesis. As I am understanding the terms “non-naturalistic realism” and “naturalistic realism,” then, it is an essential part of both views that the realist thesis is true. In contrast, it is a further substantive normative question for the expressivist whether or not realism is true. This, I think, marks an important contrast between the views and explains part of the rationale for using the term “realism” in defining the metaethical views.²⁶

²⁵ I here differ with those who claim that a “realist” about X believes that X cannot be reductively explained in terms of anything else and, instead, holds that X is a basic part of the fundamental reduction base for everything else that exists. For use of the term “realist” in this way, see (Fine 2001), (Wedgwood 2009), and (Pettit 1991).

²⁶ It should be noted that, given how I understand the term “expressivism,” all the expressivist is committed to *in virtue of her expressivism* is the denial of descriptivism in favor of an alternative pattern of explanation according to which normative judgment expresses attitudes that, at the most basic explanatory level, are desire-like or plan-like attitudes. Of course, given the motivations for adopting expressivism in metaethics in the first place – including, importantly, a desire to preserve a purely naturalistic normative metaphysics – most expressivists in metaethics will also deny the *sui generis properties* thesis. But I do not think that this metaphysical commitment is itself part of what makes a view a specifically *expressivist* one. It would be a strange view indeed that combined the *sui generis properties* thesis with expressivism. However, my point here is just that I do not think that such a view would be disqualified as a genuine expressivist view in metaethics. It should also be noted that given how I am defining realism, I am not sure that there is good reason for expressivists such as Blackburn and Gibbard who endorse realism as a substantive normative claim to qualify that their developed views are “quasi-realist” ones. Instead, I think it might be more accurate to say that their views are genuinely realist ones on which realism is vindicated by a different pattern of explanation than is adopted by most proponents of realism. This, of course, is one way of reading what the thesis of “quasi-realism” amounts to. My point then is simply that the label “quasi-realism” can be

However, my use of the terms “naturalistic realism” and “non-naturalistic realism” also has a somewhat unfortunate consequence. Since “naturalistic realism” and “non-naturalistic realism” both refer to metaethical views that endorse both descriptivism and realism, one might naturally take it that the only way to endorse the realist thesis is to embrace descriptivism. However, given how I have defined the theses in this paper, this would in fact be a mistaken inference to make. Indeed, because of this, I think that my use of the terms “non-naturalistic realism” and “naturalistic realism” is somewhat misleading. I therefore ask the reader to simply charitably grant me the use of these terms in my current way for the purposes of this paper.

1.1 The Commitments of Science-Friendly Non-naturalistic Realism.

For much of the twentieth-century, metaethical non-naturalistic realism fell out of favor. Indeed, for much of the last part of the twentieth century, philosophers often quickly dismissed the central distinctive claim of non-naturalistic realism – namely, the *sui generis* properties thesis – with little argument. For instance, in 1977, J.L. Mackie dismissively wrote that “no doubt it was an extravagance for Moore to say that ‘good’ is the name of a non-natural quality.”²⁷

Many philosophers still share the same basic sentiment with regard to non-naturalistic realism that Mackie expressed in 1977. However, at the same time, non-naturalistic realism has undergone a remarkable revival in recent years. For instance, as I said earlier, within the last decade, Derek Parfit, Jonathan Dancy, Ralph Wedgwood, David Enoch, and Russ Shafer-Landau have all defended metaethical views that – although they are often presented under different labels – are versions of the view that I am characterizing here as “non-naturalistic realism.”²⁸ In so doing, these philosophers have developed sophisticated new ways of understanding how non-naturalistic realism might be developed into a systematic metaethical view and, moreover, sophisticated new ways of arguing on its behalf.

misleading. For examples of their use of the term “quasi-realism,” see (Blackburn 1998) and (Gibbard 2003).

²⁷ (Mackie 1977, 32).

²⁸ See (Parfit 2006), (Parfit Forthcoming), (Dancy 2006), (Wedgwood 2009), (Enoch 2007), and (Shafer-Landau 2005).

One of the central claims made by many philosophers in this non-naturalistic revival is that although non-naturalistic realism asserts a fundamental divide between normative properties and the properties studied in the natural and social sciences, metaethical non-naturalism is not itself fundamentally *in conflict* with the natural and social sciences. In particular, since this has traditionally seemed the harder thing to show, such non-naturalists have sought to underscore that their versions of non-naturalism are not in conflict with the natural sciences such as physics, chemistry, and biology. The basic claim of this sort of non-naturalism is summed up well by Wedgwood in *The Nature of Normativity* when he claims that his non-naturalistic position is one in which “there is no conflict with a modern naturalistic world-view.”²⁹ Such versions of non-naturalism claim to be science-friendly in the sense of being at least compatible with our basic confidence in the idea that the natural sciences do a generally good job at studying those aspects of the world that they are concerned with. Call this the *no conflict thesis*. Proponents of non-naturalism that endorse the no conflict thesis want to make two claims. On the one hand, in order to preserve metaethical non-naturalism, they need to deny the so-called “*scientistic*” thesis that scientific claims about naturalistic properties exhaust all there is to be known about reality. The non-naturalist needs to deny scientism because, as we have seen, a core part of the non-naturalistic realist view is that there are metaphysically *sui generis* properties instantiated in our world that cannot be studied by the natural sciences. On the other hand, in order to vindicate the no conflict thesis, science-friendly non-naturalists do not want to be *anti-science*. In basic terms, science-friendly non-naturalists aim *add* to the claims that the natural sciences make about reality without denying that the sciences are essentially correct in terms of their views about the parts of reality that they study and in terms of how they study it.

For our purposes in this paper, we can specify that the no conflict thesis consists of *at least* two basic claims: 1) non-naturalistic realism is compatible with our basic confidence in the *content* of contemporary scientific theories and 2) non-naturalistic realism is compatible with our basic confidence in the methodology that the contemporary natural sciences use to generate those theories. By “content of contemporary scientific theories,” I mean all the positive claims that those theories make about how things are. By “the methodology of the contemporary natural sciences,” I mean both the concrete means by which scientists study reality and the epistemic norms employed in this way of studying reality. The natural sciences

²⁹ (Wedgwood 2009, 197).

make a great number of different claims. Furthermore, as recent work in the history and sociology of science underscores, there is a great deal of diversity in many aspects of the methodologies followed in the contemporary natural sciences. This diversity plausibly includes differences in the basic epistemic norms employed in different parts of the natural sciences.³⁰ The heart of the *no conflict thesis*, I take it, is meant to concern the potential clash between the fundamental commitments of a broadly scientific view of the world and non-naturalistic realism. Hence, for our purposes here, we can specify that the *no conflict thesis* concerns *fundamental* aspects of the content of contemporary scientific claims and methodology – aspects of these respective things that, in broad terms, we regard as crucial to our current conception of what the natural sciences hold to be true about the world as well as to our current conception of how they operate methodologically.

Putting this together, we can thus more carefully state the *no conflict thesis* as the conjunction of the following two claims:

The No Content Conflict Thesis: non-naturalistic realism is compatible with the current fundamental content of the understanding of the natural world that we get from the natural sciences.

The No Methodological Conflict Thesis: non-naturalistic realism is compatible with the current fundamental methodological commitments of the natural sciences.

Put together, these claims represent a good way of fleshing out the central claim of science-friendly versions of non-naturalism: namely, that the position involves “no conflict with a modern naturalistic world-view.”³¹

Some proponents of science-friendly non-naturalism want to make a stronger claim on their behalf as well. In *The Nature of the Normativity*, Wedgwood writes that he wants to vindicate what he sees as “a plausible *naturalistic* conception of the world” according to which “there is a sense in which the natural facts determine the fundamental nature of the world.”³² Contrast this with Shafer-Landau’s more limited aspiration to provide “a plausible metaphysics, one that is compatible with, but not exhausted by, the picture of the world as given by the natural sciences.”³³ Shafer-Landau’s aspiration here can be seen as simply a

³⁰ On this theme, see (Galison and Stump 1996) and (Daston and Galison 2007).

³¹ (Wedgwood 2009, 197).

³² (Wedgwood 2009, 10).

³³ (Shafer-Landau 2005, 4).

reflection of a broad aspiration to meet the no content conflict thesis. In contrast, Wedgwood might be seeing as adding a third thesis to what it would take to have a metaethical position that on which there is “no conflict with a modern naturalistic world-view”³⁴: namely, that it must be a metaethical position that is consistent with the thesis that “there is a sense in which the natural facts determine the fundamental nature of the world.”³⁵ Many non-naturalistic realists – especially those that favor non-causal non-naturalism – might well balk at this statement. For, when read in a strong way, it might be something close enough to scientism to undermine their basic view from the start. Given this fact, I will not take Wedgwood’s thesis here to be itself part of the no conflict thesis that all science-friendly non-naturalists would need to vindicate.

Wedgwood’s claim, however, underscores something that I *do* think unites the recent proponents of science-friendly non-naturalism: namely, even if these non-naturalists would not all agree “that there is a sense in which the natural facts determine the fundamental nature of the world,”³⁶ they nonetheless *do* want to grant the claims of the natural sciences a central place in our understanding of reality. We can thus see different non-naturalist views as giving the claims of natural sciences more or less importance in terms of our overall picture of reality. It is unclear exactly what it would be to give them enough importance to meet the bar for being “science-friendly.” At bare minimum, however, we can say that it will *at least* involve preserving a picture of reality in which *supernatural* entities do not enter the picture and which would be consistent with plausible epistemic norms that rule out justified beliefs in such things. Thus, as Shafer-Landau puts it, the claim of science-friendly non-naturalism is that one can allow that there are “genuine features of our world that remain forever outside the purview of the natural sciences,”³⁷ namely the metaphysically *sui generis* normative properties, *without* opening the door to including “all sorts of oddities”³⁸ such as ghosts, demons, deities, or other supernatural creatures.

³⁴ (Wedgwood 2009, 197).

³⁵ (Wedgwood 2009, 10).

³⁶ (Wedgwood 2009, 10).

³⁷ (Shafer-Landau 2005, 4).

³⁸ (Shafer-Landau 2005, 4).

2. Normative Explanations.

One of the central questions confronting *any* metaethical theory is whether or not normative facts – that is, the facts in virtue of which normative judgments are ultimately correct or incorrect – are themselves causally efficacious. Given our understanding of normative facts as facts concerning the instantiation of normative properties, this question can be put as follows: are normative properties themselves causally efficacious or not? This question gives us a basic way of distinguishing between two different types of non-naturalistic realist metaethical theories: *non-causal* versions that answer “no” to this question and *causal* versions that answer “yes” to it.

Most contemporary philosophers working on metaethics think that *normative judgments* are causally efficacious. To see why, consider the following two cases:

Becoming Vegetarian. Kenny has long loved eating meat. However, after reading more about animal rights, he has come to believe that eating meat is wrong. “I should become a vegetarian,” he thinks to himself. Following this, he becomes a vegetarian.

Offensive Joke. Andy used to regularly tell a joke that he has referred to as “the best joke.” However, he has recently stopped telling this joke. When asked why, he replies “well, I came to believe that it is actually offensive to some people. I therefore decided that I should no longer tell the joke.”

When read in a straightforward way, such cases suggest that our beliefs about normative facts can be causally efficacious in the exact same way as other judgments. In turn, because of this, most philosophers working in metaethics hold that *if* judgments that express beliefs in general turn out to be capable of playing a causal role, then so too will be normative judgments in particular.³⁹

The situation concerning the causal role of *normative facts*, however, is more complicated. On the one hand, many philosophers find it compelling to think that normative facts themselves can play a role in causal explanations. Consider, for instance, the two following purported historical explanations given by Nicolas Sturgeon.⁴⁰

³⁹ For an argument that normative judgments *are* epiphenomenal, see (Velleman 2009).

⁴⁰ (Sturgeon 1988).

Hitler. It is a normative fact that Hitler was morally depraved. The fact that Hitler was morally depraved explains some of the things that he did.

Slavery-Abolition. Widespread and ultimately successful opposition to plantation slavery and the slave trade arose in the 18th and 19th centuries in Europe and North America. The reason that this opposition arose at those times and places was because slavery then and there was particularly bad compared to at other times and places.

As Sturgeon argues, such historical explanations seem compelling. There will, of course, be more complicated stories to tell in each case if one wants the full story of why those events happened. Yet, it nonetheless seems that these explanations capture *part* of why certain historical events happened. If so, then this supports thinking that normative properties (of the sort we are interested in when we do ethics) can play a causal-explanatory role.

On other hand, there are also compelling arguments that we do not need such facts themselves to explain why events happen. Consider, again, the case of *Slavery-Abolition*. We might want to use the fact that *slavery is wrong* to explain the events that happened. But it also seems likely that we could explain the events that happened solely in terms of people making certain normative judgments rather than others – for instance, making the judgment that slavery is wrong. In turn, it seems probable that one might be able to explain *why* people made these judgments purely on the basis of non-normative psychological, historical, cultural, sociological, and other such facts about these people. Both Gilbert Harman and Allan Gibbard have argued that this is, in fact, true for all normative judgments in ethics. Because of this, both Harman and Gibbard deny that any *irreducibly* normative facts (at least of the sort that matter in ethics) themselves play any sort of causal role.⁴¹

Philosophers that have argued for versions of metaethical non-naturalism in general – as well as non-naturalistic realism in particular – have taken positions on both sides of this debate. On the one hand, Wedgwood and Shafer-Landau have affirmed that normative properties play a causal-explanatory role.⁴² On the other hand, Parfit and Enoch have denied that they do.⁴³ There are, of course, specific reasons why each of these philosophers have endorsed such positions given the working details of each of their theories. There are,

⁴¹ See (Harman 1986), (Harman 1997), (Gibbard 1990), and (Gibbard 2003).

⁴² (Wedgwood 2009) and (Shafer-Landau 2005).

⁴³ (Parfit Forthcoming) and (Enoch 2007). Note that there are some compelling reasons to think that Parfit, like Scanlon, is best read as pursuing a version of the quietist non-naturalist strategy that I outlined at the start. I leave this possibility to the side for the purposes of this paper.

however, general issues that all non-naturalistic realists face that put pressure for taking one direction or another.

In the next section, I will explain what I think the pressure is for non-naturalistic realists to *deny* that normative properties play a causal role. In turn, I will then discuss the resulting sort of non-naturalistic realist position and argue that non-naturalists have yet to show how it can make good on the *no conflict thesis*. I will start with this non-causal position because I believe this position is more intuitively motivated by the core commitments of non-naturalistic realism as such. Indeed, if I am right, then one major motivation for the non-naturalistic realist to grant that normative properties *do* play a causal role stems from the perceived failure of the more intuitive position for the non-naturalistic realist.

3. *Non-causal Non-naturalism.*

In order to understand the pressure for the non-naturalistic realist to *deny* that normative properties play a causal role, consider again the *sui generis* properties thesis that defines non-naturalistic realism. As David Enoch puts it, the core intuition behind what I am calling the *sui generis* properties thesis is that normative properties are just “too different” to be reduced to *any* naturalistic property.⁴⁴ This brings out that in affirming the *sui generis* properties thesis, non-naturalists mean to assert a fundamental divide between the basic normative properties and *all* the properties that we study in the natural sciences. This divide is meant to be different from whatever the divide is between different sorts of naturalistic properties: for instance, whatever the divide is between the *biological* and *chemical* properties. According to the non-naturalist, the biological and chemical properties are, in some basic sense, continuous: they are both instances of the sort of *naturalistic* properties that we study in the natural sciences. In contrast, the *normative* properties are meant to be fundamentally different in *kind*: they are meant to be metaphysically *sui generis*.⁴⁵

⁴⁴ (Enoch 2007, 244 n.47).

⁴⁵ There is an important complication here in the case of Wedgwood. Wedgwood argues that “the intentional is normative: it is impossible to give an account of the intentional mental properties and relations without mentioning normative properties and relations” (Wedgwood 2009, 201). At the same time, he argues that the reverse is true for the relationship between the normative and the intentional mental properties – and, moreover, that the intentional mental properties and the normative properties are themselves irreducible. Thus, he ends up with a position according to which the normative and intentional mental properties are necessarily mutually constituting: the essence of each sort of irreducible property makes ineliminable reference to the other. Given that the normative properties involve ineliminable reference to another set of properties, this *prima facie* presents a

How, though, are we to make good on this idea of these properties being fundamentally different in kind from the properties in the natural sciences? This is one of the most important questions that all non-naturalistic realists face in explaining the coherence of their position. In order to appreciate the basic difficulty that the non-naturalist faces on this front, first consider that we have many clear examples of naturalistic properties in our world, e.g. *being round* or *weighing five pounds*. In contrast, it is not clear that we have *any* non-controversial examples of non-naturalistic properties of the precise sort that non-naturalist realists have in mind.

Most of us *do* have some intuitive grasp of properties instantiated by *super*-natural things such as deities, demons, and ghosts. But this is decidedly *not* the sort of property that non-naturalistic realists such as Moore, Parfit, Wedgwood, Enoch, and Shafer-Landau have in mind when they assert that that normative properties are non-naturalistic. The problem, though, is that we lack core paradigmatic examples *of the same strength* as these super-natural examples that we could refer to in order to help make clear that we at least have some intuitive idea of what a non-naturalistic property is. For instance, many of us do start with the idea that some things – such as, perhaps, propositions, numbers, and rules – are *abstract objects* of a sort that cannot be directly studied by the sciences. However, it is notoriously difficult to explain what sorts of things abstract objects are and how to best make sense of them. Furthermore, many philosophers in fact reject that the purported core examples of abstract objects should even be understood as examples of abstract objects in the first place. Thus, it would not be clear exactly what the non-naturalistic realist in metaethics is asserting if she were to assert that normative properties in ethics are like properties of numbers. Moreover, the problem for the non-naturalist in metaethics is that even if there are abstract objects, it is unclear whether the properties of things such as numbers (assuming that they are abstract objects) will be a good model for capturing the sort of thing the non-naturalistic realist in metaethics has in mind. Because of this, one problem for the non-naturalistic realist in metaethics is that we lack a precise sense of what it would take to be a non-naturalistic,

problem for thinking of the idea of normative properties as *sui generis* – something Wedgwood himself insists the normative properties are. For our purposes here, we can say that, if Wedgwood's position could be made coherent, we can think of the core non-naturalist commitment to the *sui generis* properties thesis to be as follows: either normative properties are metaphysically *sui generis*, or they are an ineliminable part of a metaphysically *sui generis* mutually-constitutive property cluster that includes the normative properties. I here draw on (McPherson Manuscript) for this way of responding to the complication raised by Wedgwood's position.

irreducibly normative property in the first place. Call this the *metaphysical mystery problem*. Because of the metaphysical mystery problem, it is therefore crucial that the non-naturalist give some account of what it is that purportedly makes normative properties *sui generis*. Otherwise, non-naturalistic realism will have no defense against one of the most straightforward objections to the view that has been pressed ever since Moore first introduced it: namely, that non-naturalistic properties are mysterious at best and a metaphysical fantasy at worst.⁴⁶

One clear way of responding to this problem is as follows: whereas naturalistic properties play a causal-explanatory role in events that happen in space and time, non-naturalistic normative properties do not. Since many standard naturalistic properties do play a role in casual explanations then, if there were some properties that played no causal role at all, perhaps this would be good grounds for thinking these other properties were fundamentally different in kind from *all* naturalistic properties.⁴⁷ Hence, one good reason for the non-naturalist to deny that normative properties are causally efficacious – and hence to affirm what we can call *non-causal non-naturalism* – is that it allows the non-naturalist to respond to one of the core challenges confronting all forms of non-naturalistic realism: namely, how to explain why the *sui generis* properties thesis is true and, moreover, what exactly it is for it to be true in the first place. Since this response would still only amount to a negative characterization of what non-naturalistic properties are *not*, it would only provide the start of an account of what non-naturalistic properties *are* that would explain why the *sui generis* properties thesis is true. Nonetheless, it would help provide the start of an answer to this problem.

There are a number of serious problems that one faces in trying to vindicate a science-friendly version of *non-causal non-naturalism*. In what follows, I will focus on two of them. The first problem is that it is difficult to see how creatures like us could come to learn about these purported non-causal, irreducibly normative properties. Call this the *epistemic access problem*. The second problem concerns the methodological grounds that the non-

⁴⁶ See, for instance, (Gibbard 2003), (Ayer 1952), and (Harman 2000).

⁴⁷ It should be noted that this strategy would prove to be a dead end if one already believed that certain naturalistic properties could never play a role in causal explanations at all. Given the sorts of standard properties that we get in the natural sciences, it seems there is good reason to be skeptical of this thesis – and hence to at least think that there is something *prima facie* plausible about this proposed way for the non-naturalistic realist to explain in what sense the normative properties are *sui generis*.

naturalist has for positing the existence of such properties in the first place. If such properties are not needed to help in our best causal-explanatory accounts, then we need to know on what grounds they are posited. Call this the *ontological justification* problem. Non-causal non-naturalists, I will argue, have yet to provide scientifically viable answers to these problems.⁴⁸

Let's start with the *epistemic access problem*. In order to appreciate the problem for the non-causal non-naturalist here, consider that one way that we often secure epistemic access with something is by coming into *causal* contact with it. For instance, I come to learn that there are rocks in my backyard because I have seen and felt them. More generally, we think that we are justified in thinking that a subject S has epistemic access to P when we can explain how P is either part of the best casual-explanatory account of why S believes that P or else is analytically entailed by that account. Indeed, partly in response to the Gettier cases about knowledge, some have thought this idea about the grounds for epistemic contact can be used to ground a general constraint on when a subject S can be said to have knowledge that P.⁴⁹ The problem for the non-causal non-naturalist is that such a standard story about epistemic access is ruled out for the case of non-causal non-naturalistic properties. The non-naturalist thus owes a story about how subjects can come to learn about – and, furthermore, in some cases have knowledge about – non-naturalistic properties.

What sort of means do we have for learning about such properties? Earlier in the century, some non-naturalists seemed to hold that we have a special faculty of intuition for identifying such properties. For instance, this position is often attributed (perhaps unfairly) to W.D. Ross.⁵⁰ The problem with this sort of proposal is that – as many pointed out – we simply have no scientific grounds for thinking that creatures like us actually have such a faculty. More recently, then, when contemporary non-causal non-naturalists have put forward accounts of how we come to have epistemic access to non-naturalistic normative properties, they have not explicitly invoked anything such as a special faculty for detecting

⁴⁸ Another key problem that non-causal non-naturalists face is closely related to the *epistemic access problem*. This is that it is difficult to see how creatures like us could have come to acquire the means to refer to such properties in a natural language. Call this the *semantic access problem*. For a critical discussion of this semantic access problem confronting both causal and non-causal versions of non-naturalism alike, see (Jackson 1998), (Smith 1994), and (Railton 2005). Given the constraints of this paper, I leave the semantic access problem to the side.

⁴⁹ See, for instance, (Dretske 1981).

⁵⁰ (Ross 1930).

these properties. Rather, they have claimed that we gain access to these properties through *synthetic a priori* reasoning. This type of reasoning, non-naturalists often claim, is something we already know we need to posit in order to account for our epistemic access to many important properties. In particular, they often point out, it is something we need to account for our epistemic access to properties in mathematics involving abstract objects such as numbers. Pushing further, the non-causal non-naturalist can say the following: physical systems like us – and, indeed, physical systems such as computers – are constituted in such a way that we can come to learn about such abstract objects. Yet we do not take this to mean that we have a special faculty for doing so. What further problem remains?

There are three main problems with the line of reasoning that we just laid out. Two involve the appeal to mathematics. The third involves the initial general appeal to *synthetic a priori* knowledge.

In terms of the appeal to mathematics, the first problem is that it is not clear that there is in fact a strong analogy between the subject matters of ethics and mathematics. While the domains might be analogous in certain respects, it also seems that they are strongly disanalogous in others. For instance, one way they seem to be different is that although most of us have some intuitive sense of how counting marbles *might* be the sort of thing to help get us epistemic access to numbers, we don't start with any intuitive sense whatsoever of what physical processes would do the work in the ethics case. Furthermore, in contrast to the method of proofs in mathematics, there is much less agreement in ethics about what sort of disciplined method we have that secures us epistemic access to the subject matter. This is reflected in the fact that, even among those that grant that the method of ethics can in some sense be described as one of "reflective equilibrium," there is widespread controversy about what exactly that method involves.⁵¹ This controversy reflects a wider range of views about methodology in ethics than one sees in mathematics.⁵² Such differences make it likely that even if one were to come up with a successful story about the

⁵¹ For instance, compare the underlying views on methodology in (Kagan 1989) and (Kamm 2007).

⁵² However, it should be noted that there *are* important methodological debates in mathematics. Perhaps most importantly, even if there is relatively stable agreement about methodology for mathematical proofs, this is not the case for establishing new axioms. Instead, there is a lively debate about when we would be justified in introducing new axioms. On this issue, see (Maddy 2000) and (Easwaran 2008). The metaethical non-naturalist might appeal to this fact about debate in mathematics about axioms in order to challenge my claim that there is more methodological disagreement in ethics than in math. Thanks to Kenny Easwaran for pressing me on this point.

metaphysics and epistemology of mathematics, it would not necessarily carry over to the domain of ethics as well.

A second problem with the appeal to an analogy with mathematics is that there is just as much debate in the philosophy of mathematics as there is in metaethics about how to make sense of our epistemic access to abstract properties that appear to play no causal role.⁵³ Indeed, in part because of this issue, there is a wide range of views currently on offer about the basic ontology of numbers. It therefore seems unwarranted to cite a particular view about the epistemology and ontology of numbers to support a view in metaethics. Indeed, it is hard to see how this is any better than someone in philosophy of math doing the reverse by appeal to a contentious metaethical position such as non-naturalistic realism or expressivism. Furthermore, it should here be stressed that many of the important views in philosophy of mathematics – for instance, neo-Fregean views that attempt to derive arithmetic from logic – simply look like non-starters for generating parallel views in the case of metaethics.

The deeper problem for the non-causal non-naturalist here, however, does not have to do with the particular appeal to mathematics. Rather, it concerns the appeal to *synthetic a priori* facts in general. It might be that we need to believe in such facts to account for a wide range of things we take ourselves to be justified in believing in – perhaps, for instance, such basic things as causation. Yet, it is nonetheless notoriously difficult to give a scientifically respectable account of what makes synthetic *a priori* knowledge possible for creatures like us. Indeed, even if one does not explicitly put forward the idea of a special faculty for gaining epistemic access to synthetic *a priori* facts, it is not clear what else one can say. Most importantly, at least in the context of non-naturalistic realism, what needs to be posited is precisely some capacity we have for knowing facts that are not conceptual, not empirical, and that are somehow about properties that are a genuine part of reality at the most basic ontological level. Without some scientifically respectable story for how we have such a capacity, non-causal non-naturalism runs into a basic problem with making itself science-friendly.⁵⁴

⁵³ For the classic statement of this problem, see (Benacerraf 1973). See also (Field 1991) and (Maddy 2000).

⁵⁴ It should be also be noted that appealing to Quinean skepticism about the analytic/synthetic distinction or the *a priori*/*a posteriori* distinction is unlikely to provide much, if any, help to the non-naturalistic realist given the strongly naturalistic character of Quine's views.

Now consider the *ontological justification problem*: i.e. the problem of explaining the grounds on which the non-naturalist posits the existence of non-naturalistic normative properties in the first place. As with the *epistemic access problem*, part of the issue here for the non-causal non-naturalist stems from the fact that her basic commitments rule out an attractive way of answering this question. One of the canonical ways by which philosophers have sought to establish that something exists is by showing that its existence is part of the best explanatory account for something else that we think needs to be explained. For instance, consider the following canonical example first discussed by Harman.⁵⁵

Cloud Chamber. A physicist sees a vapor trail as she is looking into a cloud chamber. “There goes a proton,” she thinks to herself.

Harman argues that the physicists’ observation provides evidence that there actually was a proton in the cloud chamber. The reason for this, he argues, is that the best explanation for why the observation happened at all would invoke the existence of the proton itself: in other words, the best explanation would involve the fact that there *really was* a proton in the cloud chamber at the relevant time. Harman’s reasoning relies on the idea of establishing existence by *inference to the best explanation*: a type of reasoning that, even if highly contested, is one of the standard routes by which philosophers seek to establish something’s existence (including, it should be noted, in the philosophy of mathematics). The problem for the non-causal non-naturalist is that since she has already granted that the non-naturalistic properties play no causal role akin to protons in cloud chambers, then she cannot make use of this type of reasoning for establishing that normative properties exist.

There are two main types of response that non-causal non-naturalists have given in order to attempt to deal with this challenge. The first response is to expand the domain of things one is trying to explain in order to include normative facts themselves (such as the fact that *it is wrong to torture others for fun*). The second is to reject the thesis that one *needs* to use inference to best explanation (IBE) of the sort Harman appeals to in order to establish that something exists, and instead to posit other grounds for justifying ontological commitment.

The basic problem with the first option is that we need a principled way of expanding the domain of respectable explananda and, in particular, one that will not allow us

⁵⁵ (Harman 1977).

to rule in anything that we would deem to be scientifically unacceptable. For instance, imagine that one firmly believes that magical fairies exist in our world, and that one also believes they are causally isolated from everything else within it. The non-naturalist needs a way of explaining why normative facts get to count as legitimate explananda (which will in turn potentially justify positing non-naturalistic normative properties) but that facts about causally isolated magical fairies do not (which might in turn justify positing properties about supernatural beings). In short, the non-causal non-naturalist would need a principled way of expanding ontological commitment via IBE that will lead to the addition of properties in our ontology that we would regard as “adding to” but not *in conflict* with a broadly naturalistic worldview. The basic issue for the non-causal naturalist is that it is hard to see a plausible way to do this that will end up including the facts she needs – namely, normative facts.

Furthermore, there is also a deeper issue here for the non-naturalist with this first line of response, an issue that I think has been under-appreciated. As we have seen, one of the core commitments of the non-naturalist is the *sui generis thesis*. At the core of this thesis is the idea that normative properties are metaphysically different from any of the properties we study in the natural and social sciences. Now consider that it seems at least plausible to hold that one significant part of the methodology of the sciences involves using IBE to establish when we should believe in certain things rather than others. Indeed, even if it is not part of that methodology, we have already noted that it is a widely accepted way of establishing belief in certain things we might intuitively take to be quite mysterious – for instance, the protons in the cloud chamber example. If so, this raises a serious issue for the non-causal non-naturalist who wants to use IBE on what are meant to be *scientifically respectable* explananda in order to establish the existence of non-naturalistic normative properties. The issue is this: if those explananda really are scientifically respectable, what grounds does the non-naturalist then have for endorsing the *sui generis* thesis about normative properties in the first place? For if non-naturalistic normative properties are posited *following* a core part of the methodology of the natural sciences, or at least a method that many people take for granted in establishing many properties that the non-naturalist would categorize as “naturalistic,” it is then unclear why the properties that the non-naturalist has posited should be understood as *non-naturalistic* properties rather than, say, novel naturalistic properties of *an unexpected sort*. In other words, by appealing to straightforward versions of IBE in order to ground belief in

non-naturalistic normative properties, the non-casual non-naturalist ends up being stuck with a position that threatens the coherence of her basic view.

Let's now turn to the second type of response open to the non-casual non-naturalist confronted with the challenge that it seems that she cannot use IBE to ground ontological commitment about non-causal non-naturalistic properties. This is to reject the thesis that one *needs* to use inference to best explanation (IBE) of the sort that Harman appealed to in order to establish that something exists. One relatively modest option starts with the observation that the explanation was a *causal* one in the cloud chamber case. However, it seems plausible to hold that there are *non-causal* explanations that we give in many domains of inquiry – and, thus, the non-naturalist can assert that IBE should not be understood to concern solely *causal* explanations. Once it is not, perhaps this will open the door to positing any number of other things – for instance, numbers or non-naturalistic normative properties. Another, more radical option, recently put forward by David Enoch, is to argue that there are grounds other than IBE (either involving causal or non-causal explanations) for positing properties in our ontology.⁵⁶ In particular, Enoch argues that appeal to *deliberative indispensability* – that is, roughly, indispensability from the standpoint of someone engaged in practical reasoning – can ground ontological commitment just as much as appeal to *explanatory indispensability* of the sort that matters from theoretical reasoning.

It is beyond the scope of this paper to go into detail about either of these two types of proposals and whether or not, in the end, they would actually be able to deliver good grounds for positing non-causal non-naturalist properties in our ontology. It should be stressed, however, that even if such proposals provided good grounds for positing non-causal non-naturalist properties in our ontology, this will not do anything to guarantee that the results will be *science-friendly*.

Take the first option of appealing to *non-causal* explanations in IBE. Many of the paradigm cases of non-causal explanation that we think are science-friendly are ones that, even if they are not *causal* explanations, still appeal to properties that *do* play causal roles. Furthermore, those properties that do not, perhaps mathematical properties (insofar as scientific theories do in fact appeal to them), *at least* look as though they might help explain something *about* those causally efficacious properties – indeed, it is plausible to hold that is the whole reason for invoking mathematics in the sciences in the first place. Unfortunately,

⁵⁶ (Enoch 2007).

however, it seems difficult to see how non-causal non-naturalistic normative properties could help to explain the causally efficacious properties in any parallel way. Thus, even if IBE does not need to be inference to the best causal explanation, it is unclear what sort of scientifically legitimate explanatory role these normative properties could play that would parallel the explanatory roles played by such things as numbers and propositions.

Now consider the second more radical option. This is to argue that there are grounds other than IBE (either involving causal or non-causal explanations) for positing properties in our ontology. In recent work, David Enoch has suggested that these grounds could be that those properties are *indispensable* from the first-personal deliberative standpoint.⁵⁷ In rough terms, Enoch argues that deliberative indispensability is ground for ontological commitment by first arguing that the best justification for taking IBE to be ground for ontological commitment is because appeal to indispensability *in general* is good ground for ontological commitment. Enoch then argues that this idea vindicates ontological commitment derived from either explanatory indispensability or deliberative indispensability. Suppose this is so. By itself, this would do nothing to show that what we take to be indispensable from a deliberative standpoint will be things we take to be science-friendly. For instance, Enoch's strategy could – in principle – yield good grounds for positing supernatural beings such as fairies or witches or deities in our ontology – namely, if it turned out that such beliefs were deliberatively indispensable. In all likelihood, of course, Enoch's strategy would *not* yield grounds for believing in such supernatural things that play a purportedly causal role in our world but rather, at best, grounds for believing in something such as *justificatory reasons for action* (i.e. considerations or facts that could in favor of or count against certain courses of action). In turn, if negative arguments against other metaethical views were to succeed, one might then have grounds for holding that the best explanation of such justificatory reasons is non-naturalistic realism. The point, though, is this: establishing that one is justified in ontological commitment from deliberative indispensability by itself will do nothing to guarantee that the results will be science-friendly. This holds true *even if* Enoch is right that the reason why deliberative indispensability grounds ontological commitments is intimately related to the reason why explanatory indispensability does as well.

⁵⁷ (Enoch 2007).

Putting both of these proposals to the side, however, I want to focus on a general problem that *all* such proposals face – and, indeed, that I think is in fact the heart of the non-causal non-naturalist’s *ontological justification problem*. To see what this problem is, consider three basic responses we might end up with in terms of assessing how, if at all, the position of non-causal non-naturalism is supported using the methodology of the natural sciences:

1 – Non-causal non-naturalism is a well-supported scientific hypothesis that is arrived at using the methods of the natural sciences.

2 – Non-causal non-naturalism is not well-supported using solely the methods of the natural sciences. However, there are other methodologies that the non-causal non-naturalists can use to establish their view that have the following features: they are both justified methodologies and ones that in no way threaten our confidence in using the methodologies that we use in the natural sciences for studying the naturalistic world.

3 – Non-causal non-naturalism is not well-supported using solely the methods of the natural sciences. Furthermore, there are *no* other acceptable methodologies that the non-naturalist can use to establish their view except ones that would threaten our confidence in using the methodologies that we use in the natural sciences for studying the naturalistic world.

From our previous discussion so far, we know that the first response creates a serious problems for non-causal non-naturalism. This problem is that this response leaves the non-naturalistic realist without a good way of explaining how the *sui generis* properties thesis is true. The third response would also be unacceptable. For it simply amounts to saying that the no methodological conflict thesis is false. If so, then the non-causal non-naturalist should aim to show that the second type of response is true.

In order to better appreciate the difficulty of achieving this goal, we should get clear on what conditions would need to be met in order for this second type of response to be warranted. In order to identify these conditions, consider that, in everyday life, all of us make claims that rest on facts about the reliability of our own perception or on the reliability of the testimony of others. Such claims might not be *based* on using the fundamental methodology of the natural sciences. Nonetheless, it seems at least plausible to hold that these claims are not in conflict with confidence in the fundamental methodology of the natural sciences. There are, of course, deep questions here about why this is so. Yet, one good reason for thinking that these claims are not in conflict with those of the natural sciences is that, given

the content of our best scientific stories, we have good reason to hold that following the other methodologies that led to such claims will lead to reliable results. If that is so, then these claims might not be *directly* supported by the methodologies of the natural sciences, but we will have good reason to hold that the *methodologies* that led to those claims will be. Call this *indirect scientific support*. Although there might be other ways as well to show that a claim is not in conflict with confidence in the fundamental methodological commitments of the natural sciences, this brings out one paradigmatic way of doing so: namely, by showing how certain claims get *indirect scientific support*.

With this on the table, now return to the case of non-causal non-naturalism in metaethics. Given the non-naturalist's central metaethical commitments, we know that whatever response to the ontological justification problem the non-naturalist gives, it seems that her basic story of epistemic access to these properties will involve *synthetic a priori* reasoning. The problem is that, even if such reasoning does turn out to be a sort we should believe in, it currently lacks the sort of *indirect scientific support* that claims based on testimony and perception have. Thus, in staking out *ontological commitment* based on things one claims to know about through *synthetic a priori* reasoning, one is essentially doing so on methodological grounds that lack any internal support from the natural sciences. Once one does so, this opens the door to the proliferation of any number of claims that would either be in conflict with core commitments of the natural sciences or open the door to strange supernatural entities of a sort that all science-friendly non-naturalists want to deny. Thus, whatever specific method the non-naturalist appeals to in order to deal with the ontological justification problem, the non-causal non-naturalist owes a response to this basic issue of how her view can remain science-friendly in the absence of indirect scientific support.

4. *Causal non-naturalism.*

In the last section, I argued that there are some key problems that one faces in trying to make non-causal non-naturalism science-friendly and that the main responses of contemporary non-naturalists have not successfully addressed these problems. In this section, I will turn to the issues that confront the non-naturalist who grants that normative properties can be causally efficacious and who aims for the view to be science-friendly. In other words, I will turn to the issues that confront advocates of what we can call *causal non-naturalism* who also endorse *the no conflict thesis*. In short, I will argue here that advocates of

this form of non-naturalism face two main problems: 1) it is not clear why their view gets to count as a form of non-naturalism in the first place and 2) the view is in conflict with one of the core substantive ideas many have about the natural sciences, namely, the idea that the natural world is *causally closed*. Because of these problems, I argue that the causal form of non-naturalism is no better off – and perhaps worse off – than the non-causal form of non-naturalism in terms of its compatibility with our basic confidence in the natural sciences.⁵⁸

The causal form of non-naturalism has two main advantages over the non-causal form. The first advantage is that it allows the non-naturalist to – at least potentially – diffuse some of the dialectical force of the *epistemic access problem*. For if one grants that normative properties can themselves play a causal role, then these properties could at least in principle be invoked to explain why a subject forms certain beliefs as opposed to others. The second, perhaps more important, advantage is that it allows the non-naturalist a straightforward way to respond to the *ontological justification problem*. For if normative properties themselves play an ineliminable role in causal explanations of contingent naturalistic events, then the non-naturalist can appeal to straightforward forms of inference to the best explanation of naturalistic events in order to explain why one is justified in including these non-naturalistic normative properties in our ontology.

Unfortunately, this move to a causal form of non-naturalism in turn creates new, potentially even more significant issues for the aspiring science-friendly non-naturalist. The first problem is that once one grants that normative properties “play an essential role in causal explanations of certain contingent facts”⁵⁹, it makes it difficult to see in what sense these properties get to count as metaphysically *sui generis*. The non-causal non-naturalist, remember, was able to explain why normative properties are metaphysically *sui generis* by claiming that normative properties were *not* causally efficacious in this way. In so doing, she thereby had at least the start of an answer to what I called the “metaphysical mystery problem.” In contrast, the causal non-naturalist has no readily available solution. For once one grants that normative properties are part of the explanations of contingent facts – facts, in other words, of precisely the sort that the natural sciences study – it seems that one thereby grants that these properties are continuous with other naturalistic properties in a

⁵⁸ Much of what I have to say in this fourth section of the paper draws extensively on both personal communication with Hussain as well as a 2008 presentation he gave in Jerusalem entitled “Normativity and the Causal Order” (Hussain 2008). See also (Hussain Forthcoming).

⁵⁹ (Wedgwood 2009, 6).

crucial sense. Of course, if one granted this conclusion, one could still hold that the normative properties form an interesting subset of the naturalistic properties – perhaps as the biological and chemical properties are subsets of the broader class of naturalistic properties. But this would not amount to showing what the non-naturalist needs to show in order to make good on the *sui generis thesis*: namely, that the normative properties are different *in kind* from all naturalistic properties. Hence, it is not clear that causal forms of non-naturalism can be made coherent as legitimate forms of *non-naturalistic realism*.⁶⁰ At worst, the position will be a non-starter. And, at best, the causal non-naturalist faces a significant dialectical burden in explaining the coherence of this position that the non-causal non-naturalist does not face to the same degree.

For now, however, I want to put an extended discussion of this issue to the side and instead focus on a second problem with this causal form of purportedly science-friendly non-naturalism. This position states that the fundamental normative properties can play a causal role in the causation of naturalistic events.⁶¹ The problem is that this conflicts with one of the core substantive ideas many have about the natural sciences; namely, the idea that the natural world is *causally closed*. We can give a first-blush account of what this idea amounts to as follows:

Naturalistic closure: Any cause of a naturalistic event is itself a naturalistic event.

This idea is, I think, best understood as view that many of us have *about* the claims made by the natural sciences, rather than a view that is itself part of the *content* of claims made in the natural sciences. Nonetheless, I will take it that it is something most people think of as a core

⁶⁰ This worry is shared by Enoch. See (Enoch 2007, 26).

⁶¹ The phrase “in the causation of naturalistic events” is meant to underscore that the causal non-naturalist is *not* advocating a view according to which these normative properties only play a role in causing events in some non-naturalistic realm (perhaps for instance, in the normative realm). For instance, consider Wedgwood’s view. As he puts it, he grants that normative properties “enter into causal explanations of contingent facts about what happens in the world” (Wedgwood 2009, 10). These contingent facts are ones that Wedgwood explicitly grants are the same ones studied by the natural sciences – indeed, as he writes, he holds that “absolutely all contingent facts are realized in the facts of the sort that are studied by the natural sciences such as physics” (Wedgwood 2009, 3). In this way, the view in consideration is strongly disanalogous with dualist parallelism in the philosophy of mind, according to which mental events only cause other mental events and are thus causally isolated from the entire physical realm.

thesis of the so-called “modern naturalistic world-view”; and thus something that a science-friendly non-naturalism should not overturn.⁶²

There are two main reasons I think that we should accept that the thesis of naturalistic closure is part of the so-called “modern scientific worldview.” The first reason is that we think that the natural sciences could at least in principle give a complete account of the causes of purely naturalistic events – which we can here roughly designate as contingent events of the sort that natural science concerns itself with – without doing something *other* than more natural science. To think otherwise would be, in essence, to claim that the natural sciences are fundamentally limited in terms of how good a job they can do in studying precisely one of the main things they aim to study: namely, the occurrence of naturalistic events. The second reason is that we think that were the natural sciences to come up with a complete account of a naturalistic event, this account would thereby constitute an *authoritative* one in the sense that it would not simply be *one of many* complete accounts. By this, I essentially mean the following: if we develop true accounts about the natural world *other* than those offered by the natural sciences, such accounts won’t be adding any new *sufficient* causes at the basic explanatory level that the *complete* natural scientific story would leave out. Put together, then, these two basic reasons give support to thinking that naturalistic closure is a core commitment of the “modern scientific worldview.”

That the naturalistic closure thesis forms a part of our basic confidence in the natural sciences is also, I think, well-reflected by the current state of the debate about *physicalism* in the philosophy of mind. In broad terms, *physicalists* hold that mental properties can be explained in terms of basic physical properties whereas *property dualists* deny that they can be. Some dualists assert these non-physical properties could not ever be studied by the natural sciences: and hence would not could as “naturalistic” in any meaningful sense. However, those dualists who aim to make their view science-friendly – for instance, David Chalmers – explicitly *do* hold that the mental properties could be studied by the natural sciences.⁶³ In turn, since Chalmers’s view of mental properties carries over to his view of mental events,

⁶² As I noted earlier in the paper, much of what I have to say in this fourth section of the paper about the conflict between causal non-naturalism and the causal closure of the natural world draws extensively on the work of Nadeem Hussain. I here draw extensively both on personal communication with Hussain as well as his discussion of these matters in a 2008 presentation he gave in Jerusalem entitled “Normativity and the Causal Order” (Hussain 2008). See also (Hussain Forthcoming).

⁶³ (Chalmers 1996).

this means that science-friendly dualists such as Chalmers accept the naturalistic closure thesis even if they might end up denying causal closure of the *physical*. Such views might, of course, turn out to be mistaken – or perhaps even conceptually confused in some basic way. The point, however, is that Chalmers’s kind of view is a reflection that the naturalistic closure thesis is taken by many to be *common ground* in the debate between different science-friendly views about the metaphysics of mind.

Now that we have the naturalistic closure thesis on the table – as well as an understanding of why it is something that should be respected by the non-naturalistic seeking to meet the naturalistic constraint – the basic problem for causal non-naturalism can be put as follows. If I am correct that the naturalistic closure thesis is something that should be respected by the non-naturalistic seeking to meet the naturalistic constraint, then the causal non-naturalist is committed to the following three claims:

- 1) *Sui Generis Properties*. The fundamental normative properties are *irreducibly normative* properties that are metaphysically *sui generis*.
- 2) *Casually Efficacious*. The fundamental normative properties can play a causal role in the causation of naturalistic events.
- 3) *Naturalistic Closure*: Any cause of a naturalistic event is itself a naturalistic event.

These theses form an inconsistent triad. Since the non-naturalist is committed by the *sui generis* properties thesis to holding that those normative properties are fundamentally different in kind from any naturalistic property, the non-naturalist is committed to thinking that an event involving a non-naturalistic property (which we can safely say would make it a *non-naturalistic* event) causes a naturalistic event. This, however, is precisely what the naturalistic closure thesis rules out. Hence, the causal non-naturalist who aspires to vindicate the no conflict thesis is committed to an inconsistent set of claims.

In confronting this issue, Wedgwood and Shafer-Landau have sought to avoid this conclusion by arguing that causal explanations that invoke normative properties are not in conflict with those that do not.⁶⁴ In order to do so, they have not denied that both sorts of explanations are trying to give explicit explanations of the very same events. Rather, they have exploited an idea borrowed from so-called “non-reductive physicalism” in the

⁶⁴ (Wedgwood 2009) and (Shafer-Landau 2005).

philosophy of mind: namely, that mental properties are *constituted* or *realized* by base physical properties even though they are not reducible to those base physical properties. Borrowing from this, Wedgwood and Shafer-Landau claim that non-naturalistic normative properties are distinct from naturalistic properties but nonetheless are fully *constituted* and/or *realized* by them.⁶⁵ As Wedgwood puts it, this gives him a position that he takes to be perfectly compatible with a basic metaphysical naturalism “according to which absolutely all contingent facts are realized in the facts of the sort that are studied by the natural sciences such as physics.”⁶⁶

Given what we have said so far, if this claim about constitution and realization of naturalistic properties is to be made intelligible even as a *possible aid* to the non-naturalistic realist, then it must be because the closure thesis is something weaker than what I have said so far. In particular, it seems that it would need to be the following:

Modified Naturalistic Closure. any cause of a naturalistic event is either itself a naturalistic event or something that is realized and/or constituted by a naturalistic event but not reducible to a naturalistic event.

If this form of naturalistic closure were in fact the central commitment about the causal closure of the natural that is contained in our commitment to what Wedgwood calls a “modern scientific worldview,” then this would show that the causal non-naturalist can avoid conflict with the closure thesis by holding that normative properties are realized and/or constituted by naturalistic properties. There are potentially good reasons for thinking that modified naturalistic closure, rather than our initial naturalistic closure thesis, in fact represents the right way of reading the basic commitment we started with when I said that belief in the causal closure of the natural was part of our confidence in the natural sciences. Consider again the closure thesis as I initially stated it: any cause of a naturalistic event is itself a naturalistic event. Contrast this with another, weaker thesis we might hold: any naturalistic event has a sufficient naturalistic cause. This thesis is weaker than the first because it leaves open the possibility that a naturalistic event could have a non-naturalistic cause as long as it *also* has a sufficient naturalistic cause as well. Call our initial version of

⁶⁵ (Shafer-Landau 2005) and (Wedgwood 2009).

⁶⁶ (Wedgwood 2009, 3).

naturalistic closure *strong naturalistic closure* and the second, weaker thesis *weak naturalistic closure*.⁶⁷

Many philosophers think that, even if I am right that there is *prima facie* motivation for the strong naturalistic closure thesis of the sort I gave, the strong naturalistic closure thesis is simply *too strong*: it represents a *scientistic* rather than scientific understanding of the naturalistic events. In short, proponents of weak naturalistic closure might hold that, in order to have confidence in a basic naturalistic worldview, it is enough to hold that, in order to be correct as scientific explanations, scientific explanations need to invoke solely sufficient naturalistic causes but that it is simply a further *scientistic* claim to hold that there could be no other, non-naturalistic sufficient causes that could be studied by other forms of inquiry. With this proposal on the table, we can then understand the modified naturalistic closure thesis as a thesis that attempts to successfully respond to one of the core problems for those who want to defend a weak version of a closure thesis, whether about naturalistic events in general or the physical events in particular. This problem is to explain how to rule out cases of intuitively vicious types of *overdetermination* where a single event has two distinct sufficient causes coming from two radically different types of property (or, indeed, perhaps from two completely different events). At its core, the position that Wedgwood and Shafer-Landau attempt to borrow from non-reductive physicalism in the philosophy of mind is a way of explaining how this might be done.⁶⁸ Furthermore, it is a way of explaining how this might be done that is meant to preserve the core idea that “absolutely all contingent facts are realized in the facts of the sort that are studied by the natural sciences such as physics.”⁶⁹

For the purposes of this paper, there are two crucial questions about Wedgwood and Shafer-Landau’s strategy of appealing to the modified naturalistic closure thesis as a way to defend non-naturalistic realism. The first is whether or not there are in fact principled grounds for endorsing *this* thesis – or any other weakened version of our original naturalistic

⁶⁷ In distinguishing between strong and weak versions of naturalistic closure in this way, I draw on a parallel distinction between strong and weak versions of *physical* closure theses put in terms of *physical* events and causes. On this distinction, see (Crisp and Warfield 2001) and (Kim 2007).

⁶⁸ This is one of the main questions discussed about this sort of principle in the context of the debate over physicalism in the philosophy of mind where the principle is introduced as one about “physical” events rather than “naturalistic ones.” For an argument that it leads to unacceptable versions of overdetermination, see (Kim 1998) and (Kim 2007). For a defense of the idea that the types of overdetermination that such a principle leads to are ultimately acceptable, see (Crisp and Warfield 2001).

⁶⁹ (Wedgwood 2009).

closure thesis – as the right way of reading the basic commitment we started with when we said that belief in the causal closure of the natural was part of our confidence in the natural sciences. The second is whether or not appeal to the modified naturalistic closure thesis would in fact help explain the coherence of a science-friendly version of causal non-naturalistic realism. In what follows, I will argue that there are strong reasons to think that the answer to the second question is “no.” In turn, I will argue, this discussion helps bring out that there are in fact strong reasons to doubt that there is a viable weakened version of the original closure thesis that would help one defend a science-friendly version of causal non-naturalism. Because of this, I conclude that there is good reason to hold that the purportedly science-friendly causal non-naturalist will be left with the initial inconsistent triad that we started with.

In order to appreciate the reason why there are strong reasons to doubt that appeal to modified naturalistic closure thesis would in fact help explain the coherence of a science-friendly version of causal non-naturalistic realism, we first need to better appreciate why we might be attracted to such a view in the first place. Consider an explanation I might give of the following: David Ortiz hits a home run for the Red Sox and Mike Lewis, a Red Sox fan watching the game on his TV, cheers wildly. On the one hand, it seems that we could give a satisfactory explanation of why Mike Lewis cheers that simply repeats what I just said. On the other hand, on the reasonable assumption that all of the entities that I just described are constituted by collections of molecules, it looks as if one could, in principle, also give an explanation of the event solely in molecular terms – and hence never mention David Ortiz or Mike Lewis *de dicto* (that is, using the terms “David Ortiz” or “Mike Lewis”). However, it at least intuitively does not look as if the two explanations are in conflict with each other. To put it in one intuitive way, this is because the accounts look to be describing entities and events at different levels of description using different words. Because of this, nothing in the explanations are in conflict with each other; the complete ontic explanation of why the event occurred could vindicate both of them at the same time.⁷⁰

Wedgwood and Shafer-Landau want to maintain that roughly the same thing is true with the relationship between a correct explanation that references normative properties in the causal-explanatory account and a correct one that does not. In basic terms, they want to

⁷⁰ This is broadly the sort of view that is defended in (Yablo 1992). See also (Searle 1995). For critical discussion of this sort of view, see (Kim 2007).

claim that since normative properties are *realized* or *constituted* by naturalistic properties, when one invokes normative properties in a causal explanation, this is just as harmless as offering an explanation of Mike Lewis cheering by invoking Mike Lewis *de dicto* as opposed to explaining everything at the molecular level. In order to explain this, Wedgwood and Shafer-Landau both appeal to the work of non-reductive physicalists in the philosophy of mind – who, they seem to hold, show us a model for how this could be so in the case of explanations involving mental properties.⁷¹

Unfortunately, there are general reasons why this appeal to non-reductive physicalism is a dead end for anyone committed to non-naturalistic realism as I have understood it in this paper.⁷² To see why this is so, consider again the reasonable (if still potentially false) claim made by so-called “non-reductive physicalists” in the philosophy of mind that mental properties are realized or constituted by physical properties but are not reducible to them. One way in which non-reductive physicalists have explained this is by arguing that, in the first instance, mental states should be understood *functionally* (i.e. in terms of states that play a certain functional role). The basic non-reductive physicalist thought then goes as follows: since these functional roles can be realized by different physical systems, this is enough to rule out the idea that the property of having a certain mental state can be reductively explained as the property of having a certain specified physical state.⁷³

To flesh this out, consider the property of *being a stove*. Something can be a stove and operate using either gas or electric coils. Stoves are a *functional* kind: in order to be a stove, a given object needs to play a specific functional role. Thus, regardless of the specific mechanisms by which an appliance works, insofar as an appliance realizes a functional role, it thereby “realizes” and has the property of being a stove. The non-reductive physicalist claims that nothing more complicated than this is going on when a physical system “realizes” or “constitutes” a mental state (e.g. the property of being a belief). Hence, she claims, explanations that make use of mental properties are no more naturalistically suspect than those that make use of the property of *being a stove*.

⁷¹ (Wedgwood 2009, 192-198) and (Shafer-Landau 2005, 98-113).

⁷² I here draw extensively on personal communication with Nadeem Hussain and the argument of (Hussain 2008).

⁷³ For one of the best leading taxonomies of positions in the philosophy of mind that supports this idea, see (Bennett 2008). For criticism of the idea that the thesis that mental states are functional states leads to a form of physicalism worth calling “non-reductive,” see (Braddon-Mitchell and Jackson 1996).

The problem for metaethical non-naturalists such as Wedgwood and Shafer-Landau, however, is that, unlike the non-reductive physicalist in the philosophy of mind, non-naturalistic realists in metaethics cannot appeal to a pre-theoretical idea of “realization” that is exemplified in the stove case. To see why this is so, suppose that one grants that the fundamental normative properties could be identified in terms of things that play a specifiable functional role. If so, then *perhaps* the fundamental normative properties would not be identical with any *physical* property. However, such properties would then be essentially continuous with a whole range of other functionally specified properties: for instance, the property of *being a stove* or the property of *being an engine*. This will hardly be a good way of making good on the just “too different” intuition that Enoch and others take to be the core of the non-naturalist’s claim that normative properties are *sui generis*. Hence, non-naturalistic realists *cannot* grant that the fundamental normative properties could be identified in terms of things that play a specifiable functional role without again threatening the basic coherence of their own basic type of metaethical view.⁷⁴ This, I think, is usually taken to be common ground in the debate over non-naturalistic realism. This is well reflected in the fact that none of the non-naturalistic realists that we have been considering in this paper do take normative properties to be functional properties whereas certain naturalistic realists do.⁷⁵

The problem, then, for non-naturalistic realists such as Wedgwood and Shafer-Landau is this: they then need some other *non-functional* way of capturing the “realization” or “constitution” relation between naturalistic properties and the non-naturalistic normative properties that would make explanations invoking normative properties no more naturalistically suspect than those invoking stoves. Both Wedgwood and Shafer-Landau *do* have general things to say about what constitution and/or realization amounts to – and, moreover, things to say about how this relationship is more than just a basic supervenience relationship.⁷⁶ Wedgwood, for instance, cashes out the idea of one property (property *A*) being realized by another (property *B*) in terms of what he calls a “strong supervenience thesis” between the *essences* of two properties. Somewhat more precisely, on his view, “to say

⁷⁴ For a similar line of thought about why non-naturalistic realists cannot endorse a functionalist account of normative properties see (McPherson Manuscript).

⁷⁵ For instance, see (Jackson and Pettit 1996) for a naturalistic account of normative properties in terms of properties that play functional roles that are specifiable in non-normative terms.

⁷⁶ As Wedgwood puts it, the idea that the normative properties are realized by naturalistic properties is meant to *explain* why the normative properties supervene on the non-normative ones. (Wedgwood 2009, 10).

that the fact x has A is ‘realized’ in the fact that x has B is to say that it is an essential feature of property A that it strongly supervenes on properties of a certain kind, and B is the minimal non-disjunctive supervenience basis of this kind that is actually instantiated by x .⁷⁷ It would take us well beyond the scope of this paper to explain exactly what this means for Wedgwood and to evaluate its role in his broader metaethical theory. My point here, however, is not that this developed view of Wedgwood’s is a *mistaken* account of some relationship potentially worth calling a “realization” one. Rather, it is simply that when non-naturalists such as Wedgwood and Shafer-Landau appeal to the idea that “realization” and “constitution” in order to avoid the conclusion that their non-naturalism is in conflict with believing in the causal closure of the natural, they explicitly appeal to non-reductive physicalism in the philosophy of mind – an appeal that, we have just seen, in fact threatens the core of metaethical *non-naturalistic realism* as I am understand it in this paper.⁷⁸ Perhaps there *is* some other non-functionalist way of capturing the “realization” or “constitution” relation between naturalistic properties and the non-naturalistic normative properties that would a) make the direct analogy with standard non-reductive physicalism in philosophy of mind work such that explanations invoking non-naturalistic normative properties would be no more naturalistically suspect than those invoking stoves and b) not threaten the core non-naturalist idea that normative properties are metaphysically *sui generis*. The problem for those aspiring to science-friendly versions of causal non-naturalism, however, is that they have yet to deliver such an account that does all of this at once. Furthermore – unlike in the case of non-reductive physicalism in the philosophy of mind where one can appeal to the realization of functionally-defined properties as a way of explaining what the general realization relationship amounts to – there are no relatively general theoretical ideas of “realization” or “constitution” that would support thinking that there is such an account in the offing.

⁷⁷ (Wedgwood 2009, 202).

⁷⁸ Perhaps, in light of such issues, Wedgwood and Shafer-Landau would themselves ultimately choose not to be classified as non-naturalistic realists but rather as something else such as “non-reductive naturalists.” Indeed, given Wedgwood’s extended careful discussion of non-reductive naturalism as a *general* view in metaphysics, I suspect this is a potentially a good way to label at least Wedgwood’s developed metaethical view. This issue, however, is beside the point in this paper. This is because my concern is with a *broad type of view* in metaethics that Wedgwood and Shafer-Landau can at least be charitably read as exemplifying rather than on the finished details of their own metaethical views. It should be noted that, if Wedgwood and Shafer-Landau were to pursue this strategy, it would then likely make sense for Shafer-Landau to stop calling his view a “non-naturalist” one and for Wedgwood to stop calling his view a “Platonist” one.

Given this fact, I conclude that *even if* we did have principled reason for holding that the modified naturalistic closure thesis is the real commitment we have about the causal closure of the natural in light of a “modern naturalistic worldview,” we still have good reason to be skeptical that there is a coherent version of the modified naturalistic closure thesis that would help the non-naturalistic realist create a science-friendly version of her view. This, of course, does not rule out that there is some other version of a closure thesis that could do the trick here in combination with a developed version of causal non-naturalistic realism. However, it is hard to see what that would be. This is for two reasons. The first is that, as Wedgwood puts it, it seems to be a core commitment of the modern naturalistic worldview that “absolutely all contingent facts are realized in facts of the sort that are studied by the natural sciences such as physics.”⁷⁹ If this is so, it rules out views that are parallel to *interactionist dualism* in the philosophy of mind where contingent physical events are caused by non-physical mental properties that supervene on physical properties but are *not* “realized” or “constituted” by physical facts. The second reason is that even if the non-naturalistic realist were to reject Wedgwood’s view here on the grounds that it was overly-scientistic, and in turn develop a view based on interactionist dualism in the philosophy of mind, it seems that this view would leave open much more rampant forms of overdetermination than views modeled on non-reductive physicalism. Perhaps that is simply a cost of the view. Or perhaps there is nothing wrong, let alone anti-scientific, with granting such overdetermination in the first place – after all, non-reductionists about biology and chemistry might already be committed to it. If so, then perhaps the only commitment we should have about the causal closure of the natural is that any cause of a naturalistic event is either itself a naturalistic event or an event of a type that supervenes on naturalistic events. Yet, given the amount of overdetermination this would allow, and furthermore how much *overdetermination* by non-naturalistic events it would allow, I think that this proposed position would likely stray simply *too far* from Wedgwood’s “plausible *naturalistic* conception of the world” according to which “there is a sense in which the natural facts determine the fundamental nature of the world”⁸⁰ to count as “science-friendly.” In other words, putting aside the other deep problems with the view, I think it would involve *too significant* a departure from this claim to count as “science-friendly.”

⁷⁹ (Wedgwood 2009, 3).

⁸⁰ (Wedgwood 2009, 10).

Given all of this, I thus conclude that casual non-naturalistic realists have not yet given us a closure thesis that we both a) have principled reason for holding is the real commitment we have about the causal closure of the natural in light of a “modern naturalistic worldview” and b) which can be used as the basis to defend the coherence a science-friendly version of non-naturalistic realism. Therefore, even if the causal non-naturalist finds a way of explaining how a causal version of non-naturalism can be intelligible as a version of non-naturalistic realism as such, it is not clear how (if at all) causal non-naturalism can be made into a viable science-friendly metaethical view.

5. Conclusion.

We began this paper with a basic challenge faced by all those working in metaethics who do not want to end up committed to radical skepticism about our ability to form correct normative judgments: namely, the challenge of providing a theory that is at once both *practically* and *theoretically* adequate. In this paper, I have argued that non-naturalistic realism has yet to show that it is consistent with a basic confidence in the content and methodology of the natural sciences. For those of us who take this confidence to be part of our broader understanding of the world, this amounts to showing that non-naturalistic realism has failed to meet this challenge of providing an account of ethics that is at once both practically and theoretically adequate. This, I think, gives us strong reason to reject non-naturalistic realism as a metaethical view.

At the same time, however, I should stress in conclusion that I do not think that it gives us *conclusive* reason to do so. This is for the simple reason that *all* the major contemporary types of metaethical theory face daunting challenges. It might be that the problems that non-naturalistic realism faces as a type of metaethical view are more severe than the competition. But, then again, it might not be. There is no way to figure this out aside from simply *doing more work* in metaethics. The aim of this paper, then, is not to establish that non-naturalistic realism will necessarily turn out to fail in the context of overall theory-choice in metaethics. Rather, it is to establish one of the central burdens that such a view faces. If I am right, then non-naturalists have failed to meet this burden as of yet. Moreover, there is good reason to hold that they might be unable to do so given the central defining theses of the view itself.

WORKS CITED

- Ayer, A.J. 1952. *Language, Truth, and Logic*. New York: Dover.
- Benacerraf, Paul. 1973. Mathematical Truth. *Journal of Philosophy* 70 (8):661-679.
- Bennett, Karen. 2008. Exclusion, Again. In *Being Reduced*, edited by J. Hohwy and J. Kallestrup. Oxford: Oxford University Press.
- Blackburn, Simon. 1998. *Ruling Passions*. Oxford: Clarendon Press.
- Braddon-Mitchell, David, and Frank Jackson. 1996. *The philosophy of mind and cognition*. Oxford, UK ; Cambridge, MA: Blackwell Publishers.
- Chalmers, David J. 1996. *The conscious mind : in search of a fundamental theory*. New York: Oxford University Press.
- Chalmers, David J., and Frank Jackson. 2001. Conceptual Analysis and Reductive Explanation. *Philosophical Review* (110):315-61.
- Crisp, Thomas M., and Ted A. Warfield. 2001. Kim's Master Argument. *Nous* 35:2:314-316.
- Dancy, Jonathan. 2006. Nonnaturalism. In *The Oxford Handbook of Ethical Theory*, edited by D. Copp. Oxford: Oxford University Press.
- Daston, Lorraine, and Peter Galison. 2007. *Objectivity*. Cambridge, Mass.: Zone Books ; Distributed by the MIT Press.
- Dretske, F.I. . 1981. *Knowledge and the Flow of Information*. Cambridge: MIT Press.
- Easwaran, Kenny. 2008. The Role of Axioms in Mathematics. *Erkenntnis: An International Journal of Analytic Philosophy* 68 (3):381-391.
- Enoch, David. 2007. An Outline of an Argument for Robust Metanormative Realism In *Oxford Studies in Metaethics Volume 2*, edited by R. Shafer-Landau.
- Field, Hartry. 1991. *Realism, Mathematics, and Modality*. New York City: Blackwell.
- Fine, Kit. 2001. The Question of Realism. *Philosophers' Imprint* 1.
- Galison, Peter, and David J. Stump. 1996. *The Disunity of science : boundaries, contexts, and power, Writing science*. Stanford, Calif.: Stanford University Press.
- Gibbard, Allan. 1990. *Wise Choices, Apt Feelings : A Theory of Normative Judgment*. Cambridge, Mass.: Harvard University Press.
- . 2003. *Thinking How to Live*. Cambridge, Mass.: Harvard University Press.
- Harman, Gilbert. 1977. *The Nature of Morality*. Oxford: Oxford University Press.
- . 1986. Moral Explanations of Natural Facts--Can Moral Claims Be Tested against Moral Reality? *Southern Journal of Philosophy (suppl.):*57-68.
- . 1997. Ethics and Observation. In *Moral discourse and practice : some philosophical approaches*, edited by S. L. Darwall, A. Gibbard and P. A. Railton. New York: Oxford University Press.
- . 2000. *Explaining Value: And Other Essays in Moral Philosophy*, Oxford: Clarendon Press.
- Hussain, Nadeem J. Z. 2008. Normativity and the Causal Order.
- . Forthcoming. *Norms in Action: Metaethics and the Neo-Kantian Critique*.
- Hussain, Nadeem J. Z., and Nishi Shah. 2006. Misunderstanding Metaethics: Korsgaard's Rejection of Realism. In *Oxford Studies in Metaethics, Volume 1*, edited by R. Shafer-Landau.
- Jackson, Frank. 1998. *From metaphysics to ethics : a defence of conceptual analysis*. Oxford: Clarendon.
- Jackson, Frank, and Philip Pettit. 1996. Moral Functionalism, Suervenience and Reductionism. *Philosophical Quarterly* 46.
- Kagan, Shelly. 1989. *The limits of morality*. Oxford: Oxford University Press.
- Kamm, F.M. 2007. *Intricate Ethics*. Oxford: Oxford University Press.

- Kim, Jaegwon. 1998. *Mind in a physical world : an essay on the mind-body problem and mental causation*. Cambridge, Mass.: MIT Press.
- . 2007. *Physicalism, or Something near Enough*. Princeton: Princeton University Press.
- Lewis, David. 1983. New Work for a Theory of Universals. *Australasian Journal of Philosophy* 61:343-377.
- Mackie, J. L. 1977. *Ethics : inventing right and wrong*. Harmondsworth ; New York: Penguin.
- Maddy, Penelope. 2000. *Naturalism in Mathematics*. Oxford: Oxford University Press.
- Margolis, Eric, and Stephen Laurence. 1999. *Concepts : core readings*. Cambridge, Mass.: MIT Press.
- McDowell, John Henry. 1998. *Meaning, Knowledge, and Reality*. Cambridge: Harvard University Press.
- McPherson, Tristram. Forthcoming. Against Quietist Normative Realism. *Philosophical Studies*
- . Manuscript. Ethical non-naturalism and the metaphysics of supervenience. <http://www.d.umn.edu/~tristram> (Draft of 3/25/2010).
- Moore, G.E. 1993. *Principia Ethica*. Cambridge: Cambridge University Press.
- Nye, Howard. 2009. *Ethics, Fitting Attitudes, and Practical Reason: A Theory of Normative Facts*: PhD Dissertation, University of Michigan, Ann Arbor.
- Parfit, Derek. 2006. Normativity. In *Oxford Studies in Metaethics* edited by R. Shafer-Landau. New York: Oxford.
- . Forthcoming. *On What Matters*. Vol. (Draft of 7/5/2009). Oxford: Oxford University Press.
- Pettit, Philip. 1991. Realism and Response-Dependence. *Mind: A Quarterly Review of Philosophy* 100 (400):587-626.
- Plunkett, David. Forthcoming. Expressivism, Representation, and the Nature of Conceptual Analysis. *Philosophical Studies*.
- Railton, Peter. 2005. Reply to Ralph Wedgwood. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 126 (3):501-508.
- Ross, W.D. 1930. *The Right and the Good*. Oxford: Clarendon Press.
- Scanlon, T. M. 2003. Metaphysics and Morals. *Proceedings and Addresses of the American Philosophical Association* 77 (2):7-22.
- . Forthcoming. *Being Realistic About Reasons: The 2009 John Locke Lectures*. Oxford.
- Searle, John. 1995. Consciousness, The Brain, and the Connection Principle: A Reply. *Philosophy and Phenomenological Research* 55:217-232.
- Shafer-Landau, Russ. 2005. *Moral Realism: A Defense*. Oxford: Oxford University Press.
- Smith, Michael. 1994. *The Moral Problem*. Cambridge: Blackwell.
- . 2004. Moral Realism. In *Ethics and the A Priori: Selected Essays on Moral Psychology and Meta-Ethics*. Cambridge: Cambridge University Press.
- Sturgeon, Nicholas. 2009. Doubts about the Supervenience of the Evaluative. In *Oxford Studies in Metaethics*, edited by R. Shafer-Landau. Oxford: Oxford University Press.
- Sturgeon, Nicholas L. 1988. Moral Explanations. In *Essays on Moral Realism*. Ithaca: Cornell University Press.
- Velleman, J. David. 2009. *How We Get Along*. Cambridge: Cambridge University Press.
- Wedgwood, Ralph. 2009. *The Nature of Normativity*. Oxford: Oxford University Press.
- Yablo, Stephen. 1992. Mental Causation. *Philosophical Review* 101:245-80.

Chapter Two: The Metaethical Role of Evaluative Attitudes

Introduction.

Many facts about what an agent should do seem to depend on facts about what she values. For instance, suppose that Rachel is deciding whether to spend her afternoon reading in a café or watching bowling on television. What should she do? It seems natural to say that the answer will depend in part on what Rachel herself values. For instance, if Rachel generally values reading in cafés but hates bowling, then, other things being equal, she should go read in the café.

On the other hand, it seems that there are some facts about what an agent should do that hold *independently* of any facts about what she values. For instance, suppose that Louisa thinks that the most valuable thing to do in the world is to poke puppies in the eyes and watch them cry. Now imagine that Louisa wants to sneak into an animal shelter in order to poke puppies in the eyes covertly. What should she do? It seems natural to say that poking puppies in the eyes is simply not the thing Louisa should do, *regardless* of any facts about what she values.

Many people – philosophers and non-philosophers alike – find this to be an enormously compelling first-order normative thesis about what Louisa should do. More generally, it is enormously compelling to think that the prohibition on Louisa’s poking puppies in the eyes is not hostage to *any* claims about Louisa or what motivates her. To see why this is so, suppose that we replace our talk about Louisa’s *values* with talk about her desires, plans, or unreflective normative judgments about what to do. In such cases, most of us still get the intuition that Louisa should not poke puppies in the eyes. In order to have a useful way of capturing this fact, let’s call attitudes such as values, plans, desires, and unreflective normative judgments *evaluative attitudes*. Furthermore, in order to capture an important feature that such attitudes have in common, let’s stipulate that by “evaluative attitude” we mean *any* attitude that tends to produce motivation for an agent when coupled

with her prosaic beliefs.¹

With this definition of “evaluative attitudes” in hand, it appears that the case about Louisa suggests the following first-order normative thesis about the nature of normative facts in the actual world:

Ethical Agent-Attitude-Independence: there are at least some normative facts about what an agent A should do that hold independently of any facts about A’s own evaluative attitudes.²

Most people are inclined to accept ethical agent-attitude-independence in light of cases such as the one about Louisa. Furthermore, if pushed to explain why Louisa should not poke puppies in the eyes, the basic structure of the answer often strikes people as obvious – namely, because *how things are for puppies also matters* for determining what Louisa should do and, clearly, *this sort of action causes puppies pain*.

Despite its intuitive pull, the thesis of ethical agent-attitude-independence invites a number of serious philosophical challenges. Perhaps most importantly, ethical agent-attitude-independence conflicts with another normative view that many find attractive. This is the view that facts about what an agent should do cannot be ones that the agent herself would be *fundamentally alienated from* or *motivationally uninfluenced by* if she were aware of them.³ If this is so, then, when taken seriously as a normative view, this leads to the thesis that if a normative claim about what an agent should do is to be correct, then the ultimate facts that make it so must be facts at least in part about that agent’s own evaluative attitudes. If we hold fixed the idea that there are some normative facts in the actual world about what agents should do, then it thus lends support to the thesis that all normative facts about what an agent A should do are dependent on facts about A’s own evaluative attitudes. Call this thesis *ethical agent-attitude-dependence*.

How should we decide what to say in this debate over ethical agent-attitude-independence? It is a working assumption in many parts of contemporary normative ethics

¹ Which attitudes count as evaluative attitudes in this functional sense is an open question – and, importantly, will depend on one’s views in philosophy of mind and action. My use of the term “evaluative attitudes” draws on (Street 2008). It should be emphasized that I do *not* intend my use of the term “evaluative” to suggest that all evaluative attitudes concern values.

² As I qualified when I said this thesis is “about the nature of normative facts in the actual world,” I intend this to be a thesis about the normative facts *in the actual world*. By “world” I mean “possible world,” which, for our purposes here, we can understand as a complete description of the way things are.

³ For paradigmatic statements of this view, see (Williams 1981), (Railton 1986a), (Korsgaard 1996).

that, in general, first-order normative debates can and should be decided without appeal to distinctively *metaethical* views about the metaphysics, epistemology, and semantics of practical normative thought and talk.⁴ In the case of *at least* this debate, however, this assumption blinds us to the way in which our metaethical choices significantly constrain what we can reasonably say in normative ethics and vice versa. For, in this paper, I will argue that *if* one accepts a certain view in metaethics – namely, the *naturalistic realist* view, according to which, roughly, normative beliefs are genuine truth-apt beliefs solely about properties capable of being studied by the natural sciences – then one needs to accept ethical agent-attitude-dependence. Many leading naturalistic realists have believed that this is not the case. Instead, they have held that both ethical agent-attitude-independence and ethical agent-attitude-dependence are normative positions that are compatible with viable versions of this type of metaethical view. This, I argue, is an illusion. In order to maintain a viable form of naturalistic realism, one must reject ethical agent-attitude-independence.

One obvious question follows: if indeed naturalistic realism rules out ethical agent-attitude-independence, does this give us reason to reject naturalistic realism? Or does it instead give us reason to reject ethical agent-attitude-independence? My main goal in this paper is not to argue for one conclusion or the other. Rather, my main goal is to argue for the claim that forces us to make the choice between them; namely, the claim that the first-order normative thesis of ethical agent-attitude-independence needs to be abandoned if one is to develop a viable form of naturalistic realism in metaethics.

That being said, I do have an additional goal in this paper: to convince the reader that this is not an easy choice to make. This is important because some people find the thesis of ethical agent-attitude-independence so enormously compelling that they are inclined to see the main thesis of this paper as an obvious *reductio* on naturalistic realism. I think that this would be a serious mistake. For I think that naturalistic realism can in fact remain a viable type of metaethical view – indeed, that it might plausibly even turn out to be the *correct* type of metaethical view – even if it undermines ethical agent-attitude-independence. In

⁴ This assumption is explicitly stated and defended by none other than John Rawls – arguably the most influential philosopher of the second half of the twentieth-century to write on methodology in normative ethics. See (Rawls 1999b).

order to explain why this is so, I conclude the paper by sketching the form of naturalistic realism that I myself favor in light of my main argument.⁵

1. Ethical Agent-Attitude-Independence and Metaethics: What's at Issue.

The main argument of this paper is that there is an irresolvable tension between the first-order normative thesis of ethical agent-attitude-independence and the metaethical view of naturalistic realism. To see the broad outlines of what is at issue here, as well as a glimpse of why this tension is worth caring about, let me start with a useful contrast between naturalistic realism and two of the other dominant types of view in metaethics, namely, *non-naturalistic realism* and *quasi-realist expressivism*.

According to the *non-naturalistic realism* of the sort first advocated by G.E. Moore – and now defended by contemporary philosophers such as Derek Parfit, David Enoch, and Jonathan Dancy – there is a primitive normative property that only some acts have.⁶ We can have beliefs about this property, just as we have beliefs about any other property. Yet, since it is a simple non-reducible property that is intrinsically normative, it is not one that we can investigate by scientific means or use in scientific explanations. It is instead a metaphysically *sui generis* property. If some acts have this non-naturalistic property independently of any facts about an agent's attitudes, then this would be no mystery at all; the property is such

⁵ Following Street, the central debate that I am interested in is sometimes referred to as one about “mind-independence” rather than as one about “agent-attitude-independence.” See, for instance, (Street Forthcoming-a), (Street Forthcoming-b), and (Gibbard Forthcoming). Since the label “mind-independence” leaves it vague *whose* mind matters (e.g. agent vs. speaker) as well as *which aspects* of the mind matter (e.g. beliefs vs. desires), I prefer my label. My label is more cumbersome. However, I think it helps clarify what both Street and I agree is the central philosophical issue at hand for normative ethics. Street also sometimes puts the debate as one between “realist” and “anti-realist” positions in ethics (Street 2006). I think this is misleading given that there are other, more canonical ways of using these terms in metaethics that can be easily confused with this debate. It should also be noted that the debate over ethical agent-attitude-independence and ethical agent-attitude-dependence is closely related to the debate between those that accept that there can be what Bernard Williams’ calls “external” normative reasons for action and those that think that all normative reasons are what he calls “internal.” (Williams 1981). However, since this terminology has become very closely tied to Williams’s particular way of framing the issue – and since I think Street’s way of putting it does a better job of locating the key philosophical issue at hand – I will stick with my Street-inspired terminology in this paper instead of Williams’s. Finally, drawing on (Schroeder 2007), it should also be noted that there is a good sense in which all ethical agent-attitude-dependent views can be taken to be “Humean” ethical theories and all ethical agent-attitude-independent views can be taken to be “anti-Humean” ones. However, given the diversity of ways in which the term “Humean” gets used in ethics, I will not use this label in the paper.

⁶ (Moore 1993), (Parfit Forthcoming), (Enoch 2007), (Dancy 2006).

that there is no reason to expect that ethical agent-attitude independence is any more suspect than its denial.⁷

Or, to take another prominent example, consider recent versions of *quasi-realist expressivism* of the sort advocated by Simon Blackburn and Allan Gibbard.⁸ According to expressivism, the attitudes that we use normative terms to express should not be considered *in the first instance* beliefs with such-and-such propositional content, as they would on a so-called “descriptivist” account (e.g. of the sort that non-naturalistic realists endorse). Rather, these states of mind should be understood as some variety of what are sometimes called “noncognitive” attitudes: e.g. a desire, a plan, or an attitude of approval or disapproval. In turn, these attitudes might earn the right to be called “beliefs” in a sort of quasi-realist or minimalist way, as both Blackburn and Gibbard suggest. However, what is important is that, for the expressivist, it is the noncognitive attitude that is explanatorily primary and that we express when making normative judgments. As with the non-naturalistic realist, there seems to be no reason why an expressivist cannot advocate ethical agent-attitude-independence as a first-order normative thesis. In so doing, as both Gibbard and Blackburn do, the expressivist would simply be expressing something such as a plan for acting in certain ways under certain circumstances – for instance, the plan never to poke puppies in the eyes for fun regardless of what set of desires one has.⁹

Non-naturalistic realism and expressivism thus share the following feature: they appear to allow us to vindicate ethical agent-attitude-independence as a first-order normative thesis. For even though they do not themselves *settle* the issue of whether or not one should accept the first-order normative thesis of ethical agent-attitude-independence, they both have the basic resources to show that it is coherent – and, indeed, perhaps *plausible* – to do so. Both non-naturalistic realism and expressivism, however, face a number of well-known problems that have led many to reject both views. Non-naturalistic realism, for instance,

⁷ I use the term “non-naturalistic” here rather than Moore’s original “non-natural” in order to avoid confusion with the “natural” vs. “non-natural” properties distinction that is used in metaphysics (for instance, that is used when discussing so-called “grue-some” properties). See (Lewis 1983). This distinction between “non-naturalistic” and “non-natural” will become especially important later in this paper.

⁸ (Blackburn 1998) and (Gibbard 2003).

⁹ It should be noted that this does not mean that the first-order normative thesis of agent-attitude-independence *necessarily* turns out to be correct if one accepts either non-naturalism or expressivism. One could accept those metaethical positions and still argue against this thesis. Indeed, this is precisely what Street does in (Street Forthcoming-a) and (Street Forthcoming-b).

faces the difficulty of explaining a) how we could ever come to know such non-naturalistic facts, b) what exactly non-naturalistic facts are, such that they should not be dismissed as metaphysically odd, and c) why our judgments about such facts would have any motivational importance for us.¹⁰ And expressivism, for instance, faces the difficulty of explaining a) why the surface structure of normative discourse seems to participants to be one involving straightforward belief, b) what exactly it is to “express” a non-cognitive attitude, as well as c) more technical issues such as those about embedding and negation.¹¹ There is obviously widespread disagreement about how important such problems are. But they should at least give us *prima facie* reason to be interested in views that do not share these problems.

If one rejects both non-naturalistic realism and expressivism, an obvious alternative is to think that normative judgment consists in belief about some set of naturalistic facts – and hence to accept what I am calling *naturalistic realism*. And, importantly, in contrast to non-naturalistic realism and expressivism, it is not clear how exactly a naturalistic realist view in metaethics can support ethical agent-attitude-independence in the way that its rivals can. To put it another way, it is not clear that naturalistic realism can show why ethical agent-attitude-independence is coherent (let alone plausible). Unlike non-naturalism and expressivism, naturalistic realism needs to explain what naturalistic properties we are in general thinking and talking about when we engage in normative thought and talk. To be sure, there are answers here that a naturalistic realist could give that *would* leave us room to coherently endorse ethical agent-attitude-independence as a first-order normative thesis. For instance, consider the view that all normative thought and talk consists in belief about the distribution of rocks on Easter Island. Such a view is a form of naturalistic realism. And it clearly would leave us room to endorse ethical agent-attitude-independence. Such a view, however, is obviously not viable as a metaethical account. The question then is this: is there a viable account of what naturalistic properties we are thinking and talking about that would

¹⁰ For representative statements of each type of worry, see (Jackson 2000), (Ayer 1952), and (Korsgaard 2003) respectively. See (Shafer-Landau 2005) and (Parfit Forthcoming) for responses to these worries.

¹¹ For representative statements of each type of worry, see (Smith 2004b), (Jackson and Pettit 1998), and (Schroeder 2008) respectively. See (Gibbard 2003) and (Blackburn 1998) for responses to such worries. Note that there are important issues here that I am passing over about what makes something a “noncognitive” attitude – and, indeed, whether or not this is helpful terminology at all. On this latter issue, see (Gibbard 2003) and (Dreier 2004).

both leave us room to accept ethical agent-attitude-independence as well as provide the basis for a viable version of naturalistic realism in metaethics?

In this paper, I will argue that the answer is “no.” In short, I will argue that the attempt to vindicate ethical agent-attitude-independence as a naturalistic realist in metaethics undermines the ability of the naturalistic realist to meet core metaethical desiderata. Since it is beyond the scope of this paper to deal with all forms of metaethical naturalistic realism that might leave open ethical agent-attitude-independence, I will focus my work on what many take to be the leading contemporary version of such a view: the type of ideal attitude theory exemplified by the work of Michael Smith. I will then explain why I think my arguments about this specific species of metaethical naturalistic realism are instructive with respect to the genus as a whole.

2. Clarifying the Job Description for Naturalistic Realism in Metaethics.

In order to set up the argument in what follows, we first need to get clear on some preliminary points about the nature of metaethics and of naturalistic realism in particular.

Philosophers mean different things by the term “metaethics.” In this paper, I will use this term to pick out a specific philosophical project, one that many in mainstream contemporary metaethics are concerned with and that, moreover, many consider to be *the* fundamental project of metaethics. What is important in what follows is not whether the use of this term matches the use the term by all philosophers in ethical theory. Rather, what is important is *what the project is* that I pick out with this term – and, in turn, how successfully trying to engage with this project given certain theoretical commitments undermines ethical agent-attitude-independence.

According to the definition of “metaethics” that I will adopt, in contrast to the task of *normative ethics*, which aims to figure out the correct answers to normative questions, the task of *metaethics* is to make sense of what we are doing when we engage in normative thought and talk. Or, more specifically, since normative thought and talk concerns many different topics – say, for instance, the question of what to believe – metaethics is actually a subset of this broader metanormative project; namely, it is the part of this project whose task

is to make sense of practical normative thought and talk that is specifically about *what to do* (which I will henceforth just refer to as “normative thought and talk”).¹²

The task of metaethics can be broken down into two parts. On the one hand, there is what we can call the *non-fundamental* metaethical task of explaining how different parts of normative thought and talk are internally related to each other. For instance, we might want to know how talk about goodness is related to talk about obligation or reasons. On the other hand, there is what we can call the *fundamental* metaethical task of explaining how *all* of ethical thought and talk – and the ontological commitments involved in that thought and talk – are integrated into a broader conception of reality.¹³ In order to succeed at this fundamental metaethical task, a theory cannot assume certain normative thoughts as part of its explanation. Rather, it must treat the entire practice of engaging in normative thought and talk as an explanandum rather than as an explanans.¹⁴

Fundamental metaethical theories are generally defended in combination with specific non-fundamental metaethical views. For instance, some fundamental metaethical theories rely on the idea that reasons can be reductively explained in terms of value, whereas others take the opposite view. For the purposes of this paper, it will not matter which non-fundamental views are correct. Rather, what will be crucial is how the naturalistic realist makes sense of the metaphysics of what is picked out by the most explanatorily basic normative concept(s) – whatever exactly the most basic normative concept(s) turn out to be.

As I will use the term “naturalistic realism” in this paper, naturalistic realism in metaethics is the conjunction of the following three theses:

¹² For other similar glosses on the task of metaethics, see (Gibbard 1990), (Smith 1994), and (Smith 2005).

¹³ I thank Howard Nye for first suggesting to me this helpful way of drawing the distinction. For Nye’s own formulation of this distinction, see his PhD dissertation (Nye 2009).

¹⁴ I draw on (Hussain 2004, 2) in my formulation of the task of fundamental metaethics and what it must involve in order to be successful. It should be noted that this characterization is complicated by the fact that many philosophers take judgments about meaning and/or mental content to themselves be normative judgments. For reflection on this issue see (McDowell 1994), (Railton 2003), (Brandom 1994), and (Gibbard 2003). Since this might mean that *any* metaethical view would turn out to take some normative thoughts as given as an explanans, there is thus an important question whether or not the view of metaethics that I have stated is really feasible – or, at the very least, whether or not it is actually more controversial than intended. These issues become particularly complicated when we think about the nature of *metanormative* theory as a whole. Yet, for the purposes of this paper, since we are focused on *metaethics* in particular, I am going to sidestep these issues by stipulating that a view can count as a fundamental metaethical one iff it does not take any normative judgments about *what to do* or *how to revise one’s evaluative attitudes* as an explanans – thus, for now, leaving it open whether it can tacitly take normative judgments of a different sort for granted.

Naturalistic Normative Properties. The fundamental normative properties are naturalistic properties.¹⁵

Descriptivism. To make a normative judgment is to express an attitude that, at the most basic explanatory level, consists in straightforward truth-apt belief.

Realism. Some normative properties are instantiated in the actual world. This makes it the case that some possible normative judgments that concern the actual world are correct.¹⁶

In order to regiment the discussion that follows, following one standard line of contemporary thought, let's assume that the most basic notion for practical normativity is that of a *normative reason for action* (which I will henceforth just call a "normative reason"). As Scanlon puts it, this is the concept of a consideration or fact that "counts in favor of" or "calls for" a course of action in the sense that it *justifies* rather than explains that action.¹⁷ If

¹⁵ By a "fundamental" normative property, I mean a normative property that is metaphysically basic in terms of its relationship to other normative properties. More specifically, it is a normative property that can be used to explain the nature of other, further non-fundamental normative properties but that itself cannot be explained by the nature of any other normative properties. For our purposes here, it does not matter what normative properties one takes to be basic in this sense (a sense, it should be noted, that Moore was after when he called goodness a "simple" property). For instance, it won't matter whether or not one takes the class of fundamental normative properties to include the properties of *being good* and *being a normative reason* and *being a duty*, or whether one thinks only one of these properties (or another property altogether) is needed to explain the basic nature of these other properties.

¹⁶ It should be noted that I use the phrase "concern the actual world" rather than "about the actual world" to leave room for Gibbard's expressivist view according to which normative judgments express plans about what to do in the actual world rather than judgments that are strictly speaking *about* the world. It should also be noted that if realism is true, then correct normative judgments that concern the actual world will be *non-trivially* correct in virtue of the fact that there are instantiated normative properties in the actual world. In contrast, imagine that there are no normative properties instantiated in the actual world. If so, then the statement "Andy does not have a normative reason to eat the pizza" is trivially correct in the actual world.

¹⁷ (Scanlon 2000). This formulation is meant to distinguish *normative reasons* (facts that justify an agent's action whether or not that agent is aware of them) and *motivating reasons* (facts that explain an agent's action in either teleological or causal terms). On this distinction, see (Smith 1994) and (Darwall 2003). It should be noted that there is an important question of whether normative reasons should be identified with the facts themselves (the things in virtue of which propositions are true or false) or instead with the propositions. For the purposes of this paper, I am going to leave this issue to the side since the differences between these two theories does not make any difference for the core issues that I will address. See the introduction of (Schroeder 2007) for an insightful critical discussion of how to make more precise the thought that normative reasons "count in favor" of action.

so, this means that, according to naturalistic realism, the property of *being a normative reason for action* is a naturalistic property that is instantiated in the actual world.

Many non-naturalistic realists and naturalistic realists alike agree that if the property of being a normative reason (or *whatever* the most fundamental normative property is) is itself taken as primitive, then this property won't count as naturalistic. In short, this is because if there were a brute non-reducible property of simply *being a normative reason*, then this property seems as if it would be metaphysically different in kind from the other properties studied by the natural and social sciences. It seems, in other words, that it would need to be metaphysically *sui generis* as the *non-naturalistic* realist asserts. After all, none of the properties studied in the natural and social sciences look to have what J.L. Mackie called "to-be-pursuedness"¹⁸ fundamentally built into them in the way that irreducibly normative properties would seem to need. Furthermore, nowhere in the naturalistic sciences does it look as if we need to make reference to some unreduced property of "being normative" in order to explain the workings of the natural world. Therefore, the naturalistic realist seems to face the *metaphysical* task of explaining what this property is *solely* in terms of further non-normative naturalistic properties. Or, more precisely, if she wants to restrict herself to those properties that are already *widely agreed upon* to be naturalistic, then this is the metaphysical task that that she faces.¹⁹

Many philosophers, such as Frank Jackson, hold that the only way for a descriptivist to accomplish this metaphysical task will be to show that the *concepts* that we use in normative thought and talk are themselves naturalistic ones. However, while this might be so, this is not one of the stated starting commitments of the view itself. Rather, the commitment it starts with is a *metaphysical* one about the nature of normative properties.²⁰

¹⁸ See (Mackie 1977).

¹⁹ It should be noted that not *all* naturalistic realists agree that this is their metaphysical task. For instance, in parallel to *non-reductive physicalism* about mental properties, which holds that mental properties are physical properties but not reducible to any other physical properties, some naturalistic realists in metaethics hold that normative properties are naturalistic ones even though they are not reducible to any other naturalistic properties. This is, for instance, the view of (Sturgeon 2006). Against Sturgeon, I think there are arguments of the sort given by (Jackson 2000) that show that such a proposal is not consistent with the basic commitments of naturalistic realism. This is a huge issue that is beyond the scope of this paper to deal with. In order to sidestep this debate, for the purposes of this paper, one can take it as simply stipulated that, in order to qualify as a naturalistic realist, one needs to posit only *uncontroversial* naturalistic properties.

²⁰ There is an important question here of what exactly makes a property a non-normative one. For my purposes here, following (Jackson 2000) and (Gibbard 2003), I will take it that a property is a

The metaphysical task on the table is a paradigmatic instance of what Frank Jackson calls a *location problem*. In Jackson's terms, a location problem consists in figuring out how to make sense of some supposed feature of the world (feature X) solely in terms of the ingredients of some further basic account of what exists (where X is *not* itself mentioned *de dicto*).²¹ For instance, it would be a location problem to figure out how, if at all, we can make sense of what the Taj Mahal is simply in terms of the limited metaphysical ingredients we are given by physics. In our case, the feature X we are concerned with is the property of *being a normative reason*. And, insofar as we want to embrace metaphysical naturalism about practical normativity, we need to use only *non-normative naturalistic* properties in the basic account.

In order to successfully solve this location problem, we need to capture what we intuitively think of as actual *normative* thought and talk – or, at the very least, we need to capture *much* of what we intuitively think of as normative thought and talk. After all, we want to solve the location problem for the property of *being a normative reason* and not, for instance, for the property of *being a tree*. This sets up a number of important constraints on any proposed solution to our location problem. I will here mention two of the most important ones.

The first constraint stems from the fact that when we think or talk about what to do, we don't take all normative views to be on a par. For example, if Bugsy thinks that Rachel should watch bowling on TV rather than read in the café solely on the grounds that *there are green bowling balls that she would then see*, we will normally think he has made an incorrect normative claim. Indeed, we seem to presuppose that there are better or worse normative judgments *whenever* we deliberate about what to do. For if all judgments were in fact on a par, it is hard to see why we would bother deliberating about what to do rather than just picking at random. This suggests that, in order to correctly capture normative thought and talk, a metaethical view needs to capture the idea that there is *some* standard of correctness that we use when we make normative judgments. To put it another way, it needs to capture the fact that we think there is a potential gap between any token normative judgment about a topic

non-normative one iff it is capable of being fully captured by non-normative terms – that is, roughly, terms such as “rock” or “chair” that are *not* infused with so-called “ought” terminology. I will not, however, take a stand in this paper on which specific terms those are.

²¹ (Jackson 2000).

and a *correct* normative judgment about that topic. Call this the *correctness constraint*.²² A view that would fail to meet this constraint, for instance, would be one that claimed that there simply is *no* standard of correctness at all in making normative judgments. Or, to take another example, another view that would fail to meet this constraint would be one that claimed that *any* normative judgment one could make would be correct.

The second constraint stems from the fact that we find intelligible a range of first-order normative views. For instance, even if we are convinced that they are mistaken, we find it intelligible when someone claims that abortion is never permitted or that hedonic consequentialism is true. Insofar as we take all of these people to be talking about normative reasons, this means that a viable fundamental metaethical view needs to make sense of how there could be a number of different intelligible (even if radically mistaken) first-order views about normative reasons. Call this the *first-order diversity constraint*. A view that would fail to meet this constraint, for instance, would be one that claimed that it was an obvious analytic truth that normative judgment consists in judgment about how to maximize the aggregate amount of pleasure in the world.

To sum up, to succeed as a naturalistic realist view in metaethics, a view needs to:

- (A) Give an account of the content of normative thought and talk that treats the entirety of normative thought about how to live as an explanandum rather than explanans.
- (B) Accept the descriptivist thesis, the naturalistic properties thesis, and the realist thesis.
- (C) Satisfy the *correctness constraint*.
- (D) Satisfy the *first-order diversity constraint*.

3. *Smith's Ideal Attitude Theory*.

With this understanding of the aims of metaethics in hand, we can now return to the main question of this paper: is there a way to develop a viable form of naturalistic realism in metaethics that meets all these goals, at the same time that it makes room for the possibility

²² Given that I am purposefully leaving it open whether the standards of correctness here are ones involving genuine truth and falsity, this should be something that is relatively uncontroversial in contemporary metaethics.

of ethical agent-attitude-independence? According to Michael Smith, the answer here is “yes.”

The basic platitude driving Smith’s ideal attitude theory is that when an agent wonders about what she should do, one good option would be for her to ask for advice from those that are better situated than she to think about this matter. For instance, suppose that I am a novice tennis player wondering which tennis racket I should use. It would make sense is for me to ask a tennis pro who teaches tennis – not, as it were, to emulate the tennis pro by attempting to use the racket that she uses (her racket might be too advanced for me), but rather to hear what advice she has for me given that I am just starting out. As this example illustrates, our starting platitude at least initially seems perhaps best read as a first-order normative view that one could either accept or deny depending on one’s views about how to live. However, Smith suggests that this platitude can in fact be used to fix *what it is* for something to be a normative reason. In short, this is because Smith holds that it is incoherent to take oneself to be thinking about *normative reasons* as such and deny that one would have normative reason to follow the advice of an agent who is *perfectly* placed to give one advice.

Smith develops this idea as follows. To start with, he suggests that for any specific non-ideal agent A, we can conceive of another agent who is *perfectly placed* to give non-ideal A advice.²³ In rough terms, such an agent would be one whose desires were, as Smith puts it, “completely beyond reproach from the point of view of reasoned criticism.”²⁴ Call such an agent whose desires meet this condition an “ideal” or “fully rational” agent. We can then imagine such an agent taking all of non-ideal A’s own defects into account and forming desires about what to do on the supposition that she herself were in non-ideal A’s “own shoes,” so to speak. For each non-ideal agent A, the agent best suited to give A advice, claims Smith, would be one that is *in some important sense* a version of A herself – otherwise, to put it roughly, she might have desires about what non-ideal A do that do not concern promoting A’s own interests, but rather the interests of another agent. Combined with Smith’s previous thesis about the connection between advice and normative reasons, this delivers the following:

²³ As he puts it, “the core idea is that facts about our normative reasons for action ... are facts about what we would advise ourselves to do if we were perfectly placed to give ourselves advice.” (Smith 1997, 88).

²⁴ (Smith 1997, 88-89).

Smith's Account of Normative Reasons: An agent *A* has normative reason to ϕ in circumstances *C* iff an ideal version of *A* would desire that non-ideal *A* ϕ in *C*.²⁵

What should we make of this account? Can it be used to deliver a form of naturalistic realism in metaethics? And, in particular, can it be used to deliver a *viable* form of naturalistic realism that also secures the plausibility of ethical agent-attitude-independence as a first-order normative view?

In order to answer this question, we need to get a better handle on what exactly Smith's view amounts to. At first blush, Smith's theory might appear to be openly committing itself to the idea that only ethical agent-attitude-*dependent* first-order normative views are viable. This is because it explains facts about normative reasons in terms of facts about evaluative attitudes. However, what is important to note is that, for Smith, the evaluative attitudes here are ones of an agent in a hypothetical situation – namely, in the situation of being so-called “fully rational.” Thus, those evaluative attitudes might not have anything to do with the *sorts* of facts about evaluative attitudes that matter for the debate about ethical agent-attitude-independence: namely, facts about the agent's *actual* attitudes as they currently stand. Until we know what decides the standards for being fully rational – namely, whether those are further facts about an agent's own evaluative attitudes or not – it won't be clear whether his view delivers ethical agent-attitude-dependence or not.

In the second stage of developing his theory, Smith argues for an account of what it is to be fully rational, according to which this can be discovered by *a priori* investigation into what a psychology needs to be like in order to be what he calls “immune from reasoned criticism.”²⁶ Smith claims that the facts about what such an ideal psychology is like are not set by facts about the current attitudes of agents. Rather, they are set by the substantive issue of what it would be to be an agent with a properly functioning psychology.

Smith himself does not tell us exactly what this would consist in. Rather, he speculates about which norms about desire-formation might plausibly be necessary to follow in order to count as fully rational. To get a sense of what sort of norms Smith has in mind

²⁵ This is a summary of the way in which Smith himself sums up his view in (Smith 1997, 88-90). It should be noted that Smith thinks it is possible for ideal *A* to think about what *ideal A* should do. Because of this, “non-ideal *A*” in this account should strictly speaking be read as “*A* (as she actually is).” This, however, is a limit case that I will not discuss. For ease of presentation, I will therefore stick with using the phrase “non-ideal *A*” in future summaries of Smith's view.

²⁶ (Smith 1997, 91).

here, let's look at a list of such norms that he considers in his recent paper "Desires, Values, Reasons, and the Dualism of Practical Reason."²⁷ For our purposes in this paper, not much hangs on the specific details of each proposed norm, some of which are put in more complicated terms given specific debates in ethical theory that are not our main concern here. Rather, at this point, what is important to notice is that i) some of these norms are more controversial than others and ii) it is not entirely clear how to interpret each norm. In what follows, I will quote directly Smith's statement of the norm and then offer my own brief explanation of it. For our purposes here, let "RR" read "rationality requires."

R1: RR (If someone has an intrinsic desire that p and a belief that he can bring about p by bringing about q, then he has an instrumental desire that he brings about q). This is a version of the familiar means-ends principle.

R2: RR (If someone has an intrinsic desire that p, and an intrinsic desire that q, and an intrinsic desire that r, and if the objects of the desires that p and q and r cannot be distinguished from each other and from the object of the desire that s without making an arbitrary distinction, then she has an intrinsic desire that s). In essence, this norm is meant to rule out desiring one thing and not desiring another if there is no non-arbitrary difference between the objects of these desires.

R3: RR (If someone has an intrinsic desire that p, then either p itself is suitably universal, or satisfying the desire that p is consistent with satisfying desires whose contents are themselves suitably universal). This norm is meant to capture the basic idea behind the Kantian categorical imperative.

R4: $\exists p \exists q$ RR (If someone believes that p, then she has an intrinsic desire that q). This norm is meant to state that people are rationally required to have desires with certain contents if we have beliefs with certain contents. For example, one would be committed to R4 if one believed the following: if I believe that a certain episode of future agony will be my own, then I am rationally required to desire avoiding it.

R5: $\exists p$ RR (People do not desire that p). This norm is meant to state that there is at least one desire that people are rationally required *not* to have. Perhaps, for instance, one is rationally required not to have the desire to destroy the entire world for the sake of preventing one's finger from being scratched.

R6: $\exists q$ RR (People desire that q). This norm is meant to state that there is at least one desire that people are rationally required to have. Perhaps, for instance, we are rationally required to desire the promotion of happiness.

²⁷ (Smith 2009, 101-102). See also (Smith 2009b).

Smith does not commit himself to which of these norms he actually accepts – or how long the list of norms R1-RN is that we should accept. These norms are meant to be *examples* of plausible candidates for what rationality requires.

What concerns us here is not which specific candidates are correct, but rather what explanation Smith can give for what it would take for a specific candidate to be correct in the first place. In other words, the real question is why *any* norm would survive on the eventual list of R1-RN at all. What is important here is that Smith *denies* that the norms here are fixed by further facts about the evaluative attitudes of agents. This means that his view does not rule out the possibility of endorsing ethical agent-attitude-independence any more than non-naturalistic realism or expressivism.

It should be emphasized why this is so. In *The Moral Problem*, Smith claims that A has a normative reason to ϕ in C iff *all* ideal agents would converge in their desires about what non-ideal A do in C. Call this the *convergence thesis*.²⁸ At least in terms of its initial presentation by Smith in *The Moral Problem*, the convergence thesis can reasonably be read as a constraint above and beyond the norms R1-RN of what rationality requires. Otherwise, one would only need to mention *one* ideal agent rather than all ideal agents. More recently, however, Smith has suggested that the convergence thesis captures the idea that, in order for there to be normative reasons, some group of the principles in R1-RN (perhaps R3 or R4) has to be true – and, hence, the core idea is simply of a perfectly functioning psychology of an agent as such.²⁹ In what follows, for ease of presentation, I will stick with this proposed way of thinking about the convergence thesis and drop thinking of it as a constraint beyond the norms of R1-RN. Smith’s convergence thesis is often what is discussed as important about his view. Yet, no matter how one reads it, the convergence thesis is itself *not* what is important for us here. Rather, what is important is that the conditions R1-RN that Smith thinks sets the ideal conditions do not obtain because of facts about agent A’s own evaluative attitudes.

²⁸ It should be noted that Smith is sometimes misinterpreted as holding the convergence thesis only for specifically *moral* reasons. This is not so. At least in *The Moral Problem*, he holds it for all practical normative reasons. Moral reasons are there identified as normative practical reasons with a certain sort of distinctive substance or content.

²⁹ (Personal communication). Smith also proposes this basic picture in (Smith 2007) – an unpublished paper delivered at the 2007 “Norms and Analysis” conference at the University of Sydney.

The reason for this is as follows. On the one hand, certain ways in which R1-RN might be fleshed out would undermine ethical agent-attitude-independence. For instance, ethical agent-attitude-independence would be undermined if there was a norm R777 that baldly stated that rationality requires that no rational agent A desire that another agent B ϕ unless B herself desires to ϕ . On the other hand, other ways in which R1-RN might be fleshed out would vindicate ethical agent-attitude-independence. For instance, ethical agent-attitude-independence would be vindicated if R6 were fleshed out to require that all rational agents desire to promote happiness. Therefore, as long as we hold fixed the fact that which norms obtain is *not* fixed by further facts about the evaluative attitudes of agents, then there is no general reason to hold that all the norms that fully rational agents revise their desires according to are standards set by their current non-ideal evaluative attitudes. In turn, there is no general reason to think that in order for an ideal agent to desire that a non-ideal agent ϕ , the non-ideal agent must also have an evaluative attitude that also favors her ϕ -ing. Hence, given the basic schema of Smith's metaethical ideal deliberation theory, there is no general reason to deny that non-ideal agents sometimes have normative reasons that obtain because of facts independent of anything about their own evaluative attitudes. The upshot, then, is that even if Smith's style of view would not itself vindicate ethical agent-attitude-independence by itself, it does appear to allow one to argue coherently for such a view; and, indeed, given some of his suggestions for how to flesh out norms such as R6, it gives us some prima reason to hold that such arguments might in fact be plausible.

4. The Tension Between Ethical Agent-Attitude-Independence and Smith's Naturalistic Realism.

On the face of it, Smith's metaethical theory appears to be able to meet our five goals for a successful naturalistic realist theory in fundamental metaethics (stated at the end of section two), while also securing the coherence of ethical agent-attitude-independence as a first-order normative view. But does Smith's style of theory actually provide a viable form of metaethical naturalistic realism? In what follows, I will argue that it does not. Moreover, I will argue that once we appreciate the reasons why it does not, we will see that, in order to be viable, a naturalistic realist view in fundamental metaethics needs to hold the thesis that the property of *being a normative reason for A* is ultimately a property that concerns what A's evaluative attitudes favor. Somewhat more precisely, I will argue that it needs to hold the following *metaethical* thesis:

Metaethical Agent-Attitude-Dependence: to make a normative judgment that an agent A should ϕ in C is to express the belief that A's own evaluative attitudes favor A ϕ -ing in C.

This thesis analytically entails that if there are any normative facts about how given agents should act, then those facts must be ones that concern what is favored by the evaluative attitudes of the given agent in question. If we hold fixed the anti-skeptical idea that there are normative facts in the actual world, which is something that all naturalistic realists in metaethics must grant given realist thesis that is partly constitutive of their view, it thus turns out that if one accepts *metaethical* agent-attitude-dependence, then one must also accept *ethical* agent-attitude-dependence. Thus, if my argument is right that naturalistic realists need to embrace *metaethical* agent-attitude-dependence, then it will turn out that they need to embrace *ethical* agent-attitude-dependence as well.

4.1. *Satisfying the Aims of Fundamental Metaethics.*

Take the complete list of norms R1-RN (whatever those turn out to be) that Smith thinks captures the norms that one needs to follow to be in the ideal conditions for desire-formation. What is the status of these norms? More specifically, what is it about these norms that makes them fit to be a part of this list? In this section, I will argue that in order to answer this question effectively, Smith needs either a) to abandon naturalistic realism or b) to abandon ethical agent-attitude-independence. To begin to see why this is so, let's start with some things that Smith *cannot* say about the status of these norms if he holds fixed his aspiration to deliver a form of naturalistic realism.

First of all, he cannot say that they are just the norms that we have the most *normative reason* to follow – and that perhaps could then be used to explain why we have normative reason to follow other norms. That move is available to someone who advances an ideal attitude theory as a type of first-order normative view (e.g. Stephen Darwall). But it won't work for someone such as Smith who wants to solve the location problem for normative reasons and thereby provide us with the naturalistic basis for a descriptivist metaethics. For, as Darwall himself emphasizes, this first-order normative view is fully compatible with both expressivism and non-naturalistic realism.³⁰

³⁰ (Darwall 1983).

Secondly, Smith cannot just make the metaethical claim that these norms R1-RN are responsive to some brute normative facts about what norms we should have. For, if what I said earlier in this paper was correct, then positing such brute normative facts is not compatible with maintaining a purely naturalistic ontology.

Finally, and perhaps most importantly, Smith cannot say that these norms are justified because they are the norms that ideal agents would themselves desire to follow. This would make sense if what it was to be fully “ideal” is nothing other than being *fully responsive to normative reasons*. However, in giving an account of the relationship between the attitudes of ideal agents and normative reasons, one cannot give explanatory priority *both* to the property of being ideal and the property of being a normative reason. Thus, despite the fact that it might seem intuitive to say that to be so-called “ideal” is nothing other than being *fully responsive to normative reasons*, Smith cannot say this if he is going to stick with his core idea of using idealization to give a metaethical explanation of what normative reasons are. The upshot is that, in order to be successful on its own terms, a metaethical ideal attitude theory such as Smith’s needs to give an *independent* characterization of the standards for idealization that does not itself appeal to what ideal agents would desire. Given that, one cannot use the idea of idealization to explain why ideal agents should follow one set of norms rather than another. For, in so doing, an ideal attitude theory will either give explanatory priority to the idea of being a normative reason in a way that undermines the core commitment of the view, or else it will be undermined by a vicious type of explanatory circularity.

What, then, can he say? He can, of course, simply insist that the property of A having a normative reason to ϕ in circumstances C *just is* equivalent to the property of agents fully following principles R1-RN having the desire for non-ideal A to ϕ in C. Or, to say at least something, as Smith does in one his presentations of the account in *The Moral Problem*, he might say that it is *a priori* that they turn out to be equivalent given the correct analysis of the *concepts* NORMATIVE REASON or IDEAL or FULLY RATIONAL.³¹ What should we make of this basic idea?

³¹ Following one standard convention, I will henceforth use smallcaps to designate concepts. On this convention, see (Margolis and Laurence 1999).

For *any* naturalistic realist account of normative reasons, there will eventually just be some basic identity between some naturalistic properties and the property of being a normative reason. We bottom out somewhere on such an account.³²

The issue for Smith's account, then, is not that it stops *somewhere* with the identification of normative reasons with such-and-such naturalistic facts. Rather, the issue is that, given that Smith's account stops *where it does*, the account is unsatisfying as a fundamental metaethical theory.

To underscore why this is so, take the concept that Smith expresses when he uses the term "rational." In order to help us remember that this is not necessarily the *same* concept that all speakers express by the term "rational", but rather that it is a concept that plays a particular functional role in his ideal attitude theory, let's call that concept RATIONAL_{SMITH}. Suppose that Smith is right that this concept just does pick out some combination of R1-R6 or more as an *a priori* matter.³³ Call this set of norms *the rational_{Smith} norms*. Given the central role these norms play in generating our normative reasons, it seems as though there should be some story about why these norms *in particular* get to play this functional role. What, then, makes those norms in particular so special? It is unsatisfactory to say that these are the special norms because this *particular* concept RATIONAL_{SMITH} picks out some set of norms as an *a priori* matter. This is because there are *any number* of concepts that *could* play the same functional role in an ideal attitude theory here in terms of setting the ideal

³² To emphasize this point, consider what an expressivist like Gibbard says: "This property [the property of being normative] isn't picked out by the meaning of the term alone, or by its meaning plus the facts..." Indeed, he says, when it comes to normative questions, "two agents might even disagree and be factually omniscient." (Gibbard 1992, 10). In contrast, a descriptivist like Smith *cannot* say this. Rather, as a descriptivist, he needs to say the property of being a normative reason, like prosaic descriptive properties, *just is* picked out by the meaning of the term alone plus the facts. Moreover, as a descriptivist committed to ontological *naturalism*, he needs to say the property of being a normative reason is picked out by the meaning of the term alone plus the *naturalistic* facts. How broadly that is interpreted is an open question – for instance, whether the so-called "naturalistic facts" here include indexical information or not. However, the point is that, for the naturalistic realist, barring issues about indeterminacy, there can be no further reasoned disagreement about what the normative facts are once everything has been set about *what the naturalistic facts are* and what our normative terms mean.

³³ There are parts of *The Moral Problem* that seem to suggest that Smith also thinks this is *analytic* given the meaning of the term "rational." However, in more recent writing as in (Smith 2004a), he does not claim that this is so and instead only claims that it is *a priori*. This, of course, suggests that Smith thinks that some of these norms might be picked out as a *synthetic a priori* matter – which, especially given the philosophical dispositions of many of those drawn to metaethical naturalism, will strike many as questionable at best. There is thus an interesting question of what sort of vindication of the *synthetic a priori* Smith has to rely on. For now, however, we need to leave this issue to the side.

conditions in the account. So why privilege this concept? For instance, to draw on David Enoch, we can ask: why not use the concept of being *SCHMATONAL* instead?³⁴ For argument's sake, suppose that *SCHMATONAL* has the same conceptual content as *RATIONAL_{SMITH}* except for some very minor differences such that, as an *a priori* matter, *SCHMATONAL* picks out one norm that *RATIONAL_{SMITH}* does not. To underscore how close these concepts might be, we can take one of the norms that Smith considers in his work but which was not on our initial list of R1-R6. For instance, suppose this is the following:

R4good-for: RR (If someone believes that a certain episode of happiness could both feel the way that happiness does and be his own, then he desires that he enjoys that episode of happiness)³⁵

One option here is to grant that you *could* be using the concept *SCHMATONAL* to play this functional role but insist that, as a matter of fact, you *are* using the concept *RATIONAL_{SMITH}*. However, by itself, this does not seem like a satisfying answer for a fundamental metaethical account. For it still leaves it mysterious why this concept is doing the work here – not in the causal sense of how it came to play that role, but rather if there is anything that would justify its doing so. In short, if the proposal were to stop here, it would be an *entirely* arbitrary fact whether we are using one of these concepts or the other. That is a problem. For even if we need to accept that there is some degree of arbitrariness here, it contradicts the intuition that many of us have – and that Smith himself explicitly draws on in developing his theory – that practical normativity itself does not fundamentally rest on *completely* arbitrary facts. If that intuition were to be totally abandoned, most of us would start to lose a grip on why we are concerned with normative reasons in the first place *either* as engaged agents thinking about what to do or as metaethicists.

In response to this worry about arbitrariness, Smith has suggested using David Lewis' idea of a *reference magnet* to gain traction on why it is that the concept *RATIONAL_{SMITH}* matters whereas *SCHMATONAL* does not.³⁶ According to Lewis, a reference magnet is something to which our words are pulled toward referring *not* because of our own dispositions to apply the word one way or another, but rather because of *how things are*.³⁷

³⁴ (Enoch 2006).

³⁵ (Smith 2009).

³⁶ (Personal communication). Smith also hints at this approach in (Smith 2009b).

³⁷ (Lewis 1983) and (Lewis 1984).

Consider the property of *being green* as opposed to the property of *being grue*.³⁸ Many have thought that our dispositions to apply the term “green” one way or another won’t by themselves settle whether we attribute the property of being green or being grue when we use the term. Indeterminacy about reference thus looms as a threat. Reference magnetism is a proposed way of dealing with it. According to proponents of the reference magnet view, it is a feature of reference that when our pattern of usage is sufficiently close to picking out a so-called “natural” property that carves the world “at its joints,” so to speak, then this property acts as a magnet toward which the term is drawn. This is why, for instance, our term “green” picks out the property of being green and not the property of being grue.³⁹

It is unclear exactly what it means to say that one property is more “metaphysically natural” than another. But the intuitive idea is relatively straightforward. According to Lewis, metaphysically natural properties are ones that make for *genuine* metaphysical similarity between things rather than just *some* similarity. For instance, the idea of metaphysical naturalness is meant to capture that if two things share the property of *being a quark* then this has different metaphysical significance than if two things share the property of *being a quark-or-elephant*.

Drawing on Lewis’ idea that such metaphysically natural properties serve as reference magnets, Smith might grant that there are a lot of different concepts in the ballpark of the concept $RATIONAL_{SMITH}$ that could in theory play a similar functional role in an ideal attitude theory. However, he might then hold that, given the nature of reference, all of these nearby concepts are drawn to a privileged kind that serves as a reference magnet – and that, importantly, is what Smith intends to use the concept $RATIONAL_{SMITH}$ to pick out.

Suppose, for now, that we grant that a reference-magnet picture of reference akin to Lewis’ is on the right track. How much help will this be to Smith in dealing with the problem at hand? The answer, I think, is very little.

To see why this is so, consider that, in order for reference magnetism to help Smith *at all* in explaining why we pick out the $rational_{Smith}$ norms, one would need to establish a) that the $rational_{Smith}$ norms are *more* natural than the schmatational norms and b) that they are

³⁸ For our purposes here, we can follow one standard way of talking in contemporary metaphysics and stipulate that x is grue at time t iff: x is green at t and $t < t_0$ or x is blue at t and $t \geq t_0$. T_0 is some specified time in the future. This is the definition of “grue” given by (Sider 2009). It should be noted that this is *not* the original definition given by Goodman in the third chapter of (Goodman 1955).

³⁹ For more recent use of the idea of reference magnets, see (Sider 2001) and (Weatherson 2003).

sufficiently more natural than the schmational norms to warrant thinking that there is no nearby pattern of usage that would lead us to refer to the schmational norms instead. Both of these theses might be true. But it also might turn out to be that there are multiple *equally* natural groups of norms – or at least groups of norms that are not *very* different in terms of how natural they are. Indeed, given how I have identified these groups so far, it seems ungrounded to hold that the rational_{Smith} norms are vastly more natural than the schmational norms. Absent independent argument, then, there is no reason to think that one group of norms will be a sufficiently strong reference magnet to help out Smith in this case.

Moreover, in order for the idea of reference magnets to be convincing, Smith needs to tell us more about what this property of naturalness *is* anyway – or, at the very least, he needs to tell us what *type* of unity makes for metaphysical naturalness in this case. To see what I have in mind here, consider the property of being green. It seems that it is more natural than the property of being grue because of the way things are in the natural world. But what sort of feature is being rational_{Smith}? Smith here faces a dilemma.

On the one hand, that it is a property that is privileged given that it picks out a psychologically unified kind – and hence carves the *descriptive psychological* domain at its joints. This seems reasonable given that it would mirror how green is established to be more natural than grue. The problem is that, if Smith goes this route, it would then not be clear why this itself has any *normative* import. After all, to continue using the metaphors adopted by the proponents of reference magnetism, we want to use our normative terms to get at the normative rather than the psychological joints.

On the other hand, suppose that Smith holds that the property of being rational_{Smith} is metaphysically natural because it picks out a normatively unified kind – and hence carves the *normative* domain at its joints. This might get us around the problem from the previous horn of the dilemma. But it would then become less clear what advantage naturalistic realism has over non-naturalistic realism. The non-naturalistic realist, remember, insists that our normative thought and talk picks out a non-reducible property of simply *being normative*. Naturalists object that this sort of property a) looks metaphysically suspect and b) is one that we could not have reasonable epistemic access to. Now we are supposing that the naturalistic realist thinks that normative thought and talk picks out *the important normative property that carves the domain of normativity at its joints*. It is unclear why *that* property should be any less metaphysically suspect or epistemologically accessible than the one the non-

naturalist posits. Thus, putting aside whatever problems it might have for the general theory of reference, I do not think that a naturalistic realist such as Smith should use the idea of reference magnetism to explain why it is that the norms R1-RN are privileged in the correct account of normative reasons.

What else can Smith say? He might try to adopt a stronger version of our initial suggestion and say this: not only are you *in fact* using this concept RATIONAL_{SMITH}, but using that concept is *constitutive of* your taking up the normative standpoint at all. In order to have *prima facie* plausibility as a position, this so-called “normative standpoint” must be characterized not just in a trivial sense such as the standpoint of responding to normative reasons as Smith understands them. Rather, the normative standpoint must be characterized in some independent way. For now, paraphrasing Gibbard, we might say that this is the standpoint of being a *planning agent* capable of asking questions of *what to do* at all. If Smith were to hold that one *needed* to use the concept RATIONAL_{SMITH} in order to coherently ask *what to do* questions in the first place, then there would be no important question left of why one should use this concept rather than the concept of SCHMATIONAL in asking about what to do. Call this the *conditions of agency response*.

This response *would* deal with the concern about arbitrariness. However, if Smith were to adopt the conditions of agency response, then he would no longer have any room to vindicate ethical agent-attitude-independence as we have defined it. For, it would then turn out that the norms for being ideal are nothing more than the norms that one needs to follow in order to count as someone who has an evaluative attitude of a certain sort in the first place – e.g. the attitude of *planning*. In other words, on this proposal, to be ideal would be nothing other than to have revised one’s evaluative attitudes according to the norms of attitude-revision that one already accepts insofar as one has evaluative attitudes at all. If a naturalistic realist were to accept this thesis, it would amount to accepting *metaethical* agent-attitude-dependence, and, thus, as we have, it would commit the naturalistic realist *ethical* agent-attitude-dependence as well.

On the one hand, if the norms necessary for planning agency are robust enough, then the *contingent* evaluative attitudes of an agent A might themselves not matter for what the ideal version of A looks like. On the other hand, on the plausible assumption that those norms won’t themselves be so robust as to transform *any* agent’s evaluative attitudes into the very same set, and holding fixed that we do think there are facts about normative reasons in

many cases, then the contingent facts about what evaluative attitudes an agent starts with *will* themselves make a difference for what ideal A looks like. I will return to this issue later on in this paper. At this point, however, the important point is that, either way, if Smith were to accept the conditions of agency response, then his metaethical view would rule out ethical agent-attitude-independence as a first-order normative view.⁴⁰ To see why this is so, first consider that, as we defined it earlier, ethical agent-attitude-independence is the thesis that there are at least some normative facts about what an agent A should do that hold independently of *any* facts about A's evaluative attitudes. Hence, as we have defined it, ethical agent-attitude-independence includes independence from *necessary* as well as *contingent* facts about A's evaluative attitudes. Both versions of the conditions of agency response we are considering here involve the thought that *at least* necessary facts about A's own evaluative attitudes matter for settling what ideal A looks like. In turn, since the attitudes of ideal A are what generate non-ideal's normative reasons, those normative reasons themselves will thus be grounded in facts about A's own evaluative attitudes.

So, to sum up, if what I have been saying about the status of the rational_{Smith} norms R1-RN is correct, we are left with either a) a view that is not even a fundamental metaethical one at all; b) a collapse into non-naturalism; c) an unsatisfactory solution to the location problem; or d) something that relies on the conditions of agency response and hence rules out ethical agent-attitude-independence. Hence, I think we should conclude that a version of ideal attitude theory such as Smith's cannot both work as a naturalistic realist view in fundamental metaethics at the same time that it makes coherent ethical agent-attitude-independence as a first-order normative view.

⁴⁰ It should be noted here that some philosophers think that the norms constitutive of being a planning agent will turn out to be quite substantial: for instance, some Kantians such as Christine Korsgaard hold that they will include the categorical imperative. See (Korsgaard 1996) and (Korsgaard 2009). If something such as that turned out to be the case, Smith might then be able to show that normative facts are independent of any *contingent* features of agents' psychology – and hence get a type of Kantian objectivity that, despite not getting agent-attitude-independence, might nonetheless capture many of the key normative intuitions that drive people to that thesis. Others of a more Humean spirit – such as Jamie Dreier, Sharon Street, Michael Bratman, and Allan Gibbard – think that there are very few norms that one needs to follow in order to ask planning questions of *what to do*. See (Street 2008), (Gibbard 1999), (Dreier 2007), and (Bratman 1999). One worry, then, is that if the more Humean-minded philosophers are correct (which I suspect that they are), very few of the norms that Smith hypothesizes might be in R1-RN would actually remain if we went this route.

4.2. Normative Epistemology and Reflective Equilibrium.

In the last section, I argued that in order to produce a viable form of naturalistic realism in metaethics, an ideal attitude theorist such as Smith needs to take the *conditions of agency response* to the question of why certain norms are given a privileged status in an ideal attitude account – and, therefore, I claimed that such a theorist will need to give up ethical agent-attitude-independence. In order to further support my argument that naturalistic realists should abandon ethical agent-attitude-independence, and to provide some basic groundwork for how I think a successful naturalistic theory should be developed, I now want to raise two other important problems that Smith’s theory faces by *not* accepting the *conditions of agency response*. The first problem concerns normative epistemology. The second problem concerns the motivational role of normative judgment.

Let’s start with the issue about normative epistemology. Imagine that Smith is right that an agent A has normative reason to ϕ in circumstances C iff all ideal agents would desire that A ϕ in C. How are we supposed to have epistemic access to such facts about what ideal agents desire? Smith’s answer is that we use the standard method already implicit in most of normative ethics: the method of reflective equilibrium. Following Rawls’ original characterization of reflective equilibrium, Smith claims that this method can be roughly summarized as follows. We start with specific normative judgments about which we are most confident and then try to provide them with a “compelling and comprehensive underlying justification.”⁴¹ In turn, we then “amend our initial attempt to provide a comprehensive justification in the light of our specific judgments when we are more confident about our specific judgments; we amend our specific judgments when we are more confident about the comprehensive justification; and so on and so forth until, eventually, we achieve a reflective equilibrium: a point at which we deem no further adjustments to be required.”⁴²

According to Smith, normative facts are set by what agents would desire in a very specific scenario – namely, one in which those agents met Smith’s qualifications for being

⁴¹ (Smith 2005, 4).

⁴² (Smith 2005, 4). For Rawls’ characterization of the method of reflective equilibrium, see (Rawls 1999a). It should be noted that Smith’s gloss here of reflective equilibrium is in terms of so-called “narrow” rather than “wide” reflective equilibrium. For our purposes here, this decision is not crucial.

ideal by following norms R1-RN. Why should we then think that this process of reflective equilibrium is a helpful way to access the normative facts?

For instance, imagine that Jessica starts off with a set of evaluative attitudes E1. Suppose further that, by her own lights, she is more committed to or confident in a certain part of that set of attitudes – call this subset E2. She then comes up with a justification that she herself finds compelling for E2. She then amends that subset in light of the justification and ends up thinking that she should have the set of evaluative attitudes E3, amends that again and so on and so forth until she ends up thinking that she should have the set of evaluative attitudes E4. In arriving at that point, Jessica is using the norms of justification that *she herself* finds compelling. Since Smith does not think that all agents must aim to follow norms R1-RN by virtue of having evaluative attitudes at all, there is no guarantee that the norms that Jessica is using are the same ones as R1-RN.

Given what Smith thinks the normative facts are, it is thus not clear why he thinks the process of reflective equilibrium is a good one for an agent such as Jessica to use in order to figure out what to do. Not only might Jessica be using different norms from those that Smith picks out as R1-RN to guide her, but also the norms might have virtually no overlap at all (i.e. no overlap beyond those norms that *are* constitutive for being a planning agent such as, at least plausibly, R1). Thus, it is not only unclear why she should have confidence in this methodology, but, more drastically, why she should not believe that she is *totally hopeless* at gaining epistemic access to facts about normative reasons by using this method.

This worry can be fleshed out as follows. Suppose, for the sake of argument, that a metaethical theorist holds that facts about what normative reasons Jessica has are set by the evaluative attitudes of a mysterious monster that lives on a distant planet. If Jessica follows the method of reflective equilibrium as we have sketched it here, there is little reason to think that this method will help her to understand the evaluative attitudes of this mysterious monster. Hence, she will have little reason to suppose that she is not hopeless at knowing the normative facts. We do not tend to think that the attitudes of Smith's ideal agents are as opaque as the attitudes of such a mysterious monster. But why not? If Jessica does not have reason to suppose that her starting evaluative attitudes favor revising according to norms R1-RN, why should she have confidence that the method of reflective equilibrium is a useful way of gaining epistemic access to the normative facts?

What should Smith say in response to this epistemological worry? One move he might take to undercut its force is to claim that, even if we do not all follow norms R1-RN in revising our evaluative attitudes when engaged in reflective equilibrium, we do all have a *standing desire* to be rational that is stronger than our other desires – where, in turn, “rational” means having evaluative attitudes that are revised according to norms R1-RN.

There are two problems with this proposal. The first problem is that, if we in fact had such a desire, which we assume will always trump our other desires in the process of reflective equilibrium, then it is not clear why we should bother engaging in reflective equilibrium at all, rather than just asking what would satisfy that desire. The second problem is that there seems to be little reason to suppose that we would all have such a desire in the first place. To see why this is so, consider our beliefs about objects like rocks. Although there are, of course, other hypotheses that one might suggest, it is at least reasonable to hold that the reason people started to believe in rocks is that *there were rocks in the world* and believing so helped promote our ancestors’ reproductive success. As Street argues, however, the best explanation of why our ancestors started to have the evaluative attitudes that they did is not that they were tracking something independent – such as the normative joint picked out by RATIONAL_{SMITH} – but simply that having certain desires tended to lead to reproductive success and others did not. This thesis would support thinking that there are many desires we are all likely to share now – for instance, the desire to promote the well-being of our offspring, or the desire to avoid being run over by cars, but it is hard to see why it would support thinking that we all have the desire to follow the rational_{Smith} norms. This line of defense is thus a dead-end for Smith.⁴³

What then can Smith say instead? Another line that Smith could adopt is to hold that it is simply an unexplained mystery why the epistemology here tracks the metaphysics. Moreover, he can simply believe that we just happen to be lucky that they *do* line up and not worry about the skeptical worry that reflective equilibrium *might* be taking us far astray from the normative facts.⁴⁴ This response, however, will not do. Consider the familiar compelling

⁴³ It should be noted that there are ways of modifying the suggested response here that might be more probable – for instance, that we evolved a *general capacity* for reasoning which allowed us to pick up on the normative joints even if that was not the reason for the evolution of that capacity in the first place. I think that this strategy fails for the general reasons that Street sketches in (Street 2006). However, it is beyond the scope of this paper to address this here.

⁴⁴ For responses to this type of worry that are broadly similar, see (Dworkin 1996) and (Rosen 1998).

stories about why methods in the natural sciences are helpful for gaining access to facts about the natural world. Such stories are not ones that claim that we simply happen to be lucky *from the ground all the way up*. Even if they must rely on the brute anti-skeptical assumption that we are not brains in vats, these stories do not rely solely on an appeal to luck once this basic anti-skeptical assumption is in place. Shouldn't we want a similar story to be told in the normative domain if we are to continue to use the methodology of reflective equilibrium?

If so, we have reason to abandon Smith's version of ideal attitude theory. In particular, given that the cause of the tension between the proposed methodology and metaphysics is the thesis that the property of being ideal is fixed by facts that are *independent of the agent's own evaluative attitudes*, we should drop this thesis. Instead, we should think of A's normative reasons as derived from the antecedent evaluative attitudes of A herself. More specifically, I submit, we should think of A's normative reasons as derived from what evaluative attitudes she would have at the end of reflective equilibrium.⁴⁵

This does not mean that we would need to abandon ideal attitude theory *as a whole*. To see why this is so, imagine again the case where I am trying to learn to play tennis and wondering about what racket to use. The idea there was that I should ask the *advice* of the tennis instructor about which racket to use rather than seeking to *emulate* her. This was because, as a tennis pro, what she would desire for herself to use might be inappropriate for someone at my level. Similarly, we might hold that what an agent might think that *she herself* should do at the end of reflective equilibrium might be inappropriate for the version of herself to do that had not already undergone reflective equilibrium. Does that mean that there will *always* be a gap between using the method of reflective equilibrium and the actual normative metaphysics if one uses an ideal attitude theory that works on this advice model? The answer is "no." For even if we do *not* hold that the normative facts about what an agent A should do in circumstance C are identical to what A *thinks* about what she herself should do at the end of reflective equilibrium, we can still think of the normative facts as derived from the evaluative attitudes that A would have at the end of reflective equilibrium. In particular, we can hold that what makes an *ideal* version of an agent A (call her "A+") is A's

⁴⁵ I here use the phrase "an agent's normative reasons" as a helpful shorthand to refer to the normative reasons that apply to her. In everyday contexts, we often speak of "an agent's reasons" to refer to her motivating reasons for performing a given action. However, this is *not* my usage in this paper.

arrival at the end point of reflective equilibrium as characterized above. We can then take the normative facts to be, as on the basic advice model of ideal attitude theory that we get from Smith, what A+ would desire non-ideal A do in non-ideal A's actual context.⁴⁶ We might not have *easy* access to all such facts. But they certainly would not be as inaccessible as they might turn out to be on Smith's proposal – and, indeed, their level of accessibility would seem to mirror the fact that much of normative inquiry requires difficult, on-going work.

Since the way that A reaches the end of reflective equilibrium is by revising her evaluative attitudes a) following the norms that are constitutive of her being a planning agent at all and b) weighting evaluative attitudes according to how committed she herself already to those different evaluative attitudes, this proposal suggests that to be an ideal agent is nothing other than to be an agent who has revised her evaluative attitudes according to what is, in broad terms, *already entailed* by her current evaluative attitudes.⁴⁷ In contrast to a way of thinking about what it is to be an ideal agent such as Smith's where the standards for being ideal are set by standards external to the agent's own evaluative attitudes, call this an *internally idealized* version of A.

If an ideal attitude theorist such as Smith were to use *this* version of A to explain what it is to be ideal, this would mean that he would need to give up the thesis of ethical agent-attitude-independence. But, if what I have argued in this section is correct, then not doing so threatens to leave a naturalistic realist such as Smith little better off than his non-naturalistic realist foe: stuck with an account of normative metaphysics that leaves it entirely mysterious why standard normative methodology (or any other reasonable methodology we know of) tracks the normative facts.

4.3. *The Motivational Role of Normative Judgment.*

Let's now turn to the final problem that Smith's metaethical theory confronts by not thinking of what it is to be ideal in terms of facts about an agent A's evaluative attitudes. This problem concerns the apparent motivational role of an agent's normative judgments.

⁴⁶ I take this convention of "A+" from (Railton 1986b).

⁴⁷ I borrow talk of "entailment" here from (Street 2008). I use this phrase to get at, to put it in rough terms, the consequences of what one desires in the same way that we might use the phrase to get at the consequences of what one believes. There are serious questions here of how to cash out this idea: for instance, one issue is in what way the idea of entailment here is related to the idea of strict logical entailment. For the purposes of this paper, I won't take any stand on this and related issues.

To see what's at issue here, suppose that I make a prosaic descriptive judgment, such as *there is beer in my fridge*. This can bear on my motivation to act if, for instance, I am near my fridge and have the desire to drink beer. However, it might also not motivate me in any way. For instance, I might report this to Mike as an offhand remark about what is in my fridge at home in Ann Arbor when we are both currently looking at his fridge *in Manhattan*. Contrast this with a sincere normative judgment that I might make when looking at my fridge in Ann Arbor: e.g. "I really should drink the beer in my fridge." In this case, there might be lots of other psychological forces pulling against my *actually* drinking the beer – indeed, perhaps forces that pull against my actually exhibiting any actual *motivation* to get the beer at all. For instance, I might be depressed or weak-willed. But it nonetheless seems that this normative judgment puts *some* direct causal pressure on guiding me to actually drink the beer in a way that my prosaic descriptive judgments do not. Depending on one's theory of what "motivation" means, this pressure might be too weak to actually result in motivation in some cases, but we might be able to call it *motivational pull*. Although we can easily imagine prosaic descriptive judgments that lack such motivational pull, it seems that, at least initially, it is part of what it is to be a normative judgment that it exerts such motivational pressure pull on the agent who makes the judgment.

This suggests a version of the thesis that is often called "motivational judgment internalism," and that we can call "judgment internalism" for short.

Judgment Internalism. If an agent A judges that she (i.e. A herself) has a normative reason to ϕ , then, necessarily, she has some motivational pull to ϕ .⁴⁸

Judgment internalism is a controversial thesis.⁴⁹ Yet, at least at first, there *does* seem to be something quite attractive about it. For, understood in one way, judgment internalism reflects one of the most basic ideas about *normative reasons simpliciter* that many of us have: namely, that when an agent inquires into which ones she has – either by herself or with others in genuine normative discussion – she is trying to figure out the practical matter of *what to do*. Thus, facts about normative reasons are not facts that I might have no interest in whatsoever – for instance, facts about what a cult leader might order me to do or what the

⁴⁸ This is a modified version of how Darwall defines "judgment internalism" in (Darwall 1997), 308.

⁴⁹ For recent rejections of this thesis, see (Parfit 2006) and (Dancy 2003). Both Parfit and Dancy are non-naturalists. However, we could also imagine a naturalistic realist saying this as well – e.g. if they were to advance a position about normative reasons that is parallel to the so-called "Cornell Realist" view about moral properties advanced by (Brink 1989) and (Sturgeon 2006).

norms of etiquette prescribe that I do. Rather, *if I judge that I have such reasons* this *necessarily* matters to me from my own perspective on the world as an agent trying to figure out what to do.

There are many philosophers who think that judgment internalism holds in a strong form – such that the resulting “motivational pull” to ϕ always yields an actual desire to ϕ or perhaps an intention to ϕ . For instance, Gibbard endorses a version of this strong sort of judgment internalism in his recent *Thinking How to Live* where he claims that normative judgments amount to forming an intention about how one would act in actual or hypothetical circumstances.⁵⁰ In light of cases such as those involving depression and weakness of will, however, many other philosophers including Smith have wanted something weaker than a strong form of judgment internalism such as Gibbard’s. Rather, to put it in Smith’s own terms, they have wanted something such as the following: “that those who judge it right (or good or desirable or sensible) to act in a certain way desire to act in that way, at least absent practical irrationality, and that those who judge it wrong (or undesirable or stupid) to act in a certain way are averse to acting in that way, at least absent practical irrationality.”⁵¹ Given how Smith understands what it is to be rational in the rest of his theory, can he then really use his form of rational judgment internalism to make good on the initial data that we want it to capture? More specifically, can it explain the basic idea that we started with? Remember, this was the idea that, absent certain types of defeating conditions such as being weak-willed or depressed, genuine judgments about normative reasons in our actual world have a different motivational import for creatures like us than prosaic judgment about things such as beer do. This is an empirical claim about the motivational import of normative judgment in our world. Hence, let’s call this basic idea *the empirical motivational thesis*.

Suppose that when Smith says that an agent is “practically rational,” he means either an agent who is currently following the rational_{Smith} norms R1-RN or who has a desire to form her desires in accordance with the desires of an agent who does follow those norms. If such an agent were to make a normative judgment, this *would* have motivational import for her. This, though, will not get us very far in terms of making good on the underlying theoretical motivation we started with for rational judgment internalism. In particular, it

⁵⁰ (Gibbard 2003).

⁵¹ (Smith 2001).

won't get us very far in terms of explaining the empirical motivational thesis. For, as we have already seen, there is no guarantee that most agents will be practically rational in either of the above senses. Thus, rather than making room for *some* exceptions to the internal connection between an agent's normative judgment and motivational pull, it appears as if Smith's theory will make room for *many* exceptions. The connection won't just fail to hold for those who are weak-willed or depressed. In fact, it won't hold for anyone who doesn't follow norms R1-RN or who does not want to form her desires in accordance with the desires of an agent who does follow those norms. Absent further argument, that might in fact turn out to be *many of us*. Hence, it looks as though Smith's form of judgment internalism won't do much to explain the empirical motivational thesis; which, after all, was the main reason we wanted a form of judgment internalism in the first place.

As with the other problems I have identified for Smith, this problem does not stem from the peculiarities of his version of ideal attitude theory but rather from general features of his *type* of ideal attitude theory. To see this, even if you ignore entirely the convergence requirement, consider *any* view where the conditions for being ideal are set by standards external to an agent's own evaluative attitudes. In such a case, the resulting view will threaten to undercut one's ability to capture the empirical motivational thesis absent independent empirical argument for *why* we think agents have evaluative attitudes that favor revising their attitudes according to those standards.

Contrast this with how things look for our proposal from the previous section – namely, the proposal that we should retain the basic structure of an ideal attitude theory, but think of A+ as an *internally idealized* version of A. Such a view would have a fairly straightforward way of capturing the empirical motivational thesis. For, if an agent's normative judgments about what she should do are judgments about what an internally idealized version of herself would desire for her non-ideal self, then normative judgments are ones about what an agent desires for herself. Absent further argument, it is not clear that judgments about such desires will *necessarily* have motivational significance for an agent. Perhaps there is some such necessary connection. Perhaps there is not. But, insofar as one's judgments about what *oneself actually desires* do have significant motivational upshot for us much of the time, this proposal would be able to explain at least the *empirical* connection between normative judgment and motivation, where Smith's style of proposal seems to fail. Moreover, it appears to be able to do so *without* necessarily taking on any form of judgment

internalism at all – and, hence, remaining neutral on the controversial question of whether or not one’s judgments about one’s own normative reasons *necessarily* have motivational pull.

To sum up, then, we can therefore say that the modified type of judgment internalism that Smith proposes – and which many naturalistic realists find more plausible than the strong, unmodified form – will not actually be able to capture the seeming connection between normative judgment and motivation that we wanted. We thus need to 1) develop some further independent argument about how to capture a strong version of judgment internalism and/or 2) accept the idea that A+ is an *internally idealized* version of A. The first option might be possible – but it requires further argument of a sort that I have not given in this paper. Hence, absent further argument, this issue about the motivational import of normative judgment gives us one more reason to hold that A+ is an internally idealized version of A; and, hence, give up on the idea of ethical agent-attitude-independence.

5. *The Upspot for Metaethical Naturalistic Realism More Broadly.*

As anyone even vaguely familiar with the field of metaethics will be aware of, there is a range of token metaethical views for each broad type (e.g. non-naturalistic realism, expressivism, etc). It is obviously beyond the scope of this paper to consider all the leading token versions of naturalistic realism in metaethics – let alone to consider versions that might be developed. It is thus beyond the scope of this paper to go through in detail how my arguments should impact views that are quite different in kind than Smith’s ideal attitude theory. Still, I think that my arguments so far do teach us some general lessons about naturalistic realism as a type of metaethical view.

Consider the range of arguments I have given against Smith. How many of these arguments rest on peculiarities of Smith’s position or even on the fact that his view is an *ideal attitude theory* as such? Very few. To start with, consider my first argument (in section 4.1), that Smith’s view needs to accept something like the *conditions of agency response* in order to remain a satisfactory fundamental metaethical view that blocks a collapse into non-naturalism. A version of such an argument could in fact be run against *any* naturalistic realist view that identified normative properties in terms of some norms that were *optional* for a planning agent capable of thinking about what to do. Similarly, consider the second argument (in section 4.2), about the gap between Smith’s normative epistemology and his account of normative metaphysics. A version of such an argument could run against *any*

view that identified normative properties as naturalistic properties that we did not in fact care about given our own starting evaluative attitudes. Finally, consider the third argument (in section 4.3), about the motivational role of normative judgment. A version of such an argument could be run against *any* view that tried to account for this role using a form of rational judgment internalism where the conditions for being “practically rational” are not fundamentally tethered to the agent’s own evaluative attitudes. These three arguments will not have a direct impact on *all* forms of naturalistic realism. However, taken together, they will have an impact on many of the most prominent versions of naturalistic realism— and, indeed, I think on all potentially *viable* versions of it that seek to do justice to the metaethical desiderata that I listed at the end of the second section of this paper. Given this, I conclude that, if we are going to be naturalistic realists in fundamental metaethics, then we need to give up the thesis of ethical agent-attitude-independence.

6. Fully Agent-attitude-dependent Naturalistic Realism: A Preliminary Sketch and Defense.

For the sake of argument, let’s grant that this conclusion is correct. By itself, this does not tell us whether we should give up naturalistic realism or ethical agent-attitude-independence. In order to address this question, it will be crucial to ask the following: is there a way to develop a plausible form of naturalistic realism that is compatible with this conclusion? It is obviously beyond the scope of this paper to fully answer this question. However, in this last section, I want to close by sketching how I think such a theory might be best developed. In so doing, my aim is to explain why I think my main argument in this paper might be a step along the way to developing a successful form of naturalistic realism as opposed to an argument against it.

Start with the idea of the *internally idealized* version of an agent A that I first introduced in section 4.2. This is, remember, a version of A that is, to put it somewhat metaphorically, totally regulated by what her own evaluative attitudes entail. For it is not only an agent that has revised her evaluative attitudes according to the norms that are constitutive of having the target type of evaluative attitudes at all – e.g. plans – but according to the *content* of her own evaluative attitudes where those have, in turn, been weighted according to what she herself takes to be most central. If we plug this version of A into our basic advisor-type ideal attitude schema, we get the following:

Fully Agent-Attitude-Dependent Ideal Attitude Theory: A has normative reason to ϕ in C iff the fully internally idealized version of A would desire that non-ideal A ϕ in C.

In rough terms, on this sort of view, practical normativity turns out to be settled by the agent's own evaluative attitudes *all the way down*. For, on this *fully* agent-attitude-dependent metaethical view – one that endorses what we call *complete agent-attitude-dependence* at the metaethical level – it is not only that facts about an agent's evaluative attitude always matter for her normative reasons. Rather, on this view, facts about an agent's evaluative attitude are *all* that matters for her normative reasons beyond the non-normative facts that constitute the circumstances C. This is in contrast to *partially* agent-attitude-dependent views that claim that facts about an agent's evaluative attitudes always matter but that at least one other type of fact always matters as well – e.g. the distribution of rocks on Easter Island. We might also say that such a partial view amounts to endorsing a version of *incomplete* agent-attitude-dependence. Both complete and incomplete agent-attitude-dependence are theses that, in basic terms, concern the degree to which there is a connection between an agent's normative reasons and an agent's evaluative attitudes. The basic taxonomy can in turn be used to explain positions of two different sorts: positions at the substantive normative level about the correctness-conditions for normative judgments in ethics (complete versus incomplete *ethical* agent-attitude-dependence) or metaethical positions about the nature of the normative properties that ethical judgment is about (complete versus incomplete *metaethical* agent-attitude-dependence). The claim here is that naturalistic realists should endorse *metaethical* complete agent-attitude-dependence about what the normative properties *are* – a position that then entails that ethical complete agent-attitude-dependence is true (holding fixed the assumption that some agents in the actual world have normative reasons).

If my arguments in the previous sections have been right, this fully agent-attitude-dependent metaethical theory will allow us to sidestep all of the objections to Smith's theory as a form of metaethical naturalistic realism. What, though, of the theory's own chances for success? The view faces several important challenges. I will discuss some of the most important in closing.

6.1. *Whose Attitudes Really Matter?*

At the start of this paper, I claimed that the central issue was whether naturalistic realism could be reconciled with ethical agent-attitude-independence. Now, however, I am

claiming that naturalistic realists should not only reject ethical agent-attitude-independence in favor of ethical agent-attitude-dependence, but that they should adopt a *fully* agent-attitude-dependent account in normative ethics. Why is that?

Start with why *A* herself needs to be in the basic *metaethical* account. *Any* account that understands idealization in terms of current evaluative facts – whether those of the speaker, agent, or some group of agents – at least *prima facie* appears as if it will be able to deliver a viable naturalistic account of normative ontology. The problem is that if the idealized version of the agent herself is *not* mentioned in the account, then we run into trouble with the following thesis:

The Parity of Interpersonal and Intrapersonal Normative inquiry: When a speaker engages in normative conversation with an agent, it is possible for her to talk about the *very same* normative reasons that an agent herself asks about from the first-person perspective.

If the agent herself is not included in the account – say, for instance, if only the *speaker* is included – then there is no reason to suppose that this thesis is correct.⁵² Moreover, once *anyone else* other than the agent herself is included, it is unclear *why* the attitudes of other people should matter for the agent herself *from her own perspective* when she deliberates about what to do. In all likelihood, she will hold as a first-order view that the attitudes of other people matter too. But if she does not have evaluative attitudes that favor doing so herself, then there is no reason to think that when she does know about the content of such attitudes, that that knowledge will matter to her when she thinks about *what to do*. Therefore, including such other idealized versions of other subjects will leave us in trouble explaining 1) why the method of reflective equilibrium (in which the agent starts with only *her* attitudes) gets at the normative facts and also 2) why when an agent makes a normative judgment (which we are supposing is now in part about the evaluative attitudes of other speakers) this will have any interesting internal connection to her motivation. These considerations pull in favor not only of an agent-attitude-dependent account at the metaethical level, but in fact of a *fully* agent-attitude-dependent one.⁵³ In turn, assuming that there are normative facts in the

⁵² It should be noted that, for reasons that we cannot get into here, the issue is more complicated if we import new-wave versions of relativist semantics of the sort argued for by Andy Egan and John Macfarlane. See (Egan 2006) and (MacFarlane 2007). I take up these issues in my paper “The Limits of Self-Locating Relativism about Ethics”.

⁵³ The term “parity” here might lead some readers to think of the following thesis: that when a speaker makes a judgment about an agent’s normative reasons, she is concerned with the same evaluative attitudes that she is in the first-person case (namely, her own). Gibbard advocates a

actual world, this in turn counts in favor of a *fully* agent-attitude-dependent account at the substantive normative level as well.

6.2. Which Evaluative Attitudes Really Matter?

At the start of this paper, I defined the term “evaluative attitude” as follows: the term refers to *any* attitude that tends to produce motivation for an agent when coupled with her prosaic beliefs. Following Street, I claimed that the group of evaluative attitudes likely includes such things as desires, plans, attitudes of approval and disapproval, and unreflective normative judgments. However, there are important differences between each of these types of attitudes. So which ones really matter for our account? I think that this will be one of the most crucial questions for developing my type of account. Among other things, it is not even clear that all these attitudes have the same so-called “direction of fit” – or that we should even believe in all these attitudes if we stick with the core idea of the so-called “Humean Theory of Motivation” that attitudes only have *either* a so-called “world-to-mind” direction of fit or a so-called “mind-to-world” direction of fit.

For now, the important thing I want to stress is that *there are many different options* here once one accepts the basic role that *some* evaluative attitudes need to play in the metaethical account. For instance, if we use the term “desire” in the broad functionalist sense in which it is often used in the philosophy of mind to pick out any attitude with a world-to-mind direction of fit, then the prosaic use of the term “desire” is not going to pick out all the desires in this functionalist sense of having the world-to-mind direction of fit. There are a lot of other things that we might put here in this class too. These include, for instance, the more complicated types of non-cognitive attitude that have traditionally been most interesting to expressivists: such as, for instance, Blackburn’s proposed attitude of “concern” in *Ruling Passions*⁵⁴ or Gibbard’s proposed attitude of “accepting a norm” in *Wise Choices, Apt Feelings*.⁵⁵

version of this view in (Gibbard 2003). This thesis is in fact in tension with the thesis that I am here calling “The Parity of Interpersonal and Intrapersonal Normative Inquiry.”

⁵⁴ Blackburn is explicit in (Blackburn 1998) that he sees this attitude as one with this direction of fit.

⁵⁵ It is not clear whether Gibbard’s attitude of norm acceptance is a species of desire in our sense or not. On the one hand, Darwall argues in chapter seven of (Darwall 2006) that norm acceptance has a number of important logical and phenomenological features that make it incompatible with its being a species of desire. On the other hand, it is arguably the case that most of Darwall’s points could be made within a broad enough analytic functionalist account of what it takes to be a desire as in (Braddon-Mitchell and Jackson 1996). For our purposes here, what is important is not which view Gibbard actually holds but rather a view that is at least in the ballpark of what he is talking about.

What such proposals are getting at, I think, is a common idea that there are a range of noncognitive attitudes that we have – they range from mere urges to prosaic desires to states like Blackburn’s “concern” that, to put it in somewhat vague and metaphorical terms, are more central to our sense of who we are and what basic practical projects we want to pursue in the world. In order to adequately understand normative thought and talk, we will – as either expressivists or descriptivists – need to zero in on which attitudes exactly we are talking about. I acknowledge that this is an important outstanding issue for a metaethical account such as the one I have proposed. But, as my brief canvass of different so-called “noncognitive” attitudes here is meant to suggest, it is not one that I think is insurmountable – even, importantly, for someone who accepts the strict Humean division between attitudes with different directions of fit.

6.3. The Correctness Constraint Revisited.

When I introduced the aims of metaethics at the start of this paper, I said that any viable metaethical theory must vindicate the idea that there is a potential gap between any token normative judgment about a topic and the *correct* normative judgment about that topic. I called this *the correctness constraint*. Smith’s style of ideal attitude theory can meet this constraint. But, insofar as it takes normative facts to be constituted by evaluative attitudes, won’t my proposed view fail?

To see why my view *does* in fact meet the correctness constraint, consider two cases: 1) a case where a speaker makes a judgment about what another agent should do and 2) a case where an agent makes a judgment about what she herself should do. In the first case, what makes the speaker’s judgment correct or incorrect is facts about the *agent’s* evaluative attitudes. Therefore, there is a clear gap between the *speaker’s* own evaluative attitudes and the normative facts. In the second case, the situation is somewhat more complicated. Since the standards for correctness are determined by the agent’s own evaluative attitudes, there is a sense in which, as Street puts it about the basic style of view that both she and I advocate, any mistakes that an agent makes “will, in the end, be on her own terms.”⁵⁶ However, this does *not* mean that any token normative judgment is necessarily correct. This is for two reasons. The first is that I have not yet said which *type* of evaluative attitude a normative judgment is about – thus, it is not clear whether normative judgment is *about* further

⁵⁶ (Street 2008, 2).

normative judgments or about some other subset of the agent's evaluative attitudes. If it is about the latter, then there is no worry about the normative judgment being the *same* as what it is about. The second reason is that, even if the evaluative attitudes in question were to be understood as normative judgments or as proto-normative judgments, the view does not claim that what makes it correct or not depends just on *itself*. Rather, it ultimately depends on the whole group of normative judgments revised according to the norms for revising those judgments that the agent herself accepts. Thus, even if we took the evaluative attitudes in question to be normative judgments, we still preserve the important gap we were after with the correctness constraint.⁵⁷

On a related note, it should be stressed that the view that I am suggesting does not entail that an agent's normative reasons are simply up to her decision. Or, to put it another way, the view does not entail a strong form of *voluntarism* about reasons. For although there is a broad metaphorical sense in which the fully agent-attitude-dependent metaethical view of practical normativity that I am here suggesting sees facts about normative reasons as facts that the agent herself plays an important role in determining – in the sense that the agent's own evaluative attitudes constitute what normative reasons she has – it is not up to the agent to change her evaluative attitudes *at will*. Indeed, rather than getting to *choose* what evaluative attitudes one has, one rather always finds oneself having *some* evaluative attitudes and having to work from there.⁵⁸

6.4. *What About The First-Order Costs?*

At the start of this paper, I introduced the thesis of ethical agent-attitude-independence by way of an example about someone (Louisa) who most values poking puppies in the eyes to watch them cry. I claimed that most of us initially take such cases to be near conclusive evidence for ethical agent-attitude-independence. Now, however, we have arrived at a metaethical view that conclusively rules out this first-order view. There are two major worries we should have about this.

⁵⁷ This second point is emphasized in (Street 2008).

⁵⁸ I take this to be one of the crucial themes in Heidegger's account of "thrownness" in *Being and Time*: that is, roughly, his account of how a person always finds herself inhabiting a norm-infused social-historical setting that she herself did not choose but which nonetheless has structured her basic perspective on who she is and what her proper role is in the world. See (Heidegger 1967). A similar line of thought is also underscored by Railton in his account of the role of "the unchosen" in normative reasoning. See (Railton 2006).

The first is that the ethical agent-attitude-independent thesis looks very appealing to most of us – indeed, not only in light of cases such as the one about Louisa but in light of direct intuitions that many of us have about general normative principles. One can, of course, try to dispel such worries by giving broadly Kantian arguments that, given the norms that all agents need to accept in order to be engaged in thinking about what to do at all, this makes it such that agents such as Louisa are in fact conceptually impossible. Or, to take a different approach, one can hold that – as a matter of contingent empirical fact – most people whom we will ever encounter have been shaped by forces that have led them to favor the same sorts of things we ourselves favor (e.g. not causing others gratuitous pain, caring for one’s near and dear, or perhaps even the norms R1-R6 that Smith lists).⁵⁹ But, to many, these arguments will seem too implausible and/or too contingent to vindicate the sorts of intuitions that the thesis of ethical agent-attitude-independence was drawing on. This, then, is one important worry about rejecting ethical agent-attitude-independence. Whether it can be countered depends in large part on how the above types of arguments can be best developed.

The second worry is perhaps the more pressing one in the context of this paper. Given how prevalent normative judgments that rely on ethical agent-attitude-independence seem to be, doesn’t this mean that the proposed metaethical view fails to adequately meet the *first-order diversity constraint* that we introduced at the start? For doesn’t it not only entail that ethical agent-attitude-independence is *false* as a first-order normative view but ultimately *incoherent*? Many will be inclined to think that this means the view will therefore obviously fail as a theory in fundamental metaethics.

However, it is not clear to me that this is so. I said earlier that a fundamental metaethical theory needs to explain what we intuitively think of as normative thought and talk. But I also qualified this by saying it should do so *as much as possible*. It should now be clear why I said this. Any form of naturalistic realism in fundamental metaethics is going to have to rule out some seemingly coherent first-order normative views as, in fact, either patently false or not really normative views at all. Although we should of course respect those views as initial data to be made sense of by metaethics, it is an impossible task for the naturalistic realist to vindicate the coherence of *all of them*. For the foes of naturalistic realism,

⁵⁹ Korsgaard can be seen as trying to make a move of the first sort in (Korsgaard 1996), whereas Street can be seen as trying to make a move of the second sort in (Street 2008).

this will of course simply be more evidence that this type of metaethical view is not only wrong but in fact *hopelessly* wrong. Yet, for those naturalistic realists such as myself who think that abandoning either ontological naturalism or descriptivism comes at too high of a cost, I think we simply have to accept this as a feature of the right sort of metaethical view.

Furthermore, in the end, I am not sure how worried this should make us. To start with, the core idea behind the *first-order-diversity constraint* is not that the correct metaethical account must show that *all* normative views we initially thought to be coherent really are coherent. Rather, the core idea is that there is a *range* of normative views that are coherent – a fact that helps explain why we often find ourselves engaged in extended complex debates in normative ethics. My proposed naturalistic realist view supports thinking that this core idea is correct and, furthermore, supports thinking that many problems in normative ethics will be very difficult to solve. I am not sure how much more we should want. Indeed, in general, I think we should be open to the idea that much of what we intuitively think of as normative thought and talk before doing metaethics might in fact be much more conceptually confused than we initially thought – perhaps, as some error theorists have urged, picking out things that can *a priori* be ruled out as metaphysically impossible.⁶⁰ If this is so, and the proposed metaethical view turns out to be capturing only one central part of normative thought and talk – namely, the part that is not conceptually confused and that is capable of picking out properties in the world that are naturalistic, capable of being accessed through reflective equilibrium, and capable of supporting a connection between normative judgment and motivation – then this itself will be an important metaethical accomplishment. If developed in this way, the view might then turn out not to be, strictly speaking, a complete fundamental metaethical view given how I defined the task of fundamental metaethics earlier in the paper. But my view might nonetheless be the right fundamental metaethical account of the normative thought and talk that really matters.

⁶⁰ For instance, this is how Smith reads (Mackie 1977) in (Smith 2009b).

WORKS CITED

- Ayer, A.J. 1952. *Language, Truth, and Logic*. New York: Dover.
- Blackburn, Simon. 1998. *Ruling Passions*. Oxford: Clarendon Press.
- Braddon-Mitchell, David, and Frank Jackson. 1996. *The philosophy of mind and cognition*. Oxford, UK ; Cambridge, MA: Blackwell Publishers.
- Brandom, Robert. 1994. *Making it explicit : reasoning, representing, and discursive commitment*. Cambridge, Mass.: Harvard University Press.
- Bratman, Michael. 1999. *Faces of intention : selected essays on intention and agency, Cambridge studies in philosophy*. Cambridge, U.K. ; New York: Cambridge University Press.
- Brink, David. 1989. *Moral Realism and the Foundations of Ethics*. Cambridge: Cambridge University Press.
- Dancy, Jonathan. 2003. *Practical Reality*. Oxford: Oxford University Press.
- . 2006. Nonnaturalism. In *The Oxford Handbook of Ethical Theory*, edited by D. Copp. Oxford: Oxford University Press.
- Darwall, Stephen L. 1983. *Impartial reason*. Ithaca N.Y.: Cornell University Press.
- . 1997. Reasons, Motives, and the Demands of Morality: An Introduction. In *Moral discourse and practice : some philosophical approaches*, edited by S. L. Darwall, A. Gibbard and P. A. Railton. New York: Oxford University Press.
- . 2003. Desires, Reasons, and Causes. *Philosophy and Phenomenological Research* (67):436-443.
- . 2006. *The second-person standpoint : morality, respect, and accountability*. Cambridge, Mass.: Harvard University Press.
- Dreier, James. 2004. Meta-Ethics and the Problem of Creeping Minimalism. *Nous-Supplement: Philosophical Perspectives* 18:23-44.
- . 2007. Humean Doubts about The Practical Justification of Morality. In *Ethics and Practical Reason* edited by G. Cullity and B. Gaut. Oxford: Oxford University Press.
- Dworkin, Ronald. 1996. Objectivity and Truth: You'd Better Believe It. *Philosophy and Public Affairs* 25 (2):87-139.
- Egan, Andy. 2006. Epistemic Modals, Relativism, and Assertion. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 133 (1):1-22.
- Enoch, David. 2006. Agency, Shmagency: Why Normativity Won't Come from What Is Constitutive of Action. *Philosophical Review* 115 (2):169-198.
- . 2007. An Outline of an Argument for Robust Metanormative Realism In *Oxford Studies in Metaethics Volume 2*, edited by R. Shafer-Landau.
- Gibbard, Allan. 1992. Reply to Blackburn, Carson, Hill, and Railton. *Philosophy and Phenomenological Research* 52 (4):969-980.
- . 1999. Morality As Consistency in Living: Korsgaard's Kantian Lectures. *Ethics: An International Journal of Social, Political, and Legal Philosophy* 110 (1):140-164.
- . 2003. *Thinking How to Live*. Cambridge, Mass.: Harvard University Press.
- . Forthcoming. Evolved Thinkers and Moral Concepts In *Oxford Studies in Metaethics*, edited by R. Shafer-Landau. Oxford: Oxford University Press.
- Goodman, Nelson. 1955. *Fact, Fiction, and Forecast*. Cambridge: Harvard University Press.
- Heidegger, Martin. 1967. *Being and time*. Oxford,: Blackwell.
- Hussain, Nadeem J. Z. 2004. The Return of Moral Fictionalism. *Nous-Supplement: Philosophical Perspectives* 18:149-187.

- Jackson, Frank. 2000. *From metaphysics to ethics : a defence of conceptual analysis*. Oxford: Clarendon.
- Jackson, Frank, and Philip Pettit. 1998. A Problem for Expressivism. *Analysis* 58 (4):239-251.
- Korsgaard, Christine. 2009. *Self-Constitution: Agency, Identity, and Integrity*. Oxford: Oxford University Press.
- Korsgaard, Christine M. 1996. *The sources of normativity*. Cambridge ; New York: Cambridge University Press.
- . 2003. Realism and Constructivism in Twentieth-Century Moral Philosophy. In *Philosophy in America at the Turn of the Century (APA Centennial Supplement Journal of Philosophical Research)*. Charlottesville: Philosophy Documentation Center.
- Lewis, David. 1983. New Work for a Theory of Universals. *Australasian Journal of Philosophy* 61:343-377.
- . 1984. Putnam's Paradox. *Australasian Journal of Philosophy* 62:221-223.
- MacFarlane, John. 2007. Relativism and Disagreement. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 132 (1):17-31.
- Mackie, J. L. 1977. *Ethics : inventing right and wrong*. Harmondsworth ; New York: Penguin.
- Margolis, Eric, and Stephen Laurence. 1999. *Concepts : core readings*. Cambridge, Mass.: MIT Press.
- McDowell, John Henry. 1994. *Mind and world*. Cambridge, Mass.: Harvard University Press.
- Moore, G.E. 1993. *Principia Ethica*. Cambridge: Cambridge University Press.
- Nye, Howard. 2009. *Ethics, Fitting Attitudes, and Practical Reason: A Theory of Normative Facts*: PhD Dissertation, University of Michigan, Ann Arbor.
- Parfit, Derek. 2006. Normativity. In *Oxford Studies in Metaethics* edited by R. Shafer-Landau. New York: Oxford.
- . Forthcoming. *On What Matters*. Vol. (Draft of 7/5/2009). Oxford: Oxford University Press.
- Railton, Peter. 1986a. Facts and Values. *Philosophical Topics* 14:5-31.
- . 1986b. Moral Realism. *Philosophical Review* 95:163-207.
- . 2003. Noncognitivism about Rationality: Benefits, Costs, and an Alternative. In *Facts, Values, and Norms*. Cambridge: Cambridge University Press.
- . 2006. Normative Guidance. In *Oxford Studies in Metaethics* edited by R. Shafer-Landau. New York: Oxford.
- Rawls, John. 1999a. *A theory of justice*. Rev. ed. Cambridge, Mass.: Belknap Press of Harvard University Press.
- . 1999b. The Independence of Moral Theory. In *Collected papers*. Cambridge, Mass. ; London: Harvard University Press.
- Rosen, Gideon. 1998. Critical Study of Blackburn's Essays in Quasi-Realism. *Nous* 32 (3):386-405.
- Scanlon, T. M. 2000. *What We Owe To Each Other*. Cambridge: Harvard University Press.
- Schroeder, Mark. 2007. *Slaves of the Passions*. Oxford: Oxford University Press.
- . 2008. How Expressivists Can and Should Solve Their Problem with Negation. *Nous* 42 (4):573-599.
- Shafer-Landau, Russ. 2005. *Moral Realism: A Defense*. Oxford: Oxford University Press.
- Sider, Ted. 2001. *Four-Dimensionalism : An Ontology of Persistence and Time*. Oxford Clarendon Press.
- . 2009. Ontological Realism. In *Metametaphysics*, edited by D. J. Chalmers, D. Manley and R. Wasserman. Oxford: Oxford University Press.
- Smith, Michael. 1994. *The Moral Problem*. Cambridge: Blackwell.

- . 1997. In Defense of The Moral Problem: A Reply to Brink, Copp, and Sayre-McCord. *Ethics: An International Journal of Social, Political, and Legal Philosophy* 108 (1):84-119.
- . 2001. Some Not-Much-Discussed Problems for Non-Cognitivism in Ethics. *Ratio: An International Journal of Analytic Philosophy* 14 (2):93-115.
- . 2004a. *Ethics and the A Priori: Selected Essays on Moral Psychology and Meta-Ethics*. Cambridge: Cambridge Univ Press.
- . 2004b. Moral Realism. In *Ethics and the A Priori: Selected Essays on Moral Psychology and Meta-Ethics*. Cambridge: Cambridge University Press.
- . 2005. Meta-Ethics. In *The Oxford Handbook of Contemporary Philosophy* edited by F. Jackson and M. Smith. New York: Oxford Univ Pr New York.
- . 2007. Norms, Kinds, and Functions. In *Norms and Analysis Conference*. University of Sydney.
- . 2009. Desires, Values, Reasons, and the Dualism of Practical Reason. *Ratio: An International Journal of Analytic Philosophy* (XXII):98-125.
- . 2009b. Beyond the Error Theory. In *A World Without Values: Essays on John Mackie's Moral Error Theory*, edited by R. Joyce and S. Kirchin. New York City: Springer.
- Street, Sharon. 2006. A Darwinian Dilemma for Realist Theories of Value. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 127 (1):109-166.
- . 2008. Constructivism about Reasons. In *Oxford studies in metaethics*, edited by R. Shafer-Landau. Oxford: Oxford University Press.
- . Forthcoming-a. Mind-Independence Without the Mystery: Why Quasi-Realists Can't Have It Both Ways. *Noûs* <http://homepages.nyu.edu/~jrs477/streetpapers.html> (Draft of 5/1/2010).
- . Forthcoming-b. Objectivity and Truth: You'd Better Rethink It. *Philosophy and Public Affairs* <http://homepages.nyu.edu/~jrs477/streetpapers.html> (Draft of 5/1/2010).
- Sturgeon, Nicholas L. 2006. Ethical Naturalism. In *The Oxford Handbook of Ethical Theory*, edited by D. Copp. Oxford: Oxford University Press.
- Weatherston, Brian. 2003. What Good Are Counterexamples? *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* (115):1-31.
- Williams, Bernard Arthur Owen. 1981. Internal and External Reasons. In *Moral luck : philosophical papers, 1973-1980*. Cambridge Cambridgeshire ; New York: Cambridge University Press.

Chapter Three:
Ethical and Metaethical Agent-Attitude-Dependence

Introduction.

Consider the following scenario. Henry is a struggling freelance photographer. Unexpectedly, he has just been offered his dream job, a position as a staff photographer for *National Geographic*. In order to accept the job, Henry needs to sign a contract that is being delivered today and then to return it immediately. If he does not do this, then the job will go to his arch-nemesis Winston.

When the delivery truck arrives with an envelope for Henry, what should he do? It is relatively uncontroversial to say that, other things being equal, Henry should open the envelope and sign the contract.

Suppose that Henry himself, however, judges that he should *not* open the envelope. In particular, suppose that this is because one of the following two scenarios obtains:

(1) Henry falsely believes that there is an evil supernatural deity named “Satan” who has put lethal poison in the envelope. “This offer is too good to be true,” he thinks to himself. “This must be another one of Satan’s evil plots to kill me. Best not to do anything with the letter and just hope I get a *real* offer someday from *National Geographic*.”

(2) Henry has a mild dislike for opening envelopes and therefore thinks that he should not open this one. “I know that I need to open the envelope in order to sign the contract and secure my dream job,” he thinks to himself. “And I really think that it’s important to get this job and prevent Winston from getting it. But I mildly dislike opening envelopes. So, all things considered, I should not open this envelope.”

As with the initial judgment that Henry should open the envelope and sign the contract, it is relatively uncontroversial that Henry is making an incorrect normative judgment about what he should do in both of these cases. Or, more accurately, there is a so-

called “objective” sense of “should”—roughly, the sense in which what an agent should do depends on *what is actually the case*—on which this is uncontroversial.¹

If this is so, then these two scenarios suggest two different relatively uncontroversial ways in which an agent can be mistaken in her normative judgments about what to do in a given set of circumstances.

Take the first scenario. In this case, it seems that Henry’s mistake is his falsely believing that Satan put poison in the envelope. Thus, to generalize from this, we can say that one way in which an agent can be mistaken in her normative judgments is by having a false belief about the non-normative facts that make up the actual circumstances. When this sort of mistake happens let’s say that an agent makes a mistake about the relevant *non-normative facts*.

Now consider the second scenario. Henry’s mistake here is not one of straightforward factual error about the circumstances. Rather, to put it roughly, it is that Henry isn’t making normative judgments that cohere with his *own* more fundamental sense of what is valuable and worth pursuing. If he deeply values getting this job – and only has a mild dislike for opening envelopes – then there is a sense in which Henry is making a mistake *on his own terms* if he judges that he should not open the envelope. Henry’s “own terms” here are not necessarily set by any reflective normative judgment that Henry has consciously made, but also – or perhaps entirely – by some combination of his values, plans, desires, and unreflective normative judgments that structure his motivational system. Call these his *evaluative attitudes*. Furthermore, in order to capture an important feature that such attitudes have in common, let’s stipulate that by “evaluative attitude” we mean *any* attitude that tends to produce motivation for an agent when coupled with her prosaic beliefs.² With this understanding of “evaluative attitude” in hand, we can generalize from Henry’s case and say that one way in which an agent can be mistaken in her normative judgments is by failing

¹ This objective sense of “should” can be contrasted with two different so-called “subjective” senses of “should” – (1) a sense that concerns what an agent should do in light of the available evidence and (2) a sense that concerns what an agent should do in light of what she believes. The objective sense of “should” is what I am concerned with in this paper. It should be noted that on the first subjective sense of “should,” the one that concerns what an agent should do in light of the available evidence, it is also relatively uncontroversial that Henry should open the envelope.

² Which attitudes count as evaluative attitudes in this functional sense is an open question – and, importantly, will depend on one’s views in philosophy of mind and action. My use of the term “evaluative attitudes” draws on (Street 2008a). It should be emphasized that I do *not* intend my use of the term “evaluative” to suggest that all evaluative attitudes concern values.

to make normative judgments that in some sense *cohere* with her own evaluative attitudes – in particular, with those evaluative attitudes that are in some sense “fundamental” to that agent’s sense of what is valuable and worth pursuing. When this second sort of mistake happens, let’s say that an agent makes a mistake about *coherence with her own fundamental evaluative attitudes*.³

In addition to these two ways in which an agent can be wrong about what she should do, many of us think that there is another way. Many of us think that certain evaluative attitudes can simply be the wrong attitudes to have, independent of any facts about that agent. For instance, if Henry values counting blades of grass or torturing babies for fun more than anything else in the world, many of us would think that he is making a normative mistake *regardless* of how these evaluative attitudes are related to other ones that he has. If such a mistake were possible, it would mean that an agent could make an incorrect normative judgment even if she were both omniscient about the non-normative facts and made normative judgments that cohered with her most fundamental evaluative attitudes. Were this to happen, let’s say that an agent would be making a mistake about *agent-attitude-independent normative facts*.⁴

³ As we will see, there are many issues here about how to flesh this out more precisely, including, importantly, what the talk of “coherence” here might amount to. It should also be noted that one might take a different reading of this second case with Henry. For instance, one might think that if Henry really valued getting his dream job so much more than avoiding envelopes, then this would reveal itself in terms of how efficacious it is in motivating Henry’s actions. Thus, given how Henry actually thinks and behaves, one might doubt that it would be correct to attribute to Henry the evaluative attitudes I am asking the reader to imagine him as having. This touches on important issues about the attribution-conditions of specific evaluative attitudes, issue that I will return to later in this paper. For now, even if Henry’s case is a bit extreme, what is crucial is just that one agrees that at least *in some cases* an agent can fail to make a normative judgment that coheres with what are in fact her own fundamental evaluative attitudes.

⁴ I am here using the term “normative facts” in a minimal sense to refer, roughly, to the non-conceptual substantive facts *in virtue of which* some normative judgments in ethics are correct and others are not. Somewhat more specifically, I take the term “normative facts” to refer to facts about the instantiation of *normative properties* that make those judgments correct or not. In turn, I understand normative properties simply to be those features of an object in virtue of which a normative predicate in ethics such as “is a normative reason” applies to that object. This basic taxonomy for thinking about the connection between facts and properties draws on (Jackson 1998, 15-16), (Chalmers and Jackson 2001), and (Chalmers 1996). In using the terms “facts” and “properties” in this way, I do not intend to take on board any robust ideas about the nature of normative judgment or the metaphysics of normative properties. Indeed, I take my use of the terms “facts” and “properties” in this way to pick out elements in a basic picture of normative judgment that is consistent with a wide range of views on normative judgment and normative metaphysics ranging from (Parfit 2006) to (Gibbard 2003) to (Jackson 1998).

In a series of recent papers, Sharon Street has argued that it is not possible for us to make mistakes about *agent-attitude-independent normative facts*.⁵ This, she argues, is because such facts simply do not exist – or, more precisely, they *at least* do not exist in worlds, such as our own, where the content of agents’ evaluative attitudes have been heavily influenced by contingent socio-historical pressures. Some philosophers think that if there are no agent-attitude-independent normative facts, then we should be skeptical that there are *any* ethical facts in our world. Street disagrees: there are ethical facts about how agents should act, she claims, and we can be (and often are) mistaken in our judgments about them. Street’s aim is not to deny that there are ethical facts; rather, it is to make a claim about what the ethical facts are like. In particular, she argues that normative facts about what an agent should do in circumstances C are dependent *only* a) on the non-normative facts that constitute the circumstances C and b) on the facts about that agent’s own evaluative attitudes (which include both contingent and necessary facts about that agent’s evaluative attitudes). Insofar as she holds that normative facts obtain *only* because of these two such facts, Street thus argues for what we can define as *ethical agent-attitude-dependence* as opposed to *ethical agent-attitude-independence*.⁶

⁵ (Street 2006), (Street 2008b), (Street Forthcoming-b), (Street Forthcoming-c).

⁶ Following Street, the central debate that I am interested in is sometimes referred to as one about “mind-independence” rather than as one about “agent-attitude-independence.” See, for instance, (Street Forthcoming-b), (Street Forthcoming-c), and (Gibbard Forthcoming). Since the label “mind-independence” leaves it vague *whose* mind matters (e.g. agent vs. speaker) as well as *which aspects* of the mind matter (e.g. beliefs vs. desires), I prefer my label. My label is more cumbersome. However, I think it helps clarify what both Street and I agree is the central philosophical issue at hand for normative ethics. Street also sometimes puts the debate as one between “realist” and “anti-realist” positions in ethics (Street 2006). I think this is misleading given that there are other, more canonical ways of using these terms in metaethics that can be easily confused with this debate. In this paper, I will use the term “realism” in a way that is broadly inspired by the more canonical use in (Smith 2004) rather than Street’s way. In sort, I will use the term “realism” to refer to the thesis that normative properties are instantiated in the actual world such that there are correct possible normative judgments that concern the actual world. It should also be noted that the debate over ethical agent-attitude-independence and ethical agent-attitude-dependence is closely related to the debate between those that accept that there can be what Bernard Williams’ calls “external” normative reasons for action and those that think that all normative reasons are what he calls “internal.” (Williams 1981). However, since this terminology has become very closely tied to Williams’s particular way of framing the issue – and since I think Street’s way of putting it does a better job of locating the key philosophical issue at hand – I will stick with my Street-inspired terminology in this paper instead of Williams’s. Finally, drawing on (Schroeder 2007), it should also be noted that there is a good sense in which all ethical agent-attitude-dependent views can be taken to be “Humean” ethical theories and all ethical agent-attitude-independent views can be taken to be “anti-Humean” ones. However, given the diversity of ways in which the term “Humean” gets used in ethics, I will not use this label in the paper.

My use of the term “ethical agent-attitude-dependence” here – and indeed throughout this paper – differs in two ways from my use of the term in the previous paper of the dissertation, “The Metaethical Role of Evaluative Attitudes.” The first difference is that, in that previous paper, I defined “ethical agent-attitude-dependence” to include the further thesis that there are normative facts in our world of a certain sort: namely, facts about what individual agents should do. This assumption made it much easier to present the issues at hand. However, in this paper, given the different issues that I will be discussing, I will make no such assumption. The other difference is that my definition of ethical agent-attitude-dependence in this paper holds fixed the idea that *complete* ethical agent-attitude-dependence is the correct version of the view. In basic terms, in “The Metaethical Role of Evaluative Attitudes,” I defined *complete* ethical agent-attitude-dependence as the view that that normative facts about what an agent should do in circumstances C depend only a) on the non-normative facts that constitute C and b) on the facts about that agent’s own evaluative attitudes. In contrast, I defined *incomplete* ethical agent-attitude-dependence as the view that although these two facts *always* determine what an agent should do, other facts also always do so. That is: incomplete ethical agent-attitude-dependent views hold that, *in addition* to the non-normative facts about C and the facts about an agent’s evaluative attitudes, there will always be some additional type of fact that determines what an agent should do – perhaps facts about the aggregate amount of pleasure in the world. For the purposes of this paper, it will be easier if we simply take “ethical agent-attitude-dependence” to refer to *complete* ethical agent-attitude-dependence.

In turn, in light of this difference in usage, in this paper, I will use the term “ethical agent-attitude-*independence*” to refer to the denial of complete ethical agent-attitude-dependence. This is a wider use of the term “ethical agent-attitude-*independence*” than in the previous paper. In that paper, *incomplete* ethical agent-attitude-dependence was defined as an alternative to ethical agent-attitude-*independence* rather than as a specific version of it.

I ask the reader to charitably grant me the slight differences in terminology here between these two papers. These differences do not reflect a different understanding of the basic taxonomic divisions between different types of views, but rather only a choice of which labels to use to cover which sorts of views. My hope is that what might be lost in terms of continuity between the labels in the two papers is outweighed by gains in terms of ease of presentation in what follows.

Street holds that the debate over what I am here calling “ethical agent-attitude-dependence” is one of the most crucial in contemporary ethical theory. I agree. Furthermore, with Street, I think that there are general arguments that one can give in favor of ethical agent-attitude-dependence as I am defining the view in this paper: namely, as a view about what sorts of substantive facts make it the case that an agent should act one way or another. However, in this paper, I am going to put this debate to the side and instead focus on what I think is a crucial intramural debate among those drawn to ethical agent-attitude-dependence. In short, the debate is this: how should one understand the ramifications of agent-attitude-dependence for the *metaethical* question of what it is to make a normative judgment?

As I have defined it thus far, ethical agent-attitude-dependence amounts to a *substantive normative claim* in normative ethics: it is a claim about what sort of facts make it the case that an agent should act one way or another. It is thus a claim that is compatible with a wide range of theories about the basic metaphysics, epistemology, and semantics of normative thought and talk. Or, to put it another way, the substantive normative claim of ethical agent-attitude-dependence is compatible with a wide range of *metaethical* theories: theories not about which normative judgments in ethics are correct, but rather theories about how ethical thought and talk – and the ontological commitments involved in that thought and talk – should be integrated into a broader conception of reality.⁷ For instance, nothing in the thesis of *ethical agent-attitude-dependence* entails anything about whether *descriptivists* are right to hold that, at the most basic explanatory level, normative judgments express straightforward truth-apt belief or whether *expressivists* are right to deny that they do. Similarly, nothing in the thesis of *ethical agent-attitude-dependence* entails whether *naturalists* about normative properties are right that the normative property of *being the thing one should do* is *itself* a naturalistic property or whether *non-naturalists* about normative properties are right to hold that it is metaphysically *sui generis* and incapable of being studied directly by the natural sciences.

Nonetheless, many philosophers drawn to ethical agent-attitude-dependence as a substantive normative position have been drawn not only to particular claims about what normative judgments are correct, but also to particular metaethical claims, claims about how ethical thought and talk – and the ontological commitments involved in that thought and talk – should be integrated into a broader conception of reality. This, of course, might be

⁷ For a similar gloss on the central task of metaethics, see (McPherson Forthcoming, 1).

nothing more than a correlation that reflects basic personality-differences between philosophers or other psychological/sociological facts. However, it would also be good to know whether there is any deeper justificatory connection here – and, if so, what that connection is.

In the second paper of this dissertation, “The Metaethical Role of Evaluative Attitudes,” I argue for one sort of justificatory connection between the substantive normative thesis of ethical agent-attitude-dependence and metaethics: namely, I argue that there is good reason to accept ethical agent-attitude-dependence if one is committed to a specific view in metaethics that I call *naturalistic realism*. Now, in this paper, instead of thinking about how one’s metaethical position impacts the debate over ethical agent-attitude-dependence, I want to consider the reverse direction of justification. In short, I want to know whether one’s endorsement of the substantive normative position of ethical agent-attitude-dependence – and, more broadly, the arguments that lead one to it – should have any ramifications for theory-choice in metaethics; and, if so, what these ramifications are. In addressing this issue, I hope to gain a clearer understanding of how distinct types of theses about normative facts and their dependence on agents’ evaluative attitudes are related to each other; and, thereby, more broadly, to help us to understand some of the different ways in which claims in metaethics about the nature of normative judgment can interact with claims in normative ethics about which normative judgments are correct.

In order to focus my discussion, I am going to concentrate on Street’s main argument for ethical agent-attitude-dependence as a substantive normative position. Street’s argument is one of the most prominent contemporary arguments of this type, and also, I think, one of the most powerful arguments of this type. In rough terms, Street’s main argument in favor of ethical agent-attitude-dependence is as follows: she argues that in order to reconcile the view we have of ourselves when actively engaged in practical reasoning with the view we have of ourselves from a more detached “theoretical” perspective, we must endorse ethical agent-attitude-dependence. Call this Street’s *practical/theoretical argument*. What I want to know is this: if we grant that Street’s argument supports ethical agent-attitude-dependence as a substantive normative position, what significance – if any – does this have for theory-choice in metaethics?

As I have already indicated, Street’s substantive normative claim of agent-attitude-dependence is *compatible* with a wide range of metaethical theories about how to understand

normative judgment *as such*. Nonetheless, in this paper, I will argue that *if* Street's practical/theoretical argument supports ethical agent-attitude-dependence as a substantive normative thesis, then it also has significance for theory-choice in metaethics. In particular, I will argue that it lends (defeasible) support to the following broad metaethical thesis:

metaethical agent-attitude-dependence: to make a normative judgment that an agent A should ϕ in C is to express the belief that A's own evaluative attitudes favor A ϕ -ing in C.⁸

This metaethical view might strike some as a non-starter. For even if one were convinced that ethical agent-attitude-independence is *mistaken* as a substantive normative matter – which, obviously, would be no small feat – it at least seems *conceptually coherent* to hold that it is correct. For instance, it seems intelligible (if false) to think “I know that Susie is factually omniscient and that her own evaluative attitudes favor her counting blades of grass all day long, but she still should not do something so dreadfully boring and unproductive.” This, I think, is one of the most powerful objections to the sort of metaethical agent-attitude-dependent view I just sketched. It is also, I think, an instance of one of the most powerful objections to *many* theories in metaethics – namely, all metaethical theories that rule out seemingly intelligible views as ultimately conceptually confused. For, as one strand of Moore's famous so-called “open question argument” can be seen as articulating, it *prima facie* seems that there are no constraints (or *almost* no constraints) on what things one could *intelligibly* take to provide normative support for acting one way or another.⁹ Yet, I nevertheless think that it would be premature to rule out the metaethical view of agent-attitude-dependence automatically for this sort of reason. Theory-choice in metaethics (or anywhere else) should work largely by weighing the costs and benefits of different competing theories – and those theories such as expressivism and non-naturalistic realism that attempt *not* to rule out any substantive normative views as conceptually confused face many important (and daunting) challenges of their own. Progress in metaethics, I believe, is made in large part by seeing how different sorts of theories can receive *relative* dialectical

⁸ It should be stressed that not all *attitude-dependent* metaethical views are *agent-attitude-dependent* views. This is because some attitude-dependent metaethical views take the attitudes in question not to belong to the agent – for instance, they could be those of the *speaker* making a judgment about another agent or those of the members of some specified group.

⁹ For a development of Moore's open question argument in this way, see (Gibbard 2003).

advantage over other accounts. In this paper, I aim to establish that a certain broad type of metaethical view – namely, metaethical agent-attitude-dependence – has a relative dialectical advantage over other competing views: namely, it is one that receives support if Street’s practical/theoretical argument for the substantive normative thesis of ethical agent-attitude-dependence is correct. This metaethical view, I argue, receives support both a) from the conclusion of ethical agent-attitude-dependence itself as well as b) from the basic form of the practical/theoretical argument that Street uses to argue for it.

Following this argument, I will then consider how best to further develop the thesis of metaethical agent-attitude-dependence. I will argue that this thesis is best developed as a form of *naturalistic realism* in metaethics; that is, as a view according to which normative judgments express attitudes that, at the most basic explanatory level, are beliefs about naturalistic properties that are instantiated in our world. In particular, I will argue that it lends support to a naturalistic form of *agent-centered subjectivism* according to which a) normative facts are naturalistic facts of an agent’s psychology, b) some such normative facts obtain in the actual world, and c) speakers express truth-apt beliefs about such facts when they make normative judgments.

In putting forward this subjectivist position, one of my additional aims will be to explain not only why I think it is motivated by Street’s practical/theoretical argument, but also how it is related to the metaethical view that Street herself takes to be the best way of developing an agent-attitude-dependent view in metaethics: namely, the position that she calls *metaethical constructivism*. Street claims that this view is not only a metaethical position, but, moreover, one that is distinct from the dominant broad types of views in contemporary metaethics – including, most importantly, expressivism, non-naturalistic realism, and naturalistic realism. If Street’s metaethical constructivism is to be put forward as a metaethical view in the sense that I am concerned with in this paper, I will argue that it then should in fact be developed as an instance of metaethical naturalistic realism rather than as a competitor to it. More specifically, I will argue that it is best developed as precisely a version of the type of *agent-centered subjectivism* that I put forward. Such a view faces many of the same challenges that Street’s view faces when her view is conceived of as a distinctive type of metaethical view (which is how she conceives of it). However, my proposed form of agent-centered subjectivism has important advantages as a metaethical position and represents, I

argue, the best way to develop the thesis of *metaethical agent-attitude-dependence* in light of the arguments that gave rise to the thesis in the first place.

I will break up my work in this paper into three main sections. In §1, I explain Street's practical/theoretical argument in favor of ethical agent-attitude-dependence. Street presents her practical/theoretical argument in different ways – including, importantly, sometimes as an argument directly for metaethical claims as well as for the substantive normative thesis of ethical agent-attitude-dependence. However, I think that her argument is most plausibly read as one for ethical agent-attitude-dependence in the first instance – and, moreover, that her argument forms the basis for a powerful version of such an argument regardless of the interpretative issues.¹⁰ In §2, I argue that if Street's practical/theoretical argument for ethical agent-attitude-dependence is correct, then it in fact lends support to agent-attitude-dependence at both the *substantive normative* and *metaethical* levels. With this conclusion as a departure point, I then turn to the question of how best to develop an agent-attitude-dependent view in metaethics, and argue in favor of developing it as a form of naturalistic realism that is a type of agent-centered subjectivism. In §3, I argue that, despite the fact that Street puts metaethical constructivism forward as an alternative to naturalistic realism in metaethics, her view is in fact best developed as this same type of agent-centered subjectivism that I argue for in the second section.

1. Street's Practical/Theoretical Argument and Ethical Agent-Attitude-Dependence.

The main argument that Street develops in favor of ethical agent-attitude-dependence is what I am calling her *practical/theoretical argument*: she argues that if we want to successfully integrate the views that we hold when actively engaged in practical reasoning with a scientific view of the world, then we are forced to accept ethical agent-attitude-dependence. In this first section of the paper, I will explain what I take this practical/theoretical argument to be and why Street takes it to lend support to the *substantive normative thesis* of ethical agent-attitude-dependence. In so doing – and, indeed, throughout this paper – I will rely on what I think is one charitable reading of Street's argument on this front. Even if it is not consistent with everything that Street writes, this reading of Street is

¹⁰ In taking Street's practical/theoretical argument to be one for a substantive normative thesis, specifically for the one that I am calling "ethical agent-attitude-dependence," I draw on the view of this argument in (Street Forthcoming-b), (Street Forthcoming-c), and (Gibbard Forthcoming).

supported by much of her writing – and, most importantly, forms the basis for a potentially powerful argument in favor of the substantive normative thesis of ethical agent-attitude-dependence.

In her paper “Objectivity and Truth: You’d Better Rethink It,” Street introduces her basic practical/theoretical argument by way of a puzzle that all of us face in virtue of being able to occupy what she calls two different “standpoints.” On the one hand, we are beings who make normative (and evaluative) judgments.¹¹ We are thus creatures who are capable of occupying what Street calls the “practical standpoint”: the standpoint of a creature “who is in the state of mind of “valuing”¹²; or, more broadly, who takes certain things to have normative significance for deciding practical questions about *how to live* or *what to do*. She writes:

When we occupy this standpoint we understand ourselves as beings who are capable of recognizing what practical reasons we have, what we should or ought to do, what is good, valuable, and worthwhile, what is morally right and wrong, and so on. We think of ourselves, in other words, as beings whose normative judgments (many of them, anyway) are true.¹³

Creatures like us, Street claims, can’t help but take up the practical standpoint. Even if we do not consciously engage in systematic normative reasoning about how to live or what to do, we can’t help but think about what to do or help but feel that certain facts about the world “call for” or “count in favor of” one response or another.

On the other hand, we are beings who can also engage in theoretical reflection about the nature of our normative judgments and their place in the world. We are thus creatures who are capable of occupying what Street calls the “theoretical standpoint”: the standpoint of a creature whose aim is to reflect on how things *are*, rather than on how they *should* be. She writes:

When we occupy this point of view on ourselves and our values, we understand ourselves as beings who are part of the world of cause and effect, and whose normative judgments are subject to causal explanation.¹⁴

¹¹ In what follows, I will use the term “normative” in the broad sense to cover both normative terms in a narrow sense (e.g. “right,” “reason,” and “ought”) and evaluative terms (e.g. “good” and “value”).

¹²(Street Manuscript, 1).

¹³ (Street Forthcoming-c, 1).

¹⁴ (Street Forthcoming-c, 1).

Just as the practical standpoint is inescapable, so too, Street claims, is the theoretical standpoint. Or, more precisely, we might say that it is *at least* inescapable for those drawn to basic reflection about human activity and its place in the world in light of a broadly modern scientific understanding of the world. For – just as with any psychological phenomena – we are capable of seeing our normative judgments as psychological phenomena “whose causal origins and effects we may investigate and discover.”¹⁵

Street claims that our ability to take up both these standpoints gives rise to a serious puzzle. In broad terms, this puzzle is how to simultaneously believe *both* the claims we accept in the practical standpoint as well as those that we accept in the theoretical standpoint. Street calls this the *practical/theoretical* puzzle.¹⁶ In basic terms, Street’s claim that there is a puzzle here stems from the fact that from the theoretical standpoint, we can see that there are strong causal-explanatory accounts of *why* we make the normative judgments that we do that themselves make no reference to any normative facts as such.

To see how this is so, and in turn why this creates a puzzle, consider our judgments that we should not needlessly throw ourselves under falling boulders, that we should care for our children, or even that our own survival is good. One basic way to explain the causal forces that led us to make these judgments, Street claims, runs roughly as follows: evolutionary pressures helped mold us as creatures who had evaluative attitudes that favored such things – otherwise we would have died off – and, in turn, our conscious normative judgments are grounded in the fact that we have such evaluative attitudes. This, Street claims, is the core of the strongest scientific explanation we can give for why we would be creatures that would make such judgments. Importantly, this explanation does not itself invoke any normative facts *de dicto*.¹⁷ Furthermore, Street claims, the same basic story could be said for those of us in the early twenty-first century who value such things as democracy and gender equality. This is because there are, she claims, strong social scientific explanations for why I have evaluative attitudes that favor such things – explanations about the influence

¹⁵ (Street Forthcoming-c, 2).

¹⁶ (Street Forthcoming-c, 10).

¹⁷ It should be noted that, as we saw in the first chapter, Street’s idea here that this explanation is the strongest sort of explanation at hand will hardly be granted by all sides in the debates over normative explanations. See, for instance, (Sturgeon 1988).

of contingent cultural and socio-historical setting — that do not themselves invoke the actual existence of normative entities such as normative reasons. As she puts it:

When I occupy the theoretical standpoint... I understand my normative judgments — including my judgments that I have reason to exercise, to contribute to UNICEF, and so on—as having been shaped by causes such as my upbringing, cultural background, and inherited psychological tendencies; it is clear to me that had some or all of these factors been different, I wouldn't have made the same set of normative judgments that I now make.¹⁸

This raises a question: if I am aware that the content of many (or all) of my normative judgments is capable of being explained as the product of contingent forces such as social-historical setting and evolutionary forces acting on my underlying evaluative attitudes, what confidence should I have that these judgments are actually *correct* or *true*?

This question is the heart of Street's practical/theoretical puzzle. As she puts it, this puzzle is as follows:

Insofar as we regard our normative judgments as true, we must agree that there is a striking coincidence between (1) the normative judgments that are true, and (2) the normative judgments that causal forces led us to believe [are true]. This coincidence between the normative truth and what causes led us to believe is puzzling and demands some kind of explanation.¹⁹

As Street sees it, in other words, the basic practical/theoretical puzzle is to explain why there is such an overlap between (1) the normative judgments that we take to be true when occupying the practical standpoint and (2) the normative judgments that we can explain why we make from the theoretical standpoint using non-normative causal-explanatory stories.

At the outset, Street is making only minimal assumptions about how this puzzle must be tackled: in essence, all she holds is that *some* explanation is needed for this overlap. Yet, given that she thinks both the practical and theoretical standpoints are also *inescapable* in practice, she also wants to place an additional constraint on these explanations as well. This is that these explanations must result in *non-skeptical* solutions that vindicate two central beliefs we have in the practical standpoint: a) that there really are correct answers to normative questions and, b) that we are creatures who are not totally hopeless at making progress toward those answers.

¹⁸ (Street Forthcoming-c, 1).

¹⁹ (Street Forthcoming-c, 10).

In theory, one could argue that some version of the practical/theoretical puzzle arises for *all* purported normative truths – whether those truths are about how one should live or whether they are about what beliefs one should have. Street’s main concern is with practical normative truths about how to live and what to do: “truths of the form that *X* is a normative reason to *Y*, that one should or ought to *X*, that *X* is good, valuable, or worthwhile, that *X* is morally right or wrong, and so on.”²⁰ Although the focus on this topic has already been implicit in my discussion so far, it is worth stating explicitly that this basic domain of normative truths will be my focus in this paper. Moreover, in order to regiment the discussion in what follows, I will – along with Street – focus my discussion on truths about *normative reasons for action* (which I will henceforth just call “normative reasons”). As T.M. Scanlon puts it, this is the concept of a consideration or fact that “counts in favor of” or “calls for” a course of action in the sense that it *justifies* rather than explains that action.²¹ The idea of a normative reason for action is distinct from the idea of a specifically *moral* reason for action; for one might think that we have normative reasons to do things on non-moral grounds or for non-moral ends. An agent’s normative reasons for action in circumstances *C* are meant to be about what that agent has reason to do *full stop* in those circumstances. These normative reasons are, in other words, tracking practical normativity *simpliciter*.²² In this paper, I use the term “ethics” in a broad sense to cover this question about what an agent has normative reason to do *full stop*.²³

²⁰ (Street Forthcoming-c, 2). It should be noted, however, that Street is also interested in exploring whether or not a version of her practical/theoretical argument could be applied in other normative domains as well. See (Street Forthcoming-a) for her discussion of how it might be applied to normative properties in epistemology.

²¹(Scanlon 2000). This formulation is meant to distinguish *normative reasons* (facts that justify an agent’s action whether or not that agent is aware of them) and *motivating reasons* (facts that explain an agent’s action in either teleological or causal terms). On this distinction, see (Smith 1994) and (Darwall 2003). It should be noted that there is an important question of whether normative reasons should be identified with the facts themselves (the things in virtue of which propositions are true or false) or instead with the propositions. For the purposes of this paper, I am going to leave this issue to the side since the differences between these two theories does not make any difference for the core issues that I will address. See the introduction of (Schroeder 2007) for an insightful critical discussion of how to make more precise the thought that normative reasons “count in favor” of action.

²² In this paper, I use the phrase “an agent’s normative reasons” as a helpful shorthand to refer to the normative reasons that apply to that agent. In everyday contexts, we often speak of “an agent’s reasons” to refer to her motivating reasons for performing a given action. However, this is *not* my usage in this paper.

²³ My use of “ethics” in the broad sense draws on (Gibbard 1990), (Darwall 1997), (Velleman 2009), (Williams 1985), and (Street 2008b). It should be noted that the term “ethics” is sometimes used in a narrower way – namely, one where ethics is seen as fundamentally tied to morality, which is in turn

With this understanding of the practical/theoretical puzzle and normative reasons in hand, we are now in a position to turn to Street's "practical/theoretical argument." This argument aims to establish that ethical agent-attitude-independent accounts in ethics cannot successfully answer the practical/theoretical puzzle and that, because of this, we are forced to adopt an ethical agent-attitude-dependent account instead.

Street argues that the inability of ethical agent-attitude-independent accounts to answer the practical/theoretical puzzle can be seen by considering the following question: is there some connection between, on the one hand, the contingent casual forces that shaped an agent's evaluative attitudes and, on the other hand, the normative facts? For now, in order to help focus the discussion, let us follow Street's lead in "The Darwinian Dilemma" and conceive of this question as concerning, mainly, the *evolutionary* forces shaping an agent's attitudes. Street claims that the ethical agent-attitude-independent theorist faces an unacceptable dilemma in answering this question.

Suppose that the ethical agent-attitude-independent theorist holds that there is no connection between the contingent evolutionary forces that acted on shaping an agent's evaluative attitudes and the facts about what the normative truths are. This means that one denies any relation between evolutionary influences and the agent-attitude-independent normative facts such that, as Street puts it, one regards "the influence of evolutionary forces on our evaluative attitudes as no better than random with respect to the [normative] truth."²⁴ Given how undeniably strong the impact of evolutionary forces have been on shaping the content of our evaluative attitudes – and, thus, how strong their impact has been on shaping the content of our "systems of normative judgment"²⁵ – the ethical agent-attitude-independent theorist is left with one of two options. The first is to hold that since the content of one's evaluative attitudes was fundamentally shaped by forces that in no way were tied to what I am calling "agent-attitude-independent normative facts," then it is likely that we lack the ability to identify those facts at all. In other words, the first option is to embrace normative skepticism – it is to hold, as Street puts it, that "our normative judgments are in

understood as something that concerns a more specific topic than just how an agent should live *full stop*. For instance, the term "ethics" is sometimes used refer to solely to what an agent should do on specifically *moral* topics or what her specifically *moral* reasons are. I have no such restrictions in mind when I use the term "ethics" and, moreover, have nothing at all to say in this paper about the nature of morality or specifically moral normative reasons.

²⁴ (Street 2008b, 208).

²⁵ (Street 2008b, 208).

all likelihood hopelessly off track.”²⁶ Street claims that this is unacceptable – we must hold that we are *not* hopelessly off track in this way when we take up the practical standpoint. The second option is to “hold that an astonishing coincidence took place—claiming that as a matter of sheer luck, evolutionary pressures affected our evaluative attitudes in such a way that they just happened to land on or near the true normative views among all the conceptually possible ones.”²⁷ But this too is unacceptable. It makes our confidence in our ability to make correct normative judgments akin to a matter of mere faith.

Now suppose that the ethical agent-attitude-independent theorist affirms that there is a relation between evolutionary influences on our evaluative attitudes and the agent-attitude-independent normative facts on the other. This, claims Street, leads to the second horn of the dilemma. For, once the ethical agent-attitude-independent theorist accepts this thesis, she then owes us some account of what this connection is; otherwise, as Street puts it, this line will be “no more satisfactory than positing that evolutionary forces landed us on or near the truth by sheer coincidence.”²⁸ Street considers two seemingly promising accounts that are available to the agent-attitude-independent theorist in order to avoid this conclusion.

The first account is what Street calls the *Tracking Account*. According to the Tracking Account, creatures that correctly “tracked” the agent-attitude-independent normative facts had more reproductive success – and thus, “creatures with an ability to do so were selected for.”²⁹ The problem, though, says Street, is that this account is scientifically indefensible: there is simply no need to posit such additional agent-attitude-independent normative facts *at all* in order to explain why human beings tend to make normative judgments with certain content as opposed to others.³⁰ Such facts, Street claims, are superfluous from the theoretical standpoint. For one can simply give the much simpler following account instead: we make the basic normative judgments that we do because having certain evaluative attitudes – and then tending to make certain normative judgments that affirm what those attitudes favor – helped further the reproductive success of our ancestors. In turn, we are then left as creatures with much the same underlying evaluative attitudes and dispositions to make normative judgments one way or another. For instance, consider the fact that most human

²⁶ (Street 2008b, 208).

²⁷ (Street 2008b, 209).

²⁸ (Street 2008b, 209).

²⁹ (Street 2008b, 209).

³⁰ (Street 2008b, 209).

beings tend to believe that their own survival is worth pursuing. The best scientific explanation of why we make such judgments, Street claims, is not that such creatures were recognized that there were agent-attitude-independent normative facts that make survival worth pursuing. Rather, as she puts it, it is “that creatures who valued their survival tended to do what promoted it, and therefore left more descendants.”³¹

The second account is what we can call the *Indirect Tracking Account*. According to this account, it is not that creatures who correctly “tracked” agent-attitude-independent facts were selected for *because* they tracked those facts. Rather, according to this account, we evolved a general capacity of rational reflection which we then discovered we could use to discover attitude-independent truths.³² Such an account might appeal to the case of advanced mathematical reasoning for support and claim as follows: just as we likely evolved advanced mathematical reasoning by expanding on rational capacities that we evolved for reasons *other* than responding to advanced truths in mathematics that played a casual role in the development of our reasoning, so too did we evolve a parallel rational capacity in practical reasoning. The problem with this reply, claims Street, is that it is again just as implausible on basic scientific grounds as the tracking account. It is simply extravagant, she claims, to posit that we evolved such a general capacity at all when her much simpler explanation is on the table. This is, in essence, the same reply that Allan Gibbard gives to this proposal. He writes: “We’re a bunch of cobbled together heuristics that were naturally selected because, in our ancestral populations, those a little further toward these heuristics tended to reproduce more. We didn’t evolve with a faculty set generally to track rationality, whatever that turns out to be.”³³ Street concludes that her basic proposed explanation is better off than the indirect tracking account— and, by inference, better off than *any* account that the agent-attitude-independent theorist could give. Hence, she concludes, the ethical agent-attitude-independent theorist is no better off on this second horn of the dilemma than on the first horn, and therefore has no satisfying response to the practical/theoretical puzzle.

In contrast, consider what would be the case if one endorsed agent-attitude-dependence instead of agent-attitude-independence. Start with the first horn of the dilemma

³¹ (Street 2008b, 209).

³² This is, for example, the type of account that Parfit briefly proposes in (Parfit Forthcoming, 106-107).

³³ (Gibbard 1992, 978).

for the agent-attitude-independent theorist that Street discusses: the horn where the agent-attitude-independent theorist claims there is *no* connection between the contingent evolutionary forces that acted on shaping an agent's evaluative attitudes and the normative facts. Unless the ethical agent-attitude-dependent theorist wanted to adopt the implausible position that our evaluative attitudes were in fact not at all influenced by evolutionary forces, then there is no way for her to take this first horn of the dilemma.

Now consider the second horn of the dilemma that the ethical agent-attitude-independent theorist faced: namely, the horn where one affirms a connection between the evolutionary forces that acted on an agent's evaluative attitudes and the normative facts. Here, the ethical agent-attitude-dependent theorist can simply accept the basic scientific story that Street proposes without this in any way challenging her belief that there are normative facts. This is because, according to the ethical agent-attitude-dependent theorist, the normative facts are dependent only on facts that this story has already made reference to in its causal-explanatory account – namely, facts about the agent's own evaluative attitudes and the non-normative facts about what the circumstances are. The ethical agent-attitude-dependent theorist thus avoids needing to tell any additional, scientifically questionable story about how we came to have evaluative attitudes that tracked the agent-attitude-independent normative facts – whether those facts be ones about God's will, the customs of a specific society, or the attitudes of speakers making judgments about the agent. The ethical agent-attitude-dependent account, in other words, runs into no trouble with the best causal-explanatory accounts one has from the theoretical standpoint. At the same time, it allows one to affirm in practical reasoning that one is not hopeless at figuring out the normative facts – for those facts are dependent precisely on the *actual* evaluative attitudes that one already uses in guiding one's practical reasoning. Of course, as Street emphasizes, we might have difficulty in understanding the content of our own fundamental evaluative attitudes – or especially those of other agents' when we are making judgments about *their* normative reasons. Yet, this would not undermine our basic beliefs in the practical standpoint in the way that an ethical agent-attitude-independent view would: for even if it might be *difficult* for an agent to learn about her normative reasons (or those of other agents), this conclusion does not entail that we are *hopeless* at learning about those normative reasons in the way that we seem we would be if we accept ethical agent-attitude-independence.

There are many important objections to Street's practical/theoretical argument that one would need to assess in order to establish that, at the end of the day, it really does favor ethical agent-attitude-dependence. However, in the context of this paper, this is not the important issue. My concern in this paper is not whether Street's practical/theoretical argument *does* in fact establish ethical agent-attitude-dependence, or even whether her practical/theoretical argument simply establishes a dialectical advantage that the ethical agent-attitude-dependent theorist has. Rather, in this paper, my concern is with the following issue: *if* Street's argument does provide justificatory support for the substantive normative thesis of ethical agent-attitude-dependence, then what import – if any – does this have for theory-choice in metaethics? With an understanding of Street's basic practical/theoretical argument in hand, we are now in a position to turn to this question in the next section.

2. *Metaethical Agent-Attitude-Dependence and Agent-Value Subjectivism.*

In the last section, I presented Street's practical/theoretical argument and explained why she thinks it lends support to the substantive normative position that I am calling *ethical agent-attitude-dependence*. However, when presenting her practical/theoretical argument, Street herself sometimes does so in a way that – at least *prima facie* – makes it appear that she also takes it to yield a metaethical conclusion as well.³⁴ Street's transition between these two types of claims is not always clear – let alone is it clear why she takes them to be warranted. This is an issue because, it is not clear that a successful version of Street's practical/theoretical argument would itself directly yield a metaethical conclusion in the same way that it would yield a substantive normative one.³⁵ Nonetheless, I think Street is right to suspect that her practical/theoretical argument has distinctively *metaethical* import. In this section, I will argue that *if* Street's practical/theoretical argument supports ethical agent-attitude-dependence, then it also has significance for theory-choice in metaethics. In particular, I will argue that it lends (defeasible) support to the following metaethical thesis about the nature and content of normative judgment:

Metaethical Agent-Attitude-Dependence: to make a normative judgment that an agent A should ϕ in C is to express the belief that A's own evaluative attitudes favor A ϕ -ing in C.

³⁴ For instance, in (Street 2006) Street uses the practical/theoretical argument to attack a range of positions that are usually understood to be distinctively metaethical ones.

³⁵ This point is emphasized in (Gibbard Forthcoming).

This thesis, I will argue, is in no way *entailed* by accepting Street’s conclusion in favor of *ethical* agent-attitude-dependence on the basis of her practical/theoretical argument. However, I will nonetheless argue that, when coupled with plausible views about theory-choice in metaethics, metaethical agent-attitude-dependence view does receive some important justificatory support from both a) the conclusion of ethical agent-attitude-dependence as well as b) the basic form of the practical/theoretical argument that Street uses to argue for it.³⁶ Furthermore, I will argue that once we see the reasons why metaethical agent-attitude-dependence receives such relative dialectical support, I will argue that we can then see good reason to develop metaethical agent-attitude-dependence in a certain way – namely, as a thoroughly naturalistic metaethical view that I will call *agent-value-subjectivism*.³⁷

³⁶ The fact that the substantive normative thesis of agent-attitude-dependence is distinct from *metaethical agent-attitude-dependence* is obscured in some of Street’s recent discussion about what I am calling “agent-attitude-dependence.” For, in many parts of her work, Street claims that there is a *single* issue about how agents’ evaluative attitudes are related to normative facts that can *either* be seen as a substantive normative or metaethical issue depending on one’s own metaethical commitments. See, for instance, (Street 2009) and (Street Forthcoming-b). This single debate, Street claims, is essentially the debate between agent-attitude-independence and agent-attitude-dependence. In many cases Street calls this the debate between “realism” and “antirealism” in ethical theory. See, for instance, (Street Forthcoming-c) and (Street 2006). I think that this way of talking is misleading for two reasons. The first is that the labels “realism” and “antirealism” already have other, more canonical uses in metaethics. This is a point that Street herself seems to herself now hold as well. See (Street 2009, 295 n9). In this paper, I will use “realism” in the following way: namely, to refer to different sides of a debate about whether or not normative properties are instantiated in the actual world (realists say yes; antirealists say no). The reason that Street’s way of putting things is misleading is that even if it is true that the debate Street is focusing on – a debate over when normative facts obtain – is one that will be interpreted differently depending on one’s metaethical commitments, this does *not* entail that there is no way of distinguishing whether Street herself (or anyone else) intends the thesis to be advanced as substantive normative thesis or a metaethical one. The issues at hand, I claim, are made clearer by simply stating that ethical agent-attitude-independence and –dependence are substantive normative claims, and then simply stating metaethical agent-attitude-independence and –dependence as different claims.

³⁷ The term “subjectivist” gets used in ethics and metaethics in a wide variety of different ways. Because of this, my use of the term here is somewhat unfortunate. For instance, the term “subjectivism” is sometimes used to refer to views that hold that when speakers make normative statements, this involves them stating what their *own* evaluative attitudes are. This is decidedly *not* my view. The term “subjective” is also sometimes used to refer to views that deny that facts about normative reasons are objective (in some further specified sense of “objective”). The view that I will advance will not have any such implication. Or, more accurately, it will not have any such implication given many of the standard uses of the term “objective” in this context – and, moreover, the uses of the term that I think should matter. As we will see, according to agent-value-subjectivism, the normative facts turn out to be every bit as objective as any other psychological fact. It should be noted that there is, however, a form of so-called “Kantian objectivity” about ethics that I think the view will need to deny when fully developed.

In order to situate my argument in this section, we first need to clarify what I mean by the term “metaethics” in the context of this discussion. “Metaethics” and “normative ethics” are terms of art in contemporary philosophy, and different philosophers mean different things by using these words. For the purposes of this paper, I am going to work with one fairly standard picture of this division that runs as follows.

The central task of *normative ethics* is to figure out the correct answers to practical normative questions – normative questions, that is, about such things as how to live and what to do (rather than, for instance, normative questions about what to believe). If we understand the term “normative facts” to refer to facts about the instantiation of normative properties – facts in virtue of which normative judgments are correct or not – then we can say that the task of normative ethics is to discover the normative facts. This includes specifying what those facts are in specific circumstances (for instance, whether or not Andy should donate more money to Oxfam) as well as specifying general features of those facts (for instance, whether all such facts can be explained as stemming from what promotes the maximum amount of utility in a given situation). In contrast, the central task of *metaethics* is to make sense of what it is to make a practical normative judgment as such – and, in particular, to understand the semantics, metaphysics, and epistemology involved in making and stating practical normative judgments. Rather than seeking to make correct practical normative judgments, it takes such judgments as its object of study and asks how we should best make sense of them given broader commitments about how to understand ourselves and the world. Thus, I think that the central metaethical task can be summed up as follows: to make sense of how ethical thought and talk – and the ontological commitments involved in that thought and talk – should be integrated into a broader conception of reality.³⁸ For my purposes in what follows, this will be how I will understand the normative ethics/metaethics distinction.³⁹

³⁸ For a similar gloss on the central task of metaethics, see (McPherson Forthcoming, 1).

³⁹ It should be stressed that, as I am using the term here, *ethics* involves the broad practical questions of how agents should live and act *full stop*. To put it another way, ethics concerns how one should answer these questions *simpliciter* – and not, for instance, how one should answer these questions from a distinctively *moral* point of view (which one might or might not take to be definitive for answering the former questions depending on how one understands the meaning of the term “moral”). In turn, when I use the term “metaethics,” I intend to keep this use of the term “ethics” fixed. Thus, when I use the term “metaethics,” I intend to signal that I am concerned with the project of taking *ethical* thought and talk as an object of inquiry rather than taking specifically *moral* thought and talk as an object of inquiry.

The task of metaethics as I have glossed it can be broken down into two parts. On the one hand, there is what we can call the *non-fundamental* metaethical task of explaining how different parts of normative thought and talk are internally related to each other. For instance, we might want to know how talk about goodness is related to talk about obligation or reasons. On the other hand, there is what we can call the *fundamental* metaethical task of explaining what it is to engage in any sort of normative thought and talk at all.⁴⁰ In order to succeed at this fundamental metaethical task, I will take it that a theory cannot assume certain normative thoughts as part of its explanation. Rather, it must treat the entire practice of engaging in normative thought and talk about how to live as an explanandum rather than as an explanans.⁴¹ In what follows, I will concern myself exclusively with *fundamental* metaethics and, furthermore, will thus often use the term to pick only this part of the broader metaethical task that I have identified.

Given the range of things that philosophers mean by the term “metaethics,” my use of the term in this way – and my gloss of the non-fundamental vs. fundamental tasks of metaethics – can be seen as partly stipulative if one wants. This, I think, does not much matter for my work in this paper. What is important is that this definition picks out a fairly

⁴⁰ I thank Howard Nye for first suggesting to me this helpful way of drawing the distinction. For Nye’s own formulation of this distinction, see his PhD dissertation (Nye 2009).

⁴¹ I draw on (Hussain 2004, 2) in my formulation of the task of fundamental metaethics and what it must involve in order to be successful. It should be noted that there are at least two complications that should be noted about my formulation of what the task of fundamental metaethics amounts to. The first complication has to do with the fact that many philosophers take judgments about meaning and/or mental content to themselves be normative judgments. For reflection on this issue see (McDowell 1994), (Railton 2003c), (Brandom 1994), and (Gibbard 2003). Since this might mean that *any* metaethical view would turn out to take some normative thoughts as given as an explanans, there is thus an important question whether or not the view of metaethics that I have stated is really feasible – or, at the very least, whether or not it is actually more controversial than intended. These issues become particularly complicated when we think about the nature of *metanormative* theory as a whole. Yet, for the purposes of this paper, since we are focused on *metaethics* in particular, I am going to sidestep these issues by stipulating that a view can count as a fundamental metaethical one iff it does not take any normative judgments about *what to do* or *how to revise one’s evaluative attitudes* as an explanans – thus, for now, leaving it open whether it can tacitly take normative judgments of a different sort for granted. The second complication concerns the fact that – although many philosophers (including Street) – hold that one can not take any normative thoughts for granted in formulating a fundamental metaethical view, not everyone agrees with this. Most importantly, in (Wedgwood 2009), Ralph Wedgwood puts forward a view that is explicitly circular in its understanding of the metaphysics of normativity – understanding the relationship between normativity and intentionality in an ultimately circular, but nonetheless purportedly illuminating, way. For the purposes of this paper, I am simply assuming along with (Street 2008a) that such circular definitions cannot work in fundamental metaethics.

specific philosophical project: one that many in mainstream contemporary metaethics consider to be *the* fundamental project of metaethics and that, more importantly, is not clearly impacted by Street's practical/theoretical argument in favor of the substantive normative thesis of ethical agent-attitude-dependence. What I want to know is how – if at all – accepting ethical agent-attitude-dependence on the basis of the practical/theoretical argument should matter for theory-choice in the project that I am here calling “fundamental metaethics.”

It is worth emphasizing at the start that fundamental metaethical theories concern a markedly different topic than the topic addressed in the thesis of ethical agent-attitude-dependence. To see why this is so, suppose that one had a complete list of the correct statements about what all agents should do in all circumstances. We could ask more or less abstract questions about the individual statements on this list. For instance, we could ask a rather specific question such as: what generalizations would be true of those statements on the list that concern promises made to others? Another, more abstract question we could ask is whether each and every statement on this list had the following feature: the fact's correctness depends on the total set of evaluative attitudes of the specific agent in question and, moreover, is dependent on nothing else other than the non-normative facts that make up the circumstances of that agent's situation. Those that embrace ethical *agent-attitude-independence* will deny that this is so. In contrast, those that embrace ethical *agent-attitude-dependence* will affirm this is so. Both of these theses, remember, concern the question of when certain normative facts obtain rather than others. More specifically, they concern the question of what facts make it the case that an agent has a normative reason to act in one way or another.

Now contrast this issue with another one. Take our complete list of the correct normative statements. Suppose one asks *not* “what generalizations are true about these normative statements?” but rather “what is it to think that a normative statement *is* a correct normative statement in the first place?” Unlike the first question, this second question is what I am calling a fundamental *metaethical* one. For, unlike the first question, it does not concern the substantive normative question of *which normative judgments are correct* but rather the question of how ethical thought and talk – and the ontological commitments involved in that thought and talk – should be integrated into a broader conception of reality.

Because ethical agent-attitude-dependence is itself not a metaethical theory, one can accept a wide range of leading metaethical theories and still accept agent-attitude-dependence. Perhaps most importantly, ethical agent-attitude-dependence is compatible with either side of two of the most central debates in contemporary metaethics: 1) the debate between descriptivists and expressivists about which mental states are explanatorily primary in accounting for normative judgment and 2) the debate between non-naturalists and naturalists about whether or not the basic normative property (suppose, for instance, the property of being a normative reason) is metaphysically *sui generis* or not.

Take the first debate. Roughly speaking, *descriptivists* hold that, at the most basic explanatory level, normative judgments are much the same as any prosaic judgment: they express a belief-like state that aims to represent correctly how things are. In other words, descriptivists hold that normative judgments express attitudes that, at the most basic explanatory level, are belief-like attitudes with a “mind-to-world” direction of fit: they are attitudes whose content should be adjusted to fit the way the world is. In contrast, *expressivists* hold that normative judgments express attitudes that, at the most basic explanatory level, are desire-like attitudes with a “world-to-mind” direction of fit: they are attitudes with content that aims to adjust how the world is. One could adopt either basic position in this debate and still coherently claim that ethical agent-attitude-dependence is a correct normative thesis: all that would change is one’s theory of what mental state one is in when one does so.

Now take the second debate. *Non-naturalists* about normative properties hold that the basic normative property (suppose, for instance, the property of *being a normative reason*) is a metaphysically *sui generis* property fundamentally different in kind from any of the so-called “naturalistic” properties that we are capable of studying using scientific means. In contrast, *naturalists* about normative properties hold that the property of being a normative reason is itself a naturalistic property – either because it can be reductively explained in terms of further naturalistic properties (*reductive naturalism*) or because it is a basic naturalistic property itself (*non-reductive naturalism*).⁴² As with the previous debate about expressivism, one could

⁴² It should be noted that the terms “non-naturalism” and “naturalism” can also be applied to metaethical views about the nature of normative *concepts* in addition to views about normative *properties*. In *Thinking How to Live*, Gibbard argues in favor of non-naturalism about normative concepts, but denies non-naturalism about normative properties (Gibbard 2003). I am here using the terms to refer specifically to views about normative properties.

fall on either side of this basic debate about normative properties and still coherently claim that ethical agent-attitude-dependence correctly identifies the conditions for when the property of being a normative reason obtains. The only thing that would change is one's theory of *what sort* of property thereby obtains. There is thus no reason why it is conceptually incoherent for a non-naturalist such as Parfit to say that in all circumstances an agent A should do what A's attitudes favor – for, in endorsing this, Parfit would simply be endorsing a substantive normative position that covered all reasons (paralleling how Moore embraced a global version of consequentialism as a normative position).

What all of this underscores, then, is this: one can accept the substantive normative thesis of ethical agent-attitude-dependence and still accept a range of views in fundamental metaethics about *what it is* to accept this thesis. Nonetheless, I think that if one accepts the substantive normative thesis of ethical agent-attitude-dependence – and, in particular, if one accepts it on the basis of Street's basic practical/theoretical argument – then this does give certain fundamental metaethical views a relative dialectical advantage in relationship to the competition. More specifically, when combined with certain attractive theses about theory-choice in metaethics, I will argue that it lends important (if defeasible) support to the thesis that I earlier introduced as *metaethical agent-attitude-dependence*. To repeat, this is the following view:

Metaethical Agent-Attitude-Dependence: to make a normative judgment that an agent A should ϕ in C is to express the belief that A's own evaluative attitudes favor A ϕ -ing in C.

In order to make this argument, I will begin with general reasons why the thesis of ethical agent-attitude-dependence undercuts important motivations for specific types of metaethical views. In turn, I will argue that, when combined with these general reasons, Street's practical/theoretical argument lends support to metaethical agent-attitude-dependence; and, in particular, that it lends support to developing it explicitly as a thoroughly naturalistic metaethical view that I will call *agent-value-subjectivism*.

2.1 Straightforward Incompatibilities.

Let's start with one of the most straightforward ways in which one's commitment to a substantive normative position can have significance for theory-choice in metaethics:

namely, by giving one a (defeasible) reason to reject those metaethical views that are inconsistent with that normative position.

Ethical agent-attitude-dependence is straightforwardly incompatible with a number of important metaethical theories. First of all, it is incompatible with any metaethical view that entails that ethical agent-attitude-*independence* must be true if there are any normative reasons. This would include, for instance, any descriptivist views according to which judgments about normative reasons are in part *about* agent-attitude-independent facts. This is important because, as I noted at the start of this paper, many philosophers are drawn to the idea that it is part of the very concept of NORMATIVE REASON that genuine normative reasons must be based on agent-attitude-independent normative facts. Insofar as one holds a substantive normative position to the contrary, this rules out an important range of views in metaethics. Most importantly perhaps, it rules out naturalistic descriptivist views that hold that, given the meaning of normative judgments, the truth-conditions for the statement “A has a normative reason to ϕ ” rest in part on facts about the evaluative attitudes of some other person(s) than the agent herself: for instance, the speaker, a group to which the agent and/or speaker belongs, or, more ambitiously, *all* hypothetical agents.⁴³

Secondly, insofar as ethical agent-attitude-dependence is a substantive claim about normative reasons, it is incompatible with views that claim that *all* substantive normative claims are false. Thus, the endorsement of ethical agent-attitude-dependence requires the denial of *error theory* about normative reasons: the view that there are no correct statements about normative reasons in our world because there are no normative reasons in our world. Instead, it requires the endorsement of a basic *realism* about normative judgment: the thesis that some normative properties are instantiated in the actual world such that some possible normative judgments that concern the actual world are correct.⁴⁴

⁴³ It should be noted that this point would be much harder to press if such views were combined with the sort of *relativist* semantics recently argued for by (Egan 2006) and (MacFarlane 2007). In rough terms, according to these relativist proposals, for two subjects in the same possible world, the very same proposition can be true for one subject and false for the other. It is beyond the scope of this paper to address this issue here.

⁴⁴ This use of the term “realism” captures the minimal commitment shared by those who endorse realism about a range of different self-ascribed realists mean in different domains – say, for instance, realists about beliefs, chairs, or protons. It should be stressed, however, that there are a range of things one might mean by “realism” and a range of issues that debates about “realism” seem to have concerned. On this topic, see (Railton and Rosen 2009) and (Wright 1994). It should be stressed that my use of the term “realism” here is significantly different than Street’s use of the term according to

Philosophers working in contemporary metaethics understand the status of error theory and realism in different ways. For instance, some expressivists see these theses as themselves substantive normative claims while others deny that they are. For our purposes here, however, this is largely beside the point. The point, rather, is this: ethical agent-attitude-dependence is incompatible with all views that espouse a full-blown error theory *however* that espousal is theorized. Furthermore, since at least *some* metaethical views are explicitly put forward as combining a semantic thesis with the further claim that it yields error theory in our world, this means that it is incompatible with certain metaethical views.⁴⁵

We now have two ways in which the endorsement of the position of ethical agent-attitude-dependence itself has metaethical significance: 1) by ruling out views that entail ethical agent-attitude-independence must be true if any ethical view is and 2) by ruling out full-blown versions of error theory in favor of a minimal form of realism.⁴⁶ How important are either of these conclusions, however, for any semantic claims about the *meaning* of ethical judgments about normative reasons as such? In order to explain why I think these

which “realist” views are those that endorse ethical agent-attitude-independence. It should also be noted that I use the phrase “concern the actual world” in my statement of realism rather than “about the actual world” to leave room for Gibbard’s expressivist view according to which normative judgments express plans about what to do in the actual world rather than judgments that are strictly speaking *about* the world. It should also be noted that if realism is true, then correct normative judgments that concern the actual world will be *non-trivially* correct in virtue of the fact that there are instantiated normative properties in the actual world. In contrast, imagine that there are no normative properties instantiated in the actual world. If so, then the statement “Andy does not have a normative reason to eat the pizza” is trivially correct in the actual world.

⁴⁵ It should be noted that ethical agent-attitude-dependence *is* in principle compatible with views that hold that there are no correct possible normative judgments of *certain types* that concern the actual world. In other words, ethical agent-attitude-dependence is compatible with error theories about a more limited domain of purported normative facts. For instance, since it is a claim about what facts about normative reasons need to be like, ethical agent-attitude-dependence is compatible with the idea that no agents in the actual world actually have any normative reasons. My point here is just that insofar as one takes ethical agent-attitude-dependence to itself be a non-trivial *substantive normative claim* – which, at least so far, I have been claiming it is – then the view is incompatible with a full-blown ethical error theory that holds that *all* substantive non-trivial normative claims in ethics are incorrect.

⁴⁶ It should be noted that there is no reason in principle why one should always give more credence to one’s views in normative ethics than to one’s views in metaethics. Thus, one might give – and legitimately give – more credence to other theses in metaethics than to ethical agent-attitude-dependence. In turn, those metaethical theses might either entail that error theory is true if ethical agent-attitude-dependence is true – or, more drastically, that error theory is true in every possible world remotely like our own. If so, then perhaps this will cut against ethical agent-attitude-dependence rather than ethical agent-attitude-dependence cutting against those metaethical views. For now, however, we are simply holding fixed that one is committed to ethical agent-attitude-dependence in order to explore the metaethical ramifications of being so committed.

conclusions do have semantic import, I will now turn to two attractive principles about theory-choice in semantics.

The first principle is this: semantic theories about what a given predicate F means should, in general, avoid attributing massive systematic error to those of us who use predicate F. This is not to say that we should *never* attribute systematic substantive error to speakers. More specifically, it is not to say that we should *never* hold that speakers mean things by their words such that these words turn out not to refer to anything in the actual world (as they do on error theories). The principle states a rather more limited idea – namely, that it is a defeasible *desideratum* that we should avoid committing ourselves to error theory when we can. This view is taken for granted in much work in metaethics – and, moreover, can be seen as a reflection of general views about charitable interpretation that are commonplace among both philosophers and linguists engaged in systematic work in semantics.

The second principle holds that in general, when we advance a view about what speakers mean by their words, we should have some viable empirical story for how those speakers came to mean that by their words rather than something else. Call this the *empirical grounding thesis*. The reason that this principle is plausible is that, in general, the meanings of terms in an everyday language are not the result of acts of stipulative definition – such as saying “by predicate ‘F’ I mean to refer to such-and-such objects” – but rather the result of acquiring constraints over time on the use of that predicate. If Street is right that there are no agent-attitude-independent normative facts in our world, then we have reason to believe that creatures like us *never* came into any sort of contact with any sort of agent-attitude-independent normative facts at any point in our history. If that is true, then how could our words have come to aim at referring to such things? To put it another way: what in our past practices would have led us to constrain our use of our normative predicates in such a way as to *mean* something about purportedly agent-attitude-independent normative facts? In combination with the substantive thesis of ethical agent-attitude-dependence itself, the empirical grounding thesis thus casts at least some doubt on semantic theories in metaethics that take us to mean something by our words that has us aiming to refer to agent-attitude-independent normative facts.

2.2. *Undercutting the Perceived Need to Vindicate Ethical Agent-Attitude-Independence.*

Let's now take stock of where we are. If what I have argued so far is correct, then ethical agent-attitude-dependence counts against 1) metaethical views that entail that ethical agent-attitude-*independence* must be true if any substantive normative view is (as some versions of naturalistic descriptivism do) and 2) metaethical views that yield error theory about all normative judgments in ethics. This however, still leaves open a range of metaethical views. For instance, it still leaves open the possibility that expressivism is correct. This is because expressivism does not itself entail *anything* at the substantive normative level about what makes it the case that agents have normative reasons. Furthermore, expressivism is completely neutral with respect to whether or not error theory is true: for the expressivist, that is a further substantive normative debate. Similarly, my discussion so far still leaves open the possibility that the fundamental normative properties (such as the property of being a normative reason) are themselves non-naturalistic, *sui generis* properties. Consider here the *non-naturalistic realist* position which combines a) this non-naturalist view about normative properties, b) descriptivism about ethical judgment, and c) our earlier "realist" thesis that some normative judgments are correct. Such a view is not an error theory. And, furthermore, without saying more to describe the *sui generis* normative properties, this position, like expressivism, does not itself analytically entail *anything* at the substantive normative level about what makes it the case that agents have normative reasons.

However, these two metaethical views – as well as others – are impacted by another way in which one's endorsement of the position of ethical agent-attitude-dependence has ramifications for metaethics: namely, simply by making certain metaethical views more or less appealing given broader motivations for those types of views in the first place. On this front, perhaps the most important point to have in mind is that many metaethical views are *in part* advocated because their proponents seek to make it possible for themselves to actively endorse *agent-attitude-independence* as a substantive normative thesis. For instance, this is true of many leading forms of *non-naturalistic realism* (such as the views of Derek Parfit, David Enoch, and Russ Shafer-Landau).⁴⁷ Furthermore, it is also true of many leading forms of

⁴⁷ See (Parfit Forthcoming), (Enoch 2007), and (Shafer-Landau 2005). As I use the term in this paper, "non-naturalistic realism" refers to views that combine the thesis of non-naturalism about normative properties, the thesis of descriptivism about normative judgment, and the realist thesis that normative properties are instantiated in the actual world.

naturalistic realism (such as the version advocated by Michael Smith)⁴⁸ as well as leading forms of *fictionalism* (such as the version advocated by Richard Joyce).⁴⁹ And, finally, as Street herself emphasizes, it is also true of many of the leading forms of *expressivism* (such as the views of Simon Blackburn and Allan Gibbard).⁵⁰ Of course, many of these views are motivated by something more basic as well: namely, the aim of showing that agent-attitude-independence is at least an intelligible (if still false) normative view. Yet, given that *one* motivation for these types of views is that the philosophers themselves want to endorse this substantive normative view as not only intelligible but also as *correct*, then if Street's practical/theoretical argument successfully lends support to agent-attitude-dependence – and especially if it does so (as it appears to do) using no explicitly metaethical theses that non-naturalistic realists, naturalistic realists, or expressivists would reject – then this will mean that the argument undermines at least *one* motivation for a range of metaethical views. In particular, it would undermine one motivation for seeking a metaethical view that would make it possible to affirm agent-attitude-independence at the substantive normative level. For, if Street's practical/theoretical argument is correct, that substantive normative position is incorrect.

The flipside of this is that it also undermines one prominent type of criticism of views that entail ethical agent-attitude-dependence – such as metaethical agent-attitude-dependence. To see how this is so, consider that views in metaethics are often evaluated in terms of whether or not they are extensionally correct. For instance, take a simple version of metaethical analytic hedonism – the view that it is analytic that judgment about an agent's normative reasons consists in judgment about which possible actions open to that agent would produce the most aggregate pleasure in the world. One reason that such a view is routinely rejected is that appears to deliver the wrong results about what normative reasons

⁴⁸ (Smith 1994).

⁴⁹ (Joyce 2001).

⁵⁰ (Blackburn 1998) and (Gibbard 2003). It can also be noted that, even though Blackburn and Gibbard do not take their views to analytically rule out any substantive normative views, both of them think their views do lend relative dialectical support to certain substantive normative views rather than others. A central part of Gibbard's *Reconciling our Aims*, for instance, is to show how accepting his expressivist metaethical account helps support a particular type of consequentialist theory in normative ethics (Gibbard 2008). In a similar vein, Blackburn has claimed that his metaethical view puts pressure on one's ethical commitments that seem to conflict with one's goals or happiness (Blackburn 1993, 164). Thanks to Tristram McPherson for emphasizing this point to me.

an agent has in a great number of cases. Furthermore, it seems to have a relative disadvantage here with respect to other views of a similar type: for instance, it seems to deliver the wrong results in *more* cases than a version of analytic utilitarianism that took a more sophisticated view about the good. There are, of course, other reasons that philosophers reject such views as well. For instance, expressivists such as Gibbard and Blackburn maintain that all metaethical views such as analytic hedonism that analytically rule out certain substantive normative views fail to capture the core of “what’s at issue” in competing conceptually coherent (if wildly false) normative judgments – and thus universally fail as fundamental metaethical accounts.⁵¹ Nonetheless, if one were able to show on metaethically-neutral grounds show that analytic hedonism delivered roughly what we judge to be the correct extension, then this would be a response to at least one common objection to this sort of view.

If this is so, now consider the situation with views in metaethics that entail ethical agent-attitude-dependence – or, more accurately, entail it in a world where error theory is false. Such views are often dismissed out of hand for delivering the wrong normative results. Yet, if Street’s practical/theoretical argument leads to the conclusion of ethical agent-attitude-dependence – and especially if her argument shows that agent-attitude-dependence is correct on metaethically-neutral grounds – then her argument thereby provides an important counter to this criticism of certain metaethical views. Most importantly, in providing such a counter, it thereby provides some support for those metaethical views that at least *prima facie* look to be extensionally off because they rule out ethical agent-attitude-independence as incoherent. Thus, among other metaethical views, it lends support to views that endorse the claim that to make a normative judgment that an agent A should ϕ in C is to judge that A’s own evaluative attitudes favor A ϕ -ing in C (the thesis that I am calling *metaethical agent-attitude-dependence*). This does not mean that her argument by itself shows that one should endorse such a metaethical version of agent-attitude-dependence instead of just endorsing the substantive normative thesis of ethical agent-attitude-dependence. However, in providing an important counter to a type of criticism of metaethical agent-attitude-dependence, it does lend some support to this type of metaethical view.

⁵¹ (Gibbard 2003) and (Blackburn 1998).

2.3. Extending the Practical/Theoretical Argument to Non-naturalistic Realism.

So far, in considering the metaethical significance of Street's practical/theoretical argument in favor of ethical agent-attitude-dependence, I have mainly discussed how the *conclusion* of ethical agent-attitude-dependence has ramifications for theory-choice in metaethics. In contrast, I now want to consider how Street's practical/theoretical *argument* for this position also has important metaethical ramifications – ramifications of a sort that, it should be emphasized, Street herself sometimes takes her argument to have as well.⁵² In particular, I will argue that if Street's argument is right, then it also establishes a problem for non-naturalistic realism *even though* this position is theoretically compatible with the substantive normative position of ethical agent-attitude-dependence.

According to non-naturalistic realism, a) there are irreducibly normative properties that are metaphysically *sui generis*, b) such properties are instantiated in our world, and c) normative judgment expresses attitudes that, at the most fundamental explanatory level, consist in straightforward truth-apt belief about these properties. Suppose one accepted the non-naturalist's idea that there is a metaphysically *sui generis* property of *being normative* above and beyond the properties that make it the case that this normative property obtains. Adopting what Street calls the "theoretical standpoint," we can now ask how we would have evolved to recognize – and, moreover, care about – this purported extra property. If Street is right that we can explain the content of the evaluative attitudes we have – and the systems of normative judgment that they underwrite – without making reference to such extra properties, then it seems hard to see how we would have evolved any capacity to detect such extra properties – let alone why we would have evolved in such a way as to place any practical *interest* in those extra properties whatsoever. Street's work casts doubt that there is a story to be told about how we evolved to detect such extra properties – and thus casts *at least some* doubt on positing the existence of such non-naturalistic normative properties in the first place. Furthermore, beyond this point about ontology, if we combine it with the empirical grounding thesis, this gives us reason to hold that our normative terms never aimed to refer to such properties in the first place. Thus, when combined with a plausible principle about methodology in semantics, it seems that Street's practical/theoretical argument has a

⁵² This is reflected, for instance, in her original presentation of the practical/theoretical argument in (Street 2006).

considerable impact on theory-choice in metaethics: it establishes a serious problem for non-naturalistic realism.⁵³

2.4. *Agent-Value Subjectivism.*

In this section so far, I have argued that if one endorses ethical agent-attitude-dependence – and, specifically, if one does so on the basis of Street’s practical/theoretical argument for it – then this endorsement has significance for theory-choice in metaethics. On the one hand, I have argued that it makes the following broad types of metaethical view less attractive: views that entail that, assuming error theory is false, ethical agent-attitude-independence must be true (which includes many versions of naturalistic descriptivism); any view that results in error theory for whatever reason; expressivism; fictionalism; and non-naturalistic realism. On the other hand, by undercutting the force of one significant type of criticism of a metaethical view, I have argued that it lends direct dialectical support to views that entail that ethical agent-attitude-*dependence* must be true if there are any facts about normative reasons. Finally, through the course of this argument, I have put forward some broad principles for theory-choice in semantics and metaethics – including the *empirical grounding thesis* – that, if true, should have significance for which type of theory we choose in metaethics more generally.

This line of reasoning lends dialectical support to metaethical agent-attitude-dependence. For, if the arguments about which types of metaethical view are negatively affected are correct, then two main types of metaethical view have not been negatively impacted at all: 1) versions of naturalistic realism that *do not* entail either ethical agent-attitude-independence or -dependence and 2) versions of naturalistic realism that *do* entail ethical agent-attitude-dependence. In turn, if what I have argued is correct, then versions of naturalistic realism that entail ethical agent-attitude-dependence also receive relative dialectical *support* insofar as one type of common criticism of the view would turn out to be false if Street’s argument is right (namely, the criticism that such a view is extensionally off). Thus, the type of metaethical view that I think emerges with the biggest dialectical advantage in light of Street’s argument for agent-attitude-dependence is one that holds a) descriptivism,

⁵³ Enoch agrees that Street’s practical/theoretical argument can lead to this problem for non-naturalistic realism via roughly the argumentative route I suggested. For his explanation of why – and his preferred response in defense of non-naturalistic realism – see (Enoch Forthcoming).

b) naturalism about normative properties, c) that those properties concern what an agent's evaluative attitudes favor in a given set of circumstances, and d) realism (in the sense of holding that there are some such normative properties in our world). In other words, it favors metaethical agent-attitude-dependence when it is understood as a version of *naturalistic realism*.

It is worth noting here that we can also see how such a view receives relative dialectical support from the following, further reason having to do with the empirical grounding thesis: namely, it explains how the term “normative reason” in natural language – or at least one use of the term in natural language – could refer to things for which there is a compelling empirical story about why we have cared about these things independently of our use of the term “normative reason.”⁵⁴ More generally, it would make sense of why agents would talk with each other about what normative reasons they have given agents' own practical projects of figuring out what to do. Moreover, in turn, the account would then have the resources to explain why if a speaker's statements of the form “A has a normative reason to ϕ in C” were made in many cases where the speaker knew that A's own evaluative attitudes did not favor A ϕ -ing in C, we might legitimately think that a speaker is not making a substantive normative mistake, but would in fact not be understanding the import and meaning of statements of this form.

Many important details about metaethical agent-attitude-dependence would need to be fleshed out in order to evaluate carefully how it compares to the competition in the context of overall theory-choice in metaethics – a context in which the fact that metaethical agent-attitude-dependence receives relative dialectical support from Street's practical/theoretical argument for ethical agent-attitude-dependence could obviously be swamped by more massive dialectical *disadvantages* this metaethical view has. Perhaps most

⁵⁴ I say “or at least one use of the term” because, depending on how one things terms should be individuated in a natural language, one might think that some statements where we say “an agent has a normative reason to ϕ ” get a different semantics than others despite the statements involving the very same term. I think this is a crucial issue for understanding the scope of the basic semantic thesis that I am putting forward here – that is, for understanding how many parts of natural language it actually applies to. There are different options that one might take here. One could, for instance, hold that my semantic thesis applies to a large portion of statements about “normative reasons” in a natural language but not *all* of them – just as, for instance, many theories of the semantics of epistemic modals involve combining different types of semantic accounts in order to explain the entirety of talk about epistemic modals in a natural language. See, for instance, (Egan 2006). I think that this sort of issue deserves more attention in semantic accounts of normative statements. Unfortunately, however, it is beyond the scope of this paper to talk about this issue in any depth.

importantly, there are two details of this view that would be important in the context of overall theory-choice in metaethics that have not been clarified: 1) in what sense the view is reductive or not about either normative properties or concepts and 2) which evaluative attitudes the view is talking about. In order to better explain what I think is the best way of developing metaethical agent-attitude-dependence, let me say a bit about what I think we should and should not commit to at this stage with respect to each of these issues.

In terms of the first issue, this will largely depend on what one thinks suffices for a reductive account. To see why this is so, start with the issue of normative properties. On the one hand, we have established grounds for being skeptical of non-naturalistic realist views that posit a *sui generis* property of *being normative* that, from the theoretical standpoint, is superfluous for accounting for the content of our evaluative attitudes. This skepticism should carry over to *any* views that hold that there is a brute *sui generis* property of simply *being normative*. It is possible that a non-reductive naturalistic realist *might* hold this view too but simply differ from her non-naturalistic counterparts in terms of what type of property this brute property is (they could think it is a naturalistic *sui generis* one rather than of a *sui generis* non-naturalistic kind). I believe that such a view could not be made coherent given how to best understand the meaning of “naturalism” and “non-naturalism” in metaethics: a context in which, I think, to hold that the fundamental normative property is metaphysically *sui generis* suffices to make that view a non-naturalistic one of normative properties. The point here, however, is just that *if* such a view were to be made coherent, then skepticism about a brute property of *being normative* would affect this version of non-reductive naturalistic realism as well as non-naturalistic accounts.

On the other hand, some versions of non-reductive naturalistic realism are put forward as *non-reductive* for reasons other than thinking that there is simply a brute *sui generis* property of just being normative. For instance, Richard Boyd and Nicolas Sturgeon have put forward their view in this way largely because of general views they have about reduction. In general terms, following the basic lead of Ernst Nagel, they think that, in order to reduce a claim in the special sciences to a more basic scientific claim (e.g. one in physics), one would have to specify explicit bridge laws that map the laws stated in the vocabulary of the special sciences to laws stated in the vocabulary of physics.⁵⁵ They then claim that, even though

⁵⁵ See (Boyd 1997) and (Sturgeon 2006).

ethical properties are purely naturalistic properties, one cannot supply these sort of bridge laws in the case of ethics, and thus reduction fails.

Or, to take another reason that one might have for claiming that a naturalistic view is a non-reductive one about normative properties, consider that views in philosophy of mind that hold that mental properties are defined functionally – and hence could be realized by a range of physical and/or non-physical systems – are often understood to be versions of *non-reductive physicalism* about the mind.⁵⁶ Thus, if one were attracted to this way of thinking about such views in the philosophy of mind, one might want to claim that a view counts as a form of *non-reductive naturalistic realism* in metaethics if it claims that the normative properties (perhaps because they could be defined functionally) could be realized by a range of naturalistic and/or non-naturalistic systems. Since I have given a functionalist characterization of what evaluative attitudes are in this paper, then perhaps metaethical agent-attitude-dependence should be thought of as non-reductive position about normative properties.

However, this conclusion would, I think, be premature. Even this brief discussion brings out that there are a number of different criteria philosophers use for deciding when a view about a property gets to count as reductive or not. Some philosophers, for instance, think that when a view is a functionalist one in philosophy of mind, this warrants calling it “non-reductive.” Other philosophers think it certainly does not – and, indeed, some doubt the conceptual coherence of the category of “non-reductive physicalism” in general.⁵⁷ This reflects not only substantive disagreement in philosophy of mind, but broader disagreement about how to understand what should count as “reduction” in this sort of metaphysical context more generally. Some philosophers (such as Boyd and Sturgeon) have very strict ideas of what would count as a reductive account of a property such that relatively few naturalistic properties can be reduced to the basic micro-physical properties. In contrast, other philosophers (such as Jackson and Chalmers) have a different account of reduction such that either *all* or *almost all* of the naturalistic properties can be reductively explained in terms of the basic micro-physical properties.⁵⁸

⁵⁶ See (Bennett 2008) for a taxonomy of views in philosophy in mind that reflects this way of understanding non-reductive physicalism in the philosophy of mind.

⁵⁷ See (Braddon-Mitchell and Jackson 1996) and (Jackson 1996).

⁵⁸ See (Jackson 1998) and (Chalmers 1996). Chalmers holds that the most important exception – perhaps the only exception – is qualitative mental properties.

There is a similar range of disagreement about what types of conceptual reduction are necessary for reduction at the property level – and, furthermore, about what conceptual reduction is in the first place. For instance, Block and Stalnaker hold that conceptual reduction is not necessary for claims of metaphysical reduction.⁵⁹ In contrast, Jackson and Chalmers hold that it is.⁶⁰ Furthermore, there is also the debate about what it would take for there to be a conceptual reduction in the first place – a debate where, again, there is a wide range of views.⁶¹ What all of this brings out, I think, is that whether or not a naturalistic realist view in metaethics gets to count as a so-called “reductive” one about either normative properties or concepts will depend to a great degree one’s broader views about how to understand reduction in these contexts. Furthermore, given the range of ways the term “reductive” get used in the literature, it is not even clear that the distinction between reductive and non-reductive naturalistic metaethical views tracks a particularly important distinction for the purposes of doing metaethics. Such issues are simply outside the scope of this paper to address.

That being said, it should also be stressed that they do not need to be addressed for understanding the basic metaethical view that I am putting forward here. According to the view that I am advancing, the property of being a normative reason for an agent A to ϕ in C does *not* consist in any further brute normative properties, but rather is a naturalistic property that we can identify simply by identifying properties concerning an agent’s evaluative attitudes and the non-normative facts about an agent’s circumstances. If psychological facts about the content and type of an agent’s evaluative attitudes can be understood as non-normative naturalistic facts – something that I am assuming is true – I take the metaethical view on the table to deliver much of what people often want in a “reductive” version of naturalistic realism even if the view ends up being officially classified as “non-reductive” in virtue of broader considerations about metaethical taxonomy and the nature of reduction.⁶²

Let’s now consider the second issue that I raised: namely, which evaluative attitudes matter in this account? There are two main questions involved in this issue. The first stems

⁵⁹ (Block and Stalnaker 1999).

⁶⁰ (Chalmers and Jackson 2001).

⁶¹ Contrast, for instance, the views on this topic in (Jackson 1998), (Railton 2003b), and (Gibbard 2003).

⁶² It should be stressed that this assumption about mental content is by no means uncontroversial. For arguments that facts about mental content are in part constituted by normative facts, see (Wedgwood 2009), (Davidson 2001), and (McDowell 1994).

from the fact that, given how I have defined the term “evaluative attitude” in this paper, there are a range of different *types* of evaluative attitudes that any given agent has. I earlier claimed that an attitude counts as an “evaluative attitude” if it tends to produce motivation for an agent when coupled with her beliefs. Depending on one’s broader commitments in the philosophy of mind, philosophy of action, and empirical psychology, one might take any of the following to be evaluative attitudes in the functional sense that I have defined it: desires, plans, emotions, accepted norms, the state of making a normative judgment, the state of making a “proto-normative” judgment, and unreflective evaluative tendencies. One question, then, is this: when metaethical agent-attitude-dependence says our judgments about an agent’s normative reasons consist in judgment about what an agent’s evaluative attitudes favor, which *type* of evaluative attitude matters? A second question stems from the fact that for any given type of evaluative attitude, an agent will almost always have many token instances of that type. For instance, suppose Dustin desires to drink a beer. If Dustin remotely resembles anything like a normal human being, this will not be his lone desire. Rather, he will have a whole range of desires about a whole range of topics. For instance, at the same time that he has a desire for beer, he might also have a desire to travel to Tahiti, a desire to play basketball, a desire to be able to drive home later this evening, a desire to do more to promote the common good, a desire to make his family happier, and desires about which of these desires will be most important in causing him to act one way or another. A second question, then, is this: when metaethical agent-attitude-dependence says our judgments about an agent’s normative reasons consist in belief about what an agent’s evaluative attitudes favor, which *particular* instances of a given type of evaluative attitude matter?

Let’s start with the first question about evaluative attitude *type*. How we answer this first question will, I think, be one of the crucial questions for developing metaethical agent-attitude-dependence. It is also, however, a question that is simply too big for this paper to address in any detail given how much it intersects with broader questions in philosophy of mind, philosophy of action, and psychology. That being said, I think that there a couple of broad points we can make about what type of evaluative attitude matters.

On the one hand, as I noted, if all that we want is an attitude that tends to produce motivation for an agent when coupled with her prosaic beliefs, this leaves open a range of possibilities. For instance, if one is a strict Humean about motivation, one might think that

the only attitudes that will do the trick here will be ones that are *desires* in the functional sense of being an attitude with a solely world-to-mind direction of fit.⁶³ However, one might also hold an anti-Humean view in the philosophy of action according to which the state of making a normative judgment itself motivates an agent to ϕ even in the absence of any desire that agent has to ϕ .⁶⁴ If that is right, other evaluative attitudes – such as *the state of making a normative judgment* as such – might in theory be used in the account.

On the other hand, if we want an attitude that can be the basis for a naturalistic realist account in metaethics, this adds some additional constraints. I will here mention four such constraints.

To start with, in order to succeed as a fundamental metaethical account at all, we need an account of what *the state of making a normative judgment as such* consists in. Thus, one cannot take it that we already have an account of what that state *is*. For that is precisely what is in question. This creates a problem for using certain evaluative attitudes that one might be inclined to use in developing metaethical agent-attitude-dependence. Suppose, for instance, that one had the anti-Humean view in the philosophy of action that judgments themselves motivate an agent to ϕ even in the absence of any desire that agent has to ϕ and, in turn, wanted to use the state of making a normative judgment as the evaluative attitude in developing metaethical agent-attitude-dependence. If one claims that the content of normative judgment is *further normative judgments*, using this as the evaluative attitude will either be a non-starter for the task at hand or else an ultimately circular account. For it would either leave it completely unspecified what normative judgments themselves are about or else circularly define the content of normative judgment as further normative judgments. Thus, in the context of developing a fundamental metaethical view, we will want some evaluative attitude other than “the state one is in when one makes a normative judgment.” For, in a sense, the whole debate is about what attitude (a mental state) one expresses when one *makes a normative judgment* (an act).

A second constraint on which type of evaluative attitude we should use concerns the connection between normative judgment and motivation. On the metaethical agent-attitude-dependent account I have put forward, there would be no reason to think that second- and third-personal statements by a speaker of another agent’s normative reasons would have

⁶³ (Smith 1994).

⁶⁴ See (Scanlon 2000) and (Dancy 2003).

motivational upshot for the speaker herself. It thus denies that *all* normative judgments have upshot for a speaker's motivation. This means that, given the functional definition of "evaluative attitude" that we are working with in this paper, the *state of making a normative judgment* as such would not even be an evaluative attitude.

Thirdly, insofar as we want a fundamental metaethical view to be a specifically *naturalistic realist* one, as we do if we are going to defend metaethical agent-attitude-dependence, then this adds two additional constraints about which type of evaluative attitude to use: we will want the evaluative attitude to be one that is both thoroughly naturalistic in both type and content *and* that we actually have as humans.

Finally, given the types of motivation we had for developing metaethical agent-attitude-dependence in the first place, there are also reasons to prefer using some types of evaluative attitude as opposed to others for developing metaethical agent-attitude-dependence. Consider again our initial case with Henry the aspiring photographer. The example of Henry helps underscore that when we are concerned with what an agent has normative reason to do, we are in the first instance concerned with evaluative attitudes such as the one that Henry has that favors him getting the dream job at *National Geographic* and less so with things such as mere urges or attitudes that we would label as "desires" in a prosaic sense of the term. Although some might disagree, it intuitively sounds odd to describe the evaluative attitude that Henry has in favor of becoming a photographer as a prosaic everyday "desire" he has alongside such things as a desire for beer. This type of evaluative attitude will be like prosaic desires in that token instances of it won't necessarily involve the application of specifically normative or evaluative concepts. Furthermore, the attitude will likely count as a "desire" in the broader functional sense used in analytic functionalism in the philosophy of mind and the Humean theory of motivation. Indeed, insofar as it involves at *least* a world-to-mind direction of fit, the attitude here will certainly be *desire*-like holding fixed this broader functionalist characterization of desire. But, importantly, it won't do so in virtue of being something we would label as a "desire" in a prosaic sense of the term.

Furthermore, there also seem to other important *phenomenological* differences here as well between the type of evaluative attitude we can imagine Henry having in favor of getting his dream job and an ordinary desire he has for something such as the beer in his fridge. We are all familiar with *what it is like* to experience something as "calling for" a response on our

behalf that goes beyond the basic pull of a prosaic desire or urge. The experience of taking something to have normative or evaluative significance for one's own future course of actions – and hence of certain things “calling for” certain responses – seems importantly different from the phenomenology of many other attitudes with a world-to-mind direction of fit.

With this in mind, I think it is helpful to use a term in defining metaethical agent-attitude-dependence that helps mark such facts and thus better identifies the desire-like attitudes in question. In order to do so, I suggest we use the place-holder term *valuing* for the evaluative attitude in our account.⁶⁵ This term is specific enough to mark off the evaluative attitudes in question from prosaic urges and desires. It also suggests that there might be additional constraints on consistency and coherency that would need to be met in order to attribute this attitude to an agent in the first place. At the same time, the term is broad enough not to involve any detailed commitments about how this attitude is realized in the empirical psychology of human beings or any detailed commitments to other, more specific evaluative attitudes that one might want to use to develop the theory further (for instance, by using something such as Allan Gibbard's proposed attitude of “accepting a norm”).⁶⁶ Since I purposefully want to leave those issues open for now – and instead focus on the basic schematic form of the view – the term *valuing* suits our purposes well. Such issues would, of course, need to be taken up in more detail in further developing the view. However, it should also be stressed that, even in further developing the metaethical view, there is an important limit to how much more we should say about exactly what this evaluative attitude consists in. For, insofar as it is coherent to think that other possible (or actual) non-human creatures – including ones with very different cognitive architectures that we know little about – can also have normative reasons, we should want to specify this evaluative attitude at a high enough level of abstraction in order not to rule out the coherence of such thoughts.⁶⁷

With this attitude of so-called “valuing” in hand, we can now turn to the second important unanswered question about metaethical agent-attitude-dependence that I introduced earlier: namely, which *specific* valuing of an agent matter for settling her

⁶⁵ Street also favors using this term as a place-holder in this same basic way in (Street Manuscript).

⁶⁶ For discussion of this proposed attitude of *accepting a norm*, see (Gibbard 1990), (Darwall 2006), and (Nye 2009).

⁶⁷ For a similar line of thought about the conditions of belief-attribution, see (Stalnaker 1987).

normative reasons? To see how we should answer this, consider that even if an agent such as Henry has a range of valuing about a range of topics, he does not have a mere heap of valuing. Rather, we think of some of his valuing as more important or *fundamental* to Henry's structure of valuation. For instance, Henry might value not opening envelopes in general because he dislikes the way it feels. Yet, importantly, he will also not value that as much as he values finding a fulfilling career. I think it is fairly clear which of these we think matters for Henry's normative reasons: namely, his more fundamental valuing of the career. Furthermore, if we already embrace ethical agent-attitude-dependence in light of Street's practical/theoretical argument, then I think we have a good explanation of why this valuing is more important. The explanation is *not* because we have made the substantive normative judgment that it is more important because that specific valuing of Henry's lined up with an agent-attitude-independent normative fact. Rather, it is simply because it is the one that is more crucial in Henry's own system of valuation – it is the one that Henry himself assigns (perhaps not consciously) a stronger weight to in terms of his values about which valuing should actually make a motivational impact on his life. In other words, to put it in broad terms, it is because it is the one that plays a more fundamental role in guiding Henry's actual valuations.

In order to use this conclusion in a way that helps us nail down an account of which specific instances of an agent's valuing matter for determining her normative reasons, we can now imagine all of Henry's valuing as a whole. Without specifying some weights that we should give to these different valuing, it will be impossible to tell what this whole group of valuing favor – for different token evaluative attitudes will favor different things. How then should we proceed? If we are going to stick with a *strictly* agent-attitude-dependent account, the answer is clear: we should use the weights that Henry himself (consciously or not) assigns a higher weight to in terms of what he values to actually make an on how he lives. With this in mind, we can then imagine what Henry's overall system of valuing would look like *if* it were perfectly regulated and revised according to the weights given in his own system of valuing. Call this set of valuing Henry's *internally-idealized-valuing*. In turn, even though that system will not favor one determinate thing or another in *all* circumstances, it will do so in many more circumstances than just a mere heap of valuing. And, furthermore, the things it will favor will be precisely those that are the expression of Henry's own deepest valuing. Thus, I suggest, if we are to hold that metaethical agent-attitude-dependence is

concerned with what an agent's valuing favor in a set of circumstances C, we should hold more specifically that it is concerned with what is favored by the whole set of an agent's *internally-idealized-valuing*s favor in C.

We now have some initial conclusions about which type of evaluative attitudes matter for the account, as well as which subset of these valuing in particular matter. When these conclusions are combined with our earlier discussions, then, even if we do not yet have arguments for the full details of how to best develop metaethical agent-attitude-dependence – arguments that, I have emphasized, are simply beyond the scope of this paper to give – we nonetheless have some important points of clarification that are in part motivated by the same underlying ideas that led us to metaethical agent-attitude-dependence in the first place. Putting our different conclusions so far together, we now have the following naturalistic realist view in metaethics:

Agent-Value-Subjectivism: to make the normative judgment that an agent A should ϕ in circumstances C is to express the belief that A's own internally-idealized-valuing favor ϕ -ing in C.

This view is more specific than the original thesis of “metaethical agent-attitude-dependence” that we started with. First of all, it is put forward specifically as a naturalistic realist view (one that might or might not count as “reductive” depending on how one understands the label). Second, it specifies what broad type of evaluative attitudes matter – namely, an agent's *valuing*s. Third, it specifies which specific valuing matter – namely, all of those in the set of A's own internally-idealized-valuing. While this still leaves much vague, it is, I think, the start to further developing metaethical agent-attitude-dependence. And, importantly, it is not always how people have understood that such a type of metaethical view should be developed. Indeed, in the next section, I will consider another way of developing metaethical agent-attitude-dependence – namely, Street's own preferred way – that appears to be distinct from the sort of naturalistic realist metaethical position that I am putting forward here.

However, before I do so, I want to conclude this section with two important clarifications about agent-value-subjectivism. The first concerns the status of its claim about the meaning of normative judgments. The second concerns two important ways in which normative facts are “independent” on agent-value-subjectivism.

Let's start with the first issue. When philosophers in metaethics make claims about the meaning of normative terms, they understand the status of these claims in a variety of different ways. For instance, some take the meaning to be the result of strict *analysis* of the pre-theoretic concepts we started with. Others take it to somehow be a sharpening or refinement of pre-theoretic concepts: such as on a philosophical "explication" rather than analysis or on a so-called "reforming" definition. All of these are open options for advancing agent-value-subjectivism. Which option one takes will, I think, rest on one's views about such things as the nature of conceptual analysis in general, the nature of the analytic/synthetic distinction, the nature of the *a priori*, and issues of concept individuation. In turn, these issues will help determine an important question for how one understands the basic thesis of metaethical agent-attitude-dependence: namely, whether or not it is read as presenting a strict biconditional or if it should only be read "left to right" rather than "right to left." Such issues are beyond the scope of this paper. The important point is that nothing in how I have laid out the basic schema for agent-value-subjectivism necessitates going one way or another on any of these further questions. Rather, these are further questions that matter for determining how one would develop a more specific version of the view.

Let's now turn to the second issue of clarification – namely, two important ways in which normative facts are "independent" on agent-value-subjectivism. According to agent-value-subjectivism, an agent's normative reasons are fundamentally set by what valuing she has. To some, this might suggest a voluntarist understanding of our relationship to the basic facts about normative reasons: one in which *the agent* herself actively determine or choose those facts. Agent-value-subjectivism is not such a view. For even though agent-value-subjectivism does hold that normative facts are determined by facts of an agent's psychology, this does not mean that *we ourselves* determine those facts in the sense of being able to immediately change those facts at will through consciously made decisions. Facts about the content of an agent's valuing in her actual internally-idealized-valuing are not determined by what that agent wills to be the case. Indeed, as Street's work emphasizes, the facts about the content of an agent's own evaluative attitudes (such as her valuing) are due to deeply contingent historical facts that the agent herself did not choose or create in any meaningful sense. Recognizing this, Street holds, is part of what it is involved in "coming to

terms” with the basic contingency involved in ethics.⁶⁸ The content of an agent’s valuing are something that, even if they can be *modified* by her, are *not* something that she gets to choose or “will” *all the way down*. For rather than getting to *choose* what our fundamental valuing are – ones that we might choose to have without reference to any pre-existing valuing that we already have – we rather always already find ourselves having *some* valuing and having to work from there. These are facts *of* an agent’s psychology in a way that many metaethical theories would deny. But they are not – and this is the crucial point to emphasize – facts that are in some sense created by the conscious *decisions* of the agent herself or by the *beliefs* that an agent happens to have about those facts.⁶⁹

This, then, is also part of what explains the deep possibility of error in normative reasoning and the sense in which there are normative facts we are trying to *discover*. For, at the end of the day, facts about the content of an agent’s valuing are naturalistic facts that hold independently of *anyone’s* current beliefs, prosaic desires, and judgments about those facts. They are naturalistic facts just like any other psychological facts – ones capable in principle of empirical study from the third-personal “theoretical” standpoint and just as “objective” as any other psychological facts.⁷⁰

A second important point about the “independent” nature of normative facts according to agent-value-subjectivism concerns the status of *ethical* agent-attitude-

⁶⁸ (Street Manuscript).

⁶⁹ Agent-value-subjectivism thus differs significantly from many views that end up endorsing ethical agent-attitude-dependence. For instance, contrast what I have just said against voluntarism with the view that Gibbard attributes to Korsgaard: the view that that he calls “logically constrained reflective subjectivism.” Gibbard writes that “according to this logically constrained reflective subjectivism, reflectively deciding to do something *makes* it the thing to do, so long as your policies for action are logically consistent (Gibbard 1999, 148).”

⁷⁰ It should be stressed that there are a variety of different ways the term “objective” gets used to describe views in ethics. My point here is simply that if one thinks that there is a sense in which there are objective facts about psychology, then there is no reason to deny that normative facts are objective on my view. Furthermore, since I see no reason to doubt that there is at least one good sense of the term “objective” on which there are objective facts about what an agent desires in circumstances C even if another agent has different desires in C, I see no reason to doubt that there are objective facts about what an agent’s normative reasons are even if my view leaves it open that different agents could have different normative reasons in the same circumstances. Of course, there are different uses of the term “objective” – for instance, Kantian uses – on which that would amount to a denial of practical objectivity. However, this is not the It is perhaps useful to think of the normative facts – as understood by agent-value-subjectivism – in terms of the phrase “subjectual” introduced by Railton in (Railton 2003a). In basic terms, Railton sees subjectual accounts of objectivity as ones that locate “objectivity in rules or conditions for subjects rather than a relation to external metaphysics (Railton 2003a, 104).” When read in a broad way, this could be seen as covering the sort of view that I am advocating here.

dependence itself. Is the normative judgment that ethical agent-attitude-dependence is true an agent-attitude-dependent or agent-attitude-independent fact? Or, to put it another way, is the reason that the normative judgment that “ethical-agent-attitude-dependence is correct” is itself correct due to an agent’s own attitudes? As with straightforward relativist claims about truth, it is hard to see how the thesis, if it is to be informative and coherent, could apply to itself without ever grounding out in something external to evaluative attitudes.⁷¹ Indeed, even if all of us should in fact be committed to the thesis of ethical agent-attitude-dependence in practical reasoning holding fixed the current evaluative attitudes we have – which, it should be emphasized, we should be if Street’s practical/theoretical argument is right – this is still different from saying that our current evaluative attitudes make this thesis true.

If one accepts metaethical agent-attitude-dependence, the way around this issue is straightforward: at the most basic explanatory level, what makes ethical agent-attitude-dependence true is that it is conceptually entailed by metaethical agent-attitude-dependence. In other words, the thesis of metaethical agent-attitude-dependence gives us a concept of a normative reason – which can be seen in theory as either a straight-ahead conceptual analysis or conceptual explication – which in turn sets constraints on what the normative facts about normative reasons must be like if there are any such facts. The concept *NORMATIVE REASON* alone does *not* give one an ability to derive any substantive existentially committed view – such as a view concerning facts about the instantiation of normative properties (i.e. the “normative facts”). But it does give one the ability to say something about what the facts about normative reasons are like *if* there are any such facts. Importantly, the fact that the conceptual facts entail this is not itself an attitude-dependent fact. Rather, it is simply an upshot of the concept *NORMATIVE REASON* we employ in ethical thought. The upshot of this is that, in the end, if one ends up accepting metaethical agent-attitude-dependence, then ethical agent-attitude-dependence turns out not to be a substantive normative thesis at all. Rather it is simply a statement about what the normative facts (i.e. the non-conceptual facts about the instantiation of normative properties) must be like if there are any normative facts about normative reasons – a statement, importantly, that follows solely from conceptual truths alone. In this way, affirming ethical agent-attitude-dependence can be squared with

⁷¹ For discussion of this point about relativist theories of truth, see (Boghossian 2006).

the idea that there are no substantive, non-conceptual agent-attitude-dependent normative facts.⁷²

Furthermore, my basic response here gives us a template for a response that one might give on behalf of ethical agent-attitude-dependence if one denied metaethical agent-attitude-dependence. Suppose one denies metaethical agent-attitude-dependence and, moreover, one denies that ethical agent-attitude-dependence is conceptually entailed by one's metaethical view. In such a case, ethical agent-attitude-dependence's claim that "there are no agent-attitude-independent normative facts" could still be restricted to not apply to the normative fact of ethical agent-attitude-dependence *itself*. For instance, ethical agent-attitude-dependence could be understood to be a substantive normative fact about normative facts that *concern a further normative topic*. For instance, one could specify that ethical agent-attitude-dependence – a thesis concerns that concerns correctness-conditions for normative judgments about agents' normative reasons – does not claim that there are no agent-attitude-independent normative facts at all *of any sort* but rather that there are no agent-attitude-dependent normative facts that figure into those correctness-conditions once those conditions have been specified by the thesis of ethical agent-attitude-dependence. The resulting view would, of course, face problems of a sort that metaethical agent-attitude-dependence does not face. For instance, it would leave it mysterious why this single normative fact of ethical agent-attitude-dependence obtains in the first place – a worry that could in turn make the view susceptible to Street's practical/theoretical argument when run to cover a sufficiently large enough domain of normative facts (even if it excluded normative facts about how to assign meanings to speakers' utterance, etc.). If such problems turn out to undermine the appeal of that sort of response, then so much the better for my preferred position of metaethical agent-attitude-dependence. My point here, though, is simply that even if metaethical agent-attitude-dependence is false, ethical agent-attitude-dependence could still be an informative and coherent substantive normative position.

⁷² This response differs from the one that Street gives to this same basic question about the status of ethical agent-attitude-dependence. In (Street Forthcoming-c, 36-37) she argues that the thesis of ethical agent-attitude-dependence *does* apply to itself but is nonetheless not self-defeating because, given the strength of the practical/theoretical argument, the truth of ethical agent-attitude-dependence is favored by every set of evaluative attitudes an agent might have (or at least in worlds where those attitudes have been shaped by causal forces). I find this reply unconvincing as it stands. Even if one grants that the theory is favored from every set of evaluative attitudes (either in the actual world or nearby possible worlds), this still does not settle whether the truth of ethical agent-attitude-dependence is somehow agent-attitude-dependent or not.

3. *The Relationship between Agent-Value-Subjectivism and Metaethical Constructivism.*

Let's now take stock of where we are. In the last section of this paper, I argued that *if* Street's practical/theoretical argument supports *ethical* agent-attitude-dependence, then it also has significance for theory-choice in metaethics. More specifically, I argued that it lends relative dialectical support to *metaethical* agent-attitude-dependence. In turn, I argued that metaethical agent-attitude-dependence is best developed as a naturalistic realist view in metaethics that holds the following:

Agent-Value-Subjectivism: to make the normative judgment that an agent A should ϕ in circumstances C is to express the belief that A's own internally-idealized-valuings favor ϕ -ing in C.

In her own work, Street has advanced a view that she calls *metaethical constructivism*. When this view is read as a fundamental metaethical view in the sense of "metaethics" that I have been using it in this paper, *metaethical constructivism* can be seen as Street's preferred way of developing what I am calling an "agent-attitude-dependent" view at the metaethical level. On the face of it, then, metaethical constructivism represents a distinct way of developing metaethical agent-attitude-dependence. In this section, I will consider what to make of Street's metaethical constructivism and how it relates to agent-value-subjectivism. This will allow us to further clarify agent-value-subjectivism and consider some important potential objections to developing metaethical agent-attitude-dependence as agent-value-subjectivism in the first place.

When Street puts her metaethical constructivism forward, she claims that it is distinct type of metaethical view.⁷³ By this, she means not only that it is a particular way of developing a specific species of a broad genus of metaethical view – for instance, in the way that agent-value-subjectivism is a species of the genus *naturalistic realism* – but, moreover, that it represents a distinctive genus altogether. In particular, she insists it is not best interpreted as a version of either expressivism or the descriptivist views that, in this paper, I am calling "non-naturalistic realism" and "naturalistic realism." In this section, I will argue that this way of understanding Street's metaethical constructivism is untenable. In short, I will argue that

⁷³ (Street 2008a) and (Street Forthcoming-d).

her proposed metaethical constructivism is best developed as something very similar to – if indeed not identical to –agent-value-subjectivism.

3.1 Street's Metaethical Constructivism.

Let's start with how Street explains what metaethical constructivism is. In broad terms, she claims it consists in the following thesis: “the truth of a normative claim consists in that claim's being entailed from within the practical point of view.”⁷⁴ Call this Street's *practical standpoint thesis*.

As we have seen, by a “practical standpoint” or “practical point of view” Street means the standpoint of a creature “who is in the state of mind of ‘valuing’”⁷⁵; or, more broadly, that takes certain things to have normative significance for thinking about practical questions of what to do and how to live. In the context of her discussion of metaethical constructivism, Street claims that a given agent's practical standpoint can be understood as the totality of that agent's evaluative attitudes.⁷⁶ Given this fact, the practical standpoint thesis is equivalent to saying that the truth of a normative claim consists in that claim's being entailed by the total set of evaluative attitudes that make up a particular practical point of view. Furthermore, given a) that the type of normative claims we are interested in here are claims about an agent's normative reasons and b) that Street endorses *agent-attitude-dependence*, we can thus say that the practical standpoint thesis is equivalent to the thesis that the truth of a normative judgment that an agent A should ϕ in C consists in that claim's being entailed by the total set of A's own evaluative attitudes.⁷⁷

⁷⁴ (Street Forthcoming-d, 4)

⁷⁵ (Street Manuscript, 1).

⁷⁶ (Street Forthcoming-d) and (Street 2008a).

⁷⁷ It should be noted that a different, slightly more complicated route to get to this same conclusion also suggested by Street in (Street Forthcoming-d, 32 n.16). Here, Street claims that we should start with the idea that the practical standpoint thesis is as follows: the truth of a normative claim consists in that claim's being entailed from within the practical point of view where “the practical point of view is given a formal characterization.” For Street, a “formal characterization” of the practical point of view is one that assumes nothing about the substantive content of the evaluative attitudes involved in that point of view and that instead specifies only formal features of what it is to count as having the relevant *type* of evaluative attitude at all. Then, given the strength of her practical/theoretical argument, she argues that ethical-agent-attitude-dependence is itself entailed from that formal characterization. This would thereby provide an alternative way of explaining how the particular content of an agent's evaluative attitudes come to matter for the definitions we are considering. I think my way of putting it is also consistent with much of what Street writes and is a more straightforward way of putting her view forward.

In her paper “Constructivism About Reasons,” Street proposes to develop the practical standpoint thesis into the following definition of metaethical constructivism:

According to *metaethical constructivism*, the fact that X is a reason to Y for agent A is constituted by the fact that the judgment that X is a reason to Y (for A) withstands scrutiny from the standpoint of A 's other judgments about reasons.⁷⁸

When read as a way of developing the basic practical standpoint thesis, this definition of metaethical constructivism specifies two main things beyond the practical standpoint thesis. First, it specifies that the evaluative attitudes in question are *normative judgments*. Second, it replaces the earlier discussion of a judgment being “entailed” by A 's evaluative attitudes with that of it “withstanding scrutiny” from A 's other judgments about reasons.

As stated earlier, Street's *metaethical constructivism* could charitably be read as in fact a substantive normative thesis. Its primary claim concerns what *constitutes* an agent's normative reasons. In theory, until one specifies more clearly what this constitution claim means, one could endorse this constitution claim and have any number of different fundamental metaethical views from expressivism to non-naturalistic realism. This is so because one could read this constitution claim as a substantive normative claim just as one can read a utilitarian's claim that normative reasons are constituted by what promotes the best state of affairs as a substantive normative claim. After all, many non-naturalists – for instance, Shafer-Landau and Wedgwood – are happy to talk about what “constitutes” facts about an agent's normative reasons without thinking that this in any way amounts to a claim about the basic nature of normative properties *as such*.⁷⁹

Furthermore, this substantive normative reading of the view is supported by the way in which she first presents the view in “Constructivism about Reasons.” Here, Street claims that the primary way metaethical constructivism is differentiated from more “restricted” versions of constructivism such as John Rawls' in *A Theory of Justice* is that, unlike Rawls' view, it is a “thorough-going” one that concerns not just the correctness-conditions of a certain specified subclass of normative judgments about normative reasons but rather *all* of them. Importantly, this does nothing whatsoever to mark it off as a distinctively *metaethical* thesis as I have been using the term. For many substantive normative views – for instance,

⁷⁸ (Street 2008a, 223).

⁷⁹ See (Shafer-Landau 2005) and (Wedgwood 2009).

“thorough-going” versions of utilitarianism – apply to all normative reasons. What differentiates a metaethical view is not *how many* normative reasons it concerns, but rather *what sort of claim* it is making. The claim it makes is not a normative judgment itself, but rather one about how ethical thought and talk – and the ontological commitments involved in that thought and talk – should be integrated into a broader conception of reality. Street is thus wrong to suggest that if a view such as metaethical constructivism provides “an informative account of the truth conditions of judgments about practical reasons” it is then gives us the basis for “compelling answers to all standard metaethical questions.”⁸⁰ The fact that a view gives us the truth-conditions or correctness-conditions for normative judgments is not enough to make that view a metaethical one. After all, if one thought that utilitarianism is true in every possible world – something many utilitarians in fact do think in virtue of the fact that they think the basic normative truths are necessary truths – then utilitarianism *does* give us the truth-conditions for normative judgments. But, importantly, it is still consistent with a wide range of fundamental metaethical views, including expressivism, non-naturalistic realism, and naturalistic realism.⁸¹

Nonetheless, there are compelling reasons to think that, in at least some cases, Street intends the main thesis of metaethical constructivism to be a “metaethical” one in the sense that I am using the term in this paper. First, she explicitly uses the term “metaethical” to mark it off from other views that she holds to be substantive normative ones in virtue of the different subject matter that the views are *about*. This might be because Street has a different

⁸⁰ (Street 2008a, 239).

⁸¹ A claim about the truth-conditions of normative judgment *might* entail a metaethical view if you thought the thesis applied to all reasons in all possible worlds *even* on two-dimensionalist framework of the sort advocated by (Jackson 1998) and (Chalmers 2006) according to which one considers possible worlds not just as counterfactual alternatives to the actual world (where the actual world is held fixed as the actual world we are evaluating the counterfactual possibilities from) but also as alternatives to actuality. This is for the following reason. When such a view is *combined* with the claim that concepts are individuated by the function from possible worlds to extensions that they effect – roughly the view that (Jackson 1998) advocates – then, if one showed something about the extension of the concept NORMATIVE REASON in all possible worlds, then this would be revealing of something about the conceptual content of the concept NORMATIVE REASON. I am quite sympathetic to this claim given my broader views about conceptual analysis, two-dimensionalism, and concept individuation. But, importantly, it is far from universally accepted in metaethics: certainly not by expressivists and non-naturalistic realists. It is also not a claim that Street makes and there is no reason to suppose that she shares these broader sympathies to Jackson’s views on two-dimensionalism and conceptual identity. Indeed, given that Street wants to resist claims of conceptual identity between normative and non-normative concepts in (Street 2008a, 239-242), this suggests that this basic line of thought will not be available to her.

understanding of the term that I do. But there is room to doubt this is so. For instance, consider the following quote from the paper “Constructivism About Reasons” where she introduces her position of “metaethical constructivism”:

In this paper, I work with a rough and ready understanding of the distinction between normative ethics and metaethics, according to which normative ethics investigates what reasons we have whereas metaethics investigates metaphysical, epistemological, and semantic questions about reasons and normative language.⁸²

This quote suggests that Street has a view of what metaethics consists in that is relatively similar to the one that I am working with in this paper – and thus that the main thesis of metaethical constructivism might indeed be meant to be a “metaethical” one in the sense that I am using the term in this paper. Furthermore, there is the perhaps more important fact that Street explicitly takes metaethical constructivism to be an *alternative* to standard fundamental metaethical views about the meaning and content of normative judgment: for instance, expressivism, non-naturalistic realism, and naturalistic realism. Put together, all this gives us good reason to hold that, at least some of the time, Street intends metaethical constructivism to be a fundamental metaethical view of the sort we are interested in. Suppose that this is so.

This might, of course, turn out not to be the case – and, indeed, given how little Street explicitly says about her understanding of the distinction between metaethics and normative ethics, I think it is a live option that she is at least sometimes using the term in a significantly different way than I am. However, even if Street is *not* best interpreted in the end as putting forward what I am calling a “fundamental metaethical view,” our main concern here is with how Street’s metaethical constructivism relates to agent-value-subjectivism *if* metaethical constructivism is read as a way of developing metaethical agent-attitude-dependence. Even if there are other interpretations available to us, we can then choose to read Street’s metaethical constructivism as a metaethical thesis about how ethical thought and talk – and the ontological commitments involved in that thought and talk – should be integrated into a broader conception of reality. How, though, should we understand Street’s view if we take it to be one on this topic?

⁸² (Street 2008a, 217 n.22).

Since we have already seen in section two how Street's own practical/theoretical argument casts doubt on both expressivism and non-naturalistic realism – views that Street herself also explicitly attacks on the basis of the practical/theoretical argument – I think it makes most sense to hold that her view is in fact a *naturalistic realist* way of stating metaethical agent-attitude-dependence. As I will turn to shortly, Street herself has explicit reasons for disagreeing with this. However, I think there is good reason for taking this to be the most natural way to state the view: namely, that this reading of Street's view, unlike those that treat it as a version of expressivism or non-naturalism, has *not* in any way been negatively impacted by the metaethical ramifications of the practical/theoretical argument for ethical agent-attitude-dependence. Instead, as I argued in the last section, this sort of metaethical view receives relative dialectical support from that argument. I thus think that Street's metaethical constructivism is best stated as a naturalistic realist view in metaethics that holds the following:

Modified Metaethical Constructivism: to make a normative judgment that A should ϕ in C is to express the belief that the normative judgment that A should ϕ in C withstands scrutiny from the standpoint of A's other judgments about reasons.

Our question, then, is as follows: how, if at all, does this reading of metaethical constructivism differ from agent-value-subjectivism?

In order to answer this question, there are two main issues about *metaethical constructivism* that we need to address. The first concerns Street's use of the evaluative attitude of "normative judgment" itself in her metaethical account of the content of normative judgments. The second concerns her use of the general idea of "entailment" in her account and her more specific way of spelling it out in her discussion of "withstanding scrutiny." On the face of it, both of these issues mark off metaethical constructivism as interestingly different from agent-value-subjectivism. However, I will argue that this is not the case. Thus, if I am also right that metaethical constructivism is best understood as a version of naturalistic realism, then this argument will amount to showing that metaethical constructivism is in fact best read as a version of agent-value-subjectivism.

Let's start with the first issue. In the last section of this paper, I argued *against* developing a version of metaethical agent-attitude-dependence according to which the content of normative judgments about normative reasons concerns *further normative judgments*.

This was because I argued that making use of “normative judgments” in this way would either be a non-starter for the fundamental metaethical task at hand or else an ultimately circular account of a sort that we wanted to rule out for completing the fundamental metaethical task. Furthermore, unless one specified that normative judgments expressed “*the state of making a normative judgment*,” such a view would not only be a non-starter but would seem to be making a category confusion between a type of psychological state (such as beliefs and desires) and a psychological activity (the act of making a normative judgment). On the surface, it seems as though Street might doing *exactly* what I argued against doing in terms of developing metaethical agent-attitude-dependence. However, I do not think this is the case.

To see why this is so, consider that *modified metaethical constructivism* uses the term “normative judgment” in two different places. The first is in specifying the thing we are trying to understand: the explanandum. This explanandum is what we are *all* trying to understand in fundamental metaethics: namely, normative judgment as such. The second is in specifying an explanans that we can use to understand normative judgment as such. When Street uses the term “normative judgment” in this second context, she explicitly claims that it can be understood *prior* to the idea of a judgment *about* normative reasons.⁸³ The claim is that even if we don’t understand what normative reasons are, we can identify a mental state of “taking-something-to-be-a-reason.” In turn, Street claims that we can understand the attitude of “taking-something-to-be-a-reason” as a specific type of evaluative attitude that we can identify in terms of certain logical and phenomenological features that it has. In so doing, she explicitly claims both that this attitude is desire-like, in terms of having motivational upshot for agents when combined with non-normative beliefs, but also that it is distinct from standard prosaic desires. Thus, I think, there is good *prima facie* reason to hold that Street’s second use of the term “normative judgment” refers to a type of evaluative attitude that is, at least in broad outlines, quite similar to the type of evaluative attitude that I referred to using the term “valuing.” This is, I think, reflected in the fact that, Street sometimes herself uses the exact same term “valuing” to refer to the evaluative attitude in question.⁸⁴ In turn, because of this, I think there is no good reason to hold that Street’s use

⁸³ (Street 2008a), (Street Forthcoming-d), (Street Manuscript).

⁸⁴ See, for instance, (Street Forthcoming-d) and (Street Manuscript).

of the term “normative judgment” in her understanding of metaethical constructivism itself distinguishes her metaethical view as markedly different from agent-value-subjectivism.⁸⁵

Let’s now turn to Street’s idea of something being “entailed” by the totality of an agent’s evaluative attitudes. For simplicity’s sake, let’s stick with the idea that the evaluative attitude in question is “valuing.” When Street uses the notion of “entailment” in her practical standpoint thesis, she explicitly intends it to be understood as a *non-normative* concept. She writes: “to explain this sense of entailment, we needn’t make any substantive normative assumptions—for example, about what anyone *should* or *ought* to do or infer, or about what counts as a *reason* for what; to make such assumptions would be uselessly question-begging from a metaethical point of view.”⁸⁶ In order to refrain from such question-begging, Street proposes that we think of facts about what is entailed by an agent’s valuing as simply facts about what – as a constitutive matter – must be the case for that agent to actually to have those valuing in the first place. As she puts it, according to this view, “we need only make observations about what is constitutively involved in the attitude of valuing or normative judgment itself — identifying those things such that if one fails to do them, one is not making a *mistake* of any kind, but rather is not recognizably *valuing* at all.”⁸⁷ For Street, the

⁸⁵ It should be noted that, given some of the more specific things Street wants to say about the attitude of “taking-something-to-be-a-reason,” she takes on more specific commitments about the evaluative attitude in question than I intend to with “valuing.” For instance, in her discussion of how this attitude is not only desire-like in certain ways but also importantly *belief-like* in certain ways, she seems *at least open* to the idea of there being *besires*: mental states that have *both* a “mind-to-world” and a “world-to-mind” direction of fit. I myself am skeptical that this is a coherent idea for the standard reasons given by those in favor of the Humean theory of motivation. For one of the canonical arguments against *besires*, see (Smith 1994). However, as I have stated it in this paper, the attitude of “valuing” could in theory be either a species of “desire” or “besire” in the functionalist senses of how these terms are used in philosophy of mind and action. Thus, even if Street were to take on more specific commitments about the attitude in question, such as it being a species of *besire*, these commitments can be seen as thus constituting a specific way of developing agent-value-subjectivism rather than an alternative to the view as a whole. Another way one might develop Street’s thought that the attitude of valuing is both desire-like and belief-like in certain ways is by holding that this is only when we think of the terms “belief” and “desire” in the senses we use them in natural language rather than in the specific functionalist definitions that we sometimes use for them in philosophy of mind and action. If so, then she might in the end have a view that valuing just is a species of desire in this functionalist sense. Finally, one might also develop Street’s thought that the attitude of valuing is both desire-like and belief-like by holding that valuing is a *compound state* that one only counts as instantiating if one has a certain collection of other more primitive states (e.g certain beliefs, desires, or plans). As with the view that valuing is a type of *besire*, these views can be seen as constituting specific ways of developing agent-value-subjectivism rather than alternatives to the view as a whole.

⁸⁶ (Street 2008a, 232).

⁸⁷ (Street Forthcoming-d, 7).

relationship here between valuing and things that are entailed by it, is, in basic terms like the entailment relationship between being a parent and being a child. As she writes: “Being a parent entails having children, but that does not mean that a ‘parent’ is making an error if she has no children; she’s just not a parent.”⁸⁸ Thus, for Street, the facts about what is entailed by a *particular* token valuing or of a particular agent’s valuing(s) as a whole are essentially facts about the *attribution-conditions* of that specific valuing or group of valuing(s) to that agent. They are meant to be the *minimal conditions* that an agent needs to meet in order to count as having valuing(s) with content X – without meeting those conditions, we either would think that an agent was valuing something other than X or else, as she puts it, not “recognizably *valuing* at all.”⁸⁹

In insisting that the relationship of entailment from evaluative attitudes is a constitutive rather than normative relationship, Street’s use of term “entailment” here appears to be modeled on the idea of *logical entailment*. Following standard convention, we can say that p logically entails q iff q can be logically deduced from p. Thus, if p logically entails q, this means that if we learn that p is true, then we also know that q must be true as well. In his influential paper “Logic and Reasoning” Gilbert Harman argues that such deductive logical notions as logical entailment are not inherently normative: they tell us about the necessary connections between propositions, not about how we necessarily *ought* to reason.⁹⁰ The same, Street says, is true of the technical notion of entailment that she uses in her metaethical theory.

But how exactly are we to understand Street’s notion of entailment? Is it meant *literally* to be logical entailment? If so, wouldn’t it then imply that if an agent has currently inconsistent valuing(s), then those entail everything? Perhaps then Street’s “entailment” here is just something modeled on logical entailment – something, for instance, that also brings in broader ideas about coherence that aren’t strictly a part of the idea of logical entailment. These are, I think, difficult questions with indeterminate answers given that Street has only given us a basic sketch of what entailment involves. However, some of these questions – and, importantly, the ones that I think matter for this paper – are partly answered by seeing how the idea of “entailment” connects with Street’s idea of “withstanding scrutiny.”

⁸⁸ (Street 2008a, 232).

⁸⁹ (Street Forthcoming-d, 7).

⁹⁰ (Harman 1984).

Suppose that an agent has a specific set of valuing V1 – a set that, as with those we discussed before, included a range of valuing about a range of different topics and some of which were more “fundamental” than others. The normative judgment that an agent A should ϕ in circumstances C “withstands scrutiny” from V1, proposes Street, when A can be said to actually have those specific valuing V1 *and* when the total set of valuing V1 favors ϕ -ing in C. This, she proposes, will not always be the case. In some cases, it will be because we would only attribute the attitude of *valuing* to A in the first place if she did certain things or had certain properties. And, similarly, in many cases, we would only attribute those *specific* valuing V1 to A if she did certain things or had certain properties.

Street does not say what exactly it takes to count as valuing at all and what it takes to have a particular token valuing. These are, instead, questions that she leaves open for future research. She does, however, suggest that one thing valuing as such involves is some commitment by the agent to doing what *instrumentally follows* from one’s valuing. It is unclear how exactly this should be read. Suppose we read this as somehow involving a version of the *instrumental principle* that reads as follows: one must either take the necessary means to one’s ends or else give up those ends. On the one hand, as Street herself notes, we routinely fail to either take the necessary means to achieve the ends specified in our valuing or else give up those ends. There might be a minimal disposition to follow the instrumental principle we must display in order count as valuing at all. But it certainly won’t be actually *meeting the demands of* that principle in all such cases. Otherwise, there would be no room for any *failure* to follow the norms of that principle – something we think agents can do all the time and still count as valuing. Therefore, it seems likely that it would be a dead end to understand *following* the norm of the instrumental principle a constitutive feature of valuing as such – or at least in any straightforward sense of “following” in the sense of “conforming to.”

However, I think what Street has in mind is something different from this strategy: namely, it is something about specific valuing we must attribute to an agent to count as valuing at all. Against Korsgaard, Street claims that there are almost no constraints on what content an agent’s valuing must have simply in virtue of that agent being a *valuer*. However, in suggesting that valuing constitutively entails a *commitment* to what *instrumentally follows* from one’s valuing, and not just the instantiation of a pattern of doing so, it seems reasonable to hold that Street takes it that all valuers as such must accept a version of the instrumental principle in order to count as valuers. In other words, they must *have* this particular valuing

of valuing a version of the instrumental principle – even if, as failures to follow the instrumental principle show, this will not always dictate what the actually do in practice.⁹¹

If so, this would explain how Street is able to invoke a seemingly normative principle as part of her metaethical view simply on the grounds of the fact that agents have certain mental states. In short, it would be 1) because in order to count as having that mental state, agents would have to have this particular valuing as a fundamental valuing and 2) because, as an ethical agent-attitude-dependent theorist, she holds that what an agent *should* do is a function of that agent's own fundamental valuing.

The important point here, however, is not whether Street is correct or not in the particular suggestions she gives about what is constitutively involved in having a particular valuing. The question, rather, is this: if this topic is all that the language of “entailment” is meant to pick out, what does this mean for how Street's metaethical constructivism relates to agent-value-subjectivism? In essence, what it would mean is that when a normative judgment that an agent A should ϕ in circumstances C “withstands scrutiny” from the totality of an agent's valuing (suppose set V1), the norms for what it would take to “withstand scrutiny” are coming entirely from the *content* of the particular valuing in set V1. The attribution conditions for specifying that A has that particular set of valuing V1 don't themselves deliver any extra content. Rather, they only tell us what must be the case in order for it to be true that A has that set of valuing.

In turn, as with agent-value-subjectivism, Street's metaethical constructivism needs a way of specifying what it would be for an *entire set of valuing* to favor one thing rather than another. As I noted earlier in the paper, this issue is especially important given that any agent's set of valuing will likely include many different valuing that favor different things – a fact that threatens to yield large amounts of indeterminacy about that agent's normative reasons. One option would be to use a norm that we hold fixed for all agents about which valuing should receive priority. But this, of course, would undermine the entire thrust of ethical agent-attitude-dependence. With this danger in mind, Street takes precisely the sort of route that I advocated in this sort of context when presenting agent-value-subjectivism: namely, she appeals to the different weights that an agent herself tacitly gives to her different valuing.⁹² As Street puts it, we need to appeal to the valuing of an agent that are “most

⁹¹ For a line of thought sympathetic to this claim, see (Gibbard 1999).

⁹² (Street 2008a, 232-236).

deeply *hers*.”⁹³ As she puts it, which valuing is “most deeply hers” is “a function of how strongly she holds the normative judgments [i.e. valuing] in question and how close to the center of her total web of normative judgments [i.e. valuing] they lie.”⁹⁴ In turn, we can think of further valuing either being consistent with those deepest valuing or not. For a whole set of valuing to favor ϕ -ing, it therefore seems, is for a *particular version* of this system of valuing to favor ϕ -ing: namely, the version where all of an agent’s valuing “withstand scrutiny” from the valuing of an agent’s that are “most deeply hers.”

With this proposal for which of an agent’s evaluative attitudes matter, it seems to me that the best way to develop this line of thinking is to accept the following thesis: for the purposes of explaining normative judgment, the phrase “A’s present system valuing favor ϕ -ing” is true *iff* a system of valuing that was perfectly regulated and revised according to the weights given in A’s original system of valuing favored ϕ -ing. In other words, it seems to amount to the same claim that I introduced in the context of agent-value-subjectivism: that, for the purposes of explaining normative judgment, the phrase “A’s present system valuing favor ϕ -ing” is true *iff* A’s *internally-idealized-valuing* favor ϕ -ing. If this is so, then it turns out that Street’s talk of “entailment” in the practical standpoint thesis does not signal a radically different way of developing metaethical agent-attitude-dependence. Instead, it ends up amounting to another way of saying what I argued for in the context of developing agent-value-subjectivism.

Let’s now take stock of where we are given my arguments in this section. First, if I am right, then Street’s metaethical constructivism can be understood – and, indeed, is perhaps best understood – as putting forward a view that is specifically “metaethical” in the sense that I have been using the term in this paper. Second, assuming the first conclusion is right, then Street’s metaethical constructivism is best read as a version of “metaethical naturalistic realism” in the sense that I have been using the term in this paper. Third, assuming that both the first and second conclusions are right, then it is in fact best read as another way of specifying a version of the basic metaethical view that I spelled out in the last section as “agent-value-subjectivism.” For all I have said, Street’s metaethical constructivism might be one that takes on additional commitments beyond what I have said in this paper about what the attitude of “valuing” used in the definition of agent-value-subjectivism

⁹³ (Street 2008a, 235).

⁹⁴ (Street 2008a, 235).

amounts to – and, thus, one that amounts to specific proposal for how to further develop agent-value-subjectivism. This might or might not be so. However, the important point is this: however it is further developed, Street’s metaethical constructivism is best understood as a version of metaethical naturalistic realism that essentially amounts to a version of agent-value-subjectivism.

3.2. *The Purported Contrast with Naturalistic Realism.*

As I emphasized earlier in this paper, Street herself explicitly rejects reading her metaethical constructivism as a version of what I am calling “naturalistic realism.” Moreover, there are reasons to be worried about this reading of Street’s view that she herself does not much discuss. In order to defend my argument that Street’s metaethical constructivism is best read as a version of metaethical naturalistic realism that essentially amounts to a version of agent-value-subjectivism, let’s now turn to some of the most important issues facing my reading of Street’s metaethical constructivism as a version of naturalistic realism.

Let’s start with the reason that Street herself explicitly gives against reading her metaethical constructivism as a version of what I am calling metaethical “naturalistic realism.” This reason concerns the difference between the epistemological status of facts about what entails what (which we can call *entailment facts*) and psychological facts of the sort that are often referred to in metaethical naturalistic realist theories. In order to bring out this contrast, Street imagines a contrast between her view and a naturalistic version of *ideal attitude theory* in metaethics.⁹⁵

In basic terms, according to naturalistic versions of ideal attitude theory in metaethics, to judge that an agent A should ϕ in circumstances C is to express the belief that some agent (or group of agents) would have an attitude that favored A ϕ -ing in C in ideal conditions. Different versions of ideal attitude theory differ in at least three important respects. First, they differ on the question of which agents matter (e.g. an impartial observer vs. agent A herself vs. all agents). Second, they differ on the question of what type of attitude in ideal conditions we are interested in (e.g. a desire vs. a belief vs. a plan). Third, they differ on what constitutes “ideal” conditions (e.g. it is left as an unexplained primitive state vs. it is spelled out as being in a state that is described using further normative terms vs. it is spelled

⁹⁵ (Street Forthcoming-d, 14-19).

out as being in a state that is specified as such-and-such naturalistic state that is specified in non-normative terms).

For the purposes of explaining why her view is not an instance of naturalistic realism, Street supposes that the ideal attitude theory in question specifies what it is to be in “ideal” conditions as being in such-and-such naturalistic state that is specified in non-normative terms. If so, then, as Street argues, it would be in principle possible (even if impractical in reality) to study what attitudes agents would have in those circumstances using the empirical methods of the natural and/or social sciences. Thus, if the attitudes of some agents fix what normative reasons an agent A has – say, for the sake of example, the desires of everyone in that agent’s A cultural group – then it would be in principle possible (even if impractical in reality) to study *through scientific means* what normative reasons an agent A has. This is because such facts turn out to be a complicated type of empirical fact. In contrast, Street claims, one could never study *entailment facts* through the methods of the sciences because entailment facts simply are not empirical facts amenable to such study. In turn, she claims that since *what it is* to have a thoroughly naturalistic understanding of normative properties is to hold that these properties are amenable to such study, the practical standpoint thesis departs from a through-going naturalistic view of normative properties – and, thus, cannot be accurately characterized as a version of what I am calling “metaethical naturalistic realism” in this paper.

I agree with Street that there are important differences between the methods we use to investigate entailment facts and those that we use to investigate empirical facts, such as facts about the content of an agent’s evaluative attitudes (assuming, for now, that those are empirical facts). However, I disagree that this difference is important for understanding the relationship between Street’s metaethical constructivism and naturalistic realism in metaethics.

To see why this is so, remember that, for Street, facts about what valuing entails are not meant to be normative facts. We now know that they are not meant to be empirical facts. It thus seems to make sense to think of them as *a priori* ones. But what kind? Are they supposed to be synthetic or analytic? In order to answer this question, remember that, for Street, they are meant to be facts that simply fall out of what it is constitutively involved in having a particular valuing in virtue of our very concept of *valuing*. These facts are meant to be non-mysterious facts that we can access in much the same way that we can access *any*

facts about what is constitutively involved in counting as an instance of some kind K – for instance, facts about what is constitutively involved in *being a parent*. Thus, it makes sense to think of them as *a priori* facts that are true in virtue of facts about the content of our concepts such as VALUING or PARENT – in combination, perhaps, with holding fixed the logical constants. In short, entailment facts thus look best read as analytic *a priori* facts.

Yet, if this is so, it creates no problem at all for classifying her metaethical constructivism as a version of naturalistic realism. *Many* of the leading naturalistic realist views in metaethics take for granted that there are some facts about conceptual content that are analytic *a priori* and that matter for ethics. For instance, Frank Jackson – who is one of the paradigmatic examples of a naturalistic realist in metaethics – puts such facts at the front and center of his theory.⁹⁶ Indeed, others who do not, such as Peter Railton, still often have *some* way of understanding what analytic *a priori* facts might be and how they might play a role in a theory of normativity.⁹⁷ Naturalistic realists will differ greatly on how to understand the status of these *a priori* facts – for instance, Railton takes a broadly Quinean approach to the *a priori* whereas Jackson does not.⁹⁸ Some of the reason for these differences stem from concerns about how to integrate *a priori* facts into a broadly naturalistic worldview – concerns that have animated much of twentieth-century philosophy in the wake of logical positivism. However, the point here is that, for the purposes of *metaethics*, this is best understood as a further debate. For if it is not, then almost no one who is currently classified as a “naturalist” about normative properties – including paradigmatic examples such as Jackson and Railton – is accurately categorized. Thus, the fact that Street’s view appeals to analytic *a priori* facts in combination with empirical facts about the content of agent’s evaluative attitudes should in no way disqualify her view from being understood as an instance of naturalistic realism.

Indeed, as my way of putting forward agent-value-subjectivism forward suggests, this view – and, hence, Street’s own metaethical constructivism if my argument earlier in this section is correct – can in fact be read as an instance of naturalistic ideal attitude theory of precisely the sort that Street *contrasts* her view with. Remember, in its most schematic form, an ideal attitude theory in metaethics states that to judge that an agent A should ϕ in

⁹⁶ (Jackson 1998).

⁹⁷ (Railton 2001).

⁹⁸ Contrast (Railton 2001) and (Jackson 2001).

circumstances C is to express the belief that some agent (or group of agents) would have an attitude that favored A ϕ -ing in C in ideal conditions. Now remember that, according to agent-value-subjectivism, to make the normative judgment that an agent A should ϕ in circumstances C is to express the belief that A's own internally-idealized-valuings favor ϕ -ing in C. This view can be seen as a way of specifying a particular way of developing an ideal attitude theory. First, on the question of which agents matter, agent-value-subjectivism specifies that it is *only the agent A herself*. Second, on the question of what type of attitude in ideal conditions we are interested, agent-value-subjectivism specifies that they are *valuings*. Third, on the question of what constitutes "ideal" conditions, agent-value-subjectivism specifies that they are the conditions of the agent's own valuings being *internally-idealized* in the sense of those attitudes being perfectly regulated and revised according to the weights given by that agent's own system of valuings. Such a view might differ significantly from many standard versions of ideal attitude theory – for instance, by allowing the conditions of idealization to essentially be determined by contingent facts about the content of an agent's own evaluative attitudes. Yet, it is still a version of ideal attitude theory. Moreover, if my argument in the previous paper of this dissertation "The Metaethical Role of Evaluative Attitudes" is on the right track, then such a version of ideal attitude theory might in fact be the one that *any* ideal attitude theorist committed to metaethical naturalistic realism should in fact adopt.

3.3. Other Issues for My Reading of Street as a Naturalistic Realist.

Let's now turn to three other issues that confront my argument that Street's metaethical constructivism is best developed as a version of metaethical naturalistic realism.

The first issue is that unlike many naturalistic realists – say, for instance, Frank Jackson – Street appears to think that normative *concepts* are fundamentally irreducible to non-normative concepts.⁹⁹ As she puts it, "the idea of one thing's being a reason for another cannot successfully be reduced to thoroughly non-normative terms."¹⁰⁰ Furthermore, she thinks that even though an agent who has the specific evaluative attitude of *valuing* does not necessarily need to employ normative concepts in order to count as having that attitude, we nonetheless might need to employ such concepts in order to accurately identify that attitude

⁹⁹ (Street 2008a) and (Street Forthcoming-d).

¹⁰⁰ (Street 2008a, 239).

in the first place.¹⁰¹ One might think that all of this puts pressure on my understanding of metaethical constructivism as a version of naturalistic realism. In short, one might think that Street's views about normative concepts are ultimately incompatible with the view of normative concepts and/or the view of normative properties that naturalistic realism is committed to.

I think that there is some legitimate worry here – especially given how little Street says about what it means to say that normative concepts are fundamentally irreducible to non-normative concepts. Some ways of developing this claim about the irreducibility of normative concepts lead naturally to expressivism – which, for instance, Gibbard puts forward in *Thinking How to Live* partly on the basis of the idea that normative concepts are irreducibly normative because they are *sui generis*.¹⁰² Furthermore, other ways of developing this claim about the irreducibility of normative concepts lead naturally to non-naturalistic realism – which can be seen as specifying that the reason that normative concepts are irreducibly normative is that they refer to non-naturalistic irreducibly normative properties.¹⁰³ Yet, as we have seen, there are strong reasons to hold that Street rejects both such metaethical views. Furthermore, in claiming that normative concepts are fundamentally irreducible to non-normative concepts she in fact joins the ranks of other philosophers committed to naturalistic realism in metaethics including Michael Smith and Peter Railton (who also accepts a version of Street's claim that normative concepts might be necessary to employ in order to accurately pick out certain mental states).¹⁰⁴ These philosophers hold that we can say things about how normative concepts play certain functional roles in our cognition that non-normative concepts do not *without* taking on Gibbard's idea that this means that normative concepts are *sui generis* and *without* giving up a thoroughgoing naturalism about normative properties. In the end – as both Smith and Railton are well aware of – whether or not such a view will be tenable for a naturalistic realist to hold will depend in part on what sort of conceptual reduction one thinks is necessary for property reduction. Such issues are, as I said earlier in this paper, simply beyond the scope of our concern here. The point, though, is this: without *significant further argument*, one cannot rule out views such as Smith's and Railton's from being classified as “naturalistic realist” ones *on*

¹⁰¹ (Street 2008a) and (Street Manuscript).

¹⁰² (Gibbard 2003).

¹⁰³ See, for instance, (Dancy 2006) and (Parfit Forthcoming).

¹⁰⁴ (Smith 1994) and (Railton 2003b).

the basis of their views about the irreducibility of normative concepts and/or their views about the conditions for ascriptions of mental states.¹⁰⁵ If this is so, then, pending further argument, we also lack good reason for ruling out Street’s view from counting as a “naturalistic one” on those grounds.¹⁰⁶

A second issue that confronts understanding Street’s metaethical constructivism as a version of what I am calling “metaethical naturalistic realism” is that Street herself does not explicitly put metaethical constructivism forward in combination with a descriptivist claim about the attitudes expressed in *normative judgment*. According to descriptivists, to make a normative judgment is to express an attitude that, at the most basic explanatory level, should be understood as a belief-like attitude with a “mind-to-world” direction of fit. This is in contrast to expressivists who hold that to make a normative judgment is to express an attitude that, at the most basic explanatory level, should be understood as a desire-like attitude with a “world-to-mind” direction of fit. We have seen that Street rejects expressivism – and, moreover, that there are good reasons for her to do so given the metaethical significance of her practical/theoretical argument that I discussed in the second section of this paper. This, however, leaves open another possibility that I have not discussed in this paper – namely, that Street could have a non-descriptivist view that was still not expressivism. In particular, she could in theory have the view that to make a normative judgment is to express an attitude that, at the most basic explanatory level, should be understood as a *desire* with both the “mind-to-world” direction of fit of belief-like attitudes and the “world-to-mind” direction of fit of desire-like attitudes.

¹⁰⁵ See (Jackson 1998) for an argument against the idea that one successfully defend a naturalistic account of normative properties without also reducing normative concepts to non-normative concepts.

¹⁰⁶ This important fact is, I think, obscured by Korsgaard’s understanding of “realism” in (Korsgaard 2003). Korsgaard argues that, according to realists, the basic functional role of normative concepts is a representational one. She denies this. Instead, she argues that the basic functional role of normative concepts is to help us to solve problems about how to live. One *could* have a debate about what are the most “basic” functional roles that normative concepts play in our lives. But this is not the debate we are interested in here – nor is it the one that usually is taken to be settled by naturalistic realists and non-naturalistic realists who embrace the basic descriptivist idea that normative judgments express genuine beliefs that represent a part of reality. For, in order to claim that descriptivism is true about normative judgment, one can grant that the basic functional role of normative concepts is to help us to solve problems about how to live and then, in turn, explain that they do so in part by helping to represent things as being one way or another. To put it another way: claiming that normative judgments express belief is not the same as claiming that they do so *because* their primary functional role in our lives is to correctly represent facts. For a similar line of criticism of Korsgaard’s way of framing things here, see (Hussain and Shah 2006).

I think there is good reason for taking this possibility seriously. For when Street discusses what it is to be in the state of valuing, she claims that it is not only desire-like in certain ways but, also, importantly *belief-like* in certain ways.¹⁰⁷ In doing so, she seems at least open to the idea of there being *besires*. If Street were already committed to the theoretical coherence of such an attitude, then she could in theory hold that this is the same sort of attitude expressed by normative judgment itself. In the end, perhaps that is the view that Street herself intends to endorse.

However, I think there are strong reasons against taking this to be the best way to develop metaethical constructivism as Street has stated it. The first is that the idea that there *are* any attitudes with both directions of fit is a hugely controversial thesis in the philosophy of mind and action. Hence, it would be better, if possible, not to have the basic coherence of the view rest on such a controversial idea. Secondly, one of the costs of developing such a non-descriptivist view of normative judgment (whether expressivist or not) would be to explain how a non-descriptivist semantics could capture the descriptivist surface-features of normative discourse – a cost that Street herself notes for expressivism and wants to avoid. Thirdly, and perhaps most importantly, if Street were to accept the idea that normative judgments express *besires*, it would entail that the only time a speaker could make a normative judgment that correctly identifies the normative reasons of another agent would be if the speaker herself expressed an attitude that was in part desire-like – perhaps, for instance, that was in part a pro-attitude in favor of the agent acting on her normative reasons. Importantly, however, Street herself explicitly denies that we necessarily do express any pro-or-con in many cases where we make normative judgments. In particular, Street holds that when we make normative judgments about an agent with starting valuing very different from our own, we can make normative judgments about that agent’s normative reasons without in any way either approving or condemning the agent acting on those reasons.¹⁰⁸ Some philosophers will, of course, find that to be an implausible claim about normative judgment. For, according to some philosophers, having an attitude of approval towards A’s ϕ -ing in C is at the core of what *it is* to make a normative judgment that an agent A has a normative reason to ϕ in C.¹⁰⁹ The point, however, is that *if* Street is committed to

¹⁰⁷ (Street 2008a, 230-231).

¹⁰⁸ (Street 2009).

¹⁰⁹ For a strong version of this view, see (Gibbard 1990) and (Gibbard 2003).

the idea that we can make normative judgments about those agents' normative reasons without in any way either approving or condemning an agent acting on those reasons – an idea that she in fact puts forward as a central part of her theory – then it rules out reading her as holding that normative judgment involves expressing an attitude that, at the most basic explanatory level, partly (even if not exclusively) involves a world-to-mind direction of fit.

A third issue that confronts understanding Street's metaethical constructivism as a version of what I am calling metaethical "naturalistic realism" concerns the different connotations of the word "constructivism" in ethical theory. In self-consciously presenting her metaethical view as a "constructivist" one, Street intends to draw on the work of other philosophers that she takes to be part of the constructivist tradition – most importantly, John Rawls, Christine Korsgaard, and T.M. Scanlon.¹¹⁰ Moreover, she takes her view to be the best way of developing core specifically "constructivist" insights of that tradition into a systematic view in metaethics. This, I think, presents a problem for understanding Street's metaethical constructivism as a version of what I am metaethical "naturalistic realism."

First of all, *all* the major figures that Street identifies as part of the constructivist tradition – namely, Rawls, Korsgaard, and Scanlon – explicitly reject the label of metaethical naturalistic realism as well as many of the core theses associated with it. In the case of Scanlon and Rawls, this can be chiefly explained by the fact that they do not intend for their "constructivist" views to be *either* normative theories that cover all normative reasons or metaethical theories. Yet, in Korsgaard's case, it seems that one of her central commitments – and one that she takes to be a part of her constructivism – is that ethical facts show up only from the practical standpoint and are, in principle, epistemically inaccessible from the theoretical standpoint. This *prima facie* seems to stand in stark contrast to the basic commitment of naturalistic realism to the thesis that normative properties are naturalistic. And, indeed, this is something that Korsgaard herself emphasizes in contrasting her view with naturalistic realist theories.¹¹¹

A second reason that Street's understanding of her view as "constructivist" presents a problem for my interpretation of her view as a naturalistic realist metaethical one is that Rawls, Korsgaard, and Scanlon have, in different ways, expressed discontent with

¹¹⁰ (Street 2008a) and (Street Forthcoming-d).

¹¹¹ (Korsgaard 2003) and (Korsgaard 1996).

assumptions of mainstream metaethics.¹¹² This is most notable in the case of Korsgaard who claims that her constructivist position goes beyond the working assumptions of metaethics as a whole.¹¹³ Third, in part because of this, there is a widespread assumption that if there were a view worth calling “constructivist” in metaethics, it would have to be a unique type of metaethical view that somehow transcends current divisions between theories and that, at the very least, would not be a version of a metaethical view as canonical as naturalistic realism.

Finally, as the metaphor of “construction” itself perhaps suggests, there is a strong strand in Korsgaard’s work that emphasizes the voluntarist idea that we ourselves actively create the normative facts through our decisions and/or our judgments.¹¹⁴ Street picks up on this voluntarist theme when she claims that the idea of “radical choice”¹¹⁵ has “an important place in metaethical constructivism.”¹¹⁶ As my way of presenting Street’s view underscores, I do not think that the idea of “choice” is the right one to use here to capture the basic *contingency* involved in our fundamental valuations that Street has in mind. Nonetheless, in putting things in this more voluntarist language, Street *does* draw on an important strand of previous so-called “constructivist” theories. If this voluntarist strand is taken to be central to what constructivism is, then this presents an issue for my reading of Street’s metaethical constructivism as a version of the non-voluntarist, non-judgment-dependent metaethical view of agent-value-subjectivism. Moreover, it also perhaps presents a worry for reading it as a version of naturalistic realism at all.

In the end, though, I am not sure how much any of these concerns about these different associations with the word “constructivism” matter for my argument in this paper. If Street is right, then her view represents a good way of articulating a specifically *metaethical* view based on some of the core thoughts of the so-called “constructivist” tradition in ethical theory represented by Rawls, Korsgaard, Scanlon, and others. Furthermore, as she herself emphasizes, her view certainly does not represent *all* the major thoughts associated with or

¹¹² See, for instance, (Rawls 1999), (Korsgaard 2003), (Korsgaard 1996), (Scanlon 2000), and (Scanlon Forthcoming).

¹¹³ See (Korsgaard 2003).

¹¹⁴ (Korsgaard 1996). As Gibbard argues in (Gibbard 1999), Korsgaard can be read as endorsing the view that he calls “logically constrained reflective subjectivism.” Gibbard writes that “according to this logically constrained reflective subjectivism, reflectively deciding to do something *makes* it the thing to do, so long as your policies for action are logically consistent (Gibbard 1999, 148).”

¹¹⁵ (Street 2008a, 237).

¹¹⁶ (Street 2008a, 237).

put forward by these or other philosophers associated with the “constructivist” label – let alone all their thoughts that specifically concern the “constructed” nature of ethical facts.¹¹⁷ Thus, I think that it is an open question whether Street’s view should be labeled “constructivist” or not. However, *if* Street is right that her view counts as “constructivist,” then I think that we should hold that it does so primarily based on the *substance* of the view itself rather than on the fact that she *thinks* it is different from other existing metaethical views. Suppose that assumption is right and Street is right to hold that the basic substance of her view merits calling her view “constructivist.” In turn, if I am then right that Street’s view is best developed as a version of naturalistic realism given the substance of her view, then I think that there are good grounds for holding that Street’s view would still count as “constructivist” even if turned out to be a version of naturalistic realism.

Yet, even if this is so, I am not sure how much this issue should concern us. First and foremost, this is because the label of “constructivism” does not have a clear meaning in ethical theory. Once a meaning to that term is specified that explains what theses so-called “constructivists” accept, one can then ask if there could be a viable metaethical view that accords with those theses. Furthermore, one could then ask if that viable metaethical view would be a *unique* type of metaethical theory, different from standard types of metaethical view such as expressivism, non-naturalistic realism, naturalistic realism, and fictionalism. My own view is that, for any viable way of reading what the core thesis of “constructivism” amounts to, a metaethical view that captured that thesis would, in the end, fail to be a distinct type of metaethical view different from such standard types as expressivism, non-naturalistic realism, naturalistic realism, and fictionalism.¹¹⁸ Furthermore, I see no reason to hold that there is any constructivist position – whether Street’s own or some other one – that would support Street’s Korsgaardian claim that “constructivist positions in ethics may be seen as ultimately breaking down this distinction [between metaethics and normative

¹¹⁷ For instance, in (Street Forthcoming-d), Street explicitly denies a standard way of characterizing constructivism – namely, the *proceduralist* characterization according to which constructivists hold that normative facts are derived from a certain independently specified procedure. It should also be noted that, as I have emphasized in this paper, her view is also decidedly less voluntarist than the metaphor of “construction” or her occasional talk of “radical choice” might lead one to expect.

¹¹⁸ For developed arguments that support this idea, see (Enoch 2009), (Darwall, Gibbard, and Railton 1997), and (Hussain Forthcoming).

ethics].”¹¹⁹

In this section, though, I have not argued on behalf of these claims. Rather, I have argued for a claim about how to understand the substance of Street’s view. Street herself takes her metaethical constructivist view to be a way of developing what she takes to be a core insight of the constructivist tradition: namely, the general idea that the correctness of a normative judgment is always determined *relative* to a particular practical standpoint or to a set of practical standpoints.¹²⁰ In “restricted” versions of constructivism, Street claims, the standpoints in question are specified as ones where it is taken for granted that certain normative judgments are correct – such as, she thinks, in Rawls’ political philosophy. In others – such as her version – this is not the case. The broad thesis that Street identifies here – namely, that the correctness of a normative judgment is always determined *relative* to a particular practical standpoint or to a set of practical standpoints – may or may not be an insight worth calling “constructivist.” Furthermore, there may or may not be other ones worth calling “constructivist.” All that my arguments show is this: *regardless* of whether Street is right that her metaethical constructivism is a way of developing a view in metaethics that captures this core thesis that she takes to be a “constructivist” one, her actual developed view fails to represent a distinct alternative to naturalistic realism. More specifically, it ends up being best read as a version of the naturalistic realist view that I am calling “agent-value-subjectivism.”

In the end, given all of the varied connotations the label of “constructivism” has in ethical theory – not to mention in other areas of philosophy such as philosophy of mathematics – I think that those attracted to the type of agent-centered subjectivist view I develop would do best *not* to use the label “constructivism” to describe it. This is especially so given that some take it to be part of the very aim of constructivist theories in metaethics to offer a distinctive type of theory in metaethics. Yet, if one is attached to the label “constructivism” for further reasons – perhaps because one, like Street, wants to emphasize that the view *is* a good way of developing certain ideas that have been put forward as “constructivist” ones – then such considerations may outweigh the reasons against using the

¹¹⁹ (Street 2008a, 217 n.22). For developed arguments against this bold claim on behalf of constructivism, see (Hussain and Shah 2006) and (Hussain Forthcoming).

¹²⁰ (Street 2008a) and (Street Forthcoming-d).

term. At the end of the day, nothing of much philosophical substance hangs on the use of the label.

4. Conclusion.

In this paper, I have argued that if one accepts the substantive normative thesis of *ethical* agent-attitude-dependence – and, in particular, if one accepts it on the basis of something akin to Street’s practical/theoretical argument – then this has ramifications for theory-choice in metaethics. In particular, I have argued that it lends relative dialectical support to the broad view of *metaethical* agent-attitude-dependence as well as to a particular way of developing that view: namely, as a version of metaethical naturalistic realism that I have called *agent-value-subjectivism*. Following this, I then turned to one of the main other ways in which metaethical agent-attitude-dependence has been put forward in contemporary metaethics – namely, Street’s metaethical constructivism. Street claims that this view represents a distinctive type of metaethical view that cannot be reduced to a naturalistic realist position such as agent-value-subjectivism. I argued that this reading of metaethical constructivism is untenable and that, in fact, Street’s view is best developed as an instance of agent-value-subjectivism rather than as an alternative to it.

Progress in metaethics is made in different ways on many different fronts. It is therefore important to bear in mind how this paper has attempted to make progress in metaethics and how it has not. In this paper, I have *not* given a sustained argument that something akin to Street’s practical/theoretical argument for ethical agent-attitude-dependence is correct. Nor have I given an argument about how agent-value-subjectivism should be assessed relative to other competing views in metaethics – let alone that agent-value-subjectivism is the correct view in metaethics. Therefore, I have not put forward anything like a straightforward argument for agent-value-subjectivism. Rather, what I have done was to explain how a certain position in normative ethics (namely, *ethical* agent-attitude-dependence) can give relative dialectical support to a position in metaethics (namely, *metaethical* agent-attitude-dependence) and then to explore how to best develop this metaethical position (namely, as agent-value-subjectivism). Insofar as I take Street’s argument for ethical agent-attitude-dependence to be a strong one, I think that there are strong reasons for taking seriously the metaethical position of agent-value-subjectivism that I put forward. Furthermore, I think that there are other independent arguments one can give

on its behalf that I have not put forward in this paper. These include arguments in favor of naturalistic realist positions in metaethics in general as well as those that favor specific features of agent-value subjectivism. Yet, in this paper, I have not given such arguments.

In order to further evaluate the metaethical prospects of agent-value-subjectivism, such further arguments would obviously need to be given. Furthermore, crucial aspects of the view would need to be spelled out with more detail and central objections to the view would need to be addressed. Some of the most pressing issues here include the following:

The Attitude of Valuing. In this paper, I put forward the attitude of “valuing” in very schematic form. This had specific dialectical advantages in keeping us focused on the basic structure of the view rather than on details that different agent-value-subjectivists can disagree on. What more, though, can be said about what this attitude of valuing is? And is there a particular species of the general attitude of valuing that we need to focus on more directly in accounting for judgments about normative reasons? How much more should we say on this front for the purposes of doing metaethics?

Favoring. In this paper, I have frequently talked about evaluative attitudes “favoring” one course of action or another. What more can be said about what it is for this to be the case?

Coherence Between Different Valuing. In this paper, I have frequently talked about some evaluative attitudes (such as valuing) cohering with other ones. What more can be said about what this relationship of cohering amounts to?

Centrality of Different Valuing. In this paper, I have frequently talked about some valuing being more central or fundamental to an agent’s system of valuing. What more can be said about what it is for this to be the case?

Non-Fundamental Metaethics. In addition to the concept of a NORMATIVE REASON FOR ACTION, there are a whole host of other normative and evaluative concepts that we use in ethics. These include concepts ranging from GOODNESS to MORAL OBLIGATION. What, if anything, can agent-value-subjectivism say about how to understand such concepts? More generally, if

agent-value-subjectivism is meant to be a fundamental metaethical view, are there viable non-fundamental metaethical accounts that would explain how this could be so?

Non-Ethical Metanormative Theory. In addition to using normative and evaluative concepts in ethics, we do so in *many* other domains as well: for instance, in domains that concern epistemology, aesthetics, political philosophy, philosophy of language, and philosophy of science. Can agent-value-subjectivism or some other metaethical agent-attitude-dependent view be used to cover these other domains? If so, how?

The Seeming Intelligibility of Ethical Agent-Attitude-Independence. Even if ethical agent-attitude-dependence is correct, it certainly seems *prima facie* intelligible to hold that it is not. To put it another way, it seems possible to disagree about whether or not ethical agent-attitude-dependence is true without any basic conceptual incoherence or confusion. Yet, metaethical agent-attitude-dependent views such as agent-value-subjectivism imply otherwise. How can metaethical agent-attitude-dependent theorists respond to this objection?

Wildly Counterintuitive Results. There may be good reasons to suppose that many of the agents in our world will have fundamental valuing close to our own – and certainly for thinking that it is unlikely that certain fundamental valuing will be very prevalent (for instance, those favoring throwing oneself under falling boulders at all times). Yet, it is nonetheless conceptually possible – and, indeed, one might hold highly likely as an empirical matter – that there are agents with fundamental valuing quite different from our own. It is thus possible to imagine that there are agents whose systems of valuing favor such things as *poking puppies in the eyes for fun* or *torturing human beings for fun*. Don't metaethical agent-attitude-dependent views such as agent-value-subjectivism deliver wildly counterintuitive results in such cases?

Speaker-Endorsement. Many philosophers are drawn to the idea that when a speaker makes a normative judgment about *another* agent's normative reasons, this involves a crucial element of speaker-endorsement. For instance, many would find it counterintuitive to say “Betsy has a normative reason to ϕ , but I do not support her acting on her reasons. In fact, I think we

should prevent her from doing so.” If agent-value-subjectivism gives up on the idea of speaker-endorsement, what then can be said to explain the force of such intuitions?

Determinacy. The sets of valuings that agents have in the real world might not be mere heaps of valuings. But they certainly are not internally-idealized-valuings. There might be any number of functions for getting from the non-idealized to the idealized set. Will there be enough determinacy in the content and weights of the valuings of an agent’s non-idealized set to insure that there is *one unique set* of internally-idealized-valuing? If not, won’t we be forced to say that agents lack conclusive normative reasons in many cases where we think that they clearly have such normative reasons? Similarly, will there be enough determinacy in the non-idealized set of an agent’s valuings to ensure that there will ever be a stable resting point(s) that would count as the set(s) of an agent’s internally-idealized-valuings? Without answers to such questions, why should we believe that there will be any determinacy to what an agent’s normative reasons are?

Epistemic Access. According to agent-value-subjectivism, the content of an agent’s current valuings determine what that agent’s internally-idealized-valuings are. In turn, those internally-idealized-valuings determine what normative reasons that agent has. Given that an agent currently only has epistemic access to her current valuings – and, moreover, that epistemic access is far from infallible – why should that agent be confident in her epistemic access to her normative reasons? For instance, if her valuings are now only *partly* revised according to what her valuings favor, why think that this will be a good approximation of what her internally-idealized-valuings would look like at the end of the process of revision?

Making It Too Easy to Have Normative Reasons. Many philosophers have taken very seriously the idea that agents in the actual world have no normative reasons. Furthermore, many non-philosophers are often gripped by this worry as well. Insofar as metaethical agent-attitude-dependence makes facts about normative reasons rest on facts of empirical psychology of a sort that almost certainly seem to obtain, doesn’t metaethical agent-attitude-dependence make it too easy to determine that agents in our world do in fact have normative reasons? More generally, doesn’t it also just make it *too easy* for an agent to have a normative reason?

“Just too Different” Intuitions. One reason that philosophers are drawn to non-naturalist positions in metaethics is that they think that ethical thought and talk needs to be sharply distinguished from descriptive thought and talk at the most basic explanatory level. Sometimes this is put forward as a thought about ethical properties – for instance, David Enoch puts forward the idea that normative properties are just “too different” from naturalistic properties for naturalistic realism to be true.¹²¹ In other cases, this is put forward as a thought about ethical concepts – for instance, when Allan Gibbard claims that normative concepts are fundamentally different in kind from naturalistic concepts. How should metaethical agent-attitude-dependent theorists respond to such thoughts? More generally, what account of the seeming so-called “is/ought gap” can one give in response to those who take this gap to point squarely in favor of non-naturalist metaethical positions?

Normative Concepts. In this paper, I have put to one side a number of important issues about the basic nature of normative concepts that are crucial for a further development of agent-value-subjectivism. These include, for example, the issue of whether or not a normative concept is *identical* to a descriptive concept that picks out the same normative properties in all possible worlds. What are the general views we should have about the individuation of one concept from another? How do such views impact the plausibility of agent-value-subjectivism?

As I see it, these are some of the main issues that confront the further development of agent-value-subjectivism. In assessing these issues – and thus, in assessing the promise of agent-value-subjectivism – the crucial question will *not* be whether there are answers to all of these questions that everyone will find satisfying. Rather, the crucial question is what trade-offs the view needs to make in order to make itself viable – and, more specifically, how those trade-offs compare to those made by other competing metaethical views.

¹²¹ This is what Enoch calls the “just too different” intuition. (Enoch 2007, 244 n.47).

WORKS CITED

- Bennett, Karen. 2008. Exclusion, Again. In *Being Reduced*, edited by J. Hohwy and J. Kallestrup. Oxford: Oxford University Press.
- Blackburn, Simon. 1993. *Essays in Quasi-Realism*. Oxford: Oxford University Press.
- . 1998. *Ruling Passions*. Oxford: Clarendon Press.
- Block, Ned, and Robert Stalnaker. 1999. Conceptual Analysis, Dualism, and the Explanatory Gap. *Philosophical Review* 108:1-46.
- Boghossian, Paul Artin. 2006. *Fear of knowledge : against relativism and constructivism*. Oxford: Clarendon Press ; Oxford University Press.
- Boyd, Richard. 1997. How to be a Moral Realist. In *Moral discourse and practice : some philosophical approaches*, edited by S. L. Darwall, A. Gibbard and P. A. Railton. New York: Oxford University Press.
- Braddon-Mitchell, David, and Frank Jackson. 1996. *The philosophy of mind and cognition*. Oxford, UK ; Cambridge, MA: Blackwell Publishers.
- Brandom, Robert. 1994. *Making it explicit : reasoning, representing, and discursive commitment*. Cambridge, Mass.: Harvard University Press.
- Chalmers, David J. 1996. *The conscious mind : in search of a fundamental theory*. New York: Oxford University Press.
- . 2006. Two-Dimensional Semantics. In *Oxford Handbook of the Philosophy of Language*, edited by E. Lepore and B. Smith. Oxford: Oxford University Press.
- Chalmers, David J., and Frank Jackson. 2001. Conceptual Analysis and Reductive Explanation. *Philosophical Review* (110):315-61.
- Dancy, Jonathan. 2003. *Practical Reality*. Oxford: Oxford University Press.
- . 2006. Nonnaturalism. In *The Oxford Handbook of Ethical Theory*, edited by D. Copp. Oxford: Oxford University Press.
- Darwall, Stephen L. 1997. Reasons, Motives, and the Demands of Morality: An Introduction. In *Moral discourse and practice : some philosophical approaches*, edited by S. L. Darwall, A. Gibbard and P. A. Railton. New York: Oxford University Press.
- . 2003. Desires, Reasons, and Causes. *Philosophy and Phenomenological Research* (67):436-443.
- . 2006. *The second-person standpoint : morality, respect, and accountability*. Cambridge, Mass.: Harvard University Press.
- Darwall, Stephen L., Allan Gibbard, and Peter Albert Railton. 1997. Toward Fin de siecle Ethics: Some Trends. In *Moral discourse and practice : some philosophical approaches*. New York: Oxford University Press.
- Davidson, Donald. 2001. *Inquiries into truth and interpretation*. 2nd ed. Oxford: Clarendon Press.
- Egan, Andy. 2006. Epistemic Modals, Relativism, and Assertion. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 133 (1):1-22.
- Enoch, David. 2007. An Outline of an Argument for Robust Metanormative Realism In *Oxford Studies in Metaethics Volume 2*, edited by R. Shafer-Landau.
- . 2009. Can There Be a Global, Interesting, Coherent Constructivism about Practical Reason? *Philosophical Explorations: An International Journal for the Philosophy of Mind and Action* 12 (3):319-339.
- . Forthcoming. The Epistemological Challenge Metanormative Realism: How Best to Understand It, and How to Cope with It. *Philosophical Studies*.
- Gibbard, Allan. 1990. *Wise Choices, Apt Feelings : A Theory of Normative Judgment*. Cambridge, Mass.: Harvard University Press.

- . 1992. Reply to Blackburn, Carson, Hill, and Railton. *Philosophy and Phenomenological Research* 52 (4):969-980.
- . 1999. Morality As Consistency in Living: Korsgaard's Kantian Lectures. *Ethics: An International Journal of Social, Political, and Legal Philosophy* 110 (1):140-164.
- . 2003. *Thinking How to Live*. Cambridge, Mass.: Harvard University Press.
- . 2008. *Reconciling our aims : in search of bases for ethics*. Oxford ; New York: Oxford University Press.
- . Forthcoming. Evolved Thinkers and Moral Concepts In *Oxford Studies in Metaethics*, edited by R. Shafer-Landau. Oxford: Oxford University Press.
- Harman, Gilbert. 1984. Logic and Reasoning. *Synthese: An International Journal for Epistemology, Methodology and Philosophy of Science* 60:107-128.
- Hussain, Nadeem J. Z. 2004. The Return of Moral Fictionalism. *Nous-Supplement: Philosophical Perspectives* 18:149-187.
- . Forthcoming. *Norms in Action: Metaethics and the Neo-Kantian Critique*.
- Hussain, Nadeem J. Z., and Nishi Shah. 2006. Misunderstanding Metaethics: Korsgaard's Rejection of Realism. In *Oxford Studies in Metaethics, Volume 1*, edited by R. Shafer-Landau.
- Jackson, Frank. 1996. Mental Causation: The State of the Art. *Mind* 105:377-413.
- . 1998. *From metaphysics to ethics : a defence of conceptual analysis*. Oxford: Clarendon.
- . 2001. Representation, Scepticism, and the A Priori. In *New Essays on the A Priori*, edited by P. Boghossian and C. Peacocke. Oxford: Oxford University Press.
- Joyce, Richard. 2001. *The Myth of Morality*. Cambridge: Cambridge University Press.
- Korsgaard, Christine M. 1996. *The sources of normativity*. Cambridge ; New York: Cambridge University Press.
- . 2003. Realism and Constructivism in Twentieth-Century Moral Philosophy. In *Philosophy in America at the Turn of the Century (APA Centennial Supplement Journal of Philosophical Research)*. Charlottesville: Philosophy Documentation Center.
- MacFarlane, John. 2007. Relativism and Disagreement. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 132 (1):17-31.
- McDowell, John Henry. 1994. *Mind and world*. Cambridge, Mass.: Harvard University Press.
- McPherson, Tristram. Forthcoming. Against Quietist Normative Realism. *Philosophical Studies*
- Nye, Howard. 2009. *Ethics, Fitting Attitudes, and Practical Reason: A Theory of Normative Facts*: PhD Dissertation, University of Michigan, Ann Arbor.
- Parfit, Derek. 2006. Normativity. In *Oxford Studies in Metaethics* edited by R. Shafer-Landau. New York: Oxford.
- . Forthcoming. *On What Matters*. Vol. (Draft of 7/5/2009). Oxford: Oxford University Press.
- Railton, Peter. 2001. A Priori Rules: Wittgenstein on the Normativity of Logic. In *New Essays on the A Priori*, edited by P. Boghossian and C. Peacocke. Oxford: Oxford University Press.
- . 2003a. Aesthetic Value, Moral Value, and the Ambitions of Naturalism. In *Facts, Values, and Norms: Essays toward a Morality of Consequence*. Cambridge: Cambridge Univ Pr.
- . 2003b. *Facts, values, and norms : essays toward a morality of consequence*. Cambridge, U.K. ; New York: Cambridge University Press.
- . 2003c. Noncognitivism about Rationality: Benefits, Costs, and an Alternative. In *Facts, Values, and Norms*. Cambridge: Cambridge University Press.

- Railton, Peter, and Gideon Rosen. 2009. Realism. In *Blackwell Companion to Metaphysics*, edited by J. Kim, E. Sosa and G. S. Rosenkrantz. New York: Blackwell.
- Rawls, John. 1999. The Independence of Moral Theory. In *Collected papers*. Cambridge, Mass. ; London: Harvard University Press.
- Scanlon, T. M. 2000. *What We Owe To Each Other*. Cambridge: Harvard University Press.
- . Forthcoming. *Being Realistic About Reasons: The 2009 John Locke Lectures*. Oxford.
- Schroeder, Mark. 2007. *Slaves of the Passions*. Oxford: Oxford University Press.
- Shafer-Landau, Russ. 2005. *Moral Realism: A Defense*. Oxford: Oxford University Press.
- Smith, Michael. 1994. *The Moral Problem*. Cambridge: Blackwell.
- . 2004. Moral Realism. In *Ethics and the A Priori: Selected Essays on Moral Psychology and Meta-Ethics*. Cambridge: Cambridge University Press.
- Stalnaker, Robert. 1987. Semantics for Belief. *Philosophical Topics* 15:177-190.
- Street, Sharon. 2006. A Darwinian Dilemma for Realist Theories of Value. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 127 (1):109-166.
- . 2008a. Constructivism about Reasons. In *Oxford studies in metaethics*, edited by R. Shafer-Landau. Oxford: Oxford University Press.
- . 2008b. Reply to Copp: Naturalism, Normativity, and the Varieties of Realism Worth Worrying About *Philosophical Issues* 18:207-228.
- . 2009. In Defense of Future Tuesday Indifference: Ideally Coherent Eccentrics and the Contingency of What Matters. *Philosophical Issues* 19:273-298.
- . Forthcoming-a. Evolution and the Normativity of Epistemic Reasons. ed. Russ Shafer-Landau (Oxford: Clarendon Press, 2011). In *Oxford Studies in Metaethics*, edited by R. Shafer-Landau. Oxford Oxford University Press.
- . Forthcoming-b. Mind-Independence Without the Mystery: Why Quasi-Realists Can't Have It Both Ways. *Noûs* <http://homepages.nyu.edu/~jrs477/streetpapers.html> (Draft of 5/1/2010).
- . Forthcoming-c. Objectivity and Truth: You'd Better Rethink It. *Philosophy and Public Affairs* <http://homepages.nyu.edu/~jrs477/streetpapers.html> (Draft of 5/1/2010).
- . Forthcoming-d. What is Constructivism in Ethics and Metaethics? . *Philosophy Compass* <http://homepages.nyu.edu/~jrs477/streetpapers.html> (Draft of 5/1/2010).
- . Manuscript. Coming to Terms with Contingency: Humean Constructivism about Practical Reason.
- Sturgeon, Nicholas L. 2006. Ethical Naturalism. In *The Oxford Handbook of Ethical Theory*, edited by D. Copp. Oxford: Oxford University Press.
- Sturgeon, Nicholas L. 1988. Moral Explanations. In *Essays on Moral Realism*. Ithaca: Cornell University Press.
- Velleman, J. David. 2009. *How We Get Along*. Cambridge: Cambridge University Press.
- Wedgwood, Ralph. 2009. *The Nature of Normativity*. Oxford: Oxford University Press.
- Williams, Bernard Arthur Owen. 1981. Internal and External Reasons. In *Moral luck : philosophical papers, 1973-1980*. Cambridge Cambridgeshire ; New York: Cambridge University Press.
- . 1985. *Ethics and the limits of philosophy*. Cambridge, Mass.: Harvard University Press.
- Wright, Crispin. 1994. *Truth and Objectivity*. Cambridge: Harvard University Press.