

METAETHICS FOR NEO-PRAGMATISTS:
A PRAGMATIC ACCOUNT OF LINGUISTIC MEANING FOR
MORAL VOCABULARY

by
Thomas Michael Wilk

A dissertation submitted to Johns Hopkins University in conformity with the
requirements for the degree of Doctor of Philosophy

Baltimore, Maryland
October 2019

© 2019 Thomas Wilk
All rights reserved

Abstract

In this dissertation, I aim to develop and defend a novel, pragmatist approach to foundational questions about meaning, especially the meaning of deontic moral vocabulary. Drawing from expressivists and inferentialists, I argue that meaning is best explained by the various kinds of norms that govern the use of a vocabulary. Along with inferential norms, I argue we must extend our account to discursive norms that govern normative statuses required to felicitously utter certain speech-acts—*norms of authority*—and the transitions in normative statuses affected by speech-acts—*pragmatic norms*. These *structural discursive norms* differentiate discursive practices and account for distinctive features such as objectivity and motivational import that some have and others lack. The structure exhibited by a practice is then explained in terms of its utility, making it possible to see how different discursive practices are answerable to the different needs and purposes of the discursive beings who use them. I call the resulting explanatory framework a *pragmatic analysis of linguistic meaning* (PALM).

Turning my attention to moral “ought,” I argue that the structural discursive norms of moral discourse differentiate it from other objective discourses, like empirical discourse, on the one hand, and from other normative discourses, like prudential discourse, on the other. Drawing on work in evolutionary psychology and anthropology, I complete the PALM with an account of moral discourse a meta-normative practice with a meta-coordinative function. Its utility for the discursive beings who use it lies in its enabling them to remedy certain tensions and instabilities that arise in their other coordinative, normative practices in a way that minimizes the risk of domination by alpha-type free-riders, the fracturing of social groups, and individual defection from cooperative endeavors.

In the final two chapters, I leverage the account to defend a pragmatist-friendly notion of objectivity in terms of a structure of distributed epistemic authority according to which no claim within a practice is ultimately authoritative or immune from challenge and to reconcile this sense of objectivity with a persistent pressure toward a kind of relativism that restricts the standing to make moral claims to members of the relevant communities.

First Reader and Advisor: Michael Williams

Internal Readers: Hilary Bok, Dean Moyar

External Readers: Mark Lance (Philosophy, Georgetown University), William Egginton (German and Romance Languages and Literature, Johns Hopkins University)

Acknowledgements

My path to the completion of this dissertation was circuitous, and I could easily have gotten lost without the wonderful people who encouraged, supported, and challenged me along the way. I owe a great debt of gratitude to my advisor, Michael Williams, for the many hours he spent reading drafts and discussing the ideas in this dissertation. His extensive philosophical knowledge has been instrumental in shaping my thinking on the topics explored here, and his exacting expectations guided me to a much stronger dissertation than I otherwise might have produced. I would also like to extend my deepest gratitude to Mark Lance, who, many years ago, inspired my interest in the contemporary pragmatist approach I take up in this work. His feedback and encouragement along the way have been exceedingly valuable. My sincere thanks also goes to the members of my committee—Hilary Bok, Dean Moyer, and Bill Egginton—and to Steven Gross for the time they spent reading this work and the valuable suggestions they offered about its future development.

The Philosophy Department at Gettysburg College also deserves my gratitude. This is where I first encountered philosophy, and Steve Gimbel, Jennifer Hansen (now at St. Lawrence University), Kerry Walters, Dan DeNicola, and Lou Hammann have been unwavering in their support, inviting me back to teach and offering encouragement when I needed it most. Special thanks goes to Lisa Portmess for keeping me accountable when I was struggling to produce drafts of some of the central chapters and for inviting me to present some of this work to a group of Gettysburg students.

Besides Gettysburg, I've taught in philosophy departments at Goucher College, the U. S. Naval Academy, George Washington University, and George Mason University. I wish to extend my thanks to all of them for their support, both emotional and financial, but George Mason deserves special mention. Since I completed my Master's degree there, Ted Kinnamann has been a friend and mentor, and Rachel Jones has been supportive in arranging a teaching schedule that allowed me the time to complete this work. I am deeply grateful to them both.

I'd also like to extend my thanks to our department administrators, Alicia Burley, Veronica Feldkircher-Reed, and Leslie Bean, for many conversations, help in navigating institutional hurdles, and emotional support over the years; to our Research Librarian, Heidi Herr, for never once turning down a book purchase request; to my fellow Johns Hopkins graduate students, especially Josh McBee, for valuable support and feedback along the way, and to the Virginia Theological Seminary for the use of quiet library space so close to my home that facilitated the final stages of writing.

My successful completion of this project would not have been possible without the constant support of my family. My parents, Robert and Rose Ann Wilk, never

doubted me even when I doubted myself, and I cannot thank them enough for supporting and encouraging me in every possible way. I also thank my in-laws, Marty and Tom Mereen, for being a constant source of encouragement. My children, Sam and Jack, have never known a time I wasn't working on this dissertation, and I thank them for putting up with Daddy's weekends of writing, for the many hugs and distractions, and for the examples they've provided that appear in these pages.

Above all others, I owe my greatest thanks to my partner, Kate Mereen. Without Kate's constant support and companionship, I simply would not have been able to complete this dissertation. In many ways, she's worked harder than I have for this, supporting our family financially and emotionally, providing time for me to read and write, keeping my head on straight and my eye on the prize, and never once doubting that I could succeed. I would not be the person I am today without her. She is, truly, my rock.

To Katy. So much.

Think of the tools in a tool-box: there is a hammer, pliers, a saw, a screw-driver, a ruler, a glue-pot, glue, nails and screw.—The functions of words are as diverse as the functions of these objects. (And in both cases there are similarities.)

-Ludwig Wittgenstein, *Philosophical Investigations*

Table of Contents

Abstract.....	ii
Acknowledgements.....	iii
Table of Contents.....	vii
List of Figures.....	x
Introduction.....	1
1. Two Aims.....	1
2. Neo-Pragmatism.....	3
3. The Problems of Contemporary Metaethics.....	5
4. Global Anti-Representationalism and Inferentialism.....	10
5. The Return to Metaethics.....	13
6. The Objectivity-Relativism Tension.....	14
Chapter 1: Anti-Representationalist Metaethics: From Expressivism to Inferentialism ..	18
1. Introduction.....	18
2. Traditional Expressivism.....	20
3. The Frege-Geach Problem for Traditional Expressivism.....	28
3.1. Mentalism and the Problem of Composition.....	30
3.2. Logical Inconsistency and Revisionist Semantics.....	34
4. Orthodox Inferentialism.....	42
4.1. Inferentialism and the Frege-Geach Problem.....	45
5. Warren's Hybrid Expressivist Inferentialism.....	49
6. Brandom on the Explicative Function of Deontic Normative Vocabulary.....	57
7. Brandom on the Flavors of Ought.....	61
8. Conclusion.....	64
Chapter 2: Refining Inferentialism: From EMU to PALM.....	66
1. Introduction.....	66
2. The Plurality of Discursive Norms.....	70
3. Structural Discursive Norms.....	76
4. Speech Acts and Normative Functions.....	95
5. EMUs and PALMs.....	102

6. Exploring PALM.....	112
Chapter 3: A PALM for Moral "Ought"	117
1. Introduction	117
2. Moral Discourse in Context	117
3. The Pragmatic Structural Norms of Moral Discourse.....	124
4. The Authority Structure of Moral Discourse	134
4.1. The Agent-Neutral Epistemic Norms of Moral Discourse.....	135
4.2. The Agent-Relative Authority Norms of Moral Discourse.....	151
5. Conclusion.....	153
Chapter 4: The Function of Moral Discourse	156
1. Introduction	156
2. Practices and Meta-Practices.....	157
3. The Evolution of Human Cooperation and Moral Sentiments.....	165
4. The Cognitive Structure of Joint Intentional Action.....	171
5. From Cooperative Dyads to Cooperative Communities	178
6. The Free-Rider Objection.....	185
7. The Function of Moral Discourse	193
8. Conclusion.....	202
Chapter 5: Moral Objectivity	205
1. Introduction	205
2. The Traditional Notion of Objectivity	207
3. Objectivity Challenged.....	211
4. Rehabilitating Objectivity?	217
5. Brandom on Objectivity	221
6. Sorting Discursive Practices	231
7. Rorty on the Mental.....	234
8. Moral Objectivity	242
9. Religion and Morality	253
10. Conclusion.....	258
Chapter 6: Discursive Communities, Trust, and the Norms of Standing	261
1. Introduction	261
2. Norms of Holding.....	266

3. The Topography of Holding Responsible	268
4. Who Has the Standing to Hold?	273
5. Agency Threats and the Standing to Hold	278
6. Trust and the Standing to Hold	281
7. Trust and Divided Communities	284
8. From Standing to Hold to Standing to Assert	287
9. Earning Trust.....	292
10. Conclusion.....	294
Conclusion	299
Bibliography	303
Curriculum Vitae	324

List of Figures

Figure 1: Pragmatic Structure of Speech Acts	99
Figure 2: PALM Tree	111
Figure 3: Pragmatic Structure of Moral Prescriptives	134

Introduction

Philosophers of different persuasions are hypnotized by different pictures. *Literally* pictures, little diagrams that they draw in the margins of their manuscripts...

—Wilfrid Sellars,
in Amaral's Introduction to *The Metaphysics of Epistemology*

1. Two Aims

I have two mutually supporting aims in this dissertation, one situated in the philosophy of language and the other in metaethics. On the side of philosophy of language, I advance a framework for theorizing about the meanings of different kinds of vocabularies, e.g., empirical, prudential, and moral vocabularies. Drawing lessons from both expressivism and inferentialism, I develop a broadly pragmatist theoretical framework for what I call the *pragmatic analysis of linguistic meaning* (PALM). Like inferentialism, PALM treats inferential proprieties as constitutive of conceptual content. It advances the inferentialist project by attending to kinds of discursive norms that are often ignored by inferentialists but that, I'll claim, account for the distinctive features exhibited by different vocabularies, such as objectivity and motivational import that some vocabularies exhibit and others do not. I call these *structural discursive norms*. PALM then explains why a practice has the normative shape that it does in terms of its function, making it possible to see how different discursive practices are answerable to the different needs and purposes of the discursive beings who use them.

On the metaethical side, my aim is to construct a neo-pragmatist metaethics by bringing PALM to bear on moral vocabulary, particularly on moral “ought.” The resulting account captures the conceptual content of moral vocabulary in broadly inferential terms and explains the discursive structural norms that underly the inferential patterns associated with the objectivity and motivational import of moral discourse in terms of its being a meta-normative vocabulary with a meta-coordinative function. Moral discourse, I’ll argue, serves to remedy certain tensions and instabilities that arise in our other coordinative practices in a way that minimizes the risks of domination by alpha-type free-riders, the fracturing of social groups, and individual defection. This function necessitates a practice with various structural norms that distinguish it from empirical or other fact-stating discourse on the one hand and from other flavors of normative discourse— institutional and prudential “oughts,” for example—on the other.

If the arguments for PALM are successful, they lend credence to the results we derive in applying the framework to moral discourse. On the other hand, if the account of moral discourse we arrive at successfully solves or dissolves some of the persistent problems of metaethics, that’s a feather in PALM’s cap. Indeed, this is what we shall find. I’ll argue that the PALM for moral discourse avoids entangling us in ontological concerns about moral realism and anti-realism, escapes the pitfalls of traditional expressivism (especially of the Frege-Geach variety), and contains resources for better understanding the relationship between moral vocabulary and motivation. Perhaps more surprisingly for a neo-pragmatist account, I shall also argue that PALM furnishes an account of objectivity that makes moral discourse out to be fully objective but also equips us with the resources to see why our moral practices will always exhibit relativistic tendencies

and why the tension between objectivity and relativism is one we cannot escape so long as we continue leveraging the tools of moral discourse. Given the fecundity of the PALM for moral discourse, I conclude that PALM itself is a useful approach to theorizing meaning.

2. Neo-Pragmatism

With these aims in mind, I turn in the remainder of this introduction to drawing out some of the central themes and motivations of my project. PALM is an avowedly neo-pragmatist account of meaning, so I'll start by characterizing what I mean by neo-pragmatism.¹ This philosophical tradition can be understood in terms of three commitments. First, influenced by the linguistic turn in twentieth century analytic philosophy, neo-pragmatists are committed to the priority of linguistic practices. When a neo-pragmatist perceives what might be taken to be a metaphysical issue, her commitment to *linguistic priority* means that she will begin by investigating what is distinctive about the linguistic practice from which the issue arises rather than the nature of the objects or properties that practice purports to be about. She will attend to what we're *doing* when we use modal language rather than to modal properties and ask what is distinctive of evaluative discourse rather than investigate the nature of values.

Taking their cue from classical pragmatists like James and Dewey, the neo-pragmatists' second commitment is *anti-Representationalism*. This is best understood as a commitment regarding the order of explanation. Representationalists explain the proper

¹ Neo-pragmatists are themselves a varied group, but I would include among some of the more prominent figures in the relatively young tradition Richard Rorty, Hilary Putnam, Robert Brandom, Huw Price, Michael Williams, Mark Lance, Rebecca Kukla, and Susan Haack. Quine, Sellars, and Davidson, if not neo-pragmatists, at least exert strong influences over its proponents.

use of a vocabulary item in terms of its meaning and meaning in terms of semantic relations like reference and truth-value. Anti-Representationalism turns this on its head. The properties of a vocabulary's use constitute (and, in a way, explain) its meaning, and these properties are themselves given a functional explanation. They might be explained, for example, in terms of what the vocabulary item allows speakers to do that they otherwise could not or in terms of the kinds of mental states they are used to express.²

The third neo-pragmatist commitment is to *deflationary* or *minimalist* accounts of truth (and, indeed, to semantic minimalism more broadly).³ This is the point over which contemporary neo-pragmatists break with their classical counterparts. While Peirce, Dewey, and James were drawn to epistemic accounts of truth that eschewed correspondence in favor of utility, neo-pragmatists have come to see truth as nothing more than a linguistic device useful for endorsing the claims of others, generalizing over claims you could not otherwise state, or, in some cases, to express our fallibility (Rorty's *cautionary* use (Rorty & Price, 2010)). All that we can say in a theory of truth is that our use "of the truth predicate is fully captured by our commitment to the non-paradoxical instances of some appropriate equivalence schema: for example: (DQT) 'P' is true if and only if P" (Williams, 2013, p. 129). One is entitled to append the truth predicate to any proposition to which one is entitled.⁴

² For a characterization of neo-pragmatism in terms of the two commitments just discussed, see (Macarthur & Price, 2007).

³ Minimalism comes in a number of varieties, but I'll neither distinguish them nor endorse one of them here. See, for example, the redundancy theory (Ramsey, 1927), disquotationalism (Field, 1994; Horwich, 1998b, 1998a; Leeds, 1978; Quine, 1986), propositionalism (Horwich, 1998b), and prosententialism (Båve, 2009; Brandom, 1994; Grover, Camp, Jr., & Belnap, 1975; Lance, 1997).

⁴ This characterization of neo-pragmatism in terms of three commitments is due to Michael Williams, (2013, pp. 128–129).

Neo-pragmatists tell us that when we encounter a philosophical puzzle, we should begin by investigating the vocabulary that seems to be used to talk about the puzzling phenomenon, and we should ask what the users of that vocabulary are using it to do. Minimalism then comes in to block the Representationalist response that they're using it to represent how things are in the world, it prevents us from sliding down the semantic ladder from a vocabulary to the 'objects' it 'purports' to be about. PALM is neo-pragmatist in endorsing all three of these commitments; it is an attempt to work out a function-first explanation of meaning.

3. The Problems of Contemporary Metaethics

Neo-pragmatism makes contact with contemporary metaethics via the latter's broad acceptance of Representationalism. The basic Representationalist assumption is that "the *linguistic* items [of interest] 'stand for' or 'represent' something *non-linguistic*," so that the task of giving a semantics for a language is the task of systematically associating terms with their non-linguistic referents and specifying the rules for stringing them together into sentences (Price, 2013, p. 9).⁵ This means that we need to go looking for those non-linguistic items. Even if we begin by examining vocabularies we'll end up doing ontology (Price, 2013, pp. 8–10). And this is just what contemporary metaethicists have done. They begin with reflection on ordinary moral discourse and practice, investigating the semantics of moral vocabulary, the purported metaphysics of moral properties or values, the epistemic norms that seem to govern claims to moral knowledge, and the link between morality and practical reasoning or motivation. They aim to characterize this discourse and practice along these dimensions in order to understand

⁵ Also see, (Chrisman, 2015, p. 126)

what the practitioners take the realm of the moral to be like. They want to understand what commitments—especially ontological commitments—are undertaken from within the practice. Once identified, they ask whether those commitments could possibly be made good. Stepping back from moral discourse and practice and turning their attention to the broader scope of human affairs and knowledge, they ask how these moral commitments fit in. In particular, can they be countenanced by the naturalistic picture we get from the empirical sciences? Where can the objects or properties to the existence of which practitioners are committed be “placed” in this landscape? If they cannot find the right place for them, then moral discourse and practice may be in trouble; practitioners may not be able to vindicate their commitments from this broader perspective, and, so, metaethicists begin to worry that these commitments could never be fully vindicated (Darwall, Gibbard, & Railton, 1992, pp. 127–128; Price, 2013, pp. 6–8, 23–26). Something must be amiss with the practice itself.

The frenzied search for truth-makers that results from these *placement problems*⁶ is responsible in large part for the extant landscape of metaethics. Naturalistic moral realists, for starters, claim to have discovered the natural properties that moral vocabulary purports to be about, thus vindicating the commitments of the moral discursive practitioners within the Representationalist framework.⁷ Naturalistic realism, though, faces serious challenges. Sharon Street has argued, for example, that natural selection has had a “purely distorting” effect on our evaluative judgments, so there is no reason to

⁶ The term is Frank Jackson's, but Price has adopted it, (Jackson, 1997; Price, 2013, p. 5) It's worth noting that it is not only moral or value properties that give rise to placement problems. We can ask similar questions about mental, modal, mathematical, and meaning properties (the M-worlds), as well. On these four M-worlds and the problems we face in understanding them, see (Jackson, 1997; Price, 1997).

⁷ See, (Boyd, 1988; Brink, 1989; Jackson, 1998; Sayre-McCord, 1988).

think they would have evolved to accurately represent any state of affairs (Street, 2006). Others have advanced Moore's open-question argument that for any natural property *N*, which the moral realist claims as the reduction base of a moral property, say *goodness*, it makes sense to ask whether *N* is *really* good (Moore, 1903).⁸ What seems to be missing from the naturalistic realist account is a recognition that "[v]alue terms have a special function, that of commending; and so they plainly cannot be defined in terms of other words which do not perform this function; for if this is done we are deprived of the means of performing this function" (Hare, 1952, p. 91).

If naturalistic realism won't work, another option is to adopt something like the non-naturalist realism espoused by Moore. If, however, you don't want to countenance non-natural entities, you might be more drawn to moral error theory. Mackie has famously argued that the truth-makers for moral claims would have to be "qualities or relations of a very strange sort, utterly different from anything else in the universe" (Mackie, 1977).⁹ They would need to have a "categorically imperative element [that] is objectively valid," i.e., they would have to be binding on individuals without respect for their preferences and they would have to be such that pointing them out would convince individuals that this is the case. Since no "queer" properties like this exist, moral claims must be universally erroneous or, at best, true in some fiction.¹⁰ The problem with both the error theory and fictionalism is that neither can save the phenomena. They both require us to give up on some pretty potent intuitions about what moral practices are like,

⁸ For a taste of the contemporary debate see, for example, (Brink, 1989, 2001; Horgan & Timmons, 1991, 1992; Lenman, 2014).

⁹ Also see, (Joyce, 2015).

¹⁰ For attempts to work out a fictionalist account in metaethics and some of its challenges, see (Eklund, 2009; Joyce, 2007; Kalderon, 2007, 2008).

for example, that moral claims are really true or false, that there are objectively right or wrong answers to moral questions, that we have real obligations not just in some “as-if” sense, and that moral discourse is a kind of rational discourse.

Other positions in contemporary metaethics define themselves primarily in opposition to these Representationalist views, espousing one or another *locally* anti-Representationalist position. These views go by labels like emotivism¹¹, non-cognitivism¹², expressivism¹³, prescriptivism¹⁴, projectivism, and quasi-realism.^{15, 16} Because they share anti-Representationalist commitments, neo-pragmatists have been drawn to positions like these, but, as Price has argued, it’s far from clear that they have a right to them (Macarthur & Price, 2007; Price, 2013).¹⁷ These local anti-Representationalist accounts are each committed to a bifurcation thesis. For ordinary descriptive uses of language, they accept the Representationalist explanation of meaning, but for the thorny cases like moral, modal, mental, and mathematical discourse, they offer

¹¹ The classic examples of this view are (Ayer, 1952; Barnes, 1934; Stevenson, 1972). Emotivists are committed to the view that that moral claims are not speech acts of assertion but rather something more like an interjection, a boo or hooray.

¹² This is an umbrella term for all those positions that reject the idea that moral mental states are belief like. Rather, they are directive, more like desires.

¹³ Expressivism has been most ably set out by Gibbard, (Gibbard, 1990, 2003). Expressivists claim that moral assertions are really speech acts of assertion but that they serve to express non-cognitive attitudes. Constraints of consistency are imposed on their use by the relations in which those attitudes stand to one another.

¹⁴ See, (Carnap, 1935, pp. 23–24 & 29; Hare, 1952). The prescriptivist sees moral assertions as disguised prescriptions so that uttering "Murder is wrong" is really a way of saying "Don't murder." Hare advanced the position by arguing that moral claims are akin to universal prescriptions that stand in logical relations to one another. This imposes demands of consistency on their use and aims to vindicate the rationality of moral discourse.

¹⁵ The final two positions are espoused by Blackburn, (Blackburn, 1984, 1988, 1993, 2001). Projectivism is a commitment to the Humean metaphysical view that moral properties do not exist in the world itself but are projections of our moral attitudes or evaluations. Quasi-realism builds on this position (which is akin to expressivism) by working to earn the right to use realist-sounding language to express moral attitudes.

¹⁶ Non-cognitivist positions such as these are not limited to accounts of moral discourse. Sellars has espoused the position with respect to causal modalities (Sellars, 1957), Ramsey with respect to laws, causation, and probability (Ramsey, 1927), Wittgenstein with respect to necessity (Wittgenstein, 1953), and Blackburn with respect to probability, causation, and necessity (Blackburn, 1993).

¹⁷ Also see, (Williams, 2013).

a version of anti-Representationalism (Kraut, 1990). This generates two problems for neo-pragmatists who want to coopt these programs. The first arises out of their commitment to minimalism, which makes it too easy for any assertoric discourse to be truth-apt. This “creeping minimalism” threatens the bifurcation thesis, undermining the local anti-Representationalist’s ability to sort descriptive and non-descriptive vocabularies (Dreier, 2004).¹⁸ Once we accept minimalism, it’s difficult to see what distinguishes the expressivist and the realist. Since neo-pragmatists will want to maintain anti-Representationalism, the upshot is pressure toward a global version of this view (Price, 2013).¹⁹ The second problem for neo-pragmatists is that these local anti-Representationalisms are still committed to the Representationalist order of explanation. Expressivists, for example, want to offer an account of the meaning of moral vocabulary in terms of the mental states that it expresses and then to leverage that account to explain features of its use. This is the only way they can maintain logical relations between their target vocabularies and those that get the full Representationalist treatment. In a funny way, expressive use explains meaning and meaning explains the patterns of use exhibited in discourse. I develop this problem in greater detail in Chapter 1. From the neo-pragmatist perspective, this gets things the wrong way around. Patterns of use are constitutive of meaning, not explained in terms of it. In the first half of Chapter 1, I argue that even if we follow Price in endorsing global anti-Representationalism, we cannot do

¹⁸ I want to be clear that the problem here is not directly a problem for the expressivist’s positive program. Rather, the problem is that the Representationalist’s loss of monopoly on truth-aptness robs non-cognitivism of its initial motivation. This does not rob its explanatory strategy of its plausibility but, rather, pushes toward a more global anti-Representationalism. There have also been attempts to save the bifurcation thesis. See, for example, (Blackburn, 2006; Williams, 2013).

¹⁹ For a careful treatment of the local/global distinction among anti-Representationalists, see (Simpson, 2017).

so while advancing an account of meaning modeled on expressivism. A neo-pragmatist metaethics should instead be grounded in inferentialism.

4. Global Anti-Representationalism and Inferentialism

The metaethical motivation of this project is rooted in my dissatisfaction with the extant positions in contemporary metaethics, which force us to abandon plausible intuitions about moral vocabulary. The solution I propose is rooted in a more thoroughgoing pragmatism. In the second half of Chapter 1, I examine the prospects of inferentialism as an account from which to advance a globally anti-Representationalist metaethics. I argue that they are good. Inferentialism is already a well-worked out anti-Representationalist account of meaning and it has ready answers to some of the central challenges to expressivism, e.g., the Frege-Geach problem about the contribution of non-representational bits of vocabulary in embedded contexts. Moreover, it consistently adheres to the neo-pragmatist's favored order of explanation. The extant inferentialist views, however, must be framed in the right way. If they hew too closely to traditional expressivism, as does Mark Warren's account, they will fail to leverage inferentialist resources for understanding moral discourse as a rational discourse. If neo-pragmatists adopt something closer to Robert Brandom's inferentialist program, they'll do better, but they'll face a problem grasping a particularly important kind of *functional pluralism* exhibited by discursive practices.

Neo-pragmatists are committed to understanding different vocabularies in terms of their different functions, but Brandom's inferentialism offers only limited resources for doing this. He can recognize a kind of pluralism in terms of the different explicative

functions different assertoric vocabularies play.²⁰ For example, the conditional, on his account, functions to explicate doxastic inferential proprieties already implicit in ordinary descriptive vocabulary, and normative vocabulary explicates practical inferential proprieties that are implicit in our use of other vocabularies (Brandom, 2008a). What he cannot do, though, is distinguish in more fine-grained ways between different flavors of modality or normativity—say, between moral and prudential vocabularies—nor can he tie these distinctions down to the “different needs and purposes” that these vocabularies function to fulfill, “their different origins in our complex natures and relations to our physical and social environments” (Price, 2010, p. 287). This is the kind of functional pluralism that PALM aims to capture. In Chapter 2, I introduce the resources for doing this, arguing that we must expand our account of discursive norms from strictly inferential norms to also include discursive structural norms that I label *norms of authority* and *pragmatic norms*. These are norms of assessment that practitioners both institute and rely on in the practice of keeping discursive score on their interlocutors, and they serve to shape the broadly inferential proprieties of a practice. The norms of authority of a discursive practice are those that delineate what is necessary to secure entitlement to a speech-act. Some of these are *epistemic structural norms*, but others are norms of *social-institutional standing* that, for example, make it appropriate for some people but not others to request your assistance or give you an order. The pragmatic

²⁰ There are many different kinds of pluralism and even varied functional pluralisms exhibited by language. No theorist could fail to recognize that sometimes we use language to describe, sometimes to give orders, sometimes to ask questions, sometimes to emote, etc. Initially, the functional pluralism we aim to capture is the variety of uses to which we put particularly assertoric utterances, for example, the difference between what we use a moral claim to do and what we use a prudential claim to do. I will, however, also be arguing that the traditional way of capturing different kinds of speech acts with a theory of force is itself problematic. We need a more fine-grained theory of pragmatic functions that use speech to enact. Drawing on work by Kukla and Lance, I advance such an account in Chapter 2 (Kukla & Lance, 2009).

norms are a broader set of norms that define transitions from the input conditions necessary for entitlement to a speech act to the shifts in normative standing of the speaker and others effected by the speech act.

Structural discursive norms provide a target for functional explanation by distinguishing each discursive practice as the practice that it is. Which norms we enforce in a practice, I argue, is to be explained in terms of the utility that the discursive practice has for us. The analogy here is to an explanation of the structural features of a tool in terms of what it is used to do. A claw hammer has a sturdy, straight handle, a relatively weighty steel head, and a distinctive claw opposite the head because these features make it useful for basic carpentry and framing in which one needs to efficiently hammer nails into wood and, sometimes, quickly and easily pull them out. A ball-peen hammer, on the other hand, has a hemispherical peen opposite the flat head and is made with much harder steel which allows it to fulfill its functions in metalworking, where a claw hammer might chip, and a rounded peen is useful for shaping malleable materials. Linguistic practices, likewise, evolved to serve various functions in terms of which their normative structures can be, in part, explained. Like the tools in Wittgenstein's toolbox, discursive practices share some features—they all leverage the same basic inferential machinery—but what distinguishes them, their varied structural discursive norms, is what makes them useful for the varied purposes to which they're put. These norms are what facilitate functional pluralism (O'Leary-Hawthorne & Price, 1996; Price, 2010).

The metaphor of language as a set of tools useful for fulfilling a variety of needs and purposes is, I admit, a dangerous one. We are wont to think of a tool as something ready-to-hand, the hammer to be picked up to strike the nail, the knife to be wielded to

cut the rope. Language is not found in this way; it is not ready-made for us to use in the ways we will. Language is also not a tool intentionally created to fulfil a predetermined need. Many of the purposes to which we put language are not purposes we could have had prior to language. Their very intelligibility depends on our linguistic practices (Brandom, 2011, p. 80; Faerna, 2014, pp. 18–19). We forge the tool of language in the use of our linguistic capacities, i.e., those capacities for joint attention and intention and for scorekeeping that facilitate our ability to be concept mongering beings. In the process, our needs become intelligible to us, and language provides, in some cases, the means to fulfill them. Yet, there is another perspective from which we, as theorists, can recognize language as fulfilling the needs and purposes that some beings had but could not themselves conceptualize. Language, from this perspective, is not a tool that is *intentionally* constructed to better fulfill a purpose, but rather one that emerges from circumstances in the way a method for achieving some task might arise out of trial and error in the attempt. We can see language in the way that the evolutionary theorist sees other social practices that have been selected for over our evolutionary history. In Chapter 2, I advance PALM as an account that treats language in just this way. I want to replace a picture of language *copying* the world with a picture of language as a tool for *coping* with the physical and social world we inhabit.

5. The Return to Metaethics

Chapter 3 brings us squarely back into the territory of metaethics to begin applying PALM to moral discursive practices. The chapter is mostly phenomenological and descriptive, aiming to construct a sketch of the structural discursive norms of the practice of moral discourse with which we are all familiar. Along the way, however, I also

confront the question of the relationship between moral vocabulary and motivation and argue that the neo-pragmatist account afforded by PALM makes this old problem seem much more tractable.

With a sketch of the structural norms of moral discourse in hand, I turn in Chapter 4 to an account of the function of this discursive practice. Building on accounts in evolutionary psychology and anthropology, I argue that moral vocabulary is a meta-coordinative, meta-normative vocabulary. By this I mean that it was selected for because it allowed our modern human ancestors to resolve a variety of tensions and instabilities that were unavoidable in their other social-coordinative, cultural practices, which were themselves normative practices. Moral discourse offers a way to negotiate social norms that makes them answerable to the individuals who are part of a moral discursive community without making that community vulnerable to domination by alpha-type free-riders (who our modern ancestors were quite keen to suppress) or defection and splintering by individuals or sub-groups that could lead to the group being less effective in intergroup competition for resources and intergroup conflict.

6. The Objectivity-Relativism Tension

In the final two chapters I deploy the PALM for moral discourse developed in Chapters 3 and 4 to examine the question of moral objectivity. In Chapter 5, I argue that objectivity must be understood in procedural or structural terms. Following Brandom, I claim that the objectivity of conceptual norms is a matter of maintaining the possibility from each discursive perspective that anyone or everyone could turn out to be mistaken, i.e., that what is the case outruns what anyone or everyone takes to be the case (1994, Chapter 8). I argue, however, that Brandom's account fails to distinguish between

objective and non-objective discursive practices, or, better, to distinguish degrees of objectivity. He leaves us with the impression that all discourse is objective, but this doesn't comport with our ordinary ways of thinking and talking about language. In order to draw these distinctions, we need to grasp objectivity as a feature of the authority structure of discursive practices. Some practices—discourse about mental states and preferences, for example—vest authority in particular individuals. Whatever someone sincerely avows her preferences to be, they are. There is no room for a reasonable challenge to them. Others vest such authority in the community. Whether you ought to shake hands when you are introduced to someone is a matter of the customs and expectations of the community; what *we* do settles the matter. Some discursive practices, however, exhibit a pattern of distributed authority such that no one, not even the group as a whole, is in a position to settle what is the case. In these cases, any claim that is advanced is always liable to well-motivated challenges and queries, and no claim is ever accepted as absolutely settled. These are objective discursive practices in a sense of objectivity that is related to Rorty's notion of solidarity and to Longino's notion of a democratic discourse (Longino, 1990; Rorty, 1991b). They are practices that have no truck with dogmatism, that are always prepared to re-examine even widely held commitments, and that allow everyone the authority to call anyone else to defend their commitments if they can provide good reasons for the challenge.

Moral claims are objective, yet we still find ourselves reticent to admit that our judgments about what members of moral communities that are radically distinct from our own are binding on the members of those communities. We “wet” liberals seem to think that sometimes it's not our place to judge (Rorty, 1991a). In Chapter 6, I argue that this

relativistic tendency is the result of the flipside of the authority norms of moral discourse. While the *epistemic norms* of the practice do not vest authority in any individual, the *social-institutional authority norms* function to circumscribe the moral discursive community by governing who will be recognized as having the appropriate standing to engage in the practice. Not just anyone has the standing to hold others to the moral oughts that bind them, and every moral utterance contains some element of holding. As such, there are norms of standing for the practice of moral discourse itself. We can think of this as a part of the immune system of the social group whose practice it is. Allowing just anyone to engage in the practice would grant even outsiders some control over the moral obligations of members of the group, but it would not be reasonable to grant such authority to others unless there was reason to think that they would reciprocally grant that authority to the members of the group. Non-reciprocal recognition of standing to engage in moral discourse functions to institute a hierarchical relation in which the one who is recognized has leeway to exercise her will over the others.

The conditions in which a group finds itself in part determine how stringent it will be in recognizing the standing of others with whom they are engaged in common projects or who wish to voice dissent from within the group. When social groups face scarcity, threats of violence, or disease, for example, they will be less *trusting* of outsiders. Their immune system will be in high-gear, and they will likely withhold recognition from any who are not intimately enmeshed in the coordinative practices of the group. In times of abundance, when bridges of trust are built up between the group and others with which it regularly engages in joint ventures, reciprocal recognition will come more easily. The moral community, as a result, can both contract and expand.

These dual legacies of objectivity and relativism are built into the very normative structure of moral discourse. It is neither wholly objective, open, and democratic, nor only limited to carefully circumscribed communities. As a result, the tension between these two opposing forces will always be with us. “Wet” liberals are just those of us who have come to grasp this predicament.

With this plan in view, I now turn to the task of filling in the details.

Chapter 1: Anti-Representationalist Metaethics: From Expressivism to Inferentialism

1. Introduction

If you're going to be anti-Representationalist in metaethics, then you're likely going to adopt some kind of non-cognitivist expressivism, which is clearly the most developed and most popular position in the neighborhood. You'll probably accept a broadly Representationalist semantics paired locally with a commitment to explaining the meanings of normative vocabularies in terms of the non-cognitive attitudes they express. If you're a bit more radical, you might recognize the problem of maintaining a bifurcated semantics in the face of semantic minimalism and recommend a program of global expressivism that theorizes the meanings of different vocabularies in terms of what they are used to do rather than in terms of how they map onto the world.¹ If you plunge into these depths, you'll need a lifeline, and the local version of your program offers one. You might try to hold onto the idea that different bits of vocabulary express different kinds of mental states or reactive attitudes. This central thesis of expressivism, after all, does some valuable work in evading ontological commitments and making sense of the motivational oomph of moral claims, and there's a clear way forward for thinking about canonically descriptive vocabularies as expressing states of belief. If you take this half-measure,

¹ See, for example, (Price, 2013). On the problems raised by minimalism for the expressivist's bifurcation thesis, see (Dreier, 2004).

though, you'll reach a dead end, for you'll be unable to capture the benefits offered by the global position if you insist on focusing on mental states and attitudes. This is roughly the story of the first half of this chapter.

In the second half, I tell you what you should have done if you wanted to be a global anti-Representationalist. Instead of taking the expressivist tack of theorizing semantic content in terms of the mental states or reactive attitudes expressed by normative vocabularies and blowing it up into a global semantic theory, you should have started from the other end by thinking about practical deliberation and reasoning as a normatively bound social practice and asked what role the normative vocabularies we use to engage in them play in our broader discursive lives. From there, you could have argued that playing the role that they play is what constitutes the meanings of these vocabularies. Moreover, you could have claimed that the meanings so constituted are best understood in terms of the inferential norms of these practices, i.e., you could have been an inferentialist. You would still have avoided unseemly ontological commitments and made sense of the motivational oomph of moral claims, but you also would have availed yourself of the resources necessary to answer questions about other features of moral discourse, like its seeming objectivity and rationality. More importantly, this tack would not have run you aground on the shoals of the Frege-Geach problem where most expressivisms founder.

Here's my plan for the chapter. In the next section, I introduce traditional expressivism starting from its historical roots in emotivism and canvassing briefly the contemporary positions developed by Alan Gibbard and Simon Blackburn. In Section 3, I examine the Frege-Geach problem for expressivism and the traditional expressivist

method of solving it. Focusing on arguments advanced by Mark Schroeder, I conclude that, though embedding problems raised by Geach and others are not unsolvable for expressivists, the solutions are extremely costly. In Section 4, I sketch orthodox inferentialism as a non-Representationalist, use-based account of meaning and show that it avoids Frege-Geach type problems. I then turn to the task of further developing the inferentialist position. Section 5 explores a proposal by Mark Warren for capturing the expressivist insight about functional pluralism within an inferentialist framework. I argue that by aiming for too direct an integration between expressivism and inferentialism Warren cannot make sense of moral discourse as a kind of rational discourse. In Section 6, I turn to Brandom's account of the distinctive explicative function of normative vocabulary arguing that this account fares better at capturing functional pluralism in a descriptive way but lacks the resources to explain why normative discursive practices exhibit the distinctive features they do. This sets up the task of Chapter 2, which builds on orthodox inferentialism to develop the explanatory resources of PALM.

2. Traditional Expressivism

Traditional expressivism can be characterized by commitment to two theses:

Negative Thesis of Expressivism (NTE): Not all assertoric discourse should be understood in terms of what it purports to represent, i.e., not all assertion is descriptive.

Positive Thesis of Expressivism (PTE): Some assertoric discourses are best understood in terms of the mental states they are used to express.

NTE is motivated by the recognition of two central problems of metaethical Representationalism, i.e., the proto-theoretical commitment to explaining meaning in terms of truth and reference. The first is that some of the properties that moral assertions

purport to represent are not to be found in either the natural or social world. Moral claims are binding on everyone whether or not they take themselves to be engaged in a moral practice. The trouble is that when we take a sideways on view of things, examining our claim making practices and the states of affairs those claims purport to be about from, say, a broadly scientific perspective, we don't discover anything that our moral claims could represent. From this perspective, there are states of affairs that seem to require particular actions of particular individuals. The left guard ought to pick up the blitzing linebacker when that play is called, the bank employee ought to wear business attire to work, and the driver ought to stop when the light is red, but in all of these kinds of cases, the normative requirement on the individual is a product of her engaging in some particular practice and, so, making herself subject to the rules or norms of that practice. We tend to think, however, that even the rational knave or the ideally coherent eccentric is bound by moral norms even if they claim to reject moral practices. These norms strike us as universally or objectively binding on behavior, but without appeal to some practice whose norms they are, we can't seem to make sense of them from this sideways on perspective.²

The second problem for Representationalism is that moral claims play a peculiarly motivational role. Not only must moral oughts be universally binding, they must also link up to our motivational psychology so that when we recognize we morally ought to do something, we feel some compunction to do it. Only if they do so can they

² This is essentially Mackie's argument from queerness. When we go looking for objective normative properties that exist out there, independent of us, our attitudes, and our practices, we discover that the existence of such properties in the world would require that we posit "qualities or relations of a very strange sort, utterly different from anything else in the universe" (Mackie, 1977, p. 38). Our best science gives us no reason to countenance the positing of such "queer" properties, for values are "not part of the fabric of the world" (Mackie, 1977, p. 15). Also see, (Joyce, 2015a).

play the role of practical reasons in our deliberation, and being practical (or motivating) reasons is the only way that they can feature in explanations of action (M. Smith, 1994, Chapter 1). If, as Representationalism holds, moral claims express beliefs, which have a mind-to-world direction of fit, then they're the wrong sort of thing to play this role, for conative not cognitive attitudes—those with a world-to-mind direction of fit—motivate action.

Noting these problems, expressivists have argued that the central mistake of Representationalist semantics is to treat all assertions as if they function to describe how things are in the world. Assertions can play other roles in our linguistic economy, so, perhaps, the contents of some assertions are best understood in terms of these other roles that they play. Some assertions are prosaically descriptive. The propositional content of “Grass is green” should be understood in terms of truth and reference as the Representationalist contends. “Jeremy ought to care for his aging mother,” however, isn't primarily an attempt to describe a state of affairs. Among other things, it attributes a commitment to Jeremy, endorses a course of action, expresses commitment to chiding Jeremy if he fails to care for his mother, and attempts to motivate Jeremy to care for his mother. None of these functions straightforwardly requires us to posit some properties in the world to which the assertion refers. Instead, we might explain the content of the assertion directly in terms of how it fulfills one or more of these functions. Endorsing NTE, then, is adopting what I'll call *functional pluralism*: different vocabularies fulfill

different needs and purposes *and* the meanings of these vocabularies is best explained in terms of these functions.³

PTE is motivated by a further commitment. In order to explain an action, philosophers have usually thought that we must appeal both to a belief about the likely outcome of the action and to a desire to bring about that outcome. On this Humean psychology, beliefs are about the world; they aim to represent how things are or have a mind-to-world direction of fit. Desires, on the other hand, have a world-to-mind direction of fit. Their success conditions are satisfied if the world conforms to them whereas those of beliefs are satisfied if they conform to the world. As we've already seen, moral assertions are problematic within a Humean psychology because they seem to have both directions of fit. They seem to make an objective claim about some properties out there in the world, but they also aim to guide behavior, to bring about certain outcomes.

Expressivists take this as motivation to treat moral assertions more like expressions of desire than of belief. They give expression to non-cognitive (or conative) mental states not beliefs. In one fell swoop, this answers both of the problems faced by Representationalism. If moral claims express conative mental states, then we don't need to go looking for the objects and properties that they seemed to be about, and we have no need to worry over their motivational oomph.

This position has its historical roots in the emotivism of Ayer, Stevenson, and Barnes, which claims that moral assertions function to express an emotion and, perhaps,

³ I do not mean to deny that Representationalists also recognize a form of functional pluralism. Language can be used to do many different things, including the sorts of things I've just listed. This is usually captured in a theory of pragmatics appended to the semantics for a language. The difference I intend to mark here is where this functional pluralism is located. For expressivists, the function of a vocabulary (perhaps understood in terms of pragmatics) is part and parcel of the explanation of meaning for that vocabulary, whereas the Representationalist first constructs a theory of meaning, then explores how these semantic contents can be put to use.

to elicit a similar emotional response in others. The proper way to theorize about the semantics of these utterances is to point out that they are conventionally used to perform a speech act that, if it is sincere, requires the speaker have a certain attitude or be in a specific emotional state, most obviously one of approval or disapproval. Importantly, emotivists argued that moral assertions do not say that the speaker is in that state, but rather function as a means of expressing that state much like an exclamation expresses surprise or excitement. "When I say 'I have a pain,' that sentence states the occurrence of a certain feeling in me: when I shout 'Oh!' in a certain way that is expressive of the occurrence in me of a certain feeling" (Barnes, 1934). This distinguishes emotivism from ordinary speaker subjectivism and enables it to avoid the serious error of treating the semantic contents of "I disapprove of murder" and "Murder is wrong" as identical.⁴ Sentences making use of positive evaluative predicates express a non-cognitive pro-attitude toward the action, situation, or thing being evaluated, while sentences employing negative evaluative predicates express con-attitudes. As such, calling some action "good" is akin to cheering for it while calling some action "bad" is akin to jeering it. This has earned emotivism the epithet of the "Boo!"/"Hooray!" theory of moral language (Ayer, 1952; Stevenson, 1972).

Contemporary versions of expressivism have evolved to address many of the pitfalls of this naive early view. The first move in this direction came from R. M. Hare (1952). Emotivists were often dismissive of the rationality of normative discourse. If it functions merely to express an individual's evaluation, then there can be nothing like moral knowledge or truth, and the practice of moral discourse could be pseudo-rational at

⁴ For a discussion of this genealogy of the position see, (Schroeder, 2010, Chapter 1).

best. Hare's prescriptivism aimed to salvage the seeming rationality of moral discourse. Moral claims that appear syntactically to be assertions are, semantically, better understood as kinds of imperatives, according to this view. Though they differ from ordinary imperatives in their universality, like imperatives, they function to direct action. Hare proposed analyzing such claims in terms of two elements: a *phrastic*, which represents a proposition or state of affairs, and a *neustic*, which, like a mood operator, indicates what speech act is being performed with the phrastic (1952, pp. 17–22). Motivated by the idea that imperatives could stand in logical relations—incompatibility, for example—with one another, his aim was to demonstrate that the logical connectives could be represented as part of the phrastic so that the standard rules of logic could be used to explain the standard inferences in any mood. The upshot would be an account of moral assertions that captures their prescriptive dimension but does not dismiss rational constraints on their use.

Two more recent attempts to make good the promises of emotivism come from Simon Blackburn (1984, 1993) and Allan Gibbard (1990, 2003). These traditional expressivists share a basic explanatory commitment: a semantics for normative language is to be constructed in terms of the conative mental states or reactive attitudes expressed by normative claims. These contemporary expressivists break from earlier emotivists in aiming to do justice to the rationality and objectivity of moral discourse. Consequently, the explanatory burden placed on mental states and the account of the expression relation is significant. It must be capable of explaining the inferential properties normally exhibited by moral discourse,

to account for inconsistency of sentences in terms of the states of mind expressed, entailment among sentences in terms of commitment

between states of mind expressed, and so on. These are the materials with which expressivists have to work (Schroeder, 2010, p. 39).

This explanatory commitment is the source of a number of problems for traditional expressivists, and most of them stem from their attempts to respond to the challenge laid down by Geach to which we'll turn in a moment.

Before we examine the traditional expressivist response to the Frege-Geach problem, it would be good to have a sense of how they characterize the mental states or attitudes central to their account. For Blackburn, these are fundamentally attitudes of “tolerating” and “hooraying,” which correspond to the object of the attitude being permissible or obligatory. These attitudes are inter-definable in terms of negation. If doing x is not obligatory, then it is permissible; if it is not permissible then doing $not-x$ is obligatory, and so on. Moreover, we can have higher-order attitudes toward these basic attitudes, which allows them to compose more complex attitudes. One could, for example, hooray making not tolerating teaching others to cheat follow upon not tolerating cheating. If one also has an attitude of non-toleration toward cheating, then one would be tied to a tree of commitments such that one had better not tolerate teaching others to cheat on pain of a kind of pragmatic inconsistency reminiscent of Moore’s paradox (Blackburn, 1988).⁵

⁵ This is a sketch of only the semantic dimension of Blackburn’s quasi-realist program. The full program also encompasses a metaphysical commitment to projectivism, i.e., the Humean thesis that moral properties are not found in the world but projections of our attitudes painted onto the world we encounter (Hume, 1751, p. 88), and a psychological commitment to non-cognitivism, i.e., that the contents of normative mental states are directive rather than belief-like. Quasi-realism itself is best understood as a philosophical program for *earning the right* to use realist seeming language—appending the truth predicate to moral claims, accepting standard logical implications involving normative claims, and even speaking of moral properties—while endorsing an anti-realist program composed of the elements just discussed (Blackburn, 1984, p. 189; Joyce, 2015b). Blackburn has outlined both fast-track and slow-track approaches to this project. On the fast-track, the quasi-realist earns the right to realist sounding language via minimalism about truth. If minimalism is right, then quasi-realists can help themselves to truth-talk, and the rest follows easily from

Gibbard, for his part, argues that judgments about rationality express an attitude of acceptance toward rules or norms so that to judge an action rational is to express acceptance of a system of norms that permits it and to judge an action irrational is to express acceptance of a system of norms that forbids it (Gibbard, 1990, p. 46). Since speakers often lack coherent systems of norms, Gibbard opts to interpret these judgments somewhat more loosely as expressing rejection of any system of norms with which the judgment is incompatible. To capture the content of normative judgments, then, he extends a possible worlds semantics. A normative judgment is represented as a world-norm pair. Predicating a normative term to an action rules out—that is, commits one to judging as irrational—any combination of a world in which an action of the relevant kind is performed (the world component) and a system of norms that permits, forbids, or requires that action (as appropriate). For example, the judgment that harming oneself is irrational is represented as a the set of world-norm pairs that have self-harm (under some description) as the world component and any set of norms that permits or requires self-harm as the norm component.⁶ Gibbard’s subsequent work has made some amendments to this account, most notably replacing the norm component of world-norm pairs with what he thinks of as contingency plans. As such, a normative claim expresses the acceptance of a plan of action given a state of affairs. He calls these fact-prac worlds (Gibbard, 2003).

there. The slow-track approach, on the other hand, earns the right to realist language piecemeal and provides security to the fast-track approach (Blackburn, 1993, pp. 184–186 & 197). Also see, (van Roojen, 2016).

⁶ Moral judgments, for Gibbard, are a subset of judgments about rationality. To judge A to be morally impermissible is to judge it to be rational to blame someone for doing A. Furthermore, to take A to be rationally blameworthy is to take it to be rational to feel guilty for doing A (Gibbard, 1990, p. 45).

Blackburn and Gibbard's positions both exhibit the central move that marks the shift from emotivism to contemporary expressivism. Rather than characterizing the meaning of normative claims in terms of the speech acts they enact—acts of booing or hooraying—expressivists have taken the force of a normative claim to indicate that it expresses a different kind of mental state or attitude than do ordinary descriptive claims. The meaning of the claim is to be understood in terms of the attitude it expresses. This means that the speech act performed could be one of assertion, but the content is to be understood in non-Representationalist terms, since the mental state constitutive of that content plays a different role in our psychology than do beliefs.

3. The Frege-Geach Problem for Traditional Expressivism

The Frege-Geach or embedding problem for traditional expressivism is a problem about how expressivists can account for the meanings of their target sentences in embedded contexts in which those sentences are not asserted.⁷ According to traditional expressivists, the meaning of a sentence is determined at least in part by the non-cognitive attitude (or stance⁸) it is used to express. So, for example, the meaning of “One ought not to lie” is given by its use to express a negative attitude toward lying. Yet, such an analysis gives us no grip on the meaning of the sentence in contexts in which it is not asserted, for in such contexts the sentence is not used to express such an attitude. When one asserts, “If one ought not to lie, then one ought not to get one's little brother to lie,”

⁷ It is perhaps better to say problems. It's a delicate matter to state the problem exactly, and it may be best to think of it as a cluster of problems about embedding in various contexts that are deeply related but not exactly identical. In any case, I take what I say here to represent central issues in play. At the very least, the way I put the issue here captures the way it is most commonly presented.

⁸ This is Blackburn's preferred term. The essence of a “stance” lies in its practical import or pressure toward action, though the particular feel or phenomenology of a stance may vary according to cultural contexts and, perhaps, other factors. There is, he says, no need for the stance to feel like a desire (Blackburn, 1993, p. 169).

though one very well may be committed to condemning lying, one is not condemning it in this instance. The sentence could be felicitously uttered by someone who embraces the telling of lies.

The standard way of raising the problem is through examination of a *modus ponens* argument.⁹ Consider:

- (1) If one ought not to tell a lie, then one ought not to get little brother to tell a lie.
- (2) One ought not to tell a lie.
- (3) Therefore, one ought not to get little brother to tell a lie.

According to traditional expressivists, the meaning of (2) is given by the negative attitude toward lying that it expresses. The attitude expressed is essentially directive, having a world-to-mind direction of fit. If this is right, how are we to understand “One ought not to tell a lie” as it appears in (1), where it isn’t being used to express a negative attitude toward lying? Since we can’t understand it in terms of the expression of a negative attitude, the contribution that “One ought not to tell a lie” makes to (1) differs from what it means in (2), so this *modus ponens* argument is an equivocation. The argument does not equivocate, so something must be amiss with the traditional expressivist understanding of (2).¹⁰

⁹ The standard way is owed to Geach, who attributed the insight to Frege (1960), but Searle also raised this as an instance of what he called the “speech act fallacy” (1962).

¹⁰ It is worth noting that Geach’s original presentation of this case was not intended as a challenge for non-cognitivists to meet but rather as conclusive evidence that normative claims do mean the same thing in embedded and unembedded contexts. After all, we can say the argument is *logically valid* only if they mean the same thing in both contexts. As such, Geach took himself to have shown that noncognitivism is centrally committed to something false (Schroeder, 2010, p. 19). The problem can be stated another way. Neither the embedded nor unembedded occurrence is truth-apt, according to the expressivist. The consequent of the conditional detaches only if the antecedent is satisfied. If, however, the antecedent of the conditional is an embedded non-representational bit of language, the antecedent cannot be satisfied since it is not truth-apt. As such, *modus ponens* seems not to be a valid form of inference in moral discourse according to the non-cognitivist (Tebben, 2015, pp. 10–11).

Saying what exactly is amiss is tricky. We might try to clarify by noting why moral cognitivists are not subject to this objection. Cognitivists, one might think, face a similar problem since they say in (2) “One ought not to tell a lie” is used to *assert* that one ought not to tell a lie, but in (1) it is not used to assert anything at all. As such, the argument equivocates. But this isn’t right. The reason is that cognitivists distinguish between force and content. The content of the sentence is the proposition with which it is associated whether it is asserted, wished, desired, feared, attributed to another, conditionally asserted, etc.. Geach dubbed this the Frege Point: "A thought may have just the same content whether you assent to its truth or not; a proposition may occur in discourse now asserted, now unasserted, and yet be recognizably the same proposition" (1965, p. 449). The proposition associated with “One ought not to tell a lie” is the same in (1) and (2), and this is all that matters for the validity of the argument. Expressivists, however, cannot distinguish between force and content in this way. According to expressivism, the force of a sentence is partly constitutive of the mental state expressed by the claim and cannot be divorced from its propositional content. It is this elision of force and content that seems to be at the heart of the problem raised by Geach (Eklund, 2009, pp. 705–706).

3.1. Mentalism and the Problem of Composition

What are the steps that the expressivist must take in order to demonstrate that mental states can stand in the appropriate relationships to account for logical validity, inconsistency, and so on? Expressivism begins with an account of the meanings of atomic normative sentences, but the embedding problem arises when we consider how those sentences can function as constituents of more complex sentences. So, expressivists will

need an account of how the meanings of complex sentences are a function of the meanings of their constituent parts. This first step, as we shall see, already brings with it significant costs, but the task is not complete, since expressivists also need to demonstrate that the intuitive logical relations between various atomic and complex sentences can be sustained on their account. A sensible place to begin is with inconsistency. If expressivists can make sense of the inconsistency between a sentence and its negation, then they are well on their way to making sense of valid arguments. After all, a valid argument is one in which it is impossible for the premises to be true and the conclusion false. We shall see, though, that constructing an account of inconsistency for normative sentences understood in expressivist terms leads to a number of complications which force costly revisions in our semantics.

Let's begin with the problem of composition, i.e., how the meanings of complex sentences can be understood as a function of the atomic sentences of which they are composed. Traditional expressivists account for the content of normative assertions in terms of the mental states they express. For negations, conditionals, conjunctions, and disjunctions in which all of the atomic sentences are normative, the most straightforward account for traditional expressivism is to treat the meanings of these sentences as functions of the meanings of their parts. Since the meanings of their parts are constituted by the mental states they express, the logical connectives must act as functions from those mental states to a complex mental state expressed by the complex claim. This is analogous to the way that the meanings of complex descriptive claims are propositions with truth-values that are functions of the truth-values of the propositions associated with

their atomic parts.¹¹ The fundamental move, then, is to treat the logical connectives as functions on mental states rather than on propositions.

Given this commitment, what is the expressivist to do with complex sentences that have both normative and descriptive atomic parts? One option is to hold that the meanings of the logical connectives are different depending on whether the atomic sentences to which they are applied are normative or descriptive. Consider the conditional. The meaning of a conditional claim with normative atomic parts would be given as a function of the mental states expressed by those parts. The meaning of a conditional with descriptive atomic parts would be given as a truth-function of the truth-values of the propositions associated with its atomic parts. The meaning of a conditional with a descriptive antecedent and a normative consequent, then, must be given by some function from the proposition referred to by the antecedent and the mental state expressed by the consequent. This option, however, should strike us as rather implausible. If the logical connectives enact different functions—have different meanings—depending on whether they are applied to normative or descriptive sentences, then we would have to know for each sentence of the language whether it is normative or descriptive in order to compute the meanings of complex claims composed of them. But natural languages do not syntactically mark this distinction and it doesn't seem that the logical connectives are functioning differently in these different contexts.

Moreover, there is significant independent pressure toward thinking that the meanings of the logical connectives are univocal. We do not have normative-negation and

¹¹ This enacts what Schroeder calls the *Basic Expressivist Maneuver*, which is to show that their account requires the same basic explanatory resources as the Representationalist, and so is in no worse shape (2010, p. 18).

descriptive-negation in our languages, and we take it that the inconsistency between a sentence and its negation should be assured by the meaning of the negation and not by the fact that the sentence is either descriptive paired with a descriptive negation or normative paired with a normative negation (Schroeder, 2010, p. p.22). If, however, the meanings of the connectives are univocal, then the only option that remains for traditional expressivists is to maintain that the meanings of ordinary descriptive claims are also constituted by the mental states that they express. This is a rather radical conclusion to be pushed to by one's metaethical semantics, but it is the only way to maintain that the connectives enact *the same function* from the mental states expressed by the atomic sentences to the mental state expressed by the complex claim. This is the position Schroeder calls *mentalism*.

It's easy to see why mentalism is a radical break from the traditional semantics of ordinary descriptive claims. The meaning of an ordinary descriptive claim is given by the truth-value of the proposition with which it is associated. The content of "cats are mammals" is just that cats are mammals. Traditional compositional semantics then explains the propositional contents of complex claims as a function from their constituent parts given by the logical connectives that join them. The content of "cats are mammals and pheasants are birds" is just the proposition that cats are mammals and pheasants are birds and this is explained as an operation of the conjunction on the propositions that cats are mammals and that pheasants are birds. The expressivist, however, can't say this, for she is committed to the conjunction operating as a function not on propositions but on mental states. Her explanation that the propositional content of "cats are mammals and pheasants are birds" is the proposition that cats are mammals and pheasants are birds

must be that the sentence expresses the *belief* with that proposition as its content and that the content of that belief is a function of the conjunction operating not on propositions but on the mental states (also beliefs) out of which the complex belief is composed. The compositional semantics of ordinary descriptive claims must detour through a semantics populated by mental states. Expressivism's core explanatory commitments, as such, lead to a radical revision of all of compositional semantics and, perhaps, to the semantics of atomic descriptive sentences, as well (Schroeder, 2010, pp. 22–24).

I will not try to claim that mentalism is false. For all I know, it may be the correct account of the semantics of descriptive sentences. My aim is only to show that expressivists are forced into this commitment by way of their account of the semantics of normative sentences, and that this is a high price to pay. For now, though, let's accept mentalism and see what the expressivist must do in order to complete her account of the logic of attitudes in a way that makes sense of inconsistency and validity.

3.2. Logical Inconsistency and Revisionist Semantics

Logical inconsistency is usually conceived in terms of truth. Two propositions are inconsistent just in case it is not possible for both to be true at the same time.

Expressivists cannot accept this account since they are committed, at least in the first instance, to the non-truth-aptness of the contents of normative claims. Moreover, since expressivists account for the meaning of normative sentences in terms of the mental states they express, inconsistency must be a matter of incompatibility between mental states rather than between propositions. Given this starting point, what options are available to expressivists? There are essentially two. Following Hare, we can distinguish the descriptive content of a mental state—Hare's *phrastic*—from an attitude toward that

content—his *neustic*. On the first option, inconsistency is matter of having the same attitude toward inconsistent contents. I'll follow Schroeder in calling this A-type inconsistency. The second option treats it as a matter of having conflicting attitudes toward the same content. This is B-type inconsistency. I'll take up the latter option first.

Blackburn's program is emblematic of B-type accounts in taking as fundamental attitudes of 'tolerance' and 'hooraying,' which correspond to treating acts as permissible and obligatory.¹² Schroeder has forcefully argued, however, that B-type inconsistency suffers from several significant problems. Most notably, expressivists who rely on B-type inconsistency lack the resources to explain why the attitudes they favor are inconsistent. Consider, by comparison, how we explain inconsistency in beliefs (a paradigm case). Suppose that S has both the mental state of believing that *p* and of believing that *not-p*. What is wrong with S's overall mental state? Well, we might say, it is inconsistent because believing that *p* is taking *p* to be the case and believing that *not-p* is taking *not-p* to be the case, but *p* and *not-p* cannot both be the case. This is a contradiction. Since the descriptive elements are inconsistent, one goes wrong in believing them both. We cannot tell such a story about S tolerating *p* and hooraying *p*, for they are not the same attitude toward inconsistent contents but different attitudes toward the same content, so why are these two mental states to be inconsistent? This must be a kind of brute inconsistency for which traditional expressivists can offer no explanation. To be fair, we have some grip on this idea. To enjoy and abhor the same experience strikes us as holding inconsistent attitudes, for example. They are inconsistent in the sense that someone who has these

¹² This is an amendment of his work in *Spreading the Word*, in which he took as primary attitudes of 'booing' and 'hooraying'—corresponding to *impermissible* and *obligatory*. The amendment is required by recognition that the denial of *p* being obligatory is not *p* being impermissible but, rather, "tolerating $\sim p$ or allowing it as consistent with an ideal world" (1988, pp. 511–512).

attitudes toward the same object—the same action, activity, outcome—would find it difficult to plan her actions, a kind of pragmatic inconsistency. This, however, is a very slippery idea, for sometimes we enjoy (in some way) the very activities that we abhor. Think, for example, of an extraordinarily difficult climb while cycling. Moreover, it is far from clear that this kind of pragmatic inconsistency is of the same kind as the logical inconsistency that seems to be what we need in our semantic account (Schroeder, 2010, pp. 47–48).¹³

This problem for B-type inconsistency arises when we consider negations of atomic normative sentences, but the problems proliferate when we turn to complex sentences. What mental state is expressed by the conjunction of two normative sentences: one ought not to lie and one ought not to steal? Presumably, each of the atomic sentences expresses disapproval, but the mental state that results from the conjunction operating on these mental states cannot be defined in terms of disapproval. There are too many options for inserting the conjunction. The resulting mental state could be one in which S disapproves of the combination of lying and stealing together, one in which S disapproves of each lying and stealing, or one in which S disapproves of lying and S

¹³ This is the common objection raised against Blackburn’s account. He argues, for example, that someone who disapproves of lying and approves of making disapproval of getting little brother to lie follow upon disapproval of lying

must hold the consequential disapproval: he is committed to disapproving of getting little brother to lie, for if he does not his attitudes clash. He has a fractured sensibility which cannot itself be an object of approval. The ‘cannot’ here follows not (as a realist explanation would have it) because such a sensibility must be out of line with the moral facts it is trying to describe, but because such a sensibility cannot fulfill the practical purposes for which we evaluate things (1984, p. 195).

However, one who accepts p and $p \rightarrow q$ but fails to accept q doesn’t suffer from a “fractured sensibility.” She suffers from irrationality. Blackburn subsequently responds to this charge with his elaboration of commitment-theoretic semantics, but I won’t take that up here (2001, pp. 68–77).

disapproves of stealing. Schroeder takes this to show that the conjunction expresses a state that is distinct from disapproval. Moreover, the negation of the conjunction must express some state that is incompatible with the state expressed by the conjunction. Once again, this will be a primitive incompatibility, and, so, it must be a distinct state. As Schroeder remarks, it is interesting “to see just how quickly [B-type theorists] end up needing to posit an infinite list of distinct kinds of attitude to go with disapproval and tolerance—for every pair of which we will have to postulate primitive inconsistency relations” (Schroeder, 2010, p. 49).

The problem is not that the B-type theorist cannot make such an account work. Terry Horgan and Mark Timmons have, in fact, constructed just such an account. Horgan and Timmons identify both a basic non-cognitive attitude—‘ought-commitment’—and a basic cognitive attitude—‘is-commitment’—and then define the negation of each as a distinct attitude that is, by necessity, primitively incompatible with the original attitude. In similar fashion, they represent every complex sentence as expressing a distinct kind of mental state. These “logically complex commitment states” are defined in terms of their inferential roles, i.e., in terms of the incompatibility relations between different attitudes toward the same content (Horgan & Timmons, 2006). Schroeder points out, then, that

In essence Horgan and Timmons view amounts to the hypothesis that there is an unfathomably huge hierarchy of distinct kinds of mental state, together with unsupported confidence that these mental states have the right inconsistency relations with one another...But Horgan and Timmons give us no reason other than sheer optimism to believe that there really [are such states], which [have] these inferential properties (Schroeder, 2010, p. 51).

Beyond a commitment to mentalism, then, B-type theorists are saddled with the need to explain inconsistency between distinct mental states toward the same content *and* then to provide unfathomably many similar explanations for why the attitudes expressed by complex sentences stand in just the right inconsistency relations. Finally, they must

convince us that we actually have these complex mental states. Given the significant costs of a B-type account, Schroeder argues that expressivists ought to advance an A-type theory, i.e., one that relies on the same resources as the cognitivist to explain inconsistency in terms of taking the same attitude toward inconsistent states.

The central problem that drives traditional expressivists toward B-type inconsistency, according to Schroeder, is the adoption of unstructured attitudes corresponding to each normative predicate as primitives. What is needed, he thinks, is some very general noncognitive attitude that stands in the same relation to normative predicates as belief stands to descriptive predicates. Descriptive predicates “correspond to belief plus some property that is contributed by the predicate,” e.g., “The grass is green” as uttered by John expresses his belief that the grass is green. The same should hold in the normative case so that, for example, “Murder is wrong” as uttered by John expresses his *non-cognitive attitude A* that murder is **contribution of the normative predicate**. Schroeder proposes the attitude ‘being for’ to fill this role and, to fill out the account, adopts Gibbard’s analysis of wrongness in terms of blameworthiness. “Murder is wrong,” then, expresses the attitude of being for blaming for murder (Schroeder, 2010, pp. 57–58).¹⁴

The added structure in this approach allows Schroeder to solve the negation problem, i.e., the problem that there are not enough ways to negate an unstructured attitude so that their incompatibility can be explained in terms of the logic of negation.

¹⁴ Gibbard might himself best be understood as endorsing an A-type account of inconsistency with only one relevant attitude—acceptance—directed at norms regarding the aptness of different kinds of emotional responses. If this is right, then Gibbard avoids the problems that B-type theorists face, but he’ll still fall prey to the problems for A-type theorists that we’re about to discuss, i.e., it seems he’ll end up with problems accounting for tenses and modal claims as well as being committed to a bifurcated attitude semantics.

Consider the sentence “Jon thinks that murdering is wrong.” There are three places to insert negations:

- N1 Jon does not think that murdering is wrong.
- N2 Jon thinks that murdering is not wrong.
- N3 Jon thinks that not murdering is wrong.

If we grant that X is wrong expresses disapproval of X, then we can interpret only two of these negations.

- N1* Jon does not disapprove of murdering.
- N2* ???
- N3* Jon disapproves of not murdering.

There simply aren't enough spaces for the negation to go. This is what led the B-type theorist to adopt an attitude like tolerance as the second primitive, for notice that whatever N2* is, it must be the attitude that is inconsistent with Jon disapproving of murdering, i.e., it must be his toleration of it (Schroeder, 2010, pp. 44–45). Schroeder's attitude of 'being for,' however, offers a straightforward solution by adding more structure. In this case, we interpret “Murder is wrong” as uttered by Jon as “Jon is for blaming for murdering.” The negations, then, are:

- B1 Jon is not for blaming for murdering.
- B2 Jon is for not blaming for murdering.
- B3 Jon is for blaming for not murdering.

The mental state that is inconsistent with being for blaming for murdering is being for not blaming for murdering, i.e., the same attitude—being for—toward an inconsistent content. All that is needed to explain this incompatibility is an explanation of why the contents are incompatible—this just relies on our understanding of negation—and the claim that the attitude of 'being for' is *inconsistency-transmitting*, i.e., that, like the attitudes of belief and intention, two mental states that express the attitude toward

inconsistent contents will be inconsistent. Schroeder's solution, then, solves the negation problem for expressivism and allows for an explanation of inconsistency that leverages our already strong grasp of A-type inconsistency.

Even an expressivist account like Schroeder's, however, has significant costs. Schroeder identifies at least two that require revisions beyond mentalism to our semantics for ordinary descriptive sentences. The first arises when we consider complex sentences composed of both normative and descriptive atomic sentences. As things stand, we have a working account of the attitude expressed by normative claims (being for) and descriptive claims (belief), and Schroeder shows how these each can be developed to accounts for complex normative and descriptive claims using those same attitudes (Schroeder, 2010, p. chs. 5 & 6). However, we clearly use sentences that combine normative and descriptive components in various ways. What attitudes are expressed by these sentences? Given the arguments up to this point, Schroeder concludes that whatever it is, all sentences must express the *same* general, inconsistency-transmitting attitude. If this were not the case, then we would, once again, be forced to conclude that the logical connectives are not univocal (the problem that led to mentalism), and we would be forced back into the quagmire of B-type inconsistency (Schroeder, 2010, pp. 90–91). There seem to be, then, two options on the table. Either all sentences express beliefs or all sentences express being for. The first option gives up on expressivism, for it *just is* cognitivism to claim that all sentences express beliefs. The second option, however, is radical. It requires Schroeder to claim that beliefs really are just states of being for. This, however, is what he sets out to do. Schroeder argues that the belief that grass is green, for example, can be

glossed as a noncognitive state of being for *proceeding as if* grass is green (Schroeder, 2010, p. 94)!¹⁵

Schroeder has argued rather convincingly that the general structure of the account he proposes is the *only way* an expressivist account can succeed in meeting the challenges raised by the embedding and negation problems (Schroeder, 2010, p. 92). He's also argued that such an account will be costly. Expressivists will need to bear the commitment to mentalism, forcing a revision of the semantics for ordinary descriptive claims. Moreover, they will either need to accept an infinite hierarchy of attitudes to account for complex sentences in terms of B-type inconsistency or defend the position that all sentences, including ordinary descriptive sentences, express *the same* noncognitive attitude. They will, that is, need to be global expressivists. Even if they are successful in making such an account plausible, he has also argued that they would face additional, possibly insurmountable challenges in developing a semantics for tenses other than the present and for modal claims. Schroeder shows us that the costs of expressivism are dear; it's so expensive, in fact, that Schroeder himself rejects the position. I think this is enough to motivate turning away from any account that takes expressive use as explanatorily fundamental. In response to the problem of Representationalism, traditional expressivists focused on mental states and reactive attitudes, but we've now seen that such a focus leaves us paying too high a price when we try to recover the rationality and objectivity of moral discourse. Fortunately, there's another anti-Representationalist

¹⁵ This 'solution,' though, is not the end of Schroeder's problems, for it immediately gives rise to a new version of the negation problem. The state of "being for proceeding as if x" allows for *too many* negations, and this forces him to conclude that expressivists must adopt a bifurcated attitude semantics. Believing that *p* cannot be analyzed in terms of being in some relation to *p*, e.g., the relation of being for proceeding as if *p*. It must be analyzed in terms of being for *two* relations to *p*. It involves both being for proceeding as if *p* and being for *not* proceeding as if $\sim p$ (Schroeder, 2010, p. 98). I mention this just to show how deeply these problems run. With this proliferation of attitudes, things get very complicated very quickly!

alternative that will put us in a better place to address these problems by beginning with a focus on practical reasoning and the role it plays in our discursive lives.

4. Orthodox Inferentialism

Orthodox inferentialism is a non-Representationalist, use-based theory of meaning, but instead of expressive use, it conceives of meaning in terms of patterns of use within a linguistic community. Meaning is constituted by *how* a sentence is used, rather than by what it is used to express or to do. According to orthodox inferentialism, these patterns of use are explained, in part, by the inferential norms that govern the use of the sentence.¹⁶ As we shall see, in some cases, at least, these inferential norms are, in turn, explained in terms of what speakers use the sentence to do, i.e., in terms of what I will call its explicative/performative function.¹⁷

For the orthodox inferentialist the patterns of use that are constitutive of the meaning of a claim consist in the network of inferences in which it plays the role of either premise or conclusion (R. Brandom, 1994, 2000; Chrisman, 2008, p. 350). The meaning of “X is red,” for example, is constituted by its position in a network consisting in part of “X is crimson”—from which “X is red” follows—and “X is ripe”—which follows from “X is red” and “X is a Macintosh apple.” A few points about the kinds of inferences the orthodox inferentialist has in mind deserve special attention here. First, it is not only the logically or formally good inferences that are constitutive of meaning, but also the

¹⁶ This is part of Brandom’s normative pragmatics, which I’ll discuss in more detail in a moment.

¹⁷ It is possible on such an account that the meaning of a moral claim is, in the end, explained by its being used to give expression to a moral attitude. The idea would be to argue that meaning is constituted by inferential patterns of use, thereby avoiding the problems faced by traditional expressivism in the way I am about to explain, and that the particular inferential patterns exhibited by moral claims are the product of their being used to make explicit the moral attitudes of their speakers by standing in certain inferential relations to those attitudes. In Section 5, I examine and reject an account proposed by Mark Warren that takes this shape.

materially good inferences in which the claim plays the part of premise or conclusion. So, for example, “Washington, DC, is to the west of Baltimore” follows from “Baltimore is to the east of Washington, DC,” and this material relationship between these two claims is partly constitutive of the meaning of each. Second, material incompatibilities are also counted among the inferences, so the fact that something’s being monochromatically red precludes its being green underwrites the goodness of the inference from “X is monochromatically red” to “X is not green.” Finally, the inferences include language-entry transitions from non-inferential circumstances of appropriate application and language-exit transitions to non-inferential consequences of appropriate application. It is, for example, partly constitutive of the meaning of “X is red” that it is appropriately uttered by S when X is in S’s visual field, X is red, and S is disposed to reliably respond differentially to red objects by applying the concept ‘red’.¹⁸ To put it succinctly, meaning is inferential articulation broadly construed to include materially good inferences, incompatibilities, and language-entry and language-exit transitions.¹⁹

¹⁸ On Brandom’s account, perception functions as a language-entry transition that allows objects in our physical surroundings to come to bear on our conceptual apparatus by way of our reliable dispositions to respond to features of those objects by applying concepts to them. Brandom understands perception, in part, as the exercise of reliable dispositions to respond differentially to one’s environment. Such dispositions are had by everything from a piece of iron to a parrot to a human being, so reliably responding differentially to stimuli cannot be a sufficient condition for perception. But it seems hard to doubt that it is at least necessary. What distinguishes human perception from parrots and rusty pig iron is that for concept mongering beings perception consists in an exercise of our conceptual capacities. Genuine observers differentially respond to stimuli *by applying concepts*, that is by coming to believe, claiming, or reporting that such-and-such is the case. Applying concepts, for Brandom, is a linguistic and a normative affair. It is a matter of having an ability to follow the inferential connections that constitute the concept one is deploying. Just as importantly, though, it is taking up a position in the space of reasons, making oneself *liable* to respond to challenges to the claim one has made by giving reasons for it, *committing* oneself to what follows inferentially from one’s claims, and *entitling* oneself to those inferential consequences by way of defending the commitments one undertakes (R. Brandom, 2000, p. chap. 3, 2002).

¹⁹ Brandom dubs this version strong inferentialism, which he distinguishes from both weak inferentialism—which claims only that inferential connections are necessary but not sufficient for claims to have the meanings they do—and hyperinferentialism, which counts only the formally good inferences in which the claim plays a role as premise or conclusion as constitutive of content. (R. Brandom, 2007, pp. 163–164).

This inferential articulation is the product of certain social, norm-governed practices of language using beings, i.e., of the normative pragmatics of discourse. When we make claims, we put ourselves in positions of *liability* and *entitlement*, making ourselves targets of assessment by our interlocutors. In undertaking a *commitment* to “You ought not to lie” by asserting it, I make myself liable to well-motivated challenges, i.e., I owe a response to a well-motivated challenge that justifies my claim by appeal to a reason (or reasons) from which it follows. If I am entitled to this commitment, I am also entitled to certain consequences of it, for example, to chiding you if you tell a lie (a language-exit transition). Moreover, in uttering the claim, I *entitle* others to it and to deflecting challenges aimed their way to me. Now, of course, I can commit myself to a claim to which I am, in fact, *not* entitled either because I subsequently cannot provide sufficient reasons in its defense or because, unbeknownst to me, it is incompatible with some other claim to which I am committed. This latter case is possible because, besides the commitments I acknowledge, I also have among my commitments the inferential consequences of my acknowledged commitments, whether or not I am aware of them. I can have unacknowledged commitments because the normative statuses of commitment and entitlement are the product of the normative attitudes of my interlocutors who *treat* me as being committed and/or entitled (or not) to claims.²⁰ The norms of the practice are, as such, norms of assessment. They are binding insofar as they make one’s acts liable to the judgment of members of one’s linguistic community. The practice, then, is one in which interlocutors keep score on one another as they navigate this rich normative

²⁰ This is what Brandom calls *phenomenalism* about norms, i.e., the idea that the deontic *statuses* of participants are instituted by the deontic *attitudes* of participants (1994, pp. 32–37).

space.²¹ They track one another's commitments and entitlements both from their own perspectives and as they take them to be from those of their interlocutors—this is the *de re/de dicto* distinction for Brandom—and extract information about how things are in their shared world (1994, Chapter 8). This social distribution of normative statuses provides for the variety of perspectives that make objectivity of linguistic content possible, but exactly how is the topic of Chapter 5 (Peregrin, 2012, p. 5).

The explanandum for orthodox inferentialism is linguistic behavior. Rather than explaining this in terms of norms of assertability and propositional content cashed out in Representationalist terms or in terms of norms of assertability and the expression of mental states understood in traditional expressivist terms, the inferentialist explains it in terms of norms of inference that directly govern what one can or ought to infer from what, what utterances are incompatible with one another, and so on. The patterns of inference generated by these norms are constitutive of conceptual content, and the norms themselves are the product of the practices of discursive beings taking up normative attitudes toward one another's actions. Attitudes beget statuses beget inferential norms beget content.

4.1. Inferentialism and the Frege-Geach Problem

How does orthodox inferentialism fare with respect to the problems we've just explored for traditional expressivism? I argued that the primary mistake of the traditional expressivist is her elision of force and content. In taking the conative attitude expressed by a normative claim to be determinative of its meaning, she cannot explain its meaning in contexts in which it is not asserted (or endorsed). As such, on her account, the

²¹ This idea of scorekeeping was introduced by Lewis in a somewhat different context (Lewis, 1979).

following argument appears to be an equivocation, for the antecedent of (1) and (2) cannot mean the same thing:

- (1) If one ought not to lie, then one ought not to get little brother to lie.
- (2) One ought not to lie.
- (3) Therefore, one ought not to get little brother to lie.

The argument does not equivocate, so there must be something amiss with the expressivist account.

Orthodox inferentialism answers the Frege-Geach challenge by basing its account of the ingredient content of complex claims in norms of use rather than in mental states or attitudes expressed. Though moral claims may be associated with various practical or evaluative attitudes, their conceptual content is fixed by the totality of their inferential articulation not merely by their pragmatic significance (M. Williams, 2010, p. 325). The inferential articulation of the unembedded assertion, (2), is just the same as that of the embedded occurrence in (1). In fact, (1) just makes explicit one part of that inferential articulation of (2). In both cases, “One ought not to lie” is inferentially linked to commitments to refrain from lying, to condemn instances of lying, to feel guilt for having lied, and a whole host of other language-exit transitions. Though these commitments need not express an actual mental state of the speaker in the embedded case, given appropriate contexts of use, those inferential connections are activated, making it appropriate to attribute these commitments to the speaker and hold her accountable to them.²² In the *modus ponens* argument, for example, though (1) is not used to condemn either lying or

²² Compare the rules that constitute a piece on chess board as a pawn. Though a pawn might not actively be capturing *en passant* or even be capable of doing so in a particular position on the board, knowing that it is capable of doing so is part of understanding what it is for it to be a pawn. Likewise, knowing that, in the right contexts, “Lying is wrong” is inferentially bound up with a negative evaluative attitude toward lying is part of understanding the meaning of that sentence even if it is not actively displaying that part of its meaning in the present context (Warren, 2015).

getting little brother to lie, it makes explicit the inference from a claim the meaning of which is in part constituted by its inferential connections to acts of condemning and refraining from lying to a claim the meaning of which is in part constituted by its inferential connection to acts of condemning or refraining from getting little brother to lie.

Though orthodox inferentialism addresses Frege-Geach type worries in this way, it is, in fact, the long way around. The inferentialist account of the explicative role of logical vocabulary already sidesteps the problem. According to Brandom, the conditional in (1) is not required for the goodness of the inference from (2) to (3). The *modus ponens* argument does not underwrite the inference. Rather, it makes explicit a materially good inference already implicit in the commitment to (2). The inference from “One ought not to lie” to “One ought not to get little brother to lie” is good because of the content of the *nonlogical* vocabulary involved and need not be *completed* by the introduction of a conditional. Brandom is here taking a lesson from Sellars. “We need not treat all correct inferences as correct in virtue of their form... We can treat inferences such as that from ‘Pittsburgh is to the west of Philadelphia’ to ‘Philadelphia is to the east of Pittsburgh’... as *materially* good inferences” (2000, p. 85).²³ The conditional does not complete the inference but serves to make this inferential commitment explicit in the form of assertible (propositional) contents.²⁴ Without this explicating vocabulary, a linguistic community

²³ Sellars’ treatment can be found in (1957).

²⁴ Sellars and Brandom are both drawing here on Ryle’s treatment of hypotheticals as *inference tickets*. Knowing a hypothetical like “*if p, then q*” gives one license or warrant to “make a journey” from *p* to *q*, i.e., to draw the relevant inference, but “one can have an inference warrant without actually making any inferences and even without ever acquiring the premises from which to make them.” According to Ryle, we accumulate inference tickets in order to be equipped to draw inferences, but we only show them to satisfy others that we have the warrant we implicitly claim in the inference or for pedagogical purposes (Ryle, 1950, p. 308).

could still attribute inferential commitments and entitlements and accept and reject inferences, but the addition of the conditional allows interlocutors who disagree over the goodness of an inference to argue about whether the conditional is true, i.e., whether the inferential link it explicates is correct, by challenging entitlement to it and offering counterclaims (1994, Chapter 2, sect. 4). If the expressive function of the conditional is just to make explicit an already materially good inference, then inferentialism never actually faced the Frege-Geach problem in its canonical form.

Orthodox inferentialism offers an anti-Representationalist account of moral vocabulary that doesn't fall to Frege-Geach type problems. One might object, however, that it is itself a pretty radical position. The main objection I raised to the traditional expressivist responses to embedding problems was not that they failed but that they came at a high price. In particular, they forced the theorist into strong commitments about the semantics of ordinary descriptive vocabulary as a result of the apparatus required to supply a compositional semantics for normative vocabulary. Expressivists had to commit to mentalism and to the idea that all sentences, whether they are normative or descriptive, express the same attitude. Now I am proposing a "solution" to embedding problems that, once again, requires a complete revision of our semantics for ordinary descriptive vocabulary. How is this any better? The difference is that inferentialism is *already* a well-worked out position for the semantics of ordinary descriptive uses of language. In fact, this is the discursive practice for which it is most fully developed.²⁵ As such, what I am proposing is not a position in metaethics that rewrites our semantics for descriptive uses

²⁵ Inferentialism began as an account of the meanings of logical vocabulary, but Brandom's account is a full-fledged attempt to make inferentialism work, in the first instance, for empirical discourse. He then extends the position to account for modal and normative vocabulary.

of language but rather the extension of a different systematic approach to the semantics of ordinary descriptive uses of language to the context of metaethics.

Orthodox inferentialism looks like a promising anti-Representationalist approach that avoids the problems faced by traditional expressivism. Now note, however, that the Negative Thesis of Expressivism (NTE) remains intact in the face of the challenges to the positive theory. Nothing I have argued thus far requires us to reject the commitment to functional pluralism, i.e., that different vocabularies fulfill different needs and purposes of the beings who use them, and their meanings should be understood in terms of these functions. The arguments we've reviewed only put pressure on the traditional expressivist's attempt to theorize about functional pluralism in terms of expressive use. The question, then, is whether orthodox inferentialism, by focusing on practical deliberation rather than the mental states expressed, provides the resources both to capture the functional pluralism of natural languages and to explain it. Answering this question requires that we fill in the details of an inferentialist account.

5. Warren's Hybrid Expressivist Inferentialism

Suppose you want to capture the central insights of expressivism within an inferentialist framework. The hope is for an end-run around the embedding problem that still captures and explains functional pluralism by drawing on the deep bench of expressivist arguments. The most direct way to try to affect this integration is to treat the expression relation as an inferential relation. Expressivists have identified the mental states or attitudes expressed by normative claims. For the inferentialist, these become integral components in the inferential articulation of those claims. To understand the meaning of a moral claim, for instance, one must grasp what mental state the speaker

must be in to felicitously utter the claim and what further speech acts and actions are licensed by her claim. If this strategy is effective, we'll be able to rely on expressivist accounts of the functional pluralism of language to draw lines between different vocabularies in terms of their inferential relationships to different kinds of mental states.

Warren proposes a straightforward integration like this. His strategy is to identify a set of introduction and elimination rules for moral "ought." These are rules of inference that govern from which other claims "oughts" follow—introduction rules—and which other claims follow from "oughts"—elimination rules.²⁶ Focusing on these rules, Warren aims to identify how the relevant mental states are situated with respect to moral assertions. When an individual, *U*, asserts "one ought to ϕ ", her assertion licenses a variety of inferences or moves in the language game. For example, he says, it licenses one to infer "that *U* has attitudes in favor of ϕ -ing, including perhaps motivation to ϕ , a commitment to approve of one who ϕ 's, and the second-order belief 'one ought to be motivated to ϕ .'" Furthermore, one is licensed to infer "that one has similar attitudes and commitments against not ϕ -ing" that are appropriately related to feelings of guilt, disapproval, and disgust. These elimination rules tell us that the assertion of "one ought to ϕ " licenses inferences about the mental state of the speaker, about her further commitments, and about what actions she might be committed to performing. On the introduction side, Warren says that if *U* has the appropriate constellation of conative attitudes toward *S*'s ϕ -ing, e.g., a disposition to approve of it, to feel disgust at *S*'s failure to ϕ , to punish *S* for failing to ϕ , and to feel guilt when *U* herself fails to ϕ , then "*U* is

²⁶ In sentential logic, for example, the introduction rule for the logical conjunction " \wedge " tells us that we can infer $A \wedge B$ from the premises *A* and *B* while the elimination rule for " \wedge " tells us that we can infer either *A* or *B* or both from $A \wedge B$.

defeasibly licensed to assert that *S* ought to [ϕ].” Asserting “*S* ought to ϕ ,” on Warren’s account, is acceptable just when the speaker is in the appropriate mental state (Warren, 2015, pp. 1871–1872).

How does this direct integration approach fair? Not well, I’m afraid. It is a rather ham-fisted approach that fails to fully shift from the mental state dominated perspective of expressivism to the practical deliberation focused perspective of inferentialism. A mash-up of theories—even those with some affinity to one another—is not likely to produce the hoped-for result, and in this case, it falls far short. The view suffers two decisive problems. The first is that in giving explanatory precedence to conative mental states, the direct integration approach forces Warren to account for inferential proprieties in terms of those mental states rather than in the more direct way available to inferentialists. The second is in treating the expression relation as an inferential relation, the view leaves us with nothing useful or enlightening to say about the way that moral claims function as part of our reason-giving practices. We will be forced either to treat conative mental states as *reasons for* moral claims, which does not track our actual use, or to treat moral claims as bare expressions of the speaker’s reliable reaction to environmental stimuli, giving them the wrong role in our reason-giving practices and threatening the rationality of moral discourse.

Traditional expressivism arose as a local theory in metaethics aimed at accounting for the meanings of normative vocabularies within a broadly Representationalist semantic program. To situate their account within such a program, expressivists needed to adopt—at least in broad outline—the standard order of explanation found in Representationalist semantics. The explanandum, generally speaking, is linguistic behavior. Why do we use

words, sentences, clauses, signs, etc. in the ways that we do? Why do they exhibit the patterns of use and inferential properties that they exhibit? The Representationalist tries to answer these questions, in part, by positing a property of lexical items called “meaning.”²⁷ The meaning of a sentence explains (again, in part) the propriety of the inferences in which it plays a role as premise or conclusion, why it is appropriate to utter in some circumstances and not others, and why it exhibits features like objectivity, if it does. Meaning is eventually explained in terms of the semantic primitives of truth and reference, so the inferential features of a sentence are, in the end, explained in terms of the objects and properties that it purports to represent. The inferential relations in which ordinary empirical claims stand to one another are explained, in the end, in terms of the relationships of the objects and properties those claims are about. Traditional expressivism, of course, rejects the Representationalist appeal to truth and reference to explain meaning, yet they still have to explain linguistic behavior in a similar fashion, i.e., by appeal to something that accounts for the meanings of lexical items which can then account (again, in part) for their patterns of use. This is the role that mental states play in the expressivist account. The particular features of the mental states that are expressed must be such that they can explain why we use these sentences in the ways that we do, why they stand as they do to our motivational psychology, why they seem to be objective and universally binding, and so on.²⁸

²⁷ This is, of course, a massive oversimplification. Representationalists have posited different kinds of contents as the meanings of sentences and have developed programs for explaining some aspects of linguistic behavior in terms of context and pragmatics. For just one representative example, see (Recanati, 2004). I will not here engage this massive literature, as I only need the basic shape of a Representationalist theory to make my point that expressivists have been forced to shape their theories in a particular way.

²⁸ For example, the negative attitude associated with judging that some act is wrong cannot be simply an attitude of disfavoring it or disapproving of it. The attitude must in some way capture the sense that the act would be disavored for anyone in similar circumstances, that others should have a similar attitude toward the act, and that one takes one’s attitude toward the act to be appropriate, i.e., that one would not change

Orthodox inferentialism turns this Representationalist order of explanation on its head. Rather than accounting for patterns of use in terms of meaning, inferential articulation is taken to be constitutive of meaning. The inferentialist then goes on to explain inferential proprieties in terms of the norms that govern use. Traditional expressivists needed to appeal to *sui generis* mental states in order to explain the different patterns of use we observe in different vocabularies that are given the non-Representationalist treatment. Orthodox inferentialists, though, have no such need. They can explain these patterns *directly* in terms of the norms that produce them. That inferentialists can explain the important features of moral discourse without relying on an account of *sui generis* moral attitudes does not imply that she should not make such attitudes central in her account, but there are significant costs to doing so. One is that it introduces an unnecessary dogleg in their explanation of the meanings of moral vocabulary. Inferentialists who adopt the direct integration approach will need an account like the expressivist one to explain how these mental states shape the practice of moral discourse, and this explanation will need to proceed by way of the norms of the practice, for these are what govern which moves are appropriate. This is a taxing demand given that inferentialist explanations of these norms tend to be given in terms of the purpose or function of the practice.²⁹ There is also a more direct cost to positing *sui generis* moral attitudes. It commits inferentialists to claims about human psychology to which they need not be committed. The expressivist can't get by without moral attitudes, but inferentialists

one's attitude toward the act without good reason. Expressivists have posited complex moral attitudes that mirror these features of universality as well as higher-order attitudes toward one's first-order attitudes that might account for these features. For an example of the first tact, see Gibbard's planning attitudes in (2003); Blackburn develops a proposal of the latter kind (1993).

²⁹ This problem is a close cousin of the problem of mentalism that traditional expressivists face.

only need some generic motivational states. They do not need these states to have special features that help explain linguistic behavior. Instead, they can keep open the possibility that human beings have fairly simple motivational psychologies but that the inputs to them are complicated by the various cooperative and decision-making practices that we have evolved.

Let's turn now to the second, more pressing problem with Warren's attempt to affect the direct integration strategy. Recall that, on his account, the introduction rules for "*S* ought to ψ " are stated in terms of the attitudes that would defeasibly license its assertion by *U*.

If *U* has a disposition or commitment to approve of *S*'s ψ -ing, to disapprove or feel disgust at *S* not ψ -ing, to punish *S* for not ψ -ing, to pressure third parties to react punitively to *S*'s failure to ψ , etc., and if the practical significance of these commitments is not contingent on the desires or goals of particular agents, then *U* is *defeasibly* licensed to assert that *S* ought to ψ (2015, p. 2872).

There are two ways we might understand this move. On the first, Warren's strategy is essentially to treat the expression relation as an inferential relation. Where expressivists would say that "*S* ought to ψ " expresses approval of ψ -ing, disgust at not ψ -ing, and so on, Warren says that it follows inferentially from these attitudes. I think what is happening here is a tendency on Warren's part to focus on the affinity for assertability semantics exhibited by both expressivism and inferentialism. For the expressivist, the assertability conditions for a moral claim are given in terms of the mental state one must be in if one's assertion is to be appropriate.³⁰ This is one way of loosely cashing out the expression relation. Warren translates this into the inferentialist idiom by recasting the assertability conditions as introduction rules. The problem here is that inferentialist

³⁰ See (Schroeder, 2010, pp. 28–34).

introduction rules are not simply assertability conditions. Introduction rules track inferential relations, that is, they track which claims stand as *reasons* for which others. Assertability conditions need not do this, as can easily be seen in the expressivist case. Though it may be true that one should only say “*S* ought to ψ ” when one is in the requisite mental state, one cannot cite that one knows that one is in that mental state as a reason for one’s claim!³¹ The expression relation is not an inferential relation.

Perhaps, though, a more charitable interpretation is available. Another way to understand his account of the introduction rules for the moral ought is to take them to be of the same kind as the orthodox inferentialist’s introduction rules for an observation term. “Red” has observational uses. It can be applied *non-inferentially* to red things via reliable dispositions to respond differentially to those objects, and, though others can infer the proximity of a red object from a speaker’s report, one cannot cite one’s disposition to respond with “red” as a reason for one’s claim that the thing is red.³² In this way, the inferential patterns constitutive of the meaning of “*x* is red” are similar to those Warren takes to be constitutive of the meaning of “*S* ought to ψ .” Both are the product of a reliable disposition to respond differentially to environmental stimuli. In the former case, a scorekeeper can infer something about the environment from the report, but, in the latter, can she infer the presence of *S*’s obligation to ψ from *U*’s report? An affirmative answer would mean that moral terms function like observation terms, but it is far from clear that they do. While we are likely to take a report that *X* is red as evidence that it is,

³¹ Notice that this collapses the distinction between expressing and reporting that the traditional expressivist was at pains to draw.

³² One might cite one’s reliability in reporting on redness in the face of a well-motivated challenge to one’s claim, but this is not a reason in support of the content of the claim but rather in defense of one’s standing to make it. Moreover, there is a whole host of other issues about the possibility of coming to know that one is reliable in this way. See, for example, (R. Brandom, 1998, 2002; Kukla, 2000; Sellars, 1956).

we tend only to take a report that *S* ought to ψ as evidence that the speaker thinks this is the case and is willing to proffer reasons in support of the claim. The report is certainly not dispositive that *S* has such an obligation.³³ If, however, we take Warren's sketch as an account of moral terms as observation terms then it's not clear what kinds of reasons these could be. That someone takes themselves to be reliable at picking out instances of moral obligation is not a convincing reason that, in the present case, a moral obligation exists, but Warren's account gives us no foothold for thinking about what other claims might stand as reasons for a moral claim and, so, no way of making sense of the rationality of moral discourse. In the end, it looks like moral claims on this account can be nothing more than bare expressions of the speaker's moral attitudes, and, if that is all that they are, then we have no resources at all for explaining the interesting surface features of moral discourse.

Orthodox inferentialism is an advance over traditional expressivism as a non-Representationalist theory of meaning, but, as we should expect, it matters a great deal how we fill in the details, i.e., how we understand the inferential connections constitutive of the meaning of moral "ought." Warren's attempt shows the danger of hewing too closely to the expressivist line in an attempt to capture the functional pluralism to which

³³ There is an interesting question here about the role of perception in moral discourse and deliberation. I do not want to deny that we can perceive moral states of affairs, nor do I want to deny that some people have better attuned moral perception than others, but our dispositions to reliably respond deferentially by applying moral concepts plays a different role in moral discourse than does our disposition to respond to visual stimuli by applying empirical concepts. In the latter but not the former, for example, a reliable report can settle a dispute. The problem for Warren's account, if we adopt this understanding, is that moral claims will indicate that the speaker has responded to environing stimuli with a particular set of conative attitudes, and we will be forced either to say that those attitudes are dispositive about the moral state of affairs or they merely indicate the speaker's conative response. If the former, then moral perception is taken to be like visual perception in a way that it is clearly not like visual perception. If the latter, then we are left with moral claims being mere verbal ejaculations and have no resources for making sense of the rationality of moral discourse.

traditional expressivists rightly drew our attention. If we try to affect a direct integration of expressivism and inferentialism by treating the expression relation as an inferential relation, the outcome will force us into commitments about motivational psychology and leave us without the resources we need to explain the surface features of moral discourse.

The solution is to accept the inferentialist focus on practical deliberation and order of explanation and, from that starting point, see how we can explain functional pluralism as well as the various features of use distinctive of moral vocabulary. Inferentialism aims to explain patterns of use directly in terms of the norms governing the use of a vocabulary and explain those norms in terms of what we use that vocabulary to do. Luckily, we have a model to turn to for this sort of an inferentialist account of deontic normative vocabulary. Brandom has advanced a sketch of “oughts” in terms of their various explicative functions. Normative claims serve to explicate proprieties of practical inference, and different kinds of normative claims explicate different kinds of practical commitments. On this approach, we get functional pluralism in terms of different explicative functions, and we explain proprieties of inference in terms of the different kinds of practical commitments that are explicated. Let’s take a look.

6. Brandom on the Explicative Function of Deontic Normative Vocabulary

According to Brandom, deontic normative vocabulary has an explicative function like the conditional, but, instead of doxastic inferential commitments, it explicates practical inferential commitments. The standard approach to practical inferences is Davidson’s (1963). On his account, a *primary reason* consists in a belief/pro-attitude pair that together constitute a reason for action. If either part is left out, then we cannot understand why the action is reasonable from the perspective of the individual

performing it. A practical inference from a belief about some state of affairs to an intention to act must be bridged by a pro-attitude of some sort. So, for example, the inference from “Only getting on this train will get me to the theater before the show starts” to “I shall get on this train,” implicitly relies on the pro-attitude that could be expressed by “I prefer to see the beginning of the show.” Brandom contends that, as in the doxastic case, the original inference is not enthymematic. It is a materially good inference as it stands. The role of the evaluative claim is to make explicit the inferential connection between the doxastic premise and the practical conclusion opening up “a new venue...for resolving disagreements about what follows from what, about which claims rationalize which actions” (R. Brandom, 1994, p. 248). To take as entitlement preserving the pattern of reasoning ensconced in the inference and others related to it—for example, from “I must leave the house by seven to make it for the start of the show” to “I shall be ready by seven”—just is to attribute the relevant pro-attitude to the speaker.

Brandom generalizes this point to all deontic and evaluative claims. Thus, we should understand the inference from “Sam is a fellow human being” to “I shall not cause Sam pain unnecessarily” as materially good, and the assertion “One ought not to cause fellow human beings pain unnecessarily” not as completing the inference but as explicating an inferential commitment already implicit in taking Sam to be a fellow human being. Likewise, the inference from “I am a teacher” to “I shall be responsive to my students’ inquiries,” is materially good, and “Teachers ought to be responsive to their students’ inquiries” explicates the commitment already implicit in taking someone to be a teacher. Once this inferential propriety is made explicit, “it can be criticized, supported, refuted, in short, evaluated” (Sellars, 1962, p. 374). However, we need not be able to

make it explicit in order for the inference to be a good one. Like conditional vocabulary, deontic and evaluative vocabularies are optional, though perhaps very useful, additions to our linguistic repertoire. Greater precision in our criticism and evaluation of our implicit inferential commitments is the value added in making them assertible.³⁴

At this point we can summarize the role in reasoning played by normative and evaluative vocabulary, generally. This vocabulary serves to make explicit material proprieties of practical reasoning. A speaker is entitled to a normative or evaluative claim if she undertakes a commitment to (or endorses) the inference it explicates. Note that this is a normative rather than a psychological condition. It is of no concern what attitudes the speaker harbors or what her overall mental state is. Instead, what matters is the position she takes up in a normative space, i.e., what she is committing herself to in uttering the claim. She is entitled to utter the claim only if she is willing to undertake a commitment to the goodness of the inference it explicates. In undertaking such a commitment, she makes herself liable for the claim. She must be prepared to respond to challenges by providing reasons. On the downstream side, her utterance commits her to further claims about the action or state affairs under consideration, to praising those who act in accord with her judgment, and to sanctioning those whose actions violate it. This includes rightfully feeling pride (or similarly positive feelings) for compliance and shame for

³⁴ This discussion might leave the reader with the impression that inferentialism, as I understand it, takes reasoning—inference—to be a matter of drawing logical implications, both formal and material. This would be unfortunate, as I am not committed to this view. As Harman argued, the principles of reason are principles of belief and intention revision, not principles of logic. Inferentialism tells us that in order to understand a sentence, we must grasp the role it plays in a space of implications, but the role that it plays in the space of implications does not determine which implications we ought to draw. Whether what we have before us should be taken as a *modus ponens* or a *modus tollens* cannot be determined by the inferential structure alone. As such, when I use the term “role in reasoning” I mean only to invoke this idea of a position in the space of implications and not to further imply that the position occupied by a claim somehow requires us to reason one way or another with it (Harman, 1986). I am indebted to Michael Williams for prodding me to handle this point with greater care.

violation if she has judged herself to be in the conditions in which her normative or evaluative claim applies to her. She also comes to be committed to appropriately related inferences. One who undertakes commitment to the claim, “I prefer to see the beginning of the show,” is further committed to forming intentions to act in the variety of ways that facilitate making it to the show on time. Again, what is of importance is the normative position of the speaker.

Beginning with the explicative function of deontic normative vocabulary, we’ve now worked our way back to introduction and elimination rules, but they look quite different than the rules that Warren sketched. Rather than identify dispositions to respond with particular attitudes in the introduction rules, we’ve said what one must do to secure entitlement to a claim. One must be able to give the right sort of reasons. The conditions are normative, not psychological. We cannot say, however, precisely what will count as an entitling reason or what commitments will follow from a normative claim, for these are dependent on the context of the claim, on what one’s interlocutors will be willing to accept as a good reason, and what further commitments they will attribute to the speaker. What entitles a claim and what follows from it is something to be worked out in the messy retail business of giving and asking for reasons, not something that can be determined *a priori* by the theorist.

Now that we have Brandom’s sketch of the explicative function of deontic normative and evaluative vocabulary, we should ask whether this account has the resources to capture the key expressivist insight of functional pluralism. On first gloss, the answer seems to be yes. Where the traditional expressivist explained pluralism in terms of the use of different vocabularies to express different mental states, the orthodox

inferentialist can explain it in terms of the different explicative functions played by vocabularies that are not merely descriptive. Brandom, for example, provides accounts of deontic normative vocabulary and modal vocabulary in these terms, and we can see how these might be extended to evaluative vocabulary (Brandom, 2008a). Problems arise, though, when we consider the variety within some of these categories. Deontic normative vocabulary is not monolithically moral. Alongside the moral “ought,” we have at least the prudential, the institutional, and the “ought” of tastes, as well as other seemingly universal oughts like the epistemic.³⁵ Similarly, evaluative vocabulary comes in many flavors. Traditional expressivists could distinguish these in terms of the different mental states that they expressed, but what is the orthodox inferentialist to do?

7. Brandom on the Flavors of Ought

Brandom has a proposal for distinguishing these varieties of ought. Each flavor of ought, he claims, explicates a different kind of norm, and these “different sorts of norms correspond to different patterns of practical reasoning” (2000, p. 91). The trick, then, is to identify the pattern of reasoning associated with each so that we can go on to say, for example, that the moral “ought” is associated with pattern A, the prudential “ought” with B, and so on. A pattern of reasoning is distinguished in terms of the further commitments that can be properly attributed to the speaker. Some examples will help to clarify.

Consider the inference from “Exercising is part of a healthy lifestyle” to “Tom shall start exercising.”³⁶ Supposing that this is a materially good inference, it might be explicated by

³⁵ For a promising inferentialist account that distinguishes epistemic, semantic, and moral oughts, see (Chrisman, 2015).

³⁶ “Shall...” will be used to indicate the formation of an intention to act, i.e., a language-exit transition. As such, this is an inference from an empirical claim about the relationship between exercise and health to an intention to begin exercising.

the deontic normative claim, “Tom ought to adopt a healthy lifestyle.” In accepting the goodness of the original inference, I am implicitly committed to this normative claim. This commitment has further ramifications, for the original inference “is just one of a whole family of inferences that stand or fall together” (1994, p. 249). For example, the inferences from “Driving the speed limit is part of a healthy lifestyle” to “Tom shall drive the speed limit” and from “A balanced diet is part of a healthy lifestyle” to “Tom shall eat a balanced diet” are both underwritten by the same commitment, and, so, both must be taken to be materially good if the original is. This set of inferences (expanded indefinitely) constitutes the pattern of reasoning associated with “Tom ought to adopt a healthy lifestyle.” Whether one explicitly undertakes a commitment to this normative claim, one who endorses this pattern of reasoning implicitly attributes to Tom an *interest* in living a healthy lifestyle, and to do so is to implicitly undertake a commitment to a *prudential* “ought.” It is, for Brandom, this pattern of reasoning that distinguishes an “ought” as prudential.³⁷

To see what makes this pattern of reasoning distinctive, let’s consider two examples with which it contrasts. First, let’s look at the inference from “Iona is a police officer going to work” to “Iona shall wear a uniform.”³⁸ This inferential commitment can be explicated by, “Police officers ought (are required) to wear uniforms.” What is the

³⁷ Brandom’s example in this case is actually an instance of what I will call an “ought” of taste/preference. He dubs it both instrumental and prudential, but I think there is reason to distinguish between these two such that instrumental (or taste/preference based) practical reasoning is means-ends reasoning about fulfilling one’s desires, preferences, or tastes, about which one is ultimately authoritative whereas prudential (or interest based) practical reasoning is means-ends reasoning about satisfying one’s interests, which may be opaque to the individual and so about which she is not authoritative in the same way. I develop this idea in more detail in Chapter 5. For now, I am merely amending Brandom’s example because the question about differentiating kinds of “ought” is more easily stated in terms of prudential “oughts” rather than “oughts” of taste/preference. The latter are more readily made explicit not with “ought” but with expressions of preference, e.g., “I want (desire, prefer) to stay dry” (1994, pp. 245–249, 2000, pp. 84–89).

³⁸ This example hews more closely to Brandom’s own bank employee case.

pattern of reasoning associated with this deontic normative claim? In endorsing this inference, Brandom thinks, one will also be committed to the inference from “Iona is a police officer going to work” to “Iona shall carry a badge” and from “Iona is a police officer going to work” to “Iona shall faithfully execute her duties as prescribed by the law.” The pattern of inferences one undertakes commitment to in this case maps a set of norms that define an institutional role or status.³⁹ In undertaking a commitment to the goodness of one of these inferences, one implicitly undertakes a doxastic commitment to the claim that Iona is a police officer and, so, is implicitly committed to the goodness of the other practical inferences that follow from her having that status. While the prudential ought was associated with a pattern of reasoning that corresponded to the attribution of an interest, this *institutional* “ought” corresponds to the attribution of an institutional status (R. Brandom, 2000, pp. 90–91).

One final example. The inference from “Jerry repeating the gossip would harm someone, to no purpose” to “Jerry shall not repeat the gossip” is explicated by a normative principle such as “One ought not to harm anyone to no purpose” (R. Brandom, 1994, p. 245, 2000, p. 84). Endorsing this inference, then, commits one to other inferences such as from “Jack punching his sibling would harm someone to no purpose” to “Jack shall not punch his sibling” and from “Joe insulting his friend would harm someone to no purpose” to “Joe shall not insult his friend.” Someone who takes the pattern of reasoning to be entitlement preserving for some particular individual, that is, is committed to taking it to be entitlement preserving *for anyone* without regard to desires, preferences, or social statuses. Brandom dubs the “ought” corresponding to this pattern

³⁹ For example, police officers ought to carry a badge, faithfully execute the law, and so on.

unconditional, which, he says, may or may not correspond to the moral “ought” (2000, p. 91).

Brandom’s strategy is to identify a pattern of inferences to which one becomes committed in endorsing some particular practical inference. Each distinctive pattern is then associated with some flavor of “ought.” The recipe for doing so is fairly simple. First, determine what kinds of inferences the target vocabulary is used to explicate. Next, identify the variety within that category exhibited within natural language. Finally, determine the distinctive pattern of inferences underwritten by the explicating claims and associate them with the different flavors. This strategy seems fine to me if our aim is a taxonomy. We can clearly draw lines between different explicating vocabularies in terms of the inferential patterns they underwrite. The problem, as the recipe makes plain, is that we are doing this in an *ad hoc* way. What’s missing is some deeper explanation of these patterns of inference, the kind of explanation that the Representationalist gives by appeal to normative properties and the expressivist gives by appeal to the features of the mental states expressed. Why, we want to know, does our reasoning exhibit these patterns of inference and how are they sustained? What explains them?

8. Conclusion

Perhaps they need no explanation. Representationalism explains this vocabulary in terms of the objects and properties it is about, expressivism in terms of the states of mind expressed, but perhaps in adopting the inferentialist order of explanation it is enough to say that these patterns of reasoning are exhibited by our practices, and this is the vocabulary we use to make them explicit. Aiming at some *deeper* explanation invites a re-inflation of this vocabulary that the non-Representationalist aims to avoid. I find this

kind of response wholly unsatisfying. There is clearly a phenomenon here to explain. Nothing requires us to adopt the same kind of explanation as either the Representationalist or the traditional expressivist, but we also need not be forced into quietism. In the next chapter, I propose a novel way of advancing an explanation along inferentialist lines that should be welcome to most anti-Representationalists. My strategy fundamentally involves two steps. First, I will aim to explain the patterns of reasoning identified by Brandom in terms of the discursive norms that practitioners enforce on one another when engaged in a discursive practice. I'll identify various kinds of discursive norms—inferential, fundamental, and structural—and argue that the structural norms are those that give each discursive practice its distinctive shape, i.e., that generate the patterns of reasoning we associate with the practice. The second step, then, is to explain those norms. The explanation I advance is functionalist; the norms are explained in terms of the function of the practice they govern. The result is an explanation of the patterns of reasoning exhibited by different discursive practices in terms of what good those practices are for the beings who use them.

Chapter 2: Refining Inferentialism: From EMU to PALM

1. Introduction

In the last chapter, I argued that the Negative Thesis of Expressivism—that some vocabularies do not aim to describe the world and that meaning is best explained in terms of the various functions of these vocabularies—is on the right track but that traditional expressivist attempts to advance a positive account of these functions proved too costly. I then argued that orthodox inferentialism could avoid the problems faced by expressivists and found that it provides a ready answer to the Frege-Geach problem without the attendant costs. We were left with a problem, however. Orthodox inferentialism has the resources to make sense of meaning in terms of use and, as Brandom develops it, to draw some lines between different kinds of assertions roughly where we think they should go. It can distinguish, for example, between ordinary descriptive claims, modal claims, and deontic normative claims in terms of the ways in which the latter two function to explicate inferential proprieties of the former. Moreover, it can distinguish between different deontic categories—prudential, institutional, moral, and so on—in terms of the different pattern of reasoning with which each is associated. Herein lies the problem. These patterns of reasoning, for Brandom, seem to be brute, oddly inexplicable facts. He explains the explicative function of deontic normative vocabulary, in general, in terms of the utility to be had by making explicit practical inferential commitments so that they can

be more carefully examined, but when he turns to the patterns of reasoning exhibited by various flavors of “ought,” he leaves us wondering from whence they come.

These patterns of reasoning—like all patterns of language use—are the product of each of us taking up normative attitudes toward the linguistic acts of our peers, but why do we take up the particular attitudes that we do? What norms are we simultaneously instituting and enforcing in our treating certain patterns as entitlement preserving and others as not? And, just as centrally, why do we enforce these norms and not others? If we can answer these questions, then we can advance a complete account of functional pluralism for inferentialism that distinguishes between vocabularies in terms of the patterns of inference they exhibit and explains why they exhibit those patterns in terms of the functions of the practices that deploy these vocabularies.

The answers I’ll propose consist in a framework for a pragmatic account of linguistic meaning (PALM). PALM is an inferentialist account in that it treats the broadly inferential articulation of a claim as constitutive of its meaning, but it expands upon orthodox inferentialism by attending to additional discursive norms that differentiate discursive practices (and, thus, vocabularies) and sustain the patterns of reasoning distinctive of those practices. Moreover, these additional species of discursive norms afford us some explanatory purchase in our effort to understand why different discursive practices exhibit the patterns of reasoning that they do, for we can explain these norms in terms of the utility that the practice has for the beings whose practice it is. Like the various structural features of a tool that enable its function, these structural discursive norms enable a practice to fulfill its function.

I will argue that there are two relevant species of structural discursive norms:

pragmatic norms and norms of authority. These are best thought of as part of the normative pragmatics of a discursive practice, i.e., they are norms of assessment, guiding practitioners in tracking the commitments and entitlements of others. We implicitly enforce these norms on one another, but we likely lack explicit awareness of them. They come out of the darkness in drips and drabs as we make more and more of our discursive lives explicit.¹ They are primarily habits, patterns of behavior, that are summarized in their explication.² Moreover, when they are made explicit, it is only in a piecemeal way and with an implicit *ceteris paribus*, for they are deeply contextual. If something goes awry in the discourse—if someone flouts a convention or repeatedly violates a norm—interlocutors may try to make that norm explicit in order to call attention to the violation, to hold the speaker to the norm, or to sanction her, but even in such cases, the explicated norm is *defeasible* in ways we could not predict, i.e., we could not list all of the conditions under which it might not obtain.³

The purpose of this chapter is to motivate the introduction of these structural discursive norms into an inferentialist account of meaning and to sketch the resulting framework, i.e., PALM. In taking up this aim, I retreat for the time being from my central concern with moral discourse. PALM is a general account of meaning, so I will motivate

¹ The idea is that we are creatures who not only feel the pull of conformity but who also enforce that conformity—who are censorious—though often without being reflective about doing so. See, (Haugeland, 1982).

² Rorty captures this idea quite nicely, though in a different context, when he writes that “‘institutionalized norms’...take the form of bureaucrats and policemen, not of ‘rules of language’ and ‘criteria of rationality.’ To think otherwise is the Cartesian fallacy of seeing axioms where there are only shared habits, of viewing statements which summarize such practices as if they reported constraints enforcing such practices” (Rorty, 1991e, p. 26). We are each other’s bureaucrats and policemen instilling and enforcing various habits of behavior, and the norms I discuss in this text are just summaries of the practices in which we engage, not maxims that we explicitly adopt.

³ I am drawing here on Hart’s notion of a *defeasible concept*, i.e., one for the application of which a set of necessary and sufficient conditions cannot, in principle, be given because the set of defeating conditions need not have much in common with one another. These conditions are discovered in practice. See, (Hart, 1951; M. Williams, 2013b).

its development by drawing lessons from a variety of different speech act kinds, not only moral assertions. I take up moral discourse and deontic moral vocabulary again in Chapter 3.

The arguments that motivate PALM are subtle. I begin in the next section with a discussion of pluralism about discursive norms and their nature as part of the normative pragmatics of discourse. I close the section with an argument to motivate the introduction of epistemic norms as the first variety of structural discursive norms. In Section 3, I approach the more difficult task of motivating further norms of authority. There are two ways to approach this task. The first would be to argue that there are kinds of content—broadly inferential, but non-propositional content, for example—the inferential norms of which cannot be accounted for in terms of narrowly epistemic norms. If it could be shown, for example, that practical contents shaped by norms of authority other than epistemic norms exist, that would motivate including these other norms in PALM. Unfortunately, I can't see my way to such an argument.⁴ The second approach, the one that I take up in this chapter, is more circuitous. Rather than aim directly at contents that require additional structural discursive norms, I examine what kinds of speech acts—or speech act functions—are necessary for a practice to count as discursive and argue that some of those types of speech act functions are not governed by epistemic norms alone. I then argue that the norms that do govern those functions shape the content constitutive inferential proprieties of the vocabularies they deploy. If I am right that these additional norms shape the inferential norms of a practice, then they deserve to be counted among

⁴ Nuel Belnap at times pushes in this direction when he identifies, for example, imperatival and interrogative contents and claims that they cannot be accounted for in terms of the propositional contents of declarative speech acts along with some force operator. My argument would have to proceed along a similar line, but I haven't found it (Belnap, 1990).

the structural discursive norms.⁵

Having motivated the introduction of various structural discursive norms, in Section 4, I summarize the resulting explanatory structure drawing on frameworks developed by Michael Williams and by Rebecca Kukla and Mark Lance. These frameworks, when properly situated with respect to one another, constitute a pragmatic account of linguistic meaning that lays bare the way in which the function of a discursive practice shapes the structural norms to which we hold one another in engaging in the practice and, in the end, the inferential norms of the practice. Finally, in Section 5, I comment on some of the important features of PALM and prefigure some of the ways I'll put it to use in coming chapters.

2. The Plurality of Discursive Norms

Inferentialists have tended, quite naturally, to focus on norms of inference or conceptual norms. A commitment-preserving inference like that from “Pittsburgh is to the west of Philadelphia” to “Philadelphia is to the east of Pittsburgh” is a semantic relation between contents, but at the level of normative pragmatics, there is a corresponding inferential norm according to which, *ceteris paribus*, one who is committed to the first claim is committed to the second. Similar norms correspond to entitlement-preserving inferences and incompatibility relations (Laukötter, Prien, & Schepelmann, 2008, p. 82). It is not surprising that these norms are central to Brandom's normative pragmatics. His primary aim is an account conceptual content in terms of inferential articulation, and these norms are the ones we rely on to assess the moves that others make in the practice

⁵ This is the subtle part of the argument.

that is responsible for constituting such content.⁶

It is obvious that inferential norms are not the only ones deployed in assessing those moves. There must, for example, be a norm by which we assess whether one has incurred an obligation to justify a claim. When someone utters a claim, a scorekeeper assesses whether or not she is entitled to it, but, prior to that, she must have assessed whether it counted as a claim at all, that is, whether in producing the sounds that she did, the speaker incurred an obligation to perform certain further acts (usually the production of further claims that stand as reasons in support of the original claim) under certain conditions (such as a being confronted with a well-motivated challenge). The norm by which this is assessed is not itself an inferential norm, for there is no inference to which it corresponds, yet it shapes the discursive practice by identifying individuals as loci of responsibility for claims. There must, in addition, be a norm requiring a revision in one's commitments when one discovers one has incompatible commitments. Again, this is not an inferential norm, but we would not be engaged in the practice of giving and asking for reasons if we did not hold others to account for their incompatible beliefs. Laukötter, et al., dub these *fundamental discursive norms*, as “they have to be present in any practice that is to count as the game of giving an asking for reasons” (2008, pp. 82–83).⁷

The picture we have, then, is that there are inferential norms that are used to assess commitments, entitlements, and incompatibilities incurred by the speaker of a claim and fundamental discursive norms that are employed in the assessment of whether a speaker has undertaken a commitment (and, so, is responsible for defending it) and

⁶ It is important that we are thinking about these norms in the right way. As Brandom puts it, all of the norms with which he's concerned are “norms of assessment” that “have a grip” on practitioners in the sense “that one is liable to be assessed as having correctly or incorrectly done things” (Brandom, 2008b, p. 175).

⁷ Brandom agrees that there must be such fundamental discursive norms (2008b, p. 174).

whether a speaker must revise her commitments. I think, however, that we can identify a further class of discursive norms that have hitherto been treated either not as properly discursive or not at all. Let's begin by looking at an example of scorekeeping. A speaker, Elba, after placing a drop of liquid on her tongue, utters, "That solution is an acid." How do we as scorekeepers score this claim? First, we determine that it is a claim, i.e., that she has undertaken an obligation to provide reasons in the face of challenges. Next, we update her score *from her doxastic perspective* to reflect this new commitment. We add new commitments she has undertaken as a consequence of this commitment and her background beliefs, for example, that blue litmus paper dipped in the liquid would turn red, and we check for incompatible claims to which she is committed, for instance, that it was poured from a bottle marked "Calcium Hydroxide." At the same time, we update her score *from our own doxastic perspective*. We attribute to her commitments that follow from our own background commitments rather than from hers. For example, if we know that when zinc is submerged in an acid solution, the reaction gives off hydrogen, then we will attribute to her commitment to the claim that if zinc were submerged in *this* liquid and the solution was strong enough, the reaction would give off hydrogen. Finally, we will assess her entitlement to the claim. This is the point to which I want to draw your attention.

In a typical work on inferentialism, it can seem that whether a speaker is entitled to a claim or not is simply a matter of whether she is entitled to claims from which it follows inferentially as well as whether she has any commitments incompatible with it. I think this picture is far too sparse. When we raise the question of entitlement to a claim, we are interested in whether the speaker has done or can do the work required to secure

such entitlement. That work often consists in the production of other claims from which the initial claim follows, but not just any inferentially upstream claims will typically do. In the present example, if an interlocutor queried—“Are you sure it’s an acid?”—Elba might respond, “Of course I am, I just saw it on Esprit’s lab report.” Such a response, however, would not only be an indication of unethical behavior, it would also be seriously infelicitous. We may agree that Esprit is a magnificent chemistry student and that the contents of her lab report are close to authoritative. As such, it follows from Elba’s assertion and our background commitments that the sample is, in fact, an acid. Elba, however, has not merely claimed that the sample is an acid. In indicating that she knows this because she tasted it, she has claimed first-personal entitlement to this assertion.⁸ The way to secure such entitlement in the face of a well-motivated challenge or query is not to pass the buck to Esprit but either to fall back to a more secure perceptual claim—“It tasted sour”—or to defend one’s reliability as a reporter of sour tastes.

Offering a claim of which the claim in question is a consequence is not enough, it must be the right *kind* of claim.⁹ There are additional discursive norms that have come into play. These norms are properly discursive in the sense that they make the discursive practice the practice that it is. They are not in any sense external or ancillary to the practice in the way that norms of etiquette or morality might be. If Elba responds

⁸The case we are imagining is one in which Elba makes her report immediately after we have seen her taste the sample and in which she does nothing to dissuade the assumption that she has discovered on her own that it is an acid.

⁹In the text that immediately follows, I’m going to refer to “kinds” and “sorts” of claims without giving these terms much substance. I think there is an intuitive sense to there being different kinds of claims in the sense toward which I’m gesturing—say those that report on experiences, those that report preferences, and those that report on previous testimony—but I don’t yet have the resources to say just what differentiates between them. I develop those resources later in the chapter, and, by Chapter 5, will be able to provide more substance to this loose talk of “kinds” and “sorts.”

impolitely to the query—“My taste buds aren’t broken, you idiot!”—though her response may be inappropriate in some sense, it is still a legitimate response by the norms of empirical discourse. This is not so if she responds to a challenge to her first-personal report by deferring to Esprit’s authority. Moreover, though properly discursive, these norms are neither inferential nor fundamental. They are not inferential insofar as they do not underwrite any inferences; they work in tandem with inferential norms. They are not fundamental insofar as they are not necessary for a discursive practice to count as such. Different discursive practices may exhibit different norms of the kind in question. They must exhibit *some* of them, but none of them in particular are necessary.

The norms I’m pointing to here are employed in the assessment of whether the kind of claim advanced could possibly count as a reason for a claim that has been challenged—whether it does, in fact, count as a reason is also partly determined by the inferential norms. They also serve in the assessment of the kinds of challenges or queries that count as appropriate. These norms, for example, might rule out “But I don’t want it to be chocolate!” as a legitimate challenge to the assertion that the birthday cake is chocolate. What is empirically the case cannot be challenged on the basis of wishful thinking.¹⁰ They are, in general, norms by which we assess entitlement. As such, I shall

¹⁰ The epistemic structural norms of a practice do not settle all (or even most) questions about what constitutes an appropriate query, challenge, or response. “A statement is made and its making is a historic event, the utterance by a certain speaker or writer of certain words (a sentence) to an audience with reference to a historic situation, event or what not” (Austin, Strawson, & Cousin, 1950). We ask questions and make requests of particular people under specific circumstances. Orders are always aimed at getting *someone* to act in some particular way at some given time. The norms of interpersonal discourse and the common ground of the discussion shape to a large degree what one can felicitously utter. I should not inform you of things I already take you to know (unless I think you need reminding). I should not say things that are irrelevant to our discussion (unless I need to distract you). I should not ask a question of you in the presence of another when I know they should not be privy to the answer. These norms function in complex ways across all discursive practices, and they are highly contextual. The context of utterance does most of the work of ruling some pointed challenges beyond the pale, some queries irrelevant, and some responses sufficient “within reason, and for present intents and purposes” (1946, pp. 149–157). I have no interest in denying the signifi-

refer to them as the *epistemic* norms of a discursive practice.¹¹

The epistemic norms of a discursive practice, I now want to argue, stand in an important relationship to the inferential norms of the practice. They partly *shape* those inferential norms. By this I mean that the epistemic norms of a practice partly determine which inferential—content conferring—norms belong to it by constructing the way in which it interfaces with both practitioners and the environment.¹² Consider the epistemic norms with which we are most theoretically familiar, those of empirical discourse.

Copper melts at 1085C. It is, as such, an inferential norm of empirical discourse that one who is committed to “x is copper” is consequently committed to “x will melt at 1085C.”

This inferential norm is the product of our interactions with particular objects in our world—those composed of copper, in this case. Those objects are, as Brandom puts it, “grant[ed] a kind of authority over the correctness of our thinking [so that they] are (thereby) [the objects we are] thinking about” (Penco, 1999, p. 1). The epistemic norms of empirical discourse just are the norms that grant this particular authority to these objects. They determine that any claims that could possibly count as reasons for the commitment to copper’s melting at 1085C must terminate, at some point, in someone’s first-personal observation of this fact.¹³ It is this that makes the practice *empirical*. The constellation of

cant importance of context. My claim here is only that some classes or kinds of queries, challenges, and responses (which I will later identify in terms of their pragmatic structure) are ruled out by the epistemic structural norms of a discursive practice. They are one among many classes of discursive norms, but they are an important one.

¹¹ Brandom is certainly not blind to the existence of these epistemic norms or the role they play in our discursive practices. He is sensitive, for example, to how justifications for perceptual claims differ from justifications for other types of empirical claims. However, he does not focus much attention on these norms and, so, misses the important explanatory role I claim they play.

¹² These norms only *partly* determine the inferential norms, as there are other determinants such as the objects that claims are about and the interests, desires, and values of practitioners.

¹³ This is, of course, a vast oversimplification of the kind of scientific practice on which this claim rests, but I think it is enough to make my point clear.

epistemic norms of empirical discourse function to make that practice corporeal, lumpy, and thick, as Brandom has called it in various places, by integrating the physical environment into the space of giving and asking for reasons.

We can see another example of this shaping when we consider the “patterns of reasoning” to which Brandom drew our attention in the last chapter. Consider again this example: one who endorses the inference from “Exercising is part of a healthy lifestyle” to “Tom shall start exercising” undertakes commitments to a family of related inferences, in particular, those that can be explicated by the prudential assertion, “Tom ought to adopt a healthy lifestyle.” To undertake commitment to this family of inferences is to implicitly attribute to Tom an *interest* in maintaining good health. As such, whether the inferential norm explicated by the prudential assertion is one we should endorse is determined by whether Tom does indeed have such an interest. This, however, is something over which Tom has ultimate authority, in a complex way to be explored later. His desires, aims, purposes, and so on, determine what is in his interest, and no one is taken to be in a position to better know his desires or purposes than he can himself. As such, his word *is* the final word on whether the inferential norm making it proper to infer “Tom shall start exercising” from “Exercising is part of a healthy lifestyle” is included in our practice. The epistemic norms of the practice of prudential discourse are such that Tom is granted such authority in contrast to the norms of empirical discourse which grant no one individual or group such authority. This is what it is for epistemic norms to shape, in part, the inferential norms of a practice.

3. Structural Discursive Norms

As I noted in the introduction to this chapter, epistemic norms are not the only

structural discursive norms. They are a proper subset of the broader category of *norms of authority*, which are themselves distinguished from what I will call *pragmatic norms*.

What these varieties of norms share is a role in shaping the inferential norms of a discursive practice, where the differ is in the ways in which they function to shape inferential norms. To motivate the introduction of these additional structural norms, I propose to explore an important choice point for use-based, Neo-Pragmatist theories of meaning, in general, and inferentialism, in particular. Inferentialists, like most philosophers of language interested in an account of meaning, have tended to focus the lion's share of their theoretical attention on the speech act of assertion and the propositional content that it expresses. Brandom says that assertions "are essentially performances that can both serve as and stand in need of reasons" and that "propositional contents are essentially what can serve as both premises and conclusions of inferences" (R. Brandom, 2000, p. 189). As such, he claims, the inclusion of performances that count as assertions is both necessary and sufficient for a practice to count as a game of giving and asking for reasons, i.e., a discursive practice. He calls this position *linguistic rationalism*. Rather than privileging assertions, though, one might find reason to think that any practice of giving and *asking* for reasons must also include speech acts of various other kinds, e.g., questions or queries, hails, commands, or others that are yet to be identified. If this is the case, then we must investigate whether the use of vocabularies in these additional kinds of speech acts further shapes the inferential norms of the discursive practice. Finding that it does could lead one to conclude that propositional content does not exhaust the space of inferentially articulated conceptual contents. In rejecting linguistic rationalism in this way, we then need to ask whether the epistemic

norms we've identified are enough to characterize the proprieties of use of these additional speech act types. If they are not, then there is reason to include whatever norms *are* necessary as structural discursive norms. In this section, I address this choice point arguing that assertions are not enough. I present some reasons for thinking that we must also attend to a variety of non-assertional speech acts such as those I will call *observatives*. I then draw on a framework advanced by Rebecca Kukla and Mark Lance to elucidate the structural discursive norms needed to capture this diversity and sketch the broader notion of content that this implies.

A central tenet of Neo-Pragmatism in the philosophy of language is that semantics is not autonomous from pragmatics. The job of pragmatics is to provide a theory of language use and acts of thinking, i.e., of *sayings* and *thinkings*, while semantics supplies an account of the meaning or content of speech acts and thoughts, i.e., of what is *said* and *thought*. To say that semantics is not autonomous from pragmatics, then, is to commit to an account of content being in some way answerable to an account of use. In the case of an inferentialist semantics such as pursued here, semantic content is explained in terms of inferential role, which is explained in terms of linguistic items being suitably caught up in a linguistic practice. The normative structure of such a practice, then, is the subject of pragmatics (Wanderer, 2008, p. 97).

It may seem that this focus on pragmatics would naturally lead Neo-Pragmatists to attend to the broad variety of speech rather than to privilege one particular kind of speech act. Following Wittgenstein, they might recognize that though different uses of language wear a kind of common clothing, they are, in fact, functionally diverse and that none of these functions—say, “signifying something”—holds a place of privilege in

understanding our discursive practices (2001, sec. 6). An account of linguistic practices would then require attending to the ways in which these different kinds of speech acts are used, the norms by which they are governed, the broadly inferential relations between them, and the kinds of content that might be constituted by these broadly inferential relations. This is not, however, required by a commitment to the conceptual dependence of semantics on pragmatics, and, in fact, the most prominent Neo-Pragmatist account of meaning—Brandom’s inferentialism—has eschewed this kind of pluralism opting instead to privilege assertion (on the pragmatic side) and propositional content (on the semantic side).

Rather than account for the variety of things we do with speech from the outset, this standard approach has been committed to proceeding in three steps:

S1: An account of the normative pragmatics of linguistic practice is developed, where a practice is understood as linguistic if and only if it includes performances with the pragmatic significance of assertion (R. Brandom, 1994, p. 172).

S2: An account of the propositional content of assertions in terms of their inferential role is developed in which inferential role is explained in terms of the normative pragmatics developed in S1.

S3: An account of the normative pragmatics of speech acts other than assertion is developed leveraging the account of propositional content developed in S2.

On this approach, pragmatics does double-duty. In the first instance, its task is an account of the practices that count as discursive, and so linguistic. In the second, it proffers an account of speech acts like commands, requests, and questions—presumably by way of an account of (something like) the illocutionary force that attaches to the already explained propositional content—but this is possible only once an account of

propositional content is had on the back of the normative pragmatics of asserting.¹⁴ It is important to note, though, that functional diversity spans both of these instances. An account of the pragmatics of assertion will have to make some sense of the various things we use assertions to do, for example, to make ordinary empirical claims, modal claims, and normative claims. I reviewed Brandom's way of capturing this diversity in the previous chapter.

This, in broad outline, is the structure of Brandom's project. Contra Wittgenstein, he argues that "language has a downtown" comprised of acts of assertion (R. Brandom, 2000, pp. 14–15). In practices of asserting, conceptual contents are forged, and all other speech acts are parasitic upon these contents.¹⁵ Though he admits that the model of assertion might be "enriched by allowing various auxiliary sorts of speech acts" such as deferrals, disavowals, and queries, these are strictly unnecessary for a practice to count as discursive and, so, to confer semantic content on acts within the practice (1994, pp. 191–193).

Brandom's reasons for his *linguistic rationalism* are sometimes hard to discern. Following Davidson, he thinks that the only reason for a belief can be another belief. As such, Brandom claims, only propositional contents can stand in inferential relations as

¹⁴ In S1, pragmatics is conceived of in the somewhat idiosyncratic Neo-Pragmatist way, whereas in S3, it more closely resembles the standard division originally proposed by Charles Morris according to which pragmatics is the study of "the relation of signs to interpreters" (1938, p. 6). It is this "standard" approach to pragmatics that is often associated, for example, with Austin, Grice, Recanati, and Searle, to name just a few. For an overview of the attempts to distinguish between semantics and pragmatics, see (Szabó, 2005). Also see, (Austin, 1975; Grice, 1991; Recanati, 2004; Searle, 1969).

¹⁵ Indeed, if we look to the extraordinarily thorough index of *Making It Explicit*, we find that "imperatives" appears only to direct us to "See commands," and the entry for "commands" points only to a few instances where the notion of a command is used to explicate historical accounts of the bindingness of rules (with one exception discussed in the body text in a moment). The terms "interrogative" and "question" do not appear in the index at all, and "promises" are only discussed in explaining the undertaking of commitments or responsibilities.

both premise and conclusion and only assertions can both be reasons and stand in need of reasons (Davidson, 1986, p. 310). Given that inferential relations are fundamental to Brandom's account, this means that any practice that could count as discursive *must* include performances with the status of assertions. Moreover, he claims, "[i]t is only because some performances function as assertions that others deserve to be distinguished as *speech acts*" at all (1994, p. 172). Questions are distinguished only by their relation to possible answers, which come in the form of assertions. Commands don't simply alter what is permissible or obligatory but "do so specifically by *saying* or *describing* what is and is not appropriate" (1994, p. 172). This is a claim of necessity. We could not so much as understand these other performances *as* speech acts if they were not caught up in a practice that included performances having the significance of assertions. Brandom goes on to claim, however, that such performances are not only necessary but sufficient for a practice counting as discursive.¹⁶ His argument for the sufficiency claim is much more elusive. Why should we think that assertion alone will be enough? One thought seems to be that a practice of giving and asking for reasons requires only that we be able to perform some acts that could count as presenting the inferential grounds for our claims and other acts that could count as challenges to those claims. The former are clearly acts of assertion, the latter, we might think, must be some kind of pointed question. Brandom, however, claims assertions that are incompatible with a claim can play role of challenges (1994, p. 178).¹⁷ If he's right, assertions are sufficient for practices of giving and asking

¹⁶ See, for example, (Brandom, 2008a, p. 42), where he claims that in *Making It Explicit* he pursues the explanatory strategy of treating practices of assertion as sufficient for *autonomous discursive practices*, i.e., language games that one could play though one played no others. If they are sufficient for an autonomous discursive practice, then they must be sufficient for any practice to count as discursive.

¹⁷ The virtue he sees in this approach is that it puts entitlement to challenges on the table with entitlement to the claims they challenge. The default and challenge structure of entitlement requires that this is the case, i.e., that entitlement to a challenge is an open question, but it does not follow that challenges must come in

for reasons, but I see little reason for thinking that incompatible assertions are sufficient. Given the lack of an explicit argument for the sufficiency claim, I think it is best to treat Brandom's linguistic rationalism as a hypothesis, a reading that he seems to endorse himself.¹⁸ Linguistic rationalism, then, will stand or fall on the basis of the explanatory power of the resulting theory.

I will not test the limits of Brandom's account here. Instead, I want to explore a few arguments that aim to motivate the contrary hypothesis that assertions alone are not sufficient for a practice to count as discursive. The arguments I have in mind originate in the work of one of Brandom's Pittsburgh colleagues, Nuel Belnap. Belnap thinks that the privileging of assertions stems from the mis-education of philosophers at the hands of "those teachers of elementary logic" who equated "sentence" with "declarative sentence" and assumed that once we account for declaratives, everything else we do with language will either turn out to be something like that or a mere extension from it. He dubs this "the Declarative Fallacy" and argues that, rather than privileging assertions, all sentences—but especially interrogatives and imperatives—should be given equal time (Belnap, 1990, p. 1).¹⁹ Belnap's arguments proceed by showing just how the privileging

the form of assertions. Whether one's question—Why do you believe that?—is appropriate or well-motivated can just as readily be assessed. As such, this seems poor motivation for thinking that assertions alone can do the work of challenges. For a more thorough discussion of the need for interrogative speech acts see, (Wanderer, 2010a). Brandom discusses the default and challenge structure of entitlement in a number of places, but see, (R. Brandom, 1994, pp. 176–178), also see, (M. Williams, 2001, Chapter 13, 2015). I address this model in greater detail in Chapter 5.

¹⁸ Brandom recommends this understanding of his project in various places. For example, he says that he has adopted "a Popperian philosophical methodology: developing and defending the strongest, most easily falsifiable not-yet-falsified hypothesis" (R. Brandom, 2010, p. 316).

¹⁹ Fallacy is perhaps too strong a term for this phenomenon, especially if it is meant in the sense of an error in reasoning, but I will continue to use it for two reasons. First, I want to maintain continuity with Belnap's classic analysis. Second, another, more colloquial meaning of the term 'fallacy' is a mistaken belief, and the belief that we can fully capture meaning by focusing only on use in declarative sentences is, I believe, mistaken.

of declarative sentences, speech acts of assertion, propositional contents, truth-conditions, and inferences (narrowly construed) has led to an impoverished philosophical understanding of language. Among other deficits, he claims, we lack the resources to make sense of the meanings of words and phrases that appear primarily in non-declarative speech, e.g., the interrogative ‘*wh*’ words, or to account for the peculiar contribution that the disjunctive makes to the meaning of a sentence when it is used interrogatively, e.g., in “Is it declaratives or interrogatives that have inverted word order?” The latter, he argues, cannot be captured by the meaning given by truth-tables or assent-tables, yet we lack any resources to go beyond that on standard accounts (Belnap, 1990, p. 3). The upshot of Belnap’s investigation is that philosophers of language should take a deep breath and begin anew with an account that recognizes the importance of the magnificent variety of language from the outset. Only then will we be able to give an account of the meanings of words and sentences that is not blinkered to the contribution made by use in non-declarative contexts.

Belnap produces a plethora of examples to show that declaratives, assertions, propositions, and so on are not enough, but I want to turn to some theorists who have more recently taken up his cause. Training their sights on Brandom’s inferentialism, Kukla and Lance have argued that his privileging of assertion—a *declarative* speech act in their terminology—leaves him unable to properly account for the relationship between speakers and their shared world and between speakers themselves. The first of these problems arises because Brandom lacks a proper account of language-entry transitions as essentially individuating and owned by concretely situated, embodied speakers, the second because he does not attend to the directedness of speech acts—their *pragmatic*

voice—and, so, cannot make sense of speech acts as essentially calling for a response from others, i.e., enacting what they call the vocative function.

I will focus on the argument that Brandom’s account of language-entry transitions is insufficient. Kukla and Lance argue that observative speech acts—those that “give expression to our recognition of an empirical fact, object, or state of affairs in *observation*, and most paradigmatically in perception” (Kukla & Lance, 2009, p. 46), e.g., “Lo! A rabbit”—share much in common with declaratives but are distinguished from them in their pragmatic structure and voice, which marks one’s first-personal encounter with the world in a way that no declarative can. As such, observatives are not reducible to declaratives, yet they are a necessary component of the space of reasons. Consequently, any practice that does not include the resources to issue observatives—or *recognatives*, the broader category of speech acts that are structurally isomorphic with observatives but that are not necessarily linked to observation or perception—cannot count as discursive.

Observatives seem, at first glance, to be kinds of assertions, even if, in some cases, their propositional content is elliptical.²⁰ They are like assertions insofar as they license others to take up their content, i.e., issue a reassertion license. Whether I assert, “There’s a rabbit in the bush over there,” or I express my first-personal uptake of that fact, “Lo! A rabbit!,” I have said something that licenses the assertion by others that there is a rabbit in the bush.²¹ In both instances, not only can others reassert the content of my

²⁰ “A rabbit!” is not a proposition, but, perhaps, it is elliptical for “There is a rabbit in the bush.” This possibility is discussed later in this section.

²¹ It is important to distinguish between the grammatic and pragmatic structure of a speech act. The grammar of “There is a rabbit in the bush” is that of a declarative sentence, and, in this example, it is also functioning pragmatically as a declarative speech act—an assertion. Grammatical structure, however, is often a poor indication of pragmatic structure. In ordinary speech, we tend to use grammatically declarative sentences to issue pragmatically observative speech acts, i.e., it would be perfectly natural to say “There is a rabbit in the bush” as an expression of one’s first-personal uptake of that fact. In this discussion, however, I

claim, they can also defer to me when their reassertions are challenged. I have undertaken an obligation to defend my entitlement should a well-motivated challenge be directed either at me or at someone who reasserts my claim.²²

Yet, there is an important difference between the assertion and the observative. When I assert, “There is a rabbit in the bush over there,” I too can defer in the face of a challenge. I can appropriately deflect responsibility for defending the claim to the one who has informed me of the rabbit’s presence. This course is *not* open to me when I have uttered the observative. The obligation to respond to challenges rests solely and directly with me. The observative is, in this sense, *mine* in a way that the declarative is not. Observative speech acts are “essentially individuating[:]. . . an event of perception cannot be shared among several agents, even though several agents may perceive the same thing as a result of similar interactions with the world” (Kukla & Lance, 2010, p. 120). Moreover, perception yields new entitlements that, in an important sense, belong only to the observer. “To perceive—as opposed to just inheriting entitlement to a belief—is to be first-personally claimed by what I see, to recognize my perceptual episode as mine” (Kukla & Lance, 2010, pp. 120–121).

An observative is essentially, not accidentally, an expression of the speaker’s first-personal encounter with the world, and it is ineliminably marked as such. As Kukla and Lance explain, speech acts are inherently pragmatically voiced. Though declaratives—the standard model of speech for most theorists—are *impersonal*, generally lacking pragmatic voice or translatable “from one personal voice to another without its force

adopt the practice of treating declarative sentences as declarative speech acts and explicitly mark observatives either with “Lo!” or by the use of an exclamation point (Kukla & Lance, 2009, p. pp.45-46). I address the question of how observatives are typically marked in ordinary speech later in this section.

²² I discuss the epistemic structure of observatives in greater detail in Chapter 5.

being changed in the least” (Kukla & Lance, 2009, p. 59), this is not the case for other speech acts. Imperatives, for example, are inherently second-personal. If the speech act is not directed *at* someone (or some class) it is not an imperatival speech act, at all.

Imperatives can, of course, be translated into a different voice. “Go to bed!” directed at my son could be rephrased as “Sam ought to go to bed,” but something important is lost in translation. What was an order directed at Sam is now a speech act with a completely different pragmatic structure. It might still be used to remind Sam that he’s obliged to get to bed or even to hold him to that obligation, but it does not have the same force that the second-personally voiced imperative does (Kukla & Lance, 2009, pp. 61–62).

Observatives are similarly voiced, only in this case they are inherently first-personal. They are marked as owned by the individual who speaks them, indicating that their entitlement is had in a very particular way, from a very particular perspective, i.e., via perception and from the perspective of the speaker (Kukla & Lance, 2009, p. 60). It is important to note that what is at stake here is pragmatic not grammatical voice. We are all familiar from grade school that sentences come in first-, second-, and third-personal voices, but the distinction Kukla and Lance draw here is different. The pragmatic “voice of a speech act concerns the manner in which the agent takes up her entitlement to the speech act and strives to assign statuses to others” (Kukla & Lance, 2009, p. 61). Though grammatical voice often tracks pragmatic voice, as we’ve already seen with the declarative and the observative above, they can and do come apart.

With this grasp of Kukla and Lance’s position, I turn now to the argument that observatives have a distinct pragmatic structure and voice that cannot be translated into or accounted for in terms of declaratives. It is tempting to try to account for the

individuating nature of observatives in terms of a difference in propositional contents. The first step would be to argue that the content of an observative that does not exhibit a straightforwardly propositional content ought to be understood as the elliptical expression of some proposition. “Lo! A rabbit!” is really “Lo! There is a rabbit in the bush over there!” or “Lo! A rabbit is in the vicinity!” The second step would then account for the distinction between the observative, with the ellipsis suitably filled in, and the declarative with the same propositional content by noting that the observative is really a concatenation of two (or perhaps more) declaratives. “Lo! A rabbit!” is analyzed as “There is a rabbit in the bush over there” and “I am seeing that there is a rabbit in the bush over there,” or some other claim that flags the expressed content as the product of my own experiential uptake of the scene.

Kukla and Lance argue that each step in the proposed analysis is problematic. The choice of how to fill in the ellipsis in the first step is arbitrary. Any of a number of declaratives would be equally convincing as the completion of the elliptically expressed propositional content of the observative, and there seems no principled reason to choose one over the others. The second step suffers the same problem. The now propositionally contentful observative could be analyzed as “There is a rabbit in the bush over there” and “I am seeing that there is a rabbit in the bush over there” or “I see something” or “The reason I know there is a rabbit in the bush over there is that I am seeing it there,” and we could clearly go on. There again seems no principled reason to choose one expansion over the others (Kukla & Lance, 2009, pp. 55–56).

More importantly, any expansion of the observative to some concatenation of declaratives would still be missing something essential. Whatever expansion we propose,

it will have to capture that the entitlement one claims for the speech act is a product of one's first-personal encounter with the world. However, any declarative speech act that reports this to be the case, e.g., "I see that there is a rabbit in the bush over there," is, in effect, a third-personal report on one's own experience. It could always be translated without loss to, for example, "Tom sees that there is a rabbit in the bush over there," as spoken by someone else. This is just the *same assertion* in a different voice, but the essential first-person voice of the observative has now gone missing entirely. It is not marked as mine in the way that an observative speech act is, for anyone could be entitled to this assertion (Kukla & Lance, 2009, p. 59, 2010, p. 126).²³ There is no way, they conclude, to translate the first-personal, agent-relative episodes expressed by observatives into impersonal, agent-neutral declaratives without loss. "An account of perception as an assertion-like episode that is perspectively marked by its content is insufficient. Rather, perceptual episodes must be understood as first-personally structured, agent-relative events" (Kukla & Lance, 2010, p. 121).

Kukla and Lance urge that instead of trying to analyze observatives in terms of their propositional contents, we should recognize that they are not elliptical expressions of propositions but rather "complete, well-formed utterances that *imply* propositional truths" with the contents of the expansions considered and others, as well (Kukla & Lance, 2009, p. 55). The "non-propositional observations" that are "*expressed in observatives...ground justified declaratives*" (Kukla & Lance, 2009, p. 58). Which

²³Indeed, note that a challenge to the proposed analysis does not necessitate that I respond by retreating to a safer claim or defending my reliability. I could, in principle, defer to someone who is committed to both "There is a rabbit in the bush over there" and "Tom sees that there is a rabbit in the bush over there." Moreover, this speaker can also defer to another if her claim is challenged, and so on. This is a strange and troubling result. If challenges to an observation claim can always be deferred, then empirical claims never bottom out in encounters with the world. The entire edifice of empirical knowledge would be left "spinning in the void," to use McDowell's memorable phrase (1994).

declaratives they ground is a matter of the antecedent commitments of the speaker and scorekeeper, as well as their context.

Anticipating an objection from their fellow travelers in this Sellarsian terrain, Kukla and Lance consider that the claim that something with non-propositional content can imply a claim with propositional content is problematic. Only propositional contents can stand in inferential relations to one another, at least according to the rules of standard propositional logic. In response, they argue that we should not “presume that propositional logic is, or is structurally analogous to, the only inferential game in town” (Kukla & Lance, 2009, p. 58). To do so is of a piece with the declarative fallacy, “as it presumes that the only discursive logic is the propositional logic of the declarative, and that everything that isn’t propositionally structured must be somehow mute or inarticulate” (Kukla & Lance, 2009, p. 58). We should already see that such a privileging of the propositional is problematic insofar as we want to allow for inferences from propositionally contentful states to actions that do not themselves have propositional contents, i.e., language-exit transitions, and from causal responses to environmental stimuli to propositionally contentful states, i.e., language-entry transitions.²⁴ We should be prepared to further expand these possibilities to capture the inference-like transitions from one normative status to another as the result of normatively defined moves made by participants in a discursive practice.

Brandom has responded directly to the challenge advanced by Kukla and Lance. He “applauds” their analysis of the rich topography of the space of reasons, noting that it

²⁴In the Appendix of their book, Kukla and Lance (with Greg Restall) develop a formal system that aims to capture inference commonly-so-called as a special case of licensed normative moves between types of actions. Pursuing this proposal is not necessary for what I go on to do here, but I am, in general, sympathetic to this view.

furnishes useful resources for thinking about “the expressive roles of all sorts of locutions not treated” in *Making It Explicit (MIE)*, but he rejects the idea that their arguments mandate any “move beyond an assertion centered account” (R. Brandom, 2010, p. 319). According to Brandom, their arguments show only that practitioners must be capable of recognizing *implicitly* in practice the indexing of speech acts to particular individuals as speakers and targets. They need not have acquired the capacity to make these implicit recognitions explicit in the form of recognitive (or observative) speech acts. Moreover, Brandom argues, these capacities are already implicit in the scorekeeping model of *MIE*. Recognitives—observatives, in particular—are taken to be entitled only when they result from a reliable differential responsive disposition regarding which the scorekeeper is willing to endorse a reliability inference. In such instance, the scorekeeper both attributes entitlement and undertakes a commitment to the content of the report, and it is the basis of this attribution that marks its agent-relativity for Brandom (R. Brandom, 2010, p. p.316-319). The upshot is that he sees the contribution from Kukla and Lance as supplying some of the missing pieces that will allow for the completion of S3, but nothing more.

Brandom’s reply is an insufficient defense against the objections Kukla and Lance have raised. They do not claim that the normative functions enacted paradigmatically by observative and vocative speech acts need to be included in a practice by the explicit inclusion of speech acts that in some way wear their functions on their sleeves. They need not be explicit. Their claim is that these normative functions, which may be enacted by grammatically diverse speech acts, are necessary for engaging in reason-giving practices at all. What matters for them is not the explicit performances but rather the normative

functions enacted, often implicitly. Brandom is wrong about this already being captured in the perspectival scorekeeping structure elucidated in *MIE*. Consider observatives, again. Brandom claims that the peculiar way observatives are scored—in terms of the reliability inference—marks their agent-relativity, but the question of the reliability of the responsive dispositions of a speaker can only surface when the scorekeeper takes the speaker to be issuing a report on her first-person experience, i.e., an observative. In Brandom’s paradigm cases, this happens when speaker and scorekeeper are in the same context so that the scorekeeper can, in essence, “see” for herself that the speaker is making such a report. These conspicuous cases, however, are not the norm. Most reports of first-personal experience are made some time after one has had the experience. When we share a physical environment, it is somewhat rare that we have reason to report on what we already take another to be able to see for herself. This is not to say that this never happens. We may have need to give warning, to ostend for pedagogical purposes, to point out details one has missed, and so on, but much more common is the case where one reports on one’s day, what happened to her, what she saw or heard or smelled or felt when the hearer wasn’t present to share the experience. With respect to such reports, which are often grammatically declarative, we must have the ability to recognize that the proper way to score them is with regard to the speaker’s reliable differential responsive dispositions. To do this, we must recognize their first-personal pragmatic voice prior to the scoring of the report even if one’s language lacks resources, like “Lo!,” for making that voice explicit. We have many ways, both subtle and obvious, to mark this voice ranging from the tone and confidence with which one speaks to explicit locutions like “I heard...,” “I saw...,” “On my way to work today...,” etc.. Notice that, on pain of

resurfacing the sorts of problems just discussed, these locutions cannot be declaratives appended to an observative but must be markers of pragmatic voice.

By the lights of Kukla and Lance, then, Brandom's assertion centered account is in trouble.²⁵ At the very least, linguistic practices must include speech acts that enact the pragmatic function of observatives and, perhaps, others, as well.²⁶ This, however, does not entail that the structural discursive norms of assessment for non-assertional speech acts need to be included in PALM. It is possible, after all, that while observatives—or recognitives more broadly—are necessary for a practice to count as discursive, i.e., to be *meaningful*, the use of vocabularies in non-declarative speech acts is not *meaning constitutive*.²⁷ Kukla and Lance's focus on what Wanderer has called the "space of reasoning," i.e., on the embodied and embedded performances that count as speech acts, leads them to conclusions about meaningfulness, but they are not forthcoming about the consequences of their arguments for our understanding of the space of reasons, i.e., an account of semantic contents and the relations between them. Wanderer identifies three possible consequences we could derive:

- (i) Enriching the topography of the space of reasoning beyond the declarative has minimal impact on the topography of the space of reasons.
- (ii) Enriching the topography of the space of reasoning beyond the declarative substantively alters the topography of the space of reasons.
- (iii) We should reject this bifurcation of spaces and proffer a unified sense of the notion of the space of reasons shorn of adherence to the

²⁵It's worth noting that Brandom's linguistic rationalism faces other challenges, as well. Taylor, for example, has argued that Brandom's privileging of propositional content over what he calls disclosive content is problematic, and Wanderer has claimed that assertions cannot themselves comprise an autonomous discursive practice. They must at least be accompanied by challenges, which are, of necessity, second-personally voiced (Taylor, 2010; Wanderer, 2010a).

²⁶Kukla and Lance also argue that speech acts that they call *vocatives*, paradigmatically hails that are essentially second-personally voiced and call out for recognition, are equiprimordial with declaratives and observatives (Kukla & Lance, 2009, p. chaps. 6 & 7)

²⁷This would require attending to non-declarative acts in S1 but not in S2.

declarative fallacy (Wanderer, 2010b, p. 376).

We have come to see that observatives and declaratives are speech acts with distinct pragmatic structure, so if it can be shown that the language-entry transitions enacted by observatives are partly constitutive of semantic content, i.e., if they are a necessary part of the space of reasons not just of reasonings, then it follows that (i) is false. We must ask, then: are language-entry transitions partly constitutive of content or can an account be restricted to properly *inferential* circumstances and consequences of application? Brandom has himself argued that “under such a restriction, it is impossible to reconstruct the contents of actual concepts, except perhaps in some regions of mathematics” (1994, pp. 131–132) and has rejected the idea that mathematics could constitute an autonomous discursive practice. Why think this? The idea is that the notion of *empirical* content (or *practical* content, for that matter) gets little traction if such content is not in some way hooked up to the world it is purportedly about. It is part of the content of the concept red that it is appropriately used in non-inferentially elicited reports on the presence of red objects, not in the visual field of anyone in particular but *for someone*.²⁸ It would not be an empirical concept if it did not have this use. The case is the same for all of our empirical and practical concepts; their non-inferential circumstances and consequences of application partly determine their contents.

So, the motivation for including entry and exit transitions as partly meaning constitutive is to be able to reconstruct the meaning of empirical and practical

²⁸This is adamantly not to say that one does not have or cannot use the concept unless one can see red things. Having the ability to use a concept is having the ability to bind oneself by it, to undertake a commitment to the goodness of the inferences that follow from its appropriate application as judged by one’s interlocutors. One need not be in possession of all of the inferential connections constitutive of its meaning (indeed, one *could not*) in order to be able to do this.

vocabularies. Without including them, we could not make sense of empirical concepts like red as being used, often if not primarily, to report on the color properties of objects we encounter in the world. Grasping that “x is red” implies “x is colored” and is incompatible with “x is green”, etc., is *part of* being able to deploy the concept red, but failure to grasp the proper reporting use of red, even if one couldn’t so use it oneself, would indicate a fundamental failure to understand the concept. It just wouldn’t be the concept that it is if it didn’t have that reporting use. Similarly, practical concepts (ought, for instance) wouldn’t be the concepts they are if they weren’t reliably linked up to intentional acts via our dispositions to respond to their deployment.

Since observatives are the speech acts that enact these language-entry transitions, it follows that Kukla and Lance’s enrichment of the topography of the space of *reasoning* fundamentally alters the topography of the space of reasons, for, we cannot reconstruct the contents of empirical vocabularies without properly grasping how they are linked up to the world, and we cannot do that without understanding the special role played by our first-person encounters with that world.²⁹ Moreover, this means that we have reason for including any structural discursive norms required to explain the contribution to the space of reasons made by speech acts other than assertions. What are these structural discursive norms?

²⁹Adjudicating between (ii) and (iii) is more difficult in large part because it is not clear what it would mean to reject the bifurcation of the space of reasonings and the space of reasons. One possibility is that this is a rejection of the distinction between semantics and pragmatics. Not only is there a conceptual link between them but use and meaning must be theorized in tandem. This would be a radical break from the Neo-Pragmatist project in philosophy of language as it is commonly conceived, so I am inclined toward an account that pursues option (ii): recognition of the rich topography of the space of reasonings manifests in diversity of content in the space of reasons. Whether (ii) or (iii) wins the day, though, the resulting account must include more than assertions at the ground level. It will build the diversity of uses of speech in at the outset, and it would be quite strange if this didn’t result in a significant effect at the level of semantics.

4. Speech Acts and Normative Functions

Remaining within the framework of a broadly inferentialist semantics paired with a scorekeeping normative pragmatics, Kukla and Lance already provide the resources for extending our account in a way that makes room for these structural discursive norms and that helps us on our way to capturing the functional pluralism of discursive practices. They begin by noting that we can treat speech acts as enacting functions that take as their inputs the normative statuses of speakers and have as their outputs transformations of the normative statuses both of speakers and of others in their linguistic community. An individual speech act enacts at least one—but often more than one—normative function and strives to affect a normative transition in the statuses of the speaker and hearers. Speech acts, as such, are mappings “between the *normative statuses constitutive of entitlement to a given speech act* and the *normative changes (in the status of the speaker, or of others in the discursive community) that the act strives to produce*” (2009, p. 15, emphasis in original). Moreover, Kukla and Lance enrich the scorekeeping apparatus of *MIE* by identifying two distinct flavors of the normative statuses of commitment and entitlement. Each can be either agent-relative or agent-neutral. A normative input or output is *agent-relative* if, “*in virtue of its pragmatic structure,*” the speech act is indexed to a person or group of persons as a result of the particular normative positions they inhabit. If the input or output is “for everyone”, then it is *agent-neutral* (2009, p. 17). In this section, I briefly develop the topography that results from overlaying these two distinctions and explain how it can be leveraged to enrich PALM with the resources to incorporate speech acts other than assertion.

We can be more precise if we focus on inputs and outputs individually. Different

sorts of speech acts require different sorts of authority. What entitles me to a declarative like "This coffee was grown in Ethiopia" is different from what authorizes me to instruct my students to turn off their cell phones. Entitlement to the former is a product of the reasons that I could give if my claim was challenged. These reasons are available to anyone who is motivated to discover them, say, by reading the label on the bag. It is in this sense that the input of the speech act is agent-neutral. It is not that everyone *is* entitled to make the claim; many people haven't seen the packaging from which this particular cup was brewed. But anyone *could be* entitled to the claim, no matter what else happened to be the case with them.³⁰

What authorizes the imperatival speech act, on the other hand, is my particular social-normative standing as the teacher of this class. No one who is not the teacher of this class could be entitled to this speech act (on these same grounds). If the professor next door happens to walk in and tell my students to turn off their cell phones, her imperative would be infelicitous. She lacks the standing to issue it, and nothing she could do would secure that standing for her in the way that reading the label on the coffee could entitle her to the declarative speech act. What authorizes the imperative is a set of social and institutional facts. I have found myself placed in a normative position that empowers me to impose certain requirements upon students who enroll in my classes, and they have placed themselves in the position of being subject to such demands. It is this relationship, and not any reasons that I can produce, that authorizes my speech act.³¹ The entitlement

³⁰I qualify this strong notion of agent-neutrality in footnote 32.

³¹Which is not to say that one *could not* produce reasons that explain (or even vindicate or secure) the social-normative standing on which one's entitlement rests. Such reasons, though, are, in a certain sense, meta-linguistic. They don't secure entitlement by lending support to the content one has expressed but by demonstrating the standing that one has. A challenge to one's social-normative standing always bumps the discourse up to a meta-practice in which interlocutors can debate the question of whether the original speaker has the standing she has implicitly claimed for herself.

is, in this way, *agent-relative*. It is an entitlement that not just anyone can claim.

Furthermore, in this case, the entitlement is *agent-specific*, insofar as it is indexed to me in particular and not to a set of individuals who share my normative status.

The normative outputs of speech acts are similarly indexed. Consider again the declarative, "This coffee was grown in Ethiopia." Any utterance of this sentence will be directed at a specific person or persons, but this is a matter of the physical and social environment in which the speech act occurs, not of the pragmatic structure of the speech act itself. As a matter of structure, according to Kukla and Lance, declaratives strive for *universal* uptake. That is, as a constitutive ideal, a declarative speech act "seeks to impute the entitlement to assert this claim to the discursive community in general and demands that others allow its content to constrain their inferences and beliefs" (Kukla & Lance, 2009, p. 18). Anyone who doesn't give it the uptake it seeks is, in some very attenuated sense, defective. This defect is non-culpable, as one can't be held responsible for being out of earshot of a speaker in the next county or for not understanding the language in which the speech act was uttered, but it *is*, according to Kukla and Lance, a defect insofar as one *should* not undertake commitments that are incompatible with the commitments of others that one would recognize as entitled. This is just what it is for us to be part of a shared discursive community. I can, for example, be non-culpably in the wrong if I make an empirical claim that, unbeknownst to me, conflicts with the claims of a distinguished scientist that, should I have known of her claims, I would have taken to be entitled. In this sense, the output of a declarative is *agent-neutral*. It strives to alter the normative status of *all* other discursive practitioners, without regard to prior *normative* standing. By contrast, the output of the imperative issued to my class is *agent-relative*. It strives to

alter the normative statuses of all and only the students enrolled (a normative standing). Each of them now has the status of one who has been ordered to turn off their cell phone.³²

These distinctions can be represented in a diagram with four boxes into which speech acts can be categorized by their pragmatic structure (Figure 1) (Kukla & Lance, 2009, p. 40). Declaratives, with their agent-neutral inputs and outputs, quite naturally fall into place in Box 1. Imperatives, with their agent-relative inputs and outputs, find a home in Box 4. Boxes 2 and 3 are previously unrecognized (in some sense, unrecognizable) categories, and much of the interesting work in *'Yo!' and 'Lo!'* consists in filling in these boxes.

³²The way in which Kukla and Lance define the agent-neutrality of inputs and outputs may strike some as too strong. It is not as if just anyone could become entitled to just any empirical claim, after all. In saying that agent-neutral inputs are “for everyone,” they seem to imply that matters of training, expertise, education, etc., are irrelevant. Moreover, on the output side, it seems wrong to say that someone who has no conceivable reason for needing to know, for example, the number of blades of grass on my front lawn suffers from defect if I happen to have counted them. Nothing in my argument hangs on how strictly we treat agent-neutrality, so I will not try to say whether their definition is defensible here. I will, however, say that I am content with the following weaker formulation. Agent-neutral inputs are those that derive from the space of reasons and, as such, constitute entitlements which can be challenged by both undercutting and rebutting defeaters, whereas agent-relative inputs derive from social or institutional status and can be challenged only by calling into question the relevant status. Agent-neutral outputs are those that are not directed at or indexed to any particular individual or class.

Figure 1: Pragmatic Structure of Speech Acts

Input → Output ↓	Agent-neutral	Agent-relative
Agent-neutral	Box 1 [Neutral input, Neutral output] Declaratives	Box 2 [Relative input, Neutral output] Baptisms, marriage pronouncements, observatives
Agent-relative	Box 3 [Neutral input, Relative output] <i>Prescriptives (oughts)</i>	Box 4 [Relative input, Relative Output] Imperatives, promises, invitations, reproaches, ostensions, vocatives

One nice result of Kukla and Lance’s work is that we now have a framework for differentiating speech act kinds in terms of the normative functions they enact rather than in terms of a theory of force. A declarative (or assertion) is a speech act that takes agent-neutral normative statuses as inputs and has agent-neutral normative statuses as outputs; an imperative is agent-relative in, agent-relative out. Other speech acts have “mixed” inputs and outputs. For example, an act of baptism requires a special social standing to pull off—I can’t walk by the nursery in the hospital calling out names and expect others to now refer to those children by the names I’ve given—but at least some of its outputs are agent-neutral. Once a child is baptized—or, more generally, named—it becomes appropriate for everyone to call them by their given name. Box 3 is especially important for my purposes in the coming chapters, as I will argue that moral assertions—

prescriptives—fit here. Their inputs are agent-neutral, flowing, as it were, from the space of reasons, but at least some of their outputs are agent-relative, having particular normative consequences for the individuals whose moral statuses they identify.

Whereas a theory of force required the distinction between pragmatic force and semantic content, this framework allows that semantic—conceptual but, perhaps, non-propositional—content is a product not only of narrowly inferential proprieties but of the broader set of normative functions enacted in speech of various kinds. The task now is to integrate this insight—this framework—into PALM. To do so, I want to understand the Kukla-Lance framework as operating at the level of normative pragmatics. The functions from entitlements/authorizations to alterations of normative statuses that the framework identifies are associated with a set of scorekeeping norms. These are the norms by which we score whether it is appropriate to take practitioners’ normative statuses to have been transformed. On the input side, they tell us what kind of authority one must have in order to felicitously pull off the speech act and, as such, direct us to the appropriate way to score the speech act as entitled or not. They do not, themselves, provide a rule for determining entitlement, but rather guide the scorekeeper either toward an epistemic assessment or a social/institutional one. On the output side, they tell us what normative statuses one must have already had (if any) to appropriately be counted as the target of a speech act. I will call these the *pragmatic discursive norms* of a practice.

Let’s look at a few quick examples. The declarative “Copper melts at 1085C” can be uttered by anyone who could potentially produce reasons in its defense. Its inputs are agent-neutral, so we are to assess entitlement in epistemic terms. Its output is also agent-neutral, so we could take anyone to be obligated to take up this doxastic commitment,

though we may have many reasons for not taking various particular individuals to be obligated in this way. Next, consider a request made of Jack by Sam: “Please bring that block over here.” In this case, the inputs of the speech act are agent relative. There are various social and institutional norms that govern who can legitimately make a request of another, and these are the norms by which we should assess entitlement to this act. Moreover, the outputs are agent-relative, so we must determine whether Jack is in the social normative position to be appropriately targeted by Sam’s speech and, if he is, we must now score him as having been the target of a request to bring the block over here. He is, as such, called on to respond to the request by bringing the block or begging off for some reason or another. If he simply ignores the request, we can say that he has failed to give it appropriate uptake, i.e., to treat it as one ought to as a player in this language game. These examples give us a sense for the kinds of scorekeeping norms that the pragmatic discursive norms of a practice are.

These examples suggest the need for one more set of structural discursive norms in addition to the pragmatic and epistemic norms already identified. These *authority norms*, as I’ll call them, are the those by which we assess the speaker’s non-epistemic entitlement to a speech act. This is a diverse and difficult to summarize set of norms. They can be broadly thought of in institutional terms. The most straightforward examples are those of a superior giving orders to an inferior in some institutional hierarchy. However, reducing the entire class to institutional norms would miss their rich texture. In some cases, authority is a matter of hierarchy, in others of legal relationships, but in many instances it is a matter of loosely defined social relationships like friendships, acquaintances, neighbors, partners, business associates, and the relationship one has with

one's barber. The norms of these relationships are particular, negotiable, and sensitive to context, which makes the scoring of speech acts grounded in them a matter of interpretation.³³ Still, these are the kinds of norms that entitle one to make a demand, request help, extend an invitation, hold another accountable, suggest a course of action, give warning, and, as I shall argue, even to utter an observative. It is not my aim to settle what these norms are or give a list of any kind, I merely want to identify this class of norms, sketch some of its contours, and examine the way it plays with other structural discursive norms to shape our discursive practices.³⁴

5. EMUs and PALMs

PALM has grown into a complex structure. I have argued that there are scorekeeping norms—the norms associated with the normative pragmatics of discursive practices—of various kinds. First, I noted that inferential connections between semantic contents are underwritten by inferential norms. These are the norms on which we rely when scoring claims, tracking the commitments and entitlements of our interlocutors. Next, I claimed that these were not enough. Discursive practices also have *structural discursive norms* that serve to shape, in part, their inferential proprieties. I identified first the uncontroversial class of epistemic structural norms, i.e., those by which we assess epistemic entitlement to a claim. I then went on to argue that epistemic norms were not

³³Little and Macnamara's investigation of deontic pluralism can help to fill out this idea in important ways. There are many nuanced statuses that one might inhabit in the neighborhood of obligations and permissions each of which require different kinds of social or institutional relationships to make them appropriate and call for different kinds of responses. Ideally, the pragmatic norms of a practice would be identified in fairly fine detail along the lines they suggest (Little & Macnamara, 2008).

³⁴In a recent paper, Rebecca Kukla carefully examines one small corner in this space of authority norms having to do with the language of sexual negotiation, i.e., the invitations, acceptances, rejections, requests, consents, refusals, expressions of desire, expressions of gratitude, and so on, that are part and parcel of healthy (and unhealthy) sexual practices (Kukla, 2018, also see, 2019). This work provides an excellent model for the kind of careful investigation into discursive norms that PALM demands.

enough. They limit our analysis to the contents of declaratives, but we found we have reason to attend also to the myriad other speech acts in which we engage. To do this, we needed to include additional structural norms that allow us both to differentiate between speech acts of various kinds and to capture the ways in which their uses shape the broadly inferential norms of our practices. All along, I have claimed that the payoff of this complexity will be a way of explaining the meaning of vocabularies in terms of what they are used to do, i.e., in terms of their utility for the beings whose vocabularies they are. In this section I finalize and summarize this explanatory framework. Drawing on a metatheoretical framework for explanations of meaning in terms of use (EMUs) proposed by Michael Williams, I will offer a schematic version of PALM that makes its explanatory structure explicit and prepares the way for developing a PALM for the moral “ought” in the coming chapters.

Williams’ EMUs are intended to capture the central components and explanatory structure of explanations of meaning in terms of use advanced by theorists such as Brandom, Sellars, and Horwich. He introduces EMUs through an examination of Horwich’s Minimal Theory of truth (Horwich, 1998a), according to which “the meaning of the truth-predicate, is fully captured by our commitment to all (non-paradoxical) instances of the equivalence schema:

(MT) The proposition that P is true if and only if P” (M. Williams, 2010).

For Horwich, this follows from two prior commitments: first, that the meaning of a word is given by the fact that best explains its overall use, and, second, “that our underived endorsement of” MT is this fact for the truth-predicate. Our underived endorsement of MT, he says, is “explanatorily fundamental” (Horwich, 2001, p. 150). There is no deeper *theoretical* account that somehow explains this rule. He further holds

that instances of MT are “epistemologically fundamental;” we neither reach nor justify them in terms of something more obvious. Finally, he claims that “accepting the instances of MT is ‘the source of everything else we do with the truth predicate,’” in particular, using it as a generalizing device to endorse claims we could not possibly list seriatim (Horwich, 2001, p. 149; M. Williams, 2010). From this account, Williams extracts a metatheoretical framework that summarizes the central components of an explanation of the meaning of ‘true.’ He gives the following EMU:

- (1) (I-T): A material-inferential (intra-linguistic) component. Excepting sentences that generate paradox, the inference from 'Snow is white' to 'It is true that snow is white', and vice versa, is always good; the inference from 'Grass is green' to 'It is true that grass is green', and vice versa, is always good, and so on...
- (2) (E-T): An epistemological component. Such inferences are primitively acceptable (a priori). They are 'free' moves in the discursive game.
- (3) (EPF-T): An explicative/performative function component. As a generalizing device, truth-talk allows us to endorse or repudiate claims that we cannot explicitly state, for example, because we do not know what they are ('You can trust John: anything he tells you will be true') or because there are too many ('Every proposition of the form "p or not-p" is true').
- (4) (U-T): A practical significance component. The generalizing power accorded by “true” allows us to make explicit inferential or general assertional commitments and epistemic norms, which can then become objects of critical discussion.³⁵

(I-T) gives us the inferential role of the truth predicate as a device for disquotation, while (E-T) tells us that one is entitled to predicate truth to a proposition

³⁵This combines the EMU for “true” as initially conceived with the later revision in which the function clause is broken into two distinct clauses. See both (M. Williams, 2013a, pp. 134–135, 2013b, p. 17).

just in case one is entitled to utter the proposition itself. Together, these content-determining clauses capture the inferential proprieties that account for the meaning of ‘true’ and highlight the epistemic norms, which, as we’ve seen, serve an important explanatory function. They characterize how ‘true’ is to be used and, in so doing, explain its conceptual content, i.e., what we are saying when we deploy the concept (M. Williams, 2013b, p. 17).³⁶

(EPF-T) and (U-T) together comprise the functional clauses of the EMU. They capture both what discursive practitioners use ‘true’ to do and why they find it useful to have a vocabulary for doing this. The truth predicate is a generalizing device and, as such, allows us to endorse claims that we otherwise could not endorse either because we

³⁶There are both descriptive and normative versions of explanations of meaning in terms of use. Horwich presents a descriptivist version, but here I follow Williams in amending the EMU to provide a normativist account (M. Williams, 2013b, p. 16). Let me say a bit about why I think this is the right tack to take. On the descriptive side, use-theorists have adopted regularist and regulativist understandings of how use constitutes meaning. On the regularist account, meaning is a matter of the actual inferential significance of a claim, i.e., the inferences that the speaker (or, perhaps, members of her community) actually draws from a claim or, in some cases, is disposed to draw from it. Block, for example, understands a claim’s inferential role as its causal role in reasoning and deliberation and, in general, in the way the expression combines and interacts with other expressions so as to mediate between sensory inputs and behavioral outputs (Block, 1986, p. 93 cited in; Warren, 2018, p. 11). For the regulativist, on the other hand, meaning is a matter of the rules governing the inferential commitments that the speaker undertakes, i.e., it is a matter of what she should infer not what she does or is disposed to infer. Neither regularists nor regulativists are semantic deflationists. Both take the claim “X means the same as Y” to have a truth-maker. For the regularist, it is found either in the actual use of the expression in question or in the dispositions of the individual or community. For the regulativist, it is found in the rules governing the use of the expression, which will either be immanent in use—eliding the position with regularism—or transcendent in a way that will be problematic for naturalists. Both positions, as a result, face a problem of making sense of disagreement between individuals or between linguistic communities. Different speakers fail to mean the same thing by their uses of an expression either because they have different dispositions or its use is governed by different rules. The normativist, on the other hand, frames her account in terms of *proprieties* of use. According to normativism, “X means the same as Y” is to be understood not as a description of how things are but as a *prescription* for how to go on using these expressions, as providing “normative guidance for inferential behavior” (Lance & O’Leary-Hawthorne, 1997, p. 138). This avoids the problem of univocity for it allows us to conceive of communication as an ongoing endeavor to form a common linguistic community by holding each other to norms of inference (Warren, 2018, pp. 11–14). The proprieties of inference captured in PALM are the inferential and epistemic norms that we hold one another to as we go on building this shared linguistic community. These norms are not cashed out in terms of anything non-normative but are rather understood, in the end, as products of our taking up normative attitudes toward one another’s actions (both linguistic and otherwise).

could not list them all or because we don't yet know their content. This, according to Williams, is a useful thing to be able to do because it provides a means for more fine-grained evaluation of our inferential commitments and epistemic norms.

Williams is right to differentiate between explicative/performative function and utility and to recognize that these two dimensions of use play distinct explanatory roles. The EPF-clause characterizes what discursive moves are enabled by the vocabulary governed by the proprieties of the first two clauses. These include the explicative uses Brandom identifies for logical vocabularies, which we discussed in the last chapter, as well as descriptive uses and use in the performance of illocutionary acts of all kinds. The U-clause, which aims to say what good the practice is for those whose practice it is, requires some refinement without which we risk forgoing any explanatory advantage it promises to provide.

Williams claims that the utility of truth talk is that it allows users to explicate inferential or general assertional commitments and epistemic norms so that they can be made objects of critical examination. The idea is that when we use truth talk to undertake commitments to claims we couldn't otherwise make—for example, “Everything that the President says is true” when we cannot say *everything* that the President says either because we would run out of time or because we don't yet know what he will say—we give expression to a norm of accepting the claims we've generalized over as good reasons for other claims we might make. In making this generalization explicit, we are able to bring it under critical examination, i.e., to test its veracity. Given what we know in general about the truthfulness of the current President of the United States, the epistemic norm expressed by the truth claim is not one we should accept in our epistemic practices.

The claim is very likely false. The important point, though, is that if we were trying to assess the veracity of another's claims about a variety of facts but could not know that they took the President to be infallible, we would face a much more arduous task. The ability to make such norms explicit helps to reduce the amount of epistemic labor we need to do and makes our epistemic practices more efficient.³⁷

This makes the utility of truth talk apparent or at least puts us on the right footing to start making sense of it. The question of the utility of a practice is the question of what good it does for the beings who use it. It will not do simply to say that it allows them to do something they otherwise could not have done. We also need to know why that thing would be a good thing for them to be able to do. Why would they go in for the practice? Williams seems committed to this, but he does not foreground it. He tells us that truth talk allows us to make certain claims the object of critical examination, and we are to *assume* that this is a good thing to be able to do. The problem with leaving this point implicit is that utility plays a fundamental explanatory role. The utility of a practice explains its structural discursive norms, but to fulfill this explanatory promise, the utility needs to be properly specified.

Generalizing from the EMU for "true," we get a meta-theoretical framework for an explanation of meaning in terms of use for any vocabulary whatsoever. The explanation consists of giving the material-inferential norms for the use of the vocabulary, the epistemic character of those proprieties of material inference, and the function of the target vocabulary both within the context of our linguistic economy (what it is used to do) and in terms of its utility for the practitioners who use it (what it is good

³⁷With this little sketch I certainly do not claim to have settled disputes about the utility of truth talk. See, for example, (Price, 2003; Rorty & Price, 2010).

for). The content-determining I- and E-clauses constitute the conceptual content of the term in the sense of telling us what one says when one uses the term. They explain meaning in terms of inferential use. The functional EPF- and U-clauses explain the proprieties captured in the content determining clauses not by giving some deeper theoretical explanation of them but by telling us why it would be useful to have a practice governed by those proprieties. The functional clauses *rationalize* having a practice governed by these proprieties and, as such, serve to explain inferential patterns of use, or, to put it another way, the structural and inferential norms captured in the content determining clauses *enable* the practice to fulfill its function as identified by clauses (3) and (4). We can summarize EMU in the following way:

The Content Determining Clauses:

- (1) I-clause: The material-inferential (intra-linguistic) component captures the inferential proprieties (of the committive, permissive, and incompatibility varieties) of the target vocabulary item.
- (2) E-clause: The epistemological component captures proprieties of use concerning what one must be able to do in order to meet appropriate challenges, if anything at all, and what challenges will count as appropriate.

The Functional Clauses:

- (3) EPF-clause: The explicative/performative function component identifies the use of the target vocabulary in the sense of what it is used to do, i.e., what discursive moves it enables.
- (4) U-clause: The pragmatic significance (utility) component identifies the usefulness of the target vocabulary in the sense of what broader needs and purposes it fulfills for the beings who use it.

EMUs present us with a useful point of departure for summarizing PALM

precisely because they exhibit the same explanatory structure, but, we'll see, PALM is an expanded version of EMU. Williams' aim is not a comprehensive theory of meaning in terms of use but an account of use in the context of assertion that allows semantic minimalists to draw lines between different vocabularies roughly where traditional expressivists do and for similar reasons.³⁸ As such, in order to put EMU to use for present purposes, we need to free it of the Declarative Fallacy.³⁹

The first step to freeing EMU is the introduction of a new clause that captures the *pragmatic discursive norms* identified in the previous section. I'll call it the P-clause. Its purpose is to identify the pragmatic norms of assessment that, on the input side, determine how the speaker ought to have her entitlement to the speech act and, on the output side, the transformations of normative statuses of the speaker and others affected by the speech act. Once we introduce the P-clause we face a question about how to incorporate the attendant bifurcation in entitlement conditions represented by the distinction between *agent-neutral* and *agent-relative* inputs. This bifurcation, I think, is best represented by an actual bifurcation in the structure of PALM. Rather than a list of clauses, as EMU is presented, the bifurcated structure is better presented as a PALM tree (Figure 2). Moreover, the agent-neutral/agent-relative distinction on the input side tracks the distinction we have already made between *epistemic norms* and social or institutional *authority norms*. As such, we can think of the tree as having an epistemic branch and a

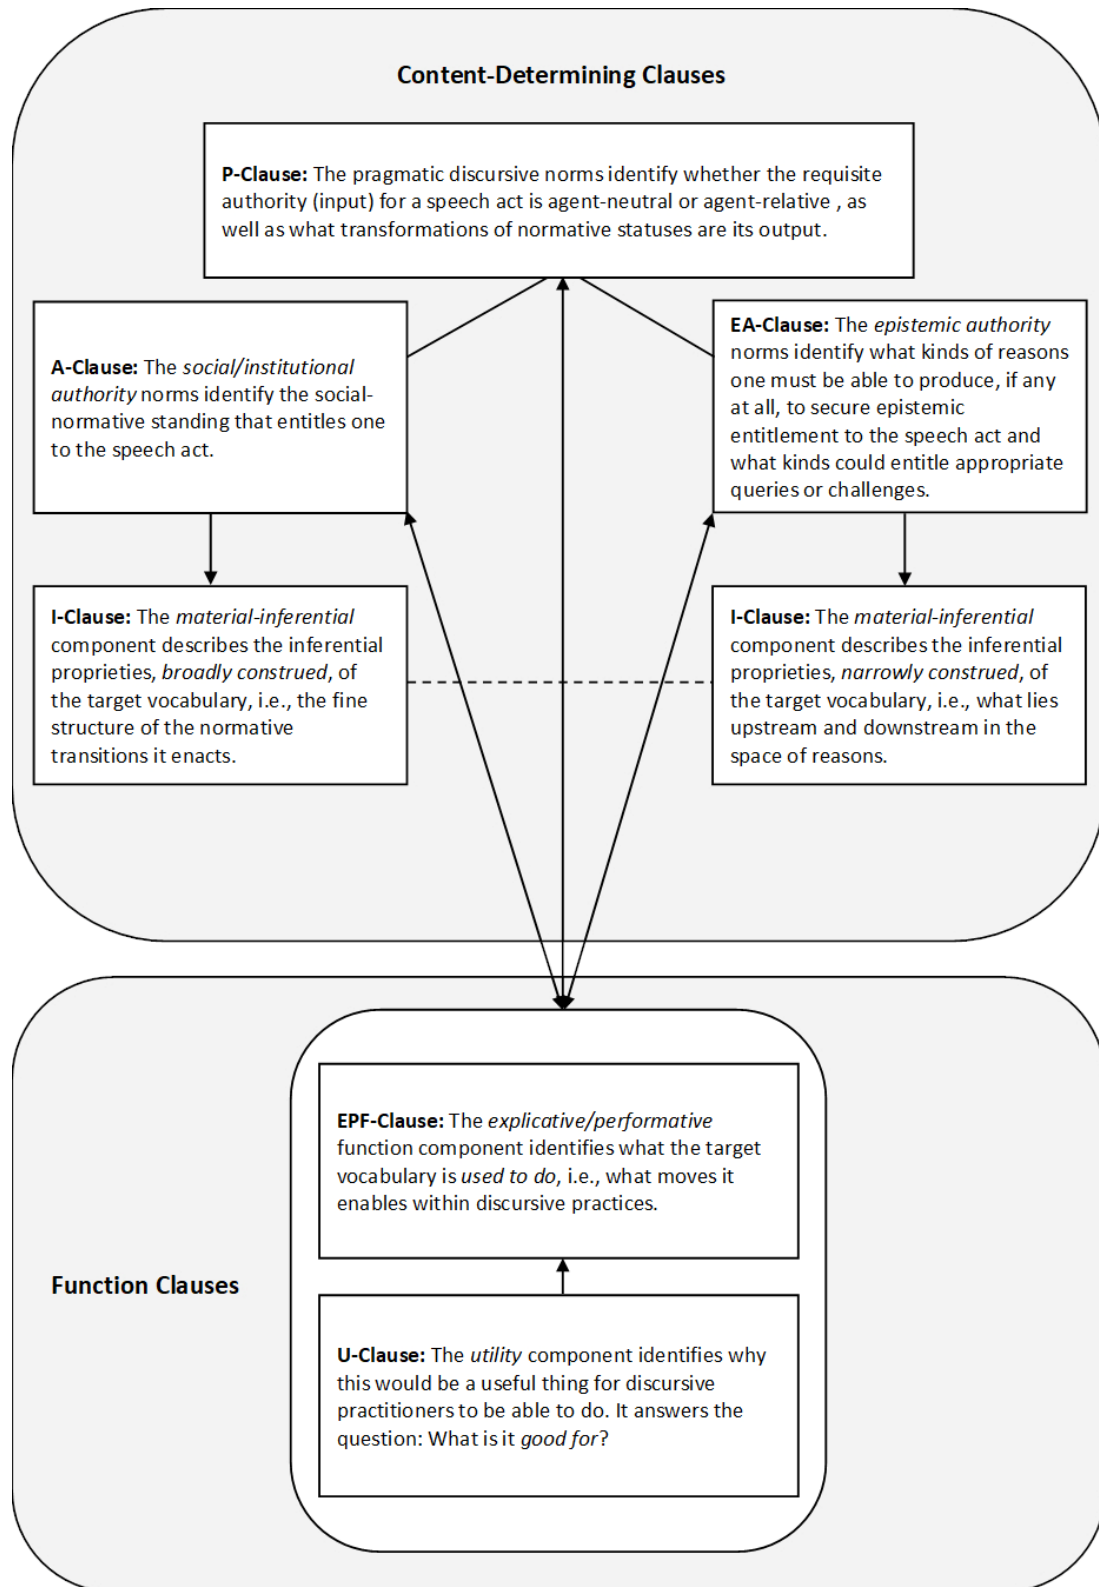
³⁸ EMU was developed as a constructive response to the work of Huw Price that attempts to bring expressivism and minimalism together in the service of global anti-Representationalism but also to the challenge posed to such a project by what Dreier has dubbed "creeping minimalism," i.e., the problem that "[m]inimalism sucks the substance out of heavy-duty metaphysical concepts...threaten[ing]to make irrealism indistinguishable from realism" (Dreier, 2004, p. 26). Also see, (Price, 2013).

³⁹ To be clear, I am not claiming that Williams himself commits the fallacy here. Given the limited scope of his project, it is open to him to expand it in various ways, some of which fall afoul of the fallacy others of which, like PALM, do not.

social/institutional authority branch.⁴⁰

⁴⁰I lack a general argument for claiming that these two distinctions track in this important way. There is something intuitive about the idea that all agent-neutral entitlement can be earned in the same way by supporting one's claims with reasons and that this is to be distinguished from agent-relative entitlements which depend, distinctively, on where one is situated in some broader normative space. I have searched for counterexamples to this hypothesis, and I've found none. The only case that I've found somewhat troubling is that of observatives, which, on this classification are not governed by what I am calling epistemic norms. I discuss this case below. Beyond this, the proof is in the pudding.

Figure 2: PALM Tree



6. Exploring PALM

Now that PALM has been properly introduced, I want to turn to clarifying some aspects of this explanatory framework. I'll begin with some comments about the diagrammatic PALM Tree. The first thing to note is the distinction between the content determining clauses and the functional clauses. The content determining clauses comprise both the structural discursive norms and the inferential norms that they shape, but, though these norms all play a role in determining content, conceptual content just is inferential articulation which is directly related only to the inferential norms. The functional clauses, as discussed, capture the explicative/performative function of a vocabulary and its utility and stand in a fundamental, bi-directional explanatory relationship to the structural discursive norms. The function clauses explain why a discursive practice—the one that uses the target vocabulary—is governed by the pragmatic and authority norms that it is in that it *rationalizes* those norms. It makes sense of why the beings whose practice it is would go in for a practice governed by these norms in just the sense that the function of a tool explains why it has the structural features it does. In the other direction, the content determining clauses can be understood as *enabling* a practice that can fulfill this function.

Next, take note of the dotted line connecting the I-Clauses of the two branches. This line is intended to indicate that these two clauses do not constitute distinct domains. The conceptual content of our vocabularies is constituted by their use both in speech acts with epistemic statuses as inputs and those with social/institutional statuses as inputs. The narrowly inferential norms—the ones that govern language-language moves, to borrow a piece of Sellars' terminology—are constitutive of propositional contents, but, as we've seen, the category of conceptual contents is not exhausted by propositional contents. This

broader notion of conceptual content includes, at least, observational contents (the non-propositional contents of observatives) and practical contents (the non-propositional contents of intentions to act as well as questions, orders, and requests which alter the normative standing of others in various ways). These varieties of content are needed to explain the fine structure of the normative transitions enacted by non-declarative speech acts. The conceptual contents of the vocabularies we study are constituted by their uses in a variety of speech acts beyond declaratives. Following Kukla and Lance, I conceive of this in terms of the *broadly* inferential proprieties governing their use.

There is one last thing to note about the structure of the diagram. The arrows from the structural discursive norms to the inferential norms are intended to represent the shaping relation previously introduced. The inferential proprieties of a discursive practice are what they are *in part* because of the pragmatic and authority norms of that practice.

Now let's turn to some more substantive commentary on PALM. Though I have been discussing PALM as an analysis of speech acts, it is better to think of it, following Kukla and Lance, as an analysis of functions enacted by speech acts. An individual speech act can enact more than one function. For example, if I say to my son, "The TV is too loud," my speech act enacts both declarative and imperative functions. It functions as a declarative with agent-neutral inputs and outputs, standing in need of reasons if challenged (though he'd better not!) and issuing a reassertion license for everyone, e.g., my wife who overhears from upstairs and repeats what I've said even if she doesn't hear the TV. However, it's also clear that my speech act is functioning as a demand with agent-relative inputs—as his father, I can tell him to turn down the volume—and outputs, as he is now obligated to turn it down. Kukla and Lance argue, in fact, that all speech acts enact

at least two functions: the one they wear on their sleeve and one that they dub the *transcendental vocative*. Every speech act, they claim, enacts a function of second-personal address, constituting a normative relationship between speaker and hearer calling for appropriate response or uptake in the form of recognition, action, or updating of one's beliefs (Kukla & Lance, 2009, Chapter 7).

Not only can individual speech acts enact more than one function, but the entitlement to an individual speech act can bridge the two branches of the PALM Tree. Authority, that is, can be partly a matter of an individual's position in some normative space and partly an epistemic matter for one in the same act. Consider, for example, a lawyer who raises an objection in court. This speech act—an act of objection—requires both that the lawyer be an appropriate official, recognized by the court, and representing a party in the case currently being heard, an agent-relative input, and that the objection be grounded in reasons, an agent-neutral, epistemic matter. If either of these conditions is unmet, it will be judged inappropriate. Warnings are also like this. Whether a warning is felicitous in part depends on whether the speaker is in a position to appropriately warn the target, e.g., only a person with proper social standing—as a friend, a parent, a spouse—can appropriately warn you not to ask a particular person on a date. The warning, however, is also assessed in terms of its epistemic propriety. It would be appropriate, after all, to ask why you shouldn't.⁴¹

Somewhat more interestingly, observatives also exhibit this duality of inputs. Observatives, on my account, primarily enact a function with agent-relative inputs—they

⁴¹ Some orders are also like this. Though in paradigmatic cases, the authority to issue a command is determined by the relevant social or institutional hierarchy, in other cases the order only stands insofar as it is *reasonable*.

are essentially individuating, as we saw earlier—and agent-neutral outputs. They give expression to one’s first-personal encounter with the world, are ineliminably first-personal in voice, but enter their contents into the space of reasons making them available *for everyone*. This puts them on the non-epistemic branch of PALM. It may be somewhat surprising that observatives count, on my account, as non-epistemic speech acts, as they are the subject of so much interest in epistemology. In part, this is merely a matter of terminology, but there is also something deeply right about it. The entitlement one has to an observative is the product of being taken to be a reliable reporter in the relevant context. Being a reliable reporter is having a particular, agent-relative normative status as a member of our epistemic community. It is a status that some lack—young children, those suffering from cognitive or perceptual impairments, habitual liars—and that cannot be earned (or recovered) by giving reasons in support of the content of one’s claim. It is right, then, that the entitling conditions for an observative are non-epistemic in the sense that I am using this term. On the other hand, observatives can face two distinct kinds of challenge or query. One might challenge the reliability as a reporter that a speaker claims for herself in issuing an observative. Such a challenge would need to be met by some evidence that one is, in fact, reliable. The other kind of query, though, is less pointed. It is a request for conceptual clarification: “Why do you think it’s a tufted titmouse?” Such a query may simply be looking for information or may be aiming to assess whether the speaker is basing her observation on anything more than a guess. If the speaker cannot respond appropriately by producing reasons for her judgment, she will be judged to lack entitlement not because she is not reliable but because she lacks facility with requisite concepts. It may, after all, be the case that she can ‘detect’ tufted titmice but doesn’t

really know what a tufted titmouse is or what, in particular, distinguishes them from other birds.

Both of these phenomena—multiple normative functions enacted by a single speech act and dual entitlement conditions for a single speech act—will play an important role in my subsequent analysis of moral assertions, or, as I shall call them, prescriptives. These speech acts, as we shall see, have agent-neutral outputs insofar as they issue reassertion licenses available for everyone, but they also enact a function with agent-relative outputs for the individual whose moral obligations or permissions they are about. Moreover, moral prescriptives have agent-neutral inputs in that they require reasons to be given in their defense, but they also require agent-relative normative standing as a member of the relevant moral community in order to enact their holding function, i.e., in order to be felicitous as attempts to bind others to their moral obligations. In the next chapter, I turn to the task of sketching the PALM for moral prescriptives, drawing out some of these interesting features.

Chapter 3: A PALM for Moral "Ought"

1. Introduction

With the pragmatic account of linguistic meaning (PALM) on the table, it's time to turn our attention to developing an account of the pragmatic and authority norms that account for the distinctive shape of moral discursive practice. This requires both a phenomenological investigation of the practice as we engage in it and an analysis of our discoveries in terms of the relevant kinds of norms. Along the way, we'll also see the first significant metaethical payoff of adopting PALM in the form of a response to the problem that Michael Smith has dubbed "*the moral problem*," i.e., the problem of how moral claims can be both truth-claims and action-guiding.

I begin in the next section with a phenomenological sketch of some central features of moral discourse and then, taking PALM as my guide, proceed to fill in first the pragmatic structural norms, in Section 3, and the authority norms, in Section 4. In the end, we'll have an account of the structural discursive norms that shape the inferential proprieties of moral discourse and that are the target of the functional explanation advanced in the next chapter.

2. Moral Discourse in Context

What is moral discourse? In the last chapter I argued that, like logical vocabularies, moral vocabulary is an explicating vocabulary. Its discursive function is to

make explicit practical inferential proprieties so that they can be examined and debated. As such, moral prescriptives will be uttered only when something has gone amiss. We operate discursively against a background of massive agreement in practice and judgment (Wittgenstein, 2001, sec. 242). Only when that agreement is threatened do we have reason to examine the propriety of practical inferences by making them explicit in the form of a prescriptive. Sometimes things go amiss in our reasoning—we find ourselves with conflicting commitments—sometimes as we engage discursively with others, and sometimes when we act. In each of these kinds of cases, it may be necessary to wheel out the machinery of moral discourse to try to elucidate what has gone wrong and, ideally, to correct it.

Let's consider a common kind of case in which we might find moral discourse put to this explicative use. Aziz, Bruce, and Carrie are discussing the upcoming blood drive to be held at their school. Carrie, taking it as given that anyone who is not prevented by some disqualifying factor like illness has a duty to give blood when there is a conveniently accessible blood drive, expects that both Aziz and Bruce are planning to donate. She forms this expectation on the basis of the background commitments she attributes to them and the fact that she takes both of them, as a result of the present conversation, to know that there is a blood drive to be conveniently held at their school on Friday. Adopting a slightly more regimented way of putting things, we can say that Carrie attributes to each implicit commitment to the propriety of the practical inference from "There is a convenient blood drive at school on Friday" to "I shall donate blood on Friday." As the conversation continues, Bruce indicates that he plans to donate, but Aziz lets it slip that he will be steering clear. Carrie, whose expectations have been disrupted,

might respond saying something like this: “Aziz, donating blood is important! You’ll be at school anyway, so why wouldn’t you do it?”¹ In saying this, Carrie makes explicit the previously implicit practical inferential propriety that she attributed to Aziz—that from the convenience of the blood drive to the intention to donate—and, furthermore, challenges him to say why he doesn’t accept its consequence. This is a significant discursive move. Not only does it challenge Aziz to present practical reasons for his lack of intent to give blood, it also opens a new discursive space that was previously unavailable to the three. Prior to making this propriety explicit, Carrie, Aziz, and Bruce could disagree about whether Aziz should give blood, but they couldn’t get at the crux of their disagreement, i.e., that Carrie and Bruce accept the now explicit inferential propriety that Aziz rejects. It is now possible for them to examine the defensibility of the inference from the convenience of the blood drive to the intention to donate by engaging in the game of giving and asking for reasons around the prescriptive that makes it explicit.

Nothing intrinsic to Carrie’s claim—say that it expresses a moral attitude, represents a moral property of the state of affairs, or aims to be categorically binding—makes it a moral claim. Rather, as I argued in the last chapter, the discursive context of a speech act—the structural norms that interlocutors enforce by being the bureaucrats and police who accept some reasons and not others—settles the kind of discourse in which

¹ In this chapter I alternate between more natural sounding, colloquial, and often more tactful and strategic expressions of moral approval and disapproval and attempts to hold others to their moral obligations and the more formulaic expressions often favored by philosophers, e.g., “S ought to φ .” The more natural language helps us to focus on the structural norms particular to moral discourse and to avoid being sidetracked by concerns about whether “So-and-so ought to such-and-such” is really a thing someone would say. Behind these tactful, contextually appropriate, and more strategically sound ways of expressing moral claims, however, there is always the deontic or evaluative assertion on which it is grounded. Carrie’s “Donating blood is important! Why wouldn’t you do it?” is just a more natural way of expressing what we might put in a more regimented way as “You (Aziz) morally ought to donate blood on Friday.” In some cases, when I aim to capture the distinctive structure of moral discourse in a regimented way, I will couch my examples in this more formulaic language.

the interlocutors are engaged. Carrie's words said in the same manner could explicate either a prudential or a moral propriety depending upon the sorts of reasons Aziz offers and the uptake that Carrie and Bruce give them. Aziz could relay that he passes out at the sight of his own blood, so that it would be imprudent for him to donate, or that he's going to be out of town making it inconvenient for him to be there. How the conversation develops, i.e., whether it takes a prudential turn, a moral turn, or goes some other way, is a matter of how the interlocutors treat these reasons. If Carrie accepts the first as defeating her prescriptive, then she has uttered a prudential prescriptive. If she accepts the second, it may yet be a moral one, but she might take the duty to donate blood to be an imperfect one such that matters of convenience bear on whether it is in force in the present case (Kant, 1785). If she rejects both and will only accept reasons that are perfectly general—having nothing to do with Aziz's interests, preferences, desires, or values but with what anyone in his circumstances ought to do—then she has uttered a moral prescriptive.

The reasons that interlocutors will accept from one another are a reflection of the discursive norms to which they hold one another. These norms generate a pattern of reasoning, which determines whether any particular bit of discourse is an instance of moral discourse and, so, whether the prescriptives advanced are moral ones.² The pattern of reasoning in the moral case is one in which the reasons given and accepted are ones

² It is worth noting an important consequence of this. There is no way to determine *a priori* which matters are moral matters. Whether some question is taken up as a moral question is determined *in situ* by the attitudes of the discursive participants mediated by the discursive norms they enforce. There may well be identifiable patterns in what kinds of matters tend to be treated as moral—for example, those in which one's actions affect the well-being of others or those that are central to an individual or group's self-understanding—but it will not be an aim of mine to say just what these patterns might be. Rather, my goal is to identify the patterns of reasoning in which we engage *when* we treat a matter as moral and to explain why we engage in this peculiar kind of practice. I'll make not predictions about when we're likely to take up that practice.

that would count as a good reason *for anyone* in the target's circumstances to act in the way prescribed. They will have nothing to do with anything particular to the individual—say her tastes or interests—or with her institutional or social status. Thus, if Carrie will not accept Aziz's reasons of queasiness and inconvenience, but would accept—or at least entertain—his claim that he is not obligated to donate blood because no one has a duty to forego his own bodily integrity in order to aid another, then Carrie's is a moral prescriptive. The aim of this chapter is a clearer picture of the norms that give moral discourse this distinctive shape.

The story about Aziz, Bruce, and Carrie is a case in which background agreement in judgment among a group of interlocutors has gone amiss. Before moving on, it will be helpful to have a rough and incomplete list of the other kinds of contexts in which things can and do go amiss, calling us to engage in moral discourse.

Personal Deliberation (PD): An individual sometimes find herself with conflicting commitments about what to do in her present circumstances or in some possible circumstances for which she is contingently planning. She thinks about what to do, provisionally accepts claims of the form “I ought to φ ”, and then reflect on reasons for and against. The output of such deliberation is usually the endorsement of a claim of the form “I ought to φ in c ” from which it follows, in circumstances c , I shall φ , i.e., an intention to act in a particular way in the specified circumstances.³

Interpersonal Deliberation (ID): We sometimes find ourselves in disagreement with one another about what to do in the present circumstances or some circumstances for which we are contingently planning, a kind of disagreement in judgment. We put forward

³ In cataloging these contexts, I am adopting the more regimented expression of prescriptives for the sake of clarity.

claims of the form “I ought to φ ,” “We ought to φ ,” or “S ought to φ ” for some S who is a participant in the deliberation. We then test those claims against reasons for and against produced by fellow participants. We test how sensitive those reasons are to changes in context, debate which empirical facts are relevant to settling the truth of the prescriptives, and try to fit the prescriptive with more general principles we accept and with our intuitions about relevantly similar moral judgments. We also make explicit and evaluate the values that are at stake in those judgments. Ideally, the output of such deliberation is the joint endorsement of a claim of the form “I ought to φ in c ,” “We ought to φ in c ,” or “S ought to φ in c ,” from which it follows that, in circumstances c , I, we, or S shall φ .

Third-personal Assessment (TPA): We often talk *about* what others—either in our community or outside of it—have done and what they ought to have done differently, a kind of disagreement in practice. We level moral assessments of their actions of the form “S (which may be plural) ought (not) to have φ -ed,” backed by reasons we take to be relevant to the circumstances in which he, she, or they acted. This assessment is not intended to directly hold S accountable for her actions. Instead, it is a means of signaling like-mindedness or membership in a moral community, a way of building community consensus about the rightness or wrongness of actions or support for sanction, a means of signaling to other members of the community what is and is not acceptable behavior, a way of reinforcing social norms, a tactic for challenging broadly accepted social norms, or a form of pedagogy. Nonetheless, should S become aware of our assessment, it would have special normative force for S that it would not for us.

Second-personal Assessment (SPA): We often talk *to* others about what they have done and what they ought to have done differently, another kind of disagreement in practice.

We level moral judgments of their behavior of the form, “You ought (not) to have φ -ed,” backed by reasons we take to be relevant to the circumstances in which they acted, though these prescriptives are often more tactfully couched. SPA is usually aimed at directly holding others accountable for their actions, eliciting recognition of norm violations, encouraging acts of contrition, arousing feelings of moral guilt, and so on, but they are not themselves a form of moral sanction.⁴ SPAs are ineliminably second-personal speech acts directed at the person whose commitments they are about. As such, they generally require the speaker to have some special relationship to the target of the assessment.

Third-personal Holding (TPH): We often talk about what others ought to do without those others being present. This is much like TPA, except that it occurs prior to the action. We are concerned with communicating what the norms that bind them require of them in these circumstances, so TPH issues in judgments of the form “S ought to φ in the present circumstances.” However, the immediate social significance of these prescriptives is not to pressure those whose commitments they are about to act in the way prescribed but rather to fulfill the same ends as SPAs. Nonetheless, should the individual whose commitments the prescriptive is about overhear, it would have special normative force for her that it does not have for us.

Second-personal Holding (SPH): We sometimes talk *to* others about what they should do and express judgments of the form “You ought (not) to φ in the present circumstances.” Like SPAs, these speech acts are ineliminably second-personal, directed

⁴ Moral sanctions, as I understand them, are speech (or physical) acts that call upon one’s agent-relative authority not only to call attention to the failure to fulfill a commitment but to *demand* corrective action either by imposing punishment of some kind or requiring some form of reparation.

at the person whose commitments they are about. They require the speaker not only to be able to provide reasons in defense of the claim, but also to stand in an appropriate social relationship to its target. SPHs are not imperatives—they do not issue a demand that the target ϕ —but function as attempts to call the target’s attention to norms that already bind her, obligations she already has and ought to recognize. They ostend these obligations but also signal the speaker’s readiness to provide reasons that they are, in fact, obligations the target has. Being second-personal speech acts, however, they also aim to add additional normative weight to these existing commitments. An individual properly targeted by a SPH who does not act in accord with it violates not only her obligation but her relationship with the speaker, as well.

In each of these contexts, moral claims are advanced to explicate an inferential propriety over which there is disagreement either in judgment or in practice so that this disagreement can be openly examined. These contexts differ, however, in that they foreground different pragmatic functions enacted by moral prescriptives such that focusing on different contexts will aid in our examination of these different functions. In SPH, for example, the aim is to elicit uptake by the target of the moral oughts that bind her so that she might act in accord with them. TPA, on the other hand, is primarily informative or pedagogical. We utter TPA prescriptives when we gossip about someone’s misdeeds or when we make an example of someone in teaching our children how not to behave.

3. The Pragmatic Structural Norms of Moral Discourse

Though different pragmatic functions dominate in different contexts of use, I’ll begin by examining what all moral prescriptives share. The first thing to note about

deontic moral assertions is something that realists have always grasped, that emotivists and other anti-realists rejected, and that contemporary expressivists and quasi-realists have tried to recapture: they are, fundamentally, truth-claims. They are the sorts of things for which we give reasons and over which we debate with the aim of discovering some answer. Though some moral disputes seem intractable, we tend to treat most moral questions—especially those we face immediately and that demand a response from us—as if they have an answer and as if it matters whether we get that answer right.⁵

Moreover, just like empirical claims, we don't think that anyone is in a position to authoritatively determine, once and for all, which moral prescriptives we must accept. There is always the possibility of discovering new reasons that bear on the question at hand and, so, of challenging the claims advanced by any individual or group. In all these respects, moral assertions behave just like any other assertions, and though others have found theoretical reasons for rejecting the truth-assessability of moral claims, they are still encumbered with the need to capture these basic patterns exhibited in their use.

Capturing these patterns of use in PALM is straightforward. Entitlement to moral prescriptives flows from the space of reasons and not from any particular, agent-relative

⁵ I tend to think that the degree to which many theorists have claimed that moral disputes are intractable is overblown. This is at least in part attributable to treating moral discourse as a kind of theoretical discourse rather than a practical one. We expect to be able to settle theoretical matters dispassionately and at some remove. We work on a problem, collect evidence, perform experiments, argue over their relevance, and so on, until we find some reason that settles the matter for everyone. We might be driven to solve a problem, but we also view solving it as a kind of intellectual exercise. Moral questions, however, are not like this. We may sit in seminar rooms debating them and spill pages and pages of ink hashing them out in journal articles, but moral matters are settled when we find ourselves forced to act. When we must decide what to do, we are much less likely to think that settling the question of what to do is impossible. We will find reasons that we accept, and we will not do so dispassionately. I think that when we focus on these cases—on the practicality of moral discourse—the intractability of moral dispute fades somewhat. We see that people do settle on what ought to be done, and then they do it. This does not mean that we will not re-evaluate their choices, that we cannot differ in our judgments, and that new disputes about these same issues cannot arise. It does mean, however, that in a particular discursive context, particular interlocutors were able to settle the question for themselves and for their purposes.

normative status held by the speaker. They are, in this way, distinct from orders, requests, invitations, and other speech acts, which require such standing, and akin to declaratives, which do not. The normative inputs of deontic normative claims are agent-neutral. In assessing entitlement to them, we assess, in the first instance, whether someone can defend these claims by giving reasons not whether they have some special standing to issue them.

This might be worrying, for some of the reasons that Representationalists like Mackie and some local anti-Representationalists like Ayer and Stevenson had for rejecting the commonsense platitude that moral claims are truth-assessable. One set of reasons is captured by Mackie's claim that, "if there were objective values, then they would be entities or qualities or relations of a very strange sort, utterly different from anything else in the universe" (Mackie, 1977, p. 38). They would have to have categorical "to-be-doneness" or "not-to-be-doneness" built right into them. Like the platonic form of the good, they would give both direction and motivation to anyone who came to know them (Mackie, 1977, p. 40). The furniture of the universe does not include such entities, but if moral claims are to be truth-claims, then it must be such entities that they purport to be truths about. Mackie's conclusion, of course, was that they are truth-claims, but they happen to be universally false. Others have argued that they are truth-claims within a kind of fiction (Joyce, 2007; Kalderon, 2007) or that they are not truth-claims at all but rather expressions of non-cognitive attitudes. In rejecting Representationalism and endorsing minimalism about truth, this concern is deflated. Moral claims are truth-assessable just in case they are grammatically declarative, and a speaker is entitled to append the truth-predicate to a moral claim just in case she is

entitled to the claim itself. The important question, then, is what entitles her to that claim? What reasons can be given its support?

A second set of reasons for thinking that moral claims could not be truth-claims stems from the thought that they are inherently motivating. Examining this challenge provides an opportunity both to defend the claim that moral prescriptives are truth-claims and to turn our attention from the normative inputs to the normative outputs of these speech acts. Most philosophers are committed to a basic division between beliefs and desires (or other motivating states). Beliefs are understood as mental states with a mind-to-world direction of fit. We seek to comport our beliefs to how things are in the world they are about. Desires, on the other hand, have a world-to-mind direction of fit; we seek to shape the world to match the ways we would like it to be. Desires are the states that motivate action, while beliefs serve a role in determining how best to fulfill our desires. Truth-claims are the sorts of things we believe, and so they cannot, on their own, motivate one to act. If moral prescriptives are merely truth-claims, then recognizing their truth, i.e., coming to believe them, could not alone account for one's being motivated to act in accord with them. Moral prescriptives, however, have often been thought to be inherently motivating, so treating them as truth-claims seems to get something very wrong about them. This kind of worry is at the heart of what Michael Smith has dubbed "*the moral problem*," i.e., the problem of how moral judgments can be both objective truth-claims and also inherently motivating or practical reason-giving (M. Smith, 1994, pp. 6–8).⁶

⁶ Smith sees three possible ways out of the problem. First, we can reject the idea that moral judgments are inherently motivating. This would be to adopt a kind of externalism about moral motivation like that advocated by Philippa Foot (Foot, 1972). Second, we can reject the idea that moral claims are truth-claims with descriptive content. This would be to advance a version of non-cognitivism ranging anywhere from the

Three tools that we've amassed so far make this old metaethical chestnut relatively easy to crack. The first is the intersecting set of distinctions between the normative inputs and outputs of speech acts and the agent-neutrality or relativity of those statuses. The second is the idea that speech acts are inherently pragmatically voiced. The third is our understanding of discursive norms as norms of assessment, i.e., part of the normative pragmatics of discourse. Prescriptives are truth-claims *about* the commitments of particular individuals (or groups). These claims, as we've seen, have agent-neutral inputs, but they are not impersonally voiced. They are often addressed second-personally to the individual whose commitments they are about—in PD, ID, SPA, and SPH—and they call for a particular kind of uptake from that individual. For anyone else, these speech acts function as ordinary declaratives. They call for third-personal recognition of a truth about someone else's commitments. So long as other conditions are met—the speaker is reliable, the hearer has no defeaters for the claim, etc.—then the hearer *epistemically* ought to build this commitment into her web of beliefs and act accordingly. For their target, however, these speech acts call for recognition “*that I, first-personally, am committed to [φ],*” which “*is essentially a matter of practically acknowledging that I am bound to [φ]*” (Kukla & Lance, 2009, p. 101 emphasis in original). Failure to give first-personal uptake to this fact, i.e., that I am so committed, is a kind of discursive failure. I literally don't understand the claim if I don't grasp its practical significance *for me*; “I don't understand what I am acknowledging.” It is not *just* a discursive failure,

emotivism of Ayer, Stevenson, and Barnes, to the prescriptivism of Hare, the expressivism of Gibbard, or the quasi-realism of Blackburn (Ayer, 1952; Barnes, 1934; Blackburn, 1984, 1988; Gibbard, 1990, 2003; Hare, 1952; Stevenson, 1972). We might also include here positions like that of Harman (Harman, 1977). Third, we could reject Humean belief-desire psychology opting instead for a kind of “besires” account. See, for example, (McDowell, 1979); the moniker “besires” comes from (Altham, 1987). For further discussion, see (Kukla & Lance, 2009, pp. 118–122).

however. “[I]f I don’t identify, first-personally, with the person whose commitment it is,” it is also a kind of deeper practical failure in the sense that, in failing to grasp *my* location in normative space, I fail to be bound by norms at all (Kukla & Lance, 2009, p. 102). If I cannot grasp that it is *me* that various norms bind, then no norms will have practical significance *for me*. I will simply fail to be engaged in the game that we all are playing.

Kukla and Lance analogize this last point to John Perry’s argument for the essential indexical. Perry argued that it is impossible to translate the practical inferential content of “I am here” into a set of propositions devoid of the indexical while still having it play the same role in the explanation of action (Perry, 1979). Similarly, Kukla and Lance claim, “no set of categorical propositions can capture our placement within the space of reasons, as bearers of particular commitments and entitlement.” I might know, for example, that the person with the 10 o’clock time slot should begin his presentation now, and that Tom has the 10 o’clock slot. If, however, I fail to grasp that *I* am Tom, the person with the 10 o’clock time slot, that bit of theoretical knowledge remains practically inert. Likewise for all commitments and entitlements: if I don’t grasp which are mine, which demand uptake from me, then these statuses will fail to make a difference for me at all (Kukla & Lance, 2009, p. 102). One who fails to exhibit this capacity not only fails to be a competent discursive practitioner, she fails to be a normatively bound agent at all.

That prescriptives call for uptake in the form of the target’s first-personal recognition of her commitments and entitlements means that they have at least some agent-relative outputs. They are speech acts that enact a function from agent-neutral inputs—their entitlement flows from the space of reasons—to agent-relative outputs that call upon a particular agent (or group) to come to recognize herself as committed to doing

what the prescriptive prescribes. As such, we have speech acts that have both directions of fit. They are truth-claims about the deontic statuses of their target—“a fact about the structure of the public world” (Kukla & Lance, 2009, p. 98)—and they have special practical significance for that target. Yet, there is nothing mysterious about these speech acts with both directions of fit, for recognizing the commitments I have is something that *should* move me to act. In fact, being able to recognize the practical significance such claims have *for their targets* is a logical condition of being able to engage in any discursive practices whatsoever. If one did not do this, one could not track the difference that doxastic commitments make to one’s particular doxastic state, i.e., one would not be able to play the game of giving and asking for reasons.

In attending to the pragmatic voice, the agent-neutrality of the inputs, and the agent-relativity of the outputs of prescriptives, we seem to have dissolved *the* moral problem, but some may feel a lingering doubt. The original problem was about how moral claims could *motivate* someone to act. We have seen that failure to give them the first-personal uptake for which they call is both a kind of discursive and practical failure, but we’ve said nothing about the link between that uptake and motivation. The residual worry, I think, stems from a misconstruction of practical motivation. Philosophers have tended to think of coming to be motivated to act as coming to have a desire for some end paired with a belief that the action is necessary to bring it about. On this view, the only way that moral claims could motivate action is if one has something akin to an antecedent desire to be moral. Lacking this desire, it seems that coming to recognize that one has a particularly moral obligation could not, by itself, motivate action. Given the broadly inferentialist framework we’re developing, it would be better to understand motivation to

act in terms of language-exit transitions rather than in terms of a psychological desire-like state. Sellars describes language-exit transitions as “learned transitions (S-R connections) in which from occupying a position [in a discursive practice] we come to behave in a way which is not a position [in that practice]” (Sellars, 2007, p. 36). To put it slightly differently, a language-exit transition is the product of a reliable disposition to respond differentially to the normative positions one inhabits in the game of giving and asking for reasons such that when one recognizes that one is the subject of an entitled, felicitous prescriptive one reliably responds by doing what is prescribed. Giving the prescription proper first-personal uptake just consists in making this transition. Having the disposition to do so inculcated in one is part and parcel of becoming a competent normative discursive practitioner in just the same way that having the disposition to respond to red objects by applying the concept red is part and parcel of coming to be a competent empirical discursive practitioner. One who lacks this disposition simply isn’t playing the game the rest of us are, but having this disposition has nothing to do with *feeling* some desire to act in any particular way. “We should not confuse acknowledgment of a practical reason with a psychological feeling of wanting to act in accordance with it, any more than we should confuse acknowledgment of a theoretical reason with a psychological feeling of wanting to draw inferences in accordance with it” (Kukla & Lance, 2009, pp. 121–122).

This formulation is on the right track, but it’s just a skosh too strong. One need not actually carry out the action in order to count as having given proper first-personal uptake to the prescriptive, i.e., as having the appropriate reliable differential responsive disposition. In fact, though first-personal practical uptake often takes the form of just

doing what one recognizes oneself as obligated to do, it just as often takes the form of feeling guilty for not doing it, of making excuses, or of suffering *akrasia*. Recognizing that one is rationally committed to φ does not always (or even usually) move one to φ . Rather, being practically motivated in this Sellarsian sense is recognizing *this* commitment as *mine*, that *I am committed to acting* in this way. This recognition manifests in that, when I do not act in the way I am obligated, I do not merely recognize a “theoretical discrepancy” between commitment and action but a “*transgression*”, a “*practical failing*” on my part to do what I ought to have done whether I felt any desire to do it or not (Kukla & Lance, 2009, pp. 121).⁷ This recognition itself constitutes proper first-personal uptake even if it does not issue in the prescribed action.

Given the pragmatic discursive norms brought to the fore here, moral prescriptives are speech acts that primarily enact discursive functions with agent-neutral inputs and agent-relative outputs, i.e., a Box-3 function in the Kukla-Lance typography from the last chapter. They are truth-claims entitlement to which derives from the space of reasons, but they have particular normative significance for those whose commitments they are about. For those practitioners, prescriptives call for uptake that recognizes the commitments *as theirs*, and such uptake is paradigmatically—though defeasibly—exhibited by acting as one is committed. In cases like PD and ID, this Box-3 function is dominant in the sense that the very point of the speech act is to enact it; it is the primary

⁷ It’s worth noting that this practical failure, not the failure of uptake, is what we normally assess as a violation of one’s obligations. Someone who fails to give first-personal uptake to the oughts that bind her—someone like the sensible knave of philosophical lore—is, at best, amoral. The ideally coherent Caligula is someone who we have failed to properly train into our linguistic practices, and our response to him should be to isolate him from our community, but when the person who can play the moral language game and give first-personal uptake to her obligations fails to act on the obligations she recognizes herself as having, the appropriate response is moral sanction. If she recognizes those norms and is responsive to them, then it is at least possible that she can be brought back into the fold.

social significance of the act. All prescriptive speech acts, however, also invariably enact a Box-1 declarative function. They are truth-claims that make some content about the obligations of particular practitioners publicly available. They issue reassertion licenses available to anyone who cares to take them up. This function dominates in acts of TPA and TPH, in which the speaker's primary concern is to communicate to others the commitments of a particular individual sometimes for pedagogical purposes, sometimes to coordinate sanction or enforcement, and oftentimes simply as gossip.⁸ Finally, we've noted that prescriptives often, but not invariably, enact a Box-4 holding function that requires the speaker to have certain agent-relative normative statuses. In cases of SPA and SPH, this Box-4 function is dominant, as the speaker's aim is to hold the target accountable to the commitments that bind her, but even in such cases, the Box-3 function of prescriptives is also relevant. These are acts that I will dub *alethic holdings*, following Kukla and Lance (2009, p. 111).⁹ They are attempts to use one's agent-relative standing to hold another to the oughts that *already* bind her and, so, the speaker must be prepared to give reasons. We'll take these up in Chapter 6 as part of a discussion of the agent-relative standing necessary to engage in moral discourse at all. We are left, then, with an analysis of the pragmatic structural norms of moral prescriptives that can be nicely

⁸ I say "simply," as this is such a common and mundane activity, but it also serves a vital role in our social practices. Gossip communicates and reinforces social norms, for example, and spreads important information about who is and is not a reliable cooperator within a community.

⁹ Note that prescriptives are not the only acts of alethic holding. We also often use Box-4 speech acts—imperatives, requests, invitations, and so on—to hold others to the oughts that bind them. In cases in which these acts are grounded in true prescriptives, they are acts of alethic holding, but when they aim to impart *new* commitments or entitlements to the target, I will follow Kukla and Lance in calling them *constatives*, even though this seems to be a slight misuse of this bit of Austinian terminology (Austin, 1962, p. 6n2). Both alethic holdings and constatives are discussed in greater detail in Chapter 6 in which I examine some dimensions of the standing necessary to engage in these acts and in moral discourse more generally (Kukla & Lance, 2009, pp. 105–113).

summed up in Figure 3. Boldface text indicates that this is the dominant function enacted by speech acts of this type.

Figure 3: Pragmatic Structure of Moral Prescriptives

Input → Output ↓	Agent-neutral	Agent-relative
Agent-neutral	Box 1 PD, ID, TPA , SPA, TPH , and SPH Declaratives	Box 2
Agent-relative	Box 3 PD, ID , TPA, SPA, TPH, and SPH Prescriptives	Box 4 PD, ID, TPA, SPA , TPH, and SPH Holdings

4. The Authority Structure of Moral Discourse

In this section, I turn to the norms of authority for moral discourse. We have seen that prescriptives enact, primarily, a Box-3 function from agent-neutral entitlements to agent-relative transformations of normative status, i.e., they are truth-claims about the commitments and entitlements of particular agents (or groups) and, as such, have special practical import for those agents (or groups). What we don't yet have is a way to differentiate moral prescriptives from other flavors of "ought" in terms of the structural discursive norms by which their use is assessed. This is the problem we raised for

Brandom at the end of Chapter 1, and the norms of authority for these various flavors of “ought” are just the tool we need to do this work.

4.1. The Agent-Neutral Epistemic Norms of Moral Discourse

Since prescriptives invariably require agent-neutral inputs, I will begin by exploring the right-hand-side of the PALM Tree for moral discourse. On the agent-neutral, epistemic side, one test of the adequacy of my account of the epistemic norms of moral discourse will be whether they can explain why moral “oughts” exhibit the pattern of reasoning that they do. Recall that Brandom distinguishes between flavors of “ought” in terms of the patterns of reasoning exhibited by each. The ought of taste is such that when Sam undertakes commitment to a taste prescriptive about Jack’s commitments he also becomes committed to a host of other inferences that are underwritten by the implicit attribution of a preference to Jack. The prudential ought is such that when Sam undertakes commitment to a prudential prescriptive about Jack, Sam comes to be committed to a whole host of other inferences underwritten by the interest he implicitly attributes to Jack. The institutional ought is such that in undertaking a commitment to an institutional prescriptive about Jack, Sam undertakes commitment to other inferences that follow from taking Jack to inhabit some particular institutional role. Finally—though there may be many other oughts of interest, I’ll limit my discussion to these—in undertaking commitment to a moral prescriptive about Jack’s commitments, Sam also becomes committed to endorsing the practical inference it explicates *for anyone* who is in circumstances relevantly like Jack’s without regard to their desires, preferences, or other

social-normative statuses, i.e., Sam takes it to be unconditionally binding.¹⁰ The challenge of this section is to say what the structural norms are that sustain this pattern of reasoning. In identifying these norms, we will also see how they differ for different varieties of normative discourse.

The epistemic structural norms of a practice are scorekeeping norms on which we rely in assessing a speaker's entitlement to a speech act that enacts a function with some agent neutral inputs, e.g., declaratives, prescriptives, questions, queries, and so on. Such assessment is couched in terms of what one must have done or be able to do in order to count as entitled or to vindicate entitlement in response to a challenge.¹¹ In general, what one must be able to do is enact some further performance, and, in most of the cases that will interest us, those further performances must be assertions. In short, one must be able to produce reasons in support of the speech acts to which one claims entitlement. The epistemic norms tell scorekeepers what sorts of performances are required and how they should be treated.

An example will help to clarify this idea. Consider Sam's claim, "There was a fire at the store today." The epistemic norms of ordinary empirical discourse tell us that Sam is default entitled to such a claim, i.e., an interlocutor ought to score him as entitled unless a well-motivated challenge arises. Now, suppose Katy responds, "Are you sure? I heard from Dory that it was a false alarm." The norms guide our scoring here in two ways. First, they tell us to score Katy's challenge as default entitled *and* as a legitimate

¹⁰ See Chapter 1, Section 7 of this dissertation for a more detailed discussion of these patterns of reasoning, including examples.

¹¹ Given the default and query structure of epistemic practice, it is the latter clause that it is usually operative. In standard cases, a speaker has default entitlement to her claims, but must be prepared to respond appropriately to well-motivated challenges or queries to vindicate that entitlement.

query to Sam's claim insofar as it addresses its content and doesn't raise a mere possibility of error. If Katy's claim were true, it would defeat Sam's entitlement to his. Second, they tell us that, if he is to vindicate his entitlement to the claim, Sam must now respond either with a reason that aims to defeat Katy's challenge—"The news reported a false alarm, but they didn't have all of the information"—or one that defends his commitment by providing good grounds for it, e.g., "Look at the pictures I took with my phone. You can see the smoke!" It would not do for him to say "I have a feeling it happened" or "I want there to have been one!" These would be the wrong kinds of reasons to vindicate his entitlement. In assessing Sam and Katy's claims and counterclaims, we were guided by the epistemic norms of the relevant discursive practice, in this case, ordinary empirical discourse.

It is not incidental that children offer wonderful examples of inappropriate reasons being brought to bear in both empirical and moral discourse. I say, "It's too late to start that project tonight," and my five-year-old son Sam responds, "But I don't want it to be late!" He does not mean only to express his disappointment, but to literally to rebut my claim that the hour is late. I say, "There are no more strawberries left," and Sam replies, "But I want more!" as if I can conjure them out of thin air and am only denying him them out of spite. Children must make a developmental leap before they come to grasp that certain kinds of reasons are not appropriate to support or challenge empirical claims; they must be inculcated into the social practice. The same is true in the moral case. For a four-year-old, "Because I want to!" seems an entirely appropriate reason for just about any action. It is only with training that they come to see otherwise.

In identifying the performances that could count as reasons for or challenges to other performances, epistemic norms function to distribute authority within a discursive practice. Consider another example. Our ordinary practices around talking and thinking about pain accord a special authority to the sincere avowals of the individual experiencing pain. In terms of the epistemic norms of pain-talk, this manifests in the fact that sincere avowals of pain are not only scored as default entitled but also as immune to challenge. Nothing could count as a well-motivated challenge to someone's sincere avowal that she feels pain. In such cases, I will say that the practice vests authority in the individual whose pain the claim is about.¹² Depending on the practice, authority may be vested in individual practitioners or in the linguistic community as a whole or it may transcend the attitudes of all practitioners. Ordinary empirical discourse, for example, does not vest authority in either individuals or the community, but rather in the objects that our claims thereby count as being about. On the other hand, discourse about what obligations one incurs by taking the vows of a religious order vests authority in the doctrines and leaders of that order.

In order to see how these norms differentiate practices by distributing authority, I will sketch, in very broad terms, oughts of taste and institutional oughts before turning to the central case of moral discourse. Suppose we go out for ice cream, and you tell me, "You ought to order the mint chocolate chip; it's particularly good here." This is a prescriptive of taste. It is true if and only if it really is the case that I ought to order the

¹² We do, of course, allow challenges to the sincerity of pain reports, and scientific and medical discourse might today or in the future diverge from ordinary pain-talk in requiring additional—perhaps neurological or biochemical—evidence to vindicate an individual's claim that she is in pain. I discuss these concerns in more detail in Chapter 5, but for present purposes, it is enough that we admit that, in general, we take people to be authoritative about what they feel, how they perceive things to be, what they think, etc..

mint chocolate chip. Whether I ought to do so depends on a number of things: whether I trust your judgment, whether I have any reason to believe that they've used sub-standard ingredients (perhaps I know of a shortage of fresh mint), and so on. Its truth, however, inescapably hangs on whether I have a preference for mint chocolate chip ice cream over other flavors at the present time in the present circumstances. It depends, that is, on my tastes. This is just an artifact of the structure of practical reasoning, which is invariably reasoning about means to achieve some end. In the case of oughts of taste, the end to be achieved is the fulfillment of the immediate preferences of the target of the prescriptive, but if the target lacks the preference on which the prescriptive is conditioned, then the prescriptive isn't true. Moreover, these preferences are something, like my own experience of pain, about which my own pronouncements are dispositive. If I say, "Yuck! Mint is disgusting!" then the prescriptive just is not true and, as such, has no grip on me. Now, of course, we know from dealing with children and the uninitiated that their tastes are not always well-informed, so their initial pronouncements may be mistaken, but once they have had a nibble or sip, the matter is settled by their response. There's just no room left for a well-motivated challenge.¹³

This feature of oughts of taste manifests in the epistemic norms governing this practice in the following way. Taste prescriptives are default entitled and are open to challenge, but there is one kind of challenge that they cannot withstand. If the individual

¹³ Relativizing to a particular time is necessary, since tastes clearly change. That I once thought broccoli disgusting and now quite enjoy it does not mean that my past self was mistaken about what his real tastes were; it means that my palate has evolved. Tastes for particular things—a fine wine, a well-crafted beer, a very dark chocolate—are, we often say, *acquired*. That very linguistic choice indicates that the neophyte's pronouncements were not a poor indicator of her *real* preferences but rather that she came to have preferences that she did not previously have. As L.A. Paul points out, tasting something for the first time can be an epistemically transformative experience in the sense that one cannot know what the phenomenology of the experience will be—and so whether it is something one prefers—until one has had the experience (Paul, 2016).

whose commitment the prescriptive is about raises a challenge that is expressed in the form of a recognitive speech act that gives expression to her own tastes or preference, and if the tastes or preferences expressed are contrary to those upon which the prescriptive is premised, that challenge is decisive. A recognitive like this enacts a function with agent-relative inputs and agent-neutral outputs and has a first-personal pragmatic voice ineliminably marking it as the speaker's first-personal recognition of some fact about her present mental state. This speech act is not just default entitled but also immune from challenge, for a speaker's claim is authoritative about her own tastes. Taste discourse vests authority over subjective preferences in the individual whose commitments the prescriptive is about by treating their recognitives about their preferences as immune from challenge.¹⁴

¹⁴ The role of individual preferences in settling matters of taste is central to debates about what has been called *faultless disagreement*. In order to salvage the idea that there really is disagreement, it seems we need to claim that the statements in question contradict one another. But if we claim that they contradict, then we seem to be committed to judging one or the other of those statements false and its speaker at fault. Salvaging what Wright has called the "Ordinary View" of faultless disagreement, i.e., that it is genuine disagreement (contradiction), faultless (no one is mistaken), and sustainable (that we can agree to disagree), seems nearly impossible (Wright, 2006, pp. 38–40). The "way forward," according to Wright, "must consist in finding the means to deny that [the speakers] must, in fact, each regard the other's view as mistaken—this despite the fact that [their] views are genuinely contradictory" (Wright, 2006, p. 52). Wright proposes a solution that he calls the "*True Relativist* exegesis of the Ordinary View." According to true relativism, "there are no absolute facts about taste—what it is true to say about taste depends upon a stance, or a set of standards, or a set of affective dispositions" (Wright, 2006, p. 52). I want to suggest, though I will not argue here, that PALM offers a novel approach to the question of faultless disagreement that is largely consistent with true relativism. The way forward is to see that expressions of taste like "Stewed rhubarb is delicious" as uttered by Tim are always taken up as relativized to some discursive practice or another. The inferential content of "Stewed rhubarb is delicious" is consistent, so that it does not contradict Crispin's claim, "Stewed rhubarb is not delicious," but what follows from these claims will depend, in part, on the ancillary commitments that are allowed to bear on it by the discursive context we take to be active. If we are engaged in helping Tim decide which desert to choose from the menu, then Tim's recognitive expression of his tastes is taken to settle the question of whether rhubarb is delicious and so what he ought to order for dessert. Crispin might disagree, but this only means that he ought not order the desert with stewed rhubarb for himself. If, however, Tim and Crispin are in the process of deciding what to put on the menu for their upcoming dinner party, the question of whether stewed rhubarb is delicious is settled neither by Tim's nor Crispin's tastes but by what they have reason to believe the preferences of their guests to be. In this case, they do disagree, and such disagreement is not faultless, for one of their expressions is better supported by the evidence, whichever it may be. Faultless disagreement plays on the ambiguity between these cases. In the

Moreover, these norms account for the pattern of reasoning observed in oughts of taste. In undertaking commitment to your prescriptive of taste, you implicitly attribute to me the relevant preference. The prescriptive presupposes that I, in fact, have the preference, and so you are committed to the whole family of inferences that follow from me having a taste for mint chocolate chip ice cream. If, however, I reject your attribution of this preference, your entitlement to the whole range of inferences collapses at once.¹⁵

Let's turn now to institutional prescriptives, i.e., those that follow from taking someone to inhabit a particular social or institutional role. This case is instructive in that it demonstrates how epistemic norms might be structured in a way that vests authority in a community rather than an individual. We each inhabit many varied social and institutional roles. We are teachers and students, parents and children, police officers and waiters, presidents and citizens. We are also friends, family members, neighbors, left tackles, goalies, and Rotarians. Each of these roles is, in the broad sense I intend, an institutional role. Each is defined by a set of norms prescribing some behaviors in some circumstances and proscribing others. We enforce these norms on one another, forming expectations on the basis of the roles we take each other to inhabit, rewarding compliance, and sanctioning violations sometimes mildly, sometimes harshly, in both deed and word. Moreover, we formulate "oughts" on the basis of these expectations. We think that a police officer ought to wear her uniform to work, that a neighbor ought not to

first, there is real disagreement about whether rhubarb is tasty, but it does not follow from this disagreement that anyone is mistaken, only that different courses of action are prescribed for Tim and Crispin. In the second case, there is again real disagreement, but in this case, there is also fault.

¹⁵ Why do we enforce these norms of authority in the practice? Why do we have these norms and not others? This is a question that takes center stage in Chapter 5, where I argue that it is not because tastes, preferences, or other mental states have some peculiar metaphysical status, nor because we have some special epistemic access to them, but, rather, because we (or our ancestors) have found that taking others verbal expressions of their mental states is the best guide to predicting their future behavior. We enforce these norms because they work for us.

take our parking space, and that the left tackle ought to have blocked the blitzing linebacker given the protection that was called. When we take these prescriptives to be true of someone, we (implicitly or explicitly) take them to inhabit the relevant role and, so, (implicitly or explicitly) attribute to them the whole family of “oughts” that follow from that habitation.¹⁶

Like that of taste prescriptives, the practice of using these institutional prescriptives has an epistemic structure with some distinctive features, though the ways in which these features manifest may vary depending on the institutional role that grounds the particular prescriptive. All such prescriptives are distinct from taste prescriptives in that authority over their truth is not vested in the individual whose commitments they are about. A neighbor who takes my parking space ought not to have done that even if she felt like doing it. Recognitives about her own preferences or desires do not defeat the prescriptive. This does not mean, however, that institutional prescriptives cannot be challenged or defeated. They are claims grounded in the space of reasons, and anyone who utters one can be challenged to defend it by providing reasons in its favor. The question, then, is what kinds of reasons will suffice?

Some institutional roles are well-defined. There are handbooks, regulations, laws, and so on, that make it the case that one who inhabits a particular role ought to act thus-and-so in such-and-such circumstances. To settle the question of when a police officer ought to wear her uniform, we consult her handbook. To decide whether the tackle really ought to pick up the blitz, we check the playbook. These institutional roles are such that when one inhabits them one grants authority to bind some part of one’s behavior either to

¹⁶ Cf. Lance’s discussion of lived normative statuses, (Lance, 2015).

those in an institutional hierarchy (the police commissioner, the offensive line coach) or to some pre-established set of rules. For the rest of us, determining which rules apply is trivial. We must determine what the relevant authority says, and that settles it.

Most of the institutional and social roles we inhabit are not so well-defined. In determining whether my neighbor really ought not to take my parking space, we appeal to the norms that define the role of neighbor. These may be quite local. For example, it may be an unwritten rule of my neighborhood but not yours that the spot directly in front of one's home is reserved for the resident. As such, it would be the case for each of my neighbors that they ought not to park in my spot. The question is how we determine whether there is such a rule however circumscribed it is. This is an empirical matter. Such an unwritten rule exists only insofar as there is some communal agreement about it, i.e., only insofar as there is some pattern of behavior among my neighbors with regard to parking behavior such that they count as following this rule, that they sanction violations of this rule, and that they would be likely to endorse this rule if it were made explicit.¹⁷ Any argument over whether my neighbor ought not to park in my spot is, in the end, about whether such a pattern exists. We can argue over whether it does on empirical grounds, but whether the pattern obtains is a matter of what we do, and we, as a community, cannot go wrong in doing it. This is a norm about which the relevant community cannot be mistaken. If we take it to be the case that one ought not to park in a neighbor's spot, then it is the case.¹⁸

¹⁷ How much agreement is necessary for this to count as a norm of our community? One is tempted to say...enough. Sellars, for one, recognizes that unanimity is unnecessary to ascribe a we-intention to a group (Sellars, 1993, p. 203)

¹⁸ As with other empirical matters, there are experts who have developed particular skills and earned special authority to speak on the norms of various communities, and we tend to give their assessments of what is required of individuals inhabiting institutional roles about which they are experts some extra weight. Longtime residents of my community are better sources on the norms of my neighborhood than are

Explicating these patterns as epistemic norms is more difficult than the previous case, but there are some general things we can say. The first thing to say is that institutional prescriptives are default entitled and always open to challenge; this is part of their nature as truth claims. Unlike taste prescriptives, however, no individual is in a position to level a decisive challenge, i.e., one that cannot itself be challenged. In general, the kinds of claims advanced in defense of or as challenges to institutional prescriptives are empirical claims about the norms of the relevant community, and the structure of authority that governs them is that of empirical discourse. We can note, however, some peculiarities. Certain kinds of recognitives again play a privileged role. These are usually couched as claims about what *we* do, i.e., first-person plural claims, that channel the norms of the community. They are, however, not merely observatives, claiming first-personal experience of what the community does, but recognitives that express the norms of the community as they have been internalized *by a member of the community*. These recognitives, while not immune from challenge, do seem to have a certain authority that mere outsider claims about the norms of the community do not. Even so, what settles the matter is not the authority of the speaker but the pattern of behavior the community exhibits, and, so, the judgment-in-action of the community as a whole, not that of any individual, is privileged.

Moral prescriptives are not merely attempts to hold others to the norms that one's community already accepts—even if they are sometimes this, too—so entitlement to a moral prescriptive cannot be vindicated by appeal to what *we*—the relevant

newcomers. Miss Manners seems to be one such authority on matters of formal etiquette and interpersonal relationships. However, even such experts can make mistakes. Their prescriptives are not immune from challenge. What constitutes such a mistake just is a misreading of what the norms of the community really are.

community—do or would accept. One cannot respond to a challenge to the moral prescriptive one has uttered by pointing to the patterns of behavior exhibited by one’s community nor can one respond by channeling the norms one has internalized in a recognitive about what *we* do.¹⁹ Whatever the social and institutional norms of a community happen to be, they are always open to moral evaluation.²⁰ It is the language we turn to in order to challenge the ways of a religious order, the laws of a nation, or the traditional practices of a community. Moral oughts, then, are distinguished from institutional oughts in that they are not grounded in empirical claims about the relevant community. The epistemic norms of the practice of moral discourse, whatever else they may be, are not such that they privilege the judgments-in-action of the community as a whole. To put it in terms of scorekeeping norms, a moral prescriptive cannot be challenge merely by appeal to an empirical claim about the practices of the community, nor can entitlement to a moral prescriptive be vindicated by such appeal.²¹ This distinguishes moral and institutional prescriptives.

Moral prescriptives are also distinct from taste prescriptives. Suppose I utter a moral prescriptive about Sam’s commitments, e.g., “Sam ought to share the last banana with his brother.” The force of this ought cannot be evaded, i.e., the prescriptive cannot be defeated, simply by Sam’s sincere utterance that he does not want to share, wants the

¹⁹ Though such attempts are often made, they are the result of the overlap between institutional and moral oughts and the ambiguity between the enforcement and revolutionary functions of moral discourse that I identify below. Moral vocabulary, in short, is often used to shore up institutional and social norms, and so it can sometimes seem appropriate to appeal to our common practices in giving reasons in support of a moral assertion.

²⁰ Though of course it is true that the more closed and authoritarian a community is, the more quickly such moral evaluation will be silenced, if anyone has the temerity to advance it.

²¹ The “merely” is not accidental. Sometimes such appeal is evidence of what we morally ought to do, sometimes conformity with the norms of the community is just what is morally called for, and many times settling empirical questions *does* settle what one morally ought to do.

whole thing for himself, or desires to see his brother cry. Sam's expression of his preferences seems to us irrelevant to settling the question of whether he really, morally ought to share. The kinds of cognitives over the truth of which Sam has absolute authority, i.e., those that settle what he ought to do to satisfy his preferences, are not scored in the same way in moral discourse. They are not accorded the same privileged authority that they are in taste discourse.

There are reasons that Sam (or others) could produce to challenge the prescriptive. He could claim that his brother really doesn't want the banana, doesn't like it, isn't feeling well, would be harmed because he is allergic, and so on.²² Or he might challenge the very idea that fair distribution of bananas is a moral good. He could, in short, produce prescriptive defeating reasons over the truth of which he has no special authority. As such, none of these reasons is immune from challenge, and so none settle the matter once and for all.

There is no way to settle *a priori* the scope of reasons relevant to a particular moral prescriptive. Any kind of claim *could* be relevant, and which are taken to be is a matter that can only be determined in context by the practitioners engaged in deliberation. Moral discourse is not unique in this way. What counts as a reason is always a matter of the context of utterance, and the background commitments, common ground, and

²² Note that, though though the target's preferences do not settle what she ought to do, it may be the case that *someone's* interests or preferences are relevant to settling the matter. If, for example, a surgeon is considering whether she ought to perform the indicated surgery, one factor that weighs decisively is whether the surgery is in the patient's best interest and, in the end, it is the patient's preferences that settle this matter. The important distinction between this case and oughts of taste is that the individual in whom authority is vested is *not* the same individual whose commitments the prescriptive is about. In this case, the patient's interests are a fact about the world that can defeat the prescriptive by demonstrating that some circumstances on which it is premised do not obtain. It is just an accident of the kind of case we are imagining that the truth about this matter is something over which some particular individual has absolute authority.

expectations of one's interlocutors are part of what constitute that context.²³ Though we cannot say in a general way what sorts of reasons will be relevant in an instance of moral discourse, beyond saying only that they will be agent-neutrally entitled, we can try to sketch the typical kinds of reasons that arise in specifically moral contexts. Setting aside the undoubtedly numerous cases in which non-moral matters must be settled, we can identify the common kinds of reasons appealed to in particularly moral circumstances. Suppose that, from the perspective of the discursive participants, all of the relevant non-moral facts have been settled. Supposing that they are still in disagreement about what ought to be done—about, say, whether the prescriptive “Aubrey ought to abort her fetus, which has been diagnosed with Tay-Sachs”—there are, I think, at least three broad classes of claims that are often brought to bear. We have understood prescriptives as having an explicative function bringing the implicit practical inferential transition from an empirical antecedent to an intention to act into the space of reasons for explicit evaluation. What needs to be evaluated is whether that transition is one to which the agent is entitled, whether this is a good inferential move. Context and background commitments that support the inference, of course, bear on this question. However, we can also address the goodness of the inference *itself*, *given* the background commitments and context. The reasons that bear on this constitute the first two classes. The first are reasons in the form of more general principles from which this particular claim follows. Call these *GP-reasons*. Tay-Sachs is a debilitating, painful, progressive neurological condition that causes extreme suffering and no hope of survival past about five years of age. As such, one might appeal to a general principle of the form, “Any fetus that is

²³ On the occasion sensitivity of reasons, see (Austin, 1946; Austin et al., 1950; Travis, 2005).

consigned to a short life of suffering ought to be aborted.” Should that not suffice—as is likely—one might then argue that a parent ought not to impose unnecessary suffering on her future child and has some obligation to prevent it. Even more generally, one might claim that all of these GP-reasons are instances of a general principle to the effect that one ought to prevent preventable suffering.

Supposing that GP-reasons don’t settle the matter, one might also appeal to the second class of reasons, which I will refer to as *value-ings*. These are recognitive expressions, i.e., first-personal expression with agent-relative inputs and agent-neutral outputs, that give voice to the values, desires, or interests of the speaker. Though I did not there label them as such, we’ve already encountered value-ings in the context of prescriptives of taste, where they were treated as authoritative. One feature that is distinctive of moral discourse is the way in which these speech acts are taken up by the practice, for any value-ing will be understood not as an expression of what the speaker values but of what the speaker takes to be objectively *valuable*. In terms of the epistemic norms of the practice, this manifests in that no value-ing is authoritative. Rather, any expression of value is always open to challenge and must be defended by reasons that could themselves be accepted by one’s interlocutors as a reason to value.²⁴ If we are engaged in moral discourse, then, the same recognitive speech act that could have settled the matter of what one prudentially ought to do can now only stand as one reason among others. It will not be treated by one’s interlocutors as settling the case.

²⁴ In reality, of course, most expressions of what is objectively valuable are attempts to win consensus, for we operate against a background of massive agreement in judgment about what is *in fact* valuable. One who is challenged to say why preventing suffering is valuable will be hard pressed to answer in any way other than exasperation with the question.

The third class of reasons that might be brought to bear on a seemingly intransigent moral disagreement is far more interesting to me. These, I think, are the reasons that, more often than others, move the needle leading to changes in belief and behavior, to what it is tempting to call moral progress. These reasons bring to the fore something we *could* see as relevant to the moral question but have not yet grasped. It may be that it is some reason of which we are aware, but we have not yet seen its relevance. Or it may be a reason that we have not yet grasped at all, i.e., something that has been, for one reason or another, epistemically unavailable to us. When we seek out a morally sagacious adviser, we hope that she will be able to spot these reasons and be a guide to moral deliberation not by aiding the systematization of our moral beliefs, i.e., in a project of making them more consistent by examination of general principles, but by bringing new reasons to bear on them. It is not just the sagacious adviser who draws our attention to these reasons. In many of the most interesting cases, those who have been or might be affected by our actions inhabit a perspective from which they can grasp reasons that we cannot. When those who have been marginalized, oppressed, victimized, or otherwise harmed but whose pain and suffering has gone hitherto unnoticed or has been dismissed as irrelevant find a voice within the dominant moral community, the recognitives they utter—their claims grounded in first-personal experience of harm—can provide important new insight. They can shed light on moral questions that a community has treated as settled, challenging institutional and moral norms alike. When these recognitives are brought to bear in a situated moral deliberation, i.e., when we are trying to decide whether to accept some prescriptive, they can sometimes provide grounds that settle the

matter for the interlocutors involved.²⁵ Though there is much to say about these recognitives and the distinctive role they play, for now I only want to note that, like all other reasons advanced in moral discourse, these are not authoritative. They can be and often are challenged. Though the speaker claims a special, first-personal authority for her utterance, it is no different from the authority claimed by someone who utters an observative from a perspective from which the relevant scene is better taken in. If I could see the getaway car drive off but you could not from your vantage point around the corner, then I am better placed to give a description of it. Moreover, there's no reason to think that only physical impediments matter. One can also be blinded by one's biases, motivations, interests, or privilege. This does not, however, mean that the claims of one who has the better perspective are unchallengeable. Once again we see that moral discourse does not vest authority in the claims of any individual.

The pattern of reasoning to which one commits in uttering a moral prescriptive is, as Brandom puts it, *unconditional* (R. Brandom, 2000, p. 91). In taking the pattern of reasoning to obligate the target to ϕ in the relevant circumstances, one is committed to taking it to obligate *anyone* in those circumstances. There is nothing about the target in particular—nothing about her preferences or her social/institutional role—that determines

²⁵ The phenomenon loosely described here is aptly seen as what happens when those who have been victims of testimonial injustice in the sense defined by Miranda Fricker find a way to be heard by the dominant moral community (2009). According to Fricker, testimonial injustice occurs when individuals in a community suffer from an undue credibility deficit usually as a result of some aspect of their identity so that their testimony is discounted. Their accounts of first-personal experience are immediately suspect in the eyes of the dominant epistemic community, and claims that they have been wronged or violated are easily dismissed. From time to time, however, marginalized speakers break into the dominant discourse. We have seen this play out in American life in the abolition movement, the women's suffrage movement, and the civil rights movement. More recently we have witnessed Black Lives Matter bringing police violence, misuse of force, and mass incarceration to public attention and the #MeToo movement challenge the silence imposed on women who have long faced sexual misconduct and assault by men in positions of power or authority.

that this practical inference is a good one. We've captured this in terms of discursive norms by showing that neither her recognitive expressions of preferences nor settled empirical claims about the relevant social or institutional norms could authoritatively defeat (or vindicate) the prescriptive. To put it another way, in moral discourse, no speaker's perspective is privileged in advance over any others and figuring out what to do is the messy retail business of weighing competing evidential and epistemic claims.

4.2. The Agent-Relative Authority Norms of Moral Discourse

When a once marginalized speaker comes to be heard in the dominant discourse, it is because she has finally been recognized to have the standing to hold those in the dominant culture to the moral oughts that bind them. She has, that is, gained the agent-relative standing required to felicitously utter a moral prescriptive to members of the dominant culture. We need now to examine these agent-relative inputs of moral prescriptives, the left-hand-side norms of authority in the PALM for moral discourse. To begin, recall that moral prescriptives are used not only in contexts of deliberation but also in contexts of assessment and holding. In particular, we sometimes utter moral prescriptives second-personally to the person whose commitments they are about either in assessment of an action they have performed (SPA) or in an attempt to hold them to their moral obligations (SPH). SPA and SPH both require something more than epistemic entitlement to a prescriptive in order to be felicitous.

Suppose that it is the case that I have a moral obligation to ensure that my children eat a nutritious diet. It follows that when my spouse, Katy, says, in a rather annoyed tone, "Tom, you shouldn't let Sam have a doughnut for dinner," she utters a true moral prescriptive in an attempt to get me to refrain from feeding him the doughnut in

recognition of my obligation to look after his health. This speech act is perfectly felicitous. As my spouse and Sam's mother, Katy has *standing* to issue such a holding. If, as Sam and I settled into our seats at the doughnut shop at dinnertime, the elderly gentleman seated across from us said, "Are you sure you want to feed him a doughnut?"²⁶ He may be quite right that I shouldn't be feeding Sam doughnuts for dinner. He's also out of line. It's no business of his what I feed my children. He lacks the agent-relative social or institutional standing to issue the prescriptive even if what he says is right. If he had said to his tablemate, "Can you believe that guy's giving his kid a doughnut for dinner?" this TPA would have been perfectly felicitous. He's entitled to the third-personal holding on the basis of the reasons he could give in its defense, but he is not entitled to the attempt to shape my behavior, as he does not stand in the proper social or institutional relationship to me.

The dominant pragmatic function enacted by SPA and SPH is a Box-4 agent-relative input to agent-relative output function. These speech acts are attempts to shape the behavior of an individual in accord with the norms that *already* bind her, but such attempts, as we've just seen, can only be pulled off by someone who the target herself recognizes as having the requisite agent-relative standing. So, what is the standing that one must have to engage in these acts of holding? This is a difficult question. It is not merely some particular institutional standing that we can pick out. It is itself something that is worked out in the ongoing engagement of agents with one another, but just how it is worked out, just when we will recognize someone as having the relevant standing with respect to us, is not something we have the resources yet to answer. For now, then, all we

²⁶ Where this speech act is understood as a more tactful way of enacting the same second-personal holding.

can say is that *some* agent-relative standing is required to felicitously pull off *some* of the pragmatic functions enacted by moral prescriptives. I fill in this part of the picture in greater detail in Chapter 6.

5. Conclusion

We now have a complete sketch of the structural discursive norms distinctive of moral discourse. Moral prescriptives primarily enact a Box-3 pragmatic function from agent-neutral inputs to agent-relative outputs. They are claims that draw entitlement from the space of reasons but that have special practical force for the individuals whose normative statuses they are claims about. They call on their target to give first-personal recognition to those normative statuses and respond appropriately. Drilling down to the agent-neutral inputs of moral prescriptives, we find that the epistemic structural norms of moral discourse are such that no individual or group perspective is privileged in advance over any other. Every moral claim—whether it is a prescriptive, a recognitive expression of value, or some other claim advanced as a reason within the context of moral discourse—is open to challenge and the speaker is encumbered with the obligation to respond to challenges by giving reasons. Moreover, neither those reasons nor any challenges are themselves immune from challenge. Interestingly, though, we've also seen that though moral claims are agent-neutrally entitled, enacting the distinctive holding function for which moral prescriptives are often used requires some special social normative standing. One must stand in the right kind of social relationship to the target to felicitously hold her to the oughts that bind her, though we have not yet made an attempt to say what constitutes such standing.

In closing this chapter, I want to take a moment to reflect on an idealization I have made in order to extract the basic normative pragmatic structure of moral discourse. Discursive practices, like all human practices, are messy works-in-progress that overlap and interpenetrate one another, often in unpredictable ways. Whether a particular prescriptive is prudential, institutional, or moral, for example, is often a matter that can itself be disputed within the practice, and so what reasons are allowed to bear on the question is more often than not itself an open question to be settled *in situ*. This question is settled within the discursive context by the structural discursive norms that interlocutors try to enforce. These are themselves sometimes made explicit and debated. We oftentimes find ourselves arguing not just over a moral issue but over *whether* it is really a moral issue or ought to be settled by institutional or legal norms or simply by settling what would be most prudent for the agent. There is no ready way to decide which questions are *really* moral questions, rather, that we engage in moral discourse indicates that we are treating a particular question as moral.

A special case of this interpenetration of discursive practices is of particular interest. It is often the case that we vacillate between settling certain issues by appeal to institutional norms and moral ones. This vacillation often manifests in the religious context. For a Catholic believer, whether abortion is permissible is, it seems, both a matter of religious teaching and morality, but the two are not easily prised apart. Our moral convictions are shaped in deep and indelible ways by our upbringing, our community, our religion, our teachers, and so on, and in many instances we take the question of what we morally ought to do to be settled by what the relevant institutions demand of us. The teaching of the church, then, settles the moral question. While

institutional and moral oughts are often entangled in this way, however, it is important to recognize that institutional authority is always open to moral appraisal, while the converse does not hold. This asymmetry shows that while institutional norms may shape our moral perspectives, moral questions are never fully settled by those institutional norms. We must recognize that, though moral prescriptives are often pressed into service in the *enforcement* of institutional norms, their use in aiming to transform those institutions—what I call their *revolutionary* use—is prior.

In unveiling the structural discursive norms of moral discourse in this chapter, I have set aside this entanglement with religious and other institutional norms, not because I think it is unimportant but because I see value in getting clear on the special role of moral discourse that comes through in its revolutionary use. One payoff of this clarity is that we come to see that, though religion has historically played an important role in shaping morality, it is not foundational to morality. In the next chapter, we turn to the question of what, if anything, is foundational to it in the form of an exploration of the function of moral discourse in evolutionary terms.

Chapter 4: The Function of Moral Discourse

1. Introduction

The pragmatic account of linguistic meaning (PALM) that I developed in Chapter 2 takes meaning to be constituted by patterns of use. Conceptual content—of both propositional and non-propositional varieties—is understood in terms of broadly inferential proprieties. These patterns of use are the product of the normative attitudes discursive practitioners exhibit toward the moves that they and their compatriots make in the practice. Those attitudes constitute a variety of discursive norms—pragmatic norms and norms of authority—that shape the inferential norms of a discursive practice. They distribute epistemic and social-normative authority within the practice and mediate the ways in which inputs to the practice—causal interactions with the world, individual preferences, emotional responses, etc.—are taken up by practitioners. Like all human practices, we have developed these over evolutionary time—sometimes purposely, but most often without design—to fulfill some needs and purposes that we have. Therefore, we can explain the normative structure of discursive practices in terms of how their pragmatic and authority norms enable them to fulfill the needs and purposes of the community whose practices they are.

In the last chapter, I set about applying the PALM framework to moral discourse by sketching its pragmatic and authority norms. In this chapter, I complete the PALM by investigating the needs and purposes fulfilled by this practice. I argue that moral

discourse is an evolutionary adaptation that resolves instabilities that arise in more basic cooperative practices under specific constraints that our ancestors faced. In particular, it provides for the integration of the interests and preferences of individual community members into the normative structure of the community's practices in a way that minimizes risks of defection, fracturing, and domination, thus making those practices answerable to individuals. In a slogan, moral discourse is meta-normative vocabulary that enables cooperation under constraint.

2. Practices and Meta-Practices

I'm going to argue that moral discourse stands as a meta-practice to various social-coordinative practices in which human communities engage. As such, I want to start by clarifying the practice/meta-practice relation. Social practices stand in all manner of relations to one another. One practice can reinforce another as the practice of advertising does for commerce. One practice can constrain another as policing constrains driving. A practice can also undermine another as lying does to promising. In some of these cases, one of the practices is parasitic on the other, but none of these is an example of one practice standing as a meta-practice to another.

Consider, on the other hand, the relation between the game played by professional basketball players on the court and the practice of the league sanctioning those games. The National Basketball Association (NBA) has changed the rules of basketball many times throughout the league's history. Some changes have been in response to the evolving athleticism of players, i.e., to things happening on the court, and some in response to demands of fans or changing technology, i.e., to things happening off the court. In all cases, the rule changes resulted from a practice distinct from the practice of

playing professional basketball. The rule changes are decided not on the court, according to the existing rules of basketball, but in the owners' meeting governed by Robert's Rules of Order. The owners' meeting in which changes to the rules of basketball are debated and decided stands as a meta-practice to the game of professional basketball played on the court. Other sports, with their various governing bodies, provide many more examples, but we can also easily see that everyday practices like driving and the political processes by which the rules of the road are set or walking and the much messier social practices in which rules of etiquette around passing in hallways, staying to the left or right on the stairs, and holding doors for others are settled stand in the same relation. In general, for any practices A and B, A is a meta-practice to B if and only if A provides for the explication, evaluation, and emendation of the norms of B (Lance, 2008). Meta-practices are the practices in which the rules of other practices are made explicit, debated, and altered.

These practices are not always as distinct as the NBA example makes it seem. In some cases, the practice and meta-practice are deeply integrated, such that the same players are engaged in each but relying on a somewhat different normative structure in changing the rules than in playing the game. This happens, for example, when a group of kids are playing a pickup game of football but make changes to the rules as the game progresses, maybe to accommodate less skilled players on one team or to avoid hazards as the evening light fades. It also happens when the owners in a condo community decide to institute new rules about the use of community space.

Practices and meta-practices may also differ in terms of the standing required to engage in them. The kids playing ball or and the members of the community all have

standing to engage in the meta-practice in which the rules are changed, but the same isn't true of NBA players. There is a distinct set of standing conditions to be acknowledged in the owners' meeting, and the players, in general, don't meet those conditions. We can also distinguish between cases in which the standing to engage in the meta-practice can only be claimed by engaging in the underlying practice and those in which it can be earned in some other way. The former are normatively closed, the latter normatively open. I will argue that moral discourse is more like the pickup game than the institutionalized structure of the NBA, and that it is relatively normatively closed insofar as one must be engaged in the underlying social-coordinative practices in order to claim standing in the meta-practice.

Sellars thought that this practice/meta-practice structure was essential to both language and thought. In "Philosophy and the Scientific Image of Man" he claimed that "anything that can properly be called conceptual thinking can occur only within a framework of conceptual thinking in terms of which it can be criticized, supported, refuted, in short, evaluated" (Sellars, 1962). Thought, properly so called, requires standards by which one can evaluate the correctness, relevance, and justification of particular thoughts. Something's counting as taking up a position in the space of implications requires already a meta-practice of evaluation of that position, without which it could be rule-governed but not rule-obeying behavior.

Sellars also seems to be committed to a stronger claim. He appears to think that any rule obeying behavior—that is, behavior *because of* a rule not just *in accord with* it—always brings in tow a meta-practice for the evaluation of the rules we are following. "Rule obeying behavior contains, in some sense, both a game and a meta-game, the latter

being the game in which belong the rules obeyed in playing the former game as a piece of rule obeying behavior” (Sellars, 1954, p. 34). For Sellars, this is what distinguishes merely conforming to rules from actually following them, what distinguishes the dance of the bee from the inscription of a linguistic sign. We can only follow rules if we can make them explicit, evaluate them, and, possibly, amend them, and this can only be done within a meta-practice.

I think that Sellars is right on this point, and the argument of this chapter exploits his idea. My thesis is that moral discourse stands as an important kind of meta-practice to various other normative practices in which we are engaged, i.e., moral vocabulary is a kind of meta-normative vocabulary.¹ The genus of meta-normative vocabularies is populated by species such as epistemic and semantic vocabulary that are often the target of great philosophical interest. They share three key characteristics. First, they are normative. They do not aim to describe, explain, or classify some episode or state of affairs but, to borrow from Sellars, to place “it in the logical space of reasons, of justifying and being able to justify what one says” or does. From a different perspective, we might say that to deploy one of these vocabularies is “to *rehearse an intention*” to treat the episode, state of affairs, or person in a particular way (Sellars, 1956, sec. 36, 1962, p. 408). Second, meta-normative vocabularies are always parasitic on some other practice. They stand in asymmetric relations to the practices for which they serve as

¹ This phrase—meta-normative vocabulary—may invoke thoughts of Brandom’s pragmatic and semantic meta-vocabularies (Brandom, 2008a, 2015). Pragmatic meta-vocabularies enable one to say what one must be able to do in order to deploy the object vocabulary, while semantic meta-vocabularies enable one to say what one says in the object vocabulary in a different vocabulary (presumably one with which we’re more theoretically comfortable, e.g., extensional rather than modal vocabulary). As should be clear, this is not the relationship I have in mind. Meta-normative vocabularies don’t explicate what one must be able to do to deploy normative vocabularies. Rather, they are themselves normative vocabularies that explicate inferential proprieties of *other* normative vocabularies.

meta-practices in that the underlying practice could be carried on without the meta-practice explicitly in place. However, the norms of the underlying practice will remain implicit in practice, for once they are explicated we are already in a meta-normative vocabulary about them. Practices in which behaviors are judged right and wrong, outcomes good and bad, and in which reasons for action can be demanded and given do not require a meta-normative vocabulary like moral discourse to carry on in their own ways. Meta-normative vocabularies, however, are introduced to solve certain kinds of problems that arise out of the underlying practices as a result of various unavoidable constraints under which practitioners engage in them. Once they are introduced, these vocabularies change the very nature of the underlying practices by providing the practitioners with resources to rationally revise the practices to fit their changing needs. They make those practices answerable to their users.

The third characteristic shared by meta-normative vocabularies is their fundamental role in explicating what Sellars called “the conceptual framework of persons,” which “is the framework in which we think of one another as sharing the community intentions which provide the ambiance of principles and standards (above all, those which make meaningful discourse and rationality itself possible) within which we live our own individual lives” (Sellars, 1962, p. 408). This ambiance of principles and standards exists as a result of practitioners taking up normative attitudes towards one another’s behaviors, but it remains ambient until they develop a practice in which those principles and standards can be made explicit. Meta-normative vocabularies like epistemic vocabulary, semantic vocabulary, and moral vocabulary explicate the principles and standards constitutive of the frameworks within which we place one another when

we evaluate each other as potential informants, speakers of languages, and members of the community responsible for our actions. In this way they form the conceptual framework of persons, the part of the manifest image that, for Sellars, must be *joined* to the scientific image of man in the world. This chapter is a step in the direction of joining this framework to the scientific image not by providing a reductively naturalistic account of moral norms but by telling a suitably naturalistic evolutionary story that situates them as the products of natural beings interacting with one another and their environment.

Though I will not argue explicitly for this claim, I also think that *all* meta-normative vocabularies are subject to the kind of functional explanation advanced in this chapter. This is not just to say that we could give a PALM for each, but to further claim that the utility clause of the PALM would be filled out by appeal to the needs of practitioners and the specific constraints under which those needs must be fulfilled. Edward Craig advances an explanation like this for epistemic vocabulary arguing that the concept of knowledge can be illuminated in terms of the basic human need for true information about one's environment along with the fact that we are very limited in terms of the information we can gather on our own. Since "the tiger that Fred can see and I can't may be after me and not Fred," it will be in my interest to have a way to tap into Fred's primary stock of beliefs as well as to have ways of tracking which of my compatriots are reliable informants (Craig, 1991, p. 11). Epistemic vocabulary enables us to do just that.²

This gives the shape of the account I present in this chapter. My aim is to do for moral discourse something like what Craig has done for epistemic discourse. I have

² On the function of epistemic vocabulary, also see (Hannon, 2015, 2018, 2019).

already sketched the explanans: the structural discursive norms of moral discourse that are responsible for its distinctive shape. Now I must explain how a practice given shape by those norms is of use to beings in the circumstances in which human beings find themselves operating under the constraints we face. While Craig tells a kind of just-so story about epistemic vocabulary, I frame my account of the underlying normative practices in evolutionary terms. Drawing on work in evolutionary psychology and anthropology, I sketch a likely evolutionary path by which human pro-sociality and related practices developed to solve coordination problems faced by our ancestors. I then argue that these accounts leave us at an unstable point, facing further problems that threaten to undermine ongoing coordination. The function of moral discourse is the solution of these problems under constraints that require limiting desertion or defection from the cooperative projects of the community. It is, as such, what I will call a meta-coordinative practice that stabilizes practices of communal coordination and cooperation.

I adopt this evolutionary approach for two reasons. First, the shifts in human sociality and coordination wrought by moral discourse itself makes its relation to the underlying normative practices opaque from our perspective. The evolutionary account I sketch clears the air, allowing us to see more plainly the predicament in which we find ourselves by focusing on the problems of our long-lost ancestors. Bringing the problem solved by moral discourse and the constraints under which it evolved into relief in this way enhances the explanatory power of the account. Second, beginning with an evolutionary account presents the possibility that the present conditions in which moral discourse is used are not always a good fit for the practice. Moral discourse evolved to solve certain problems under certain constraints, and these conditions hold to a greater or

lesser degree in different circumstances in contemporary life. In circumstances in which they hold to a high degree, moral discourse tends to function fairly well, in cases in which they do not, we find the many examples of failures to find common moral ground, to resolve moral disagreements, and so on. Beginning with the evolutionary perspective allows us to explain why such breakdowns occur and to predict the circumstances in which they are probable.³

We should not hope to explain *too* much from this evolutionary perspective. It helps us to grasp why practitioners would go in for a discursive practice governed by certain structural discursive norms. As such, it explains the structure of the practice but not the content of the claims or prescriptions in which it trades. For moral discourse, in particular, my proposed account explains the underlying structure of moral discourse—the norms by which we keep score on one another when we engage in the practice—but the content of the moral norms discovered via the practice are not determined by this structure alone. They are the products of long histories, of cultures and religions, of problems faced (some of which we solved, other to which we succumbed), of economic and power relations, and of the value-ings and evaluations of individual human beings. I do not aim to explain the great variety of moral dicta, only the structure of the practice by

³ Accounts of the function of morality in terms of an evolutionary genealogy have recently come under attack from philosophers who argue that there are too many functionally relevant discontinuities between the early human societies in which morality allegedly evolved and our contemporary human societies. Though these genealogies may explain the emergence of moral practices in earlier times, they do not shed light on its present function, for it could not serve that function in contemporary societies (Smyth, 2017). My account avoids this charge in two steps. First, in taking moral discourse rather than ‘morality’ (whatever that might be) as my explanatory target, I aim to explain a very malleable practice that lies behind a great variety of extant moral systems. Though the shape and needs of human societies have changed dramatically, there is continuity—at least in some circumstances—in the need for a practice that solves the kinds of problems I take moral discourse to solve. Second, though, I argue that the various discontinuities to which we can point serve to explain the many cases in which moral discourse fails, in which cooperation breaks down, or in which it must be secured by some other means.

which they are discovered.

Though I think that the evolutionary story I offer here is on the right track, I do not wish to wed my account of the function of moral discourse to its accuracy. Rather, I think that it stands or falls on the basis of whether it is explanatorily fruitful to see moral discourse as a meta-practice that meets particular challenges faced by beings like us who engage in different varieties of practical deliberation under various constraints and who have evolved certain pro-social dispositions and sentiments. If the story about our needs and the constraints we faced is plausible, that will be enough.

3. The Evolution of Human Cooperation and Moral Sentiments

Let us commence, then, with an account of the evolution of cooperative practices and sentiments among our early human ancestors. There is no doubt that human beings engage in kinds of social collaborative behavior that are qualitatively different from our great ape cousins. The social lives of chimpanzees and bonobos, with whom we share a common ancestor that anthropologist Christopher Boehm calls *Ancestral Pan* (Boehm, 2012a), are structured by competition for vital resources like food and mates. They engage in cooperative behaviors, but in order to understand their cooperation it must be viewed in the broader context of competition, which represents its “driving force” (Muller & Mitani, 2005, p. 278).

Chimpanzees and bonobos form coalitions and cultivate “friendships” by engaging in behaviors such as reciprocal grooming and food sharing. These friendships are maintained by a kind of “emotional reciprocity:” chimpanzee A grooms chimpanzee B out of a sense of sympathy for an individual on which he is dependent, chimpanzee B grooms A not as a result of an accounting of recent exchanges but out of sympathy for

one on whom he has now become dependent (Melis, Hare, & Tomasello, 2008, p. 951).

This interdependence based sympathy is directed toward individuals who are perceived to be fit coalition partners, i.e., those who will be dominant in fights for food, mates, dominance status, and other important resources (Tomasello, 2016, pp. 23–25). As Muller and Mitani summarize:

The most prevalent forms of cooperation among chimpanzees...are rooted in male contest competition. Chimpanzee males maintain short-term coalitions and long-term alliances to improve their dominance status within communities and defend their territories cooperatively against foreign males. Other prominent cooperative activities, such as grooming and meat sharing, relate strategically to these goals. Females are far less social than males, and do not cooperate extensively. Nevertheless, the most conspicuous examples of female cooperation also involve contest competition, as females sometimes cooperate to kill the infants of rivals (Muller & Mitani, 2005, p. 317).

Chimpanzees are prosocial only in the service of competition. As such, their sociality is not evolutionarily mysterious. They engage in low-cost prosocial behavior that directly increases fitness. Human beings, however, often engage in potentially high-cost prosocial behavior that often has no direct payoff for the individual or their kin group. We are altruistic. We act on motivations like beneficence, compassion, concern for others, and benevolence. We sacrifice ourselves—our resources, our wellbeing, even our bodies—for the benefit of others. Sometimes such sacrifices are for the good of all. Everyone, including the individual who sacrifices, will in the end be better off if justice, equality, and fairness result. Sometimes, though, these sacrifices are purely self-immolating. We act purely for the sake of the other—who is often not kin—without concern for any future benefit for ourselves. This is certainly more mysterious. Given the selfishness of evolution, why would individuals act for the sole benefit of non-kin others? Why would we choose to do something because it is 'right' or 'good' if it does not increase our own (or our kin's) fitness?

The obvious response for a naturalist is to find a reasonable, well-grounded, empirically constrained story to tell about how morally motivated actions *really are* conducive to fitness in some less direct way. Evolutionary psychologist Michael Tomasello’s “Interdependence Hypothesis,” which claims “that at some point humans created lifeways in which collaborating with others was necessary for survival and procreation,” is such a story (Tomasello, Melis, Tennie, Wyman, & Herrmann, 2012, p. 674). The Interdependence Hypothesis is a two-step account of the evolution of human moral practices, which holds that the various cognitive, communicative, and social mechanisms that underlie them were selected by the demands of obligate mutualistic collaboration. The first step is the evolution of cognitive processes to solve the kinds of coordination problems that arise in small-scale, dyadic collaborative foraging. The second step explains how further cognitive abilities and social technologies—like social norms and institutions—evolved from the cognitive capacities required for dyadic cooperation to support the kind of group-level collaboration needed to work together interdependently to compete with other social groups for limited resources and to defend one’s group from attack.

Tomasello’s focus—and ours—is on the cognitive processes that evolved to solve coordination problems generated by the imposition of obligate collaborative foraging rather than on the evolution of altruism or moral emotions. He argues, however, that the development of some basic pro-social, altruistic sentiments served to set the stage for the later development of these cognitive capacities. Around two million years ago our early human ancestors underwent a “kind of self-domestication” through which they became less aggressive and their social interactions became less dominance-based (Tomasello,

2016, pp. 42–43). This likely resulted from three intertwined selection processes mediated by partner choice: mating via pair bonding which had effects on human emotions and expanded the kin group, loosely collaborative scavenging and hunting which selected against bullies and food hogs, and collaborative childcare (alloparenting) which likely arose in connection with collaborative subsistence strategies as a kind of division of labor and which generated further interdependence relations. The result of these was a shift in human psychology and behavior through selection for more sympathetic and altruistic individuals. Our ancestors became more interdependent upon one another thus broadening sympathetic concern from immediate kin to larger, more loosely related groups. Just as in our great ape cousins, this broadened sympathetic concern would have led to more altruistic behavior.

Around the same time, the earth underwent a period of cooling and drying. The result was an expansion of open, un-forested environments. As early humans were first coming out of the forests, terrestrial monkeys, e.g., baboons, were also emerging into new environments where they likely outcompeted early *Homo* for resources. As fruit and other nutrient rich resources became scarce, our early ancestors likely turned to scavenging carcasses of large animals killed by other predators. This required the formation of coalitions to scare off those predators and ward off other scavengers while the meat was collected or consumed. By the time of *Homo heidelbergensis* (our common ancestor with Neanderthals), Tomasello claims, strong evidence suggests that our ancestors had built on these early coalitions and begun collaboratively and systematically hunting game.

This new collaborative endeavor involved a variation on the classic stag hunt

payoff matrix. In game theory, a stag hunt game has a large payoff when both parties collaborate. When only one party attempts collaboration and the other defects, the payoff for the defector is small but not zero while the payoff for the collaborator is zero. If neither party collaborates the payoff for both is small but not zero. This gives rise to a coordination problem. Stag hunt has two Nash equilibria—two points in the game at which neither player can profitably deviate given the strategy of the other—but without the ability to coordinate their strategies, the sub-optimal equilibrium is ensured by the risk of going hungry should one player hunt stag while the other defects. However, because of the increased competition for resources posed by terrestrial monkeys and other omnivores, the small payoffs available at that suboptimal equilibrium became insufficient for survival. Thus, unlike the standard stag hunt, our ancestors faced a situation in which failure to collaborate would lead not to small but acceptable payoffs but to unacceptably small ones. These unsatisfactory fallback positions made collaborative foraging *obligate*. Early humans “were interdependent with one another in much more urgent and pervasive ways than were other apes: they had to collaborate with others on a daily basis or else starve” (Tomasello, 2016, p. 44).

One important result of *Homo's* newfound interdependence was the further broadening of sympathetic concern beyond offspring and kin. In the context of obligate collaborative foraging, strong interdependence makes it so that there are direct benefits to helping both partners and *potential future partners*. Since, in this context, my partner is dependent on me for survival, she has no incentive to defect when I have provided aid and every incentive to provide aid in return. It may prove a hindrance to my immediate goals to stop and help you now. However, since my long-term success depends on your

well-being and your success in accomplishing your short-term goals, and since my long-term, overarching goals depend on (actual or possible future) mutual collaboration with you, it is in my best interest to provide aid in the present. I need you to be alive and healthy in order to collaborate in the future, which is what I need for my own survival (Tomasello, 2016, p. 46).⁴

Even if this complex bit of prudential reasoning explains why our early ancestors engaged in altruistic behavior, it is unlikely that it is what motivated them to do so. Rather, the proximate psychological mechanism driving them to look out for one another even at their own expense was likely an expansion of the sympathetic concern that they already felt for their immediate kin. This mechanism is much quicker, emotionally charged, and more reliable, and, as a result of our earlier self-domestication, it was a preadaptation already in place to be selected for. All that was necessary was to broaden its scope to all potential collaborators. As Tomasello puts it, “Because of interdependence, this sympathy...for others presumably contributes to the helper's reproductive fitness on the evolutionary level, but...the evolved proximate mechanism contains nothing about interdependence and reproductive fitness,” rather, it is experienced purely as concern for the other (Tomasello, 2016, p. 49).

⁴ Though his is a game-theoretical rather than evolutionary account of morality, Gauthier’s reasoning nicely explicates the structure of what Tomasello claims occurs at the level of selection: “[E]ach person can see the benefit, to herself, of participating with her fellows in practices requiring each to refrain from the direct endeavor to maximize her own utility, when such mutual restraint is mutually advantageous. No one, of course, can have reason to accept any unilateral constraint on her maximizing behavior; each benefits from, and only from, the constraint accepted by her fellows. But if one benefits more from a constraint on others than one loses by being constrained oneself, one may have reason to accept a practice requiring everyone, including oneself, to exhibit such a constraint. We may represent such a practice as capable of gaining unanimous agreement among rational persons who were choosing the terms on which they would interact with each other. And this agreement is the basis of morality”(Gauthier, 1998, p. 23).

4. The Cognitive Structure of Joint Intentional Action

Though the evolution of altruism and broadened sympathetic concern are important components of the development of human moral practices, they are not the centerpiece of the story I want to tell. Rather, I want to focus on the cognitive dimension of the changes spurred by obligate collaborative foraging. When early humans began engaging in cooperative endeavors, we can imagine that the mechanisms of coordination were somewhat haphazard and the results rather mixed. Some chimpanzees hunt for monkeys in groups. In most cases, a single chimpanzee would not be capable of catching a monkey; they must “work together” in some fashion. In locales with minimal forest canopy, this coordination takes the form of an every-chimp-for-itself kind of chase, but in areas where their monkey prey can move swiftly through the canopy, the chimpanzee hunters need to surround the monkey. This requires a kind of coordination. What manifests, however, is not a well-planned hunt. Rather, each hunter simply takes into account what the others will do while attempting to capture the monkey for himself. Rather than cooperating, each chimpanzee treats the others as moving monkey obstacles in their environment and triangulates its own behavior given its expectations about what the others will do. This is simply competition under another guise. The chimpanzee that captures the monkey gets the most food and only shares with those above him in the hierarchy or others who beg or forcibly demand meat. Even then, if he can, he will keep the spoils for himself, for chimpanzees lack any social mechanism for equitably sharing the spoils (Tomasello, 2016, p. 11). It is likely that early humans started with similar tactics, but because hunting became obligate and our ancestors strongly interdependent, they became more likely to share the spoils with one another to increase chances of future

success. Given that each must rely on the others to capture their prey and that the captured prey will be shared among the hunters, more reliable and efficient mechanisms of coordination could be leveraged.

Three key pieces of cognitive machinery evolved in response to this coordination problem faced by early humans: cognitive processes of *joint intention* that structured coordinative activities, a sense of *second-personal agency* that fostered concern about how others viewed one as a potential collaborator, and self-regulatory processes of *joint commitment* that provided additional motivation to follow through on collaboration. We'll look at each of these in turn.

Our ancestors could trust one another in ways that other great apes could not because they were so deeply interdependent. They could act toward a goal with the trust that their partner would act in ways to help rather than hinder. In order to act together toward *the same goal*, however, our ancestors also needed to evolve a capacity for *joint attention*. Great apes lack this capacity, but human children as young as nine months can engage in episodes of joint attention in which they do not merely attend to the same object as their adult counterpart but attend to it together with the knowledge that they are attending together. In joint attention it becomes part of the personal common ground of each individual that they know together what has been experienced (Tomasello, 2016, p. 52).⁵ Joint attention, according to Tomasello, underlies collaborative activities but also “intentional communication and language acquisition...[J]oint intentional activities were the birthplace of humans’ unique forms of cooperative communication, beginning with the natural gestures of pointing and pantomiming” (Tomasello, 2016, pp. 52–53). With

⁵ Cf. Lewis’s account of “common knowledge,” (Lewis, 1969, pp. 52–57).

the ability to jointly attend to phenomena and to communicate about how to jointly respond and the trust that compatriots will act as agreed, we have the foundations of *joint intentionality*, the human ability to join together and coordinate our behavior toward a common goal.

Joint intentionality has a dual-level structure. When agents engage together toward a common goal, each of them is both an individual agent—an “I”—in mutual collaboration with another individual—a “You”—and a component of a joint agent, “We.” The joint agent, “We,” of joint intentional action appears to be unique to humans among our great ape cousins as distinguished by its particular cognitive structure. Other great apes are capable of acting together toward a common end as the chimpanzees hunting a monkey, but when a joint agent is constituted each individual “intend[s] that ‘we’ act together jointly toward a single end, and they both know together in common ground (they both know that they both know) that this is what they both intend” (Tomasello, 2016, p. 51). Rather than merely triangulating their behavior off of one another they coordinate it with one another via convention and signaling and come to understand themselves and one another as engaged in a shared project.

Joint intentionality requires that individuals take up a perspective removed from their direct, embodied encounter with the world in the form of the perspective of the “we.” Individuals who take up this distance on their own activities while engaging in a cooperative endeavor understand that the success of their endeavor is dependent not on each individual doing what is required for her own success but on each of them fulfilling the demands of a role defined by the norms that would enable their joint success.

In joint intentional action, each individual has a particular role to play, which is

defined in terms of the ideal way of achieving their common aim. If two individuals are hunting a stag, the role of one may be to give chase while the other lies in wait with a spear. With practice and over a period of time, our ancestors would have developed a common-ground understanding of the ideal way of inhabiting each role. These role ideals generate normative standards apart from but as an extension of the individual instrumentality of each partner. They are an extension of individual instrumentality insofar as they are "fundamentally instrumental and local." They are not imposed by some external authority but, rather, contingent upon the interests and goals of the collaborators in their immediate context. Yet, they are distinct from individual instrumentality in that they are binding on individual agents only in the context of the goals of the joint agent "We." Neither individual ought to act in accord with these norms except insofar as she is engaged in joint intentional activity (Tomasello, 2016, p. 54).

Collaborative role ideals are not indexed to an individual. Either individual could, in principle, play either role in the collaborative relationship. For each individual engaged in joint intentional action, the roles are understood from a kind of "birds eye view" in which she perceives her own activities as part of the whole of the activity of the joint agent "we." "She does not peer from inside her role and perspective onto the outside of the partner and what he is doing. Rather, as she is collaborating, the individual imagines being in the partner's role and perspective, on the one hand, and also imagines how the partner is imagining her role and perspective, on the other" (Tomasello, 2016, p. 55). As such, as part of the joint agent, she understands the demands of both roles as following from the joint agent's aims independently of any characteristics of either individual agent. There may be contingent reasons for one partner giving chase rather than the other—one

may have better aim with the spear or be injured in a way that slows her down—but from the point of view of the cognitive structure of joint intentionality, dyadic cooperative roles are entirely interchangeable.

For this collaborative mechanism to reliably solve the coordination problems faced by our early ancestors, they needed to more thoroughly adapt to their new social environment by developing a sense of *second-personal agency*. Our ancestors had to learn, in particular, to choose good collaborators as partners. This involves not just the ability to judge others on the basis of their past actions and evaluations, i.e., their reputations, but also the ability to anticipate how others will evaluate you so that you can act in ways that will lead to your being chosen as a collaborative partner. In general, we needed to learn to judge and manage reputations—or cooperative identities—as reliable, trustworthy partners (Tomasello, 2016, pp. 57–63). This is an instance of a common evolutionary pattern. Adaptation is dynamic. Organisms don't simply adapt to a pre-existing ecological environment. Instead, they are “codeterminative” with their selective environments. As organisms adapt, their adaptations give rise to new “ecological design problems” that future adaptations will need to solve (Buchanan & Powell, 2018, p. 124). Adapting to obligate collaboration by evolving capacities for joint intention created a new design problem: it was now beneficial to be seen by others as a fit collaborator. Second-personal agency evolved as a tool for ensuring that this is how one would be seen. It also became the root of the highly developed evolutionary conscience that later human beings would come to possess.

The final adaptation our ancestors had to undergo was to come to see themselves as obligated to act in accord with collaborative role ideals. This was imposed by the

inherent riskiness of relying on strategic trust to govern cooperative endeavors. I trust you to go on the difficult hunt because I think I know what motivates you. I believe that you're hungry and that you believe that hunting with me is the only way you'll get to fill your belly; thus, I trust that you will be a reliable partner in the hunt. But perhaps you have reason to defect. Maybe, unbeknownst to me, you've recently had a snack, so you're much less committed to the hunt than I am. Perhaps, after we begin, you start to think that I'm less committed on the basis of my lackadaisical hunting behavior. Or maybe you just come to see the hunt as too risky. If we are relying only on strategic trust to ensure our cooperation, I'm going to be left hungry. To solve this problem, our ancestors needed to become subject to a primordial "ought," i.e., to feel themselves obligated to fulfill their collaborative role ideals even when they judged it strategically best to defect. Only in this way could collaboration become a way of life for them (Tomasello, 2016, p. 64).

The solution they hit upon, according to Tomasello, is the creation of *joint commitments*. Early humans committed to one another and together to pursue some goal, and, in doing so, they constituted the new joint agent "we" to which both took themselves to be beholden. In making this joint commitment, "I freely grant authority—legitimate authority—over 'me' to the supra-individual entity that is 'we,' and indeed, I will defer to that 'we' to the point that if you rebuke me for non-ideal behavior, I will join you in this rebuke (either overtly or in a personal feeling of guilt), judging that it is indeed deserved" (Tomasello, 2016, p. 64). I grant such authority by identifying with the "we" over the "I" because to do otherwise would risk my chances at future collaboration by revealing myself as an unreliable partner. I have also come to view myself through the lens of second-personal agency. As such, it is central to my cooperative, second-personal identity

to maintain good standing as a collaborator (Korsgaard & O'Neill, 1996). Doing so requires that I keep my commitments, and I will feel guilty if I do not (Tomasello, 2016, p. 64). Thus, these three developments are mutually reinforcing.

Our early human ancestors thus evolved cognitive capacities to navigate the dual-level structure required to manage dyadic cooperation. Faced with ecological changes that imposed conditions of obligate collaboration for survival, they developed cognitive mechanisms of joint intentionality, which required a new kind of engagement, joint agency, in which two individuals intend together, as a single joint agent “we,” to act toward some goal. This new cooperative environment selected for individuals who could manage their cooperative identities, giving rise to a sense of second-personal agency. Furthermore, collaboration could only be stable if individuals could be counted on to abide by collaborative role ideals. This selected for agents who had a sense of being beholden both to other collaborators and the joint agent, “we,” i.e., for individuals who could engage in joint commitments. There were likely a few selective mechanisms responsible for these adaptations. One was a form of partner selection. Cooperative partners—those who cultivated a reputation of being reliable, dependable, and willing to share fairly the spoils of the hunt—were more likely to survive and reproduce while non-cooperative individuals were less likely to establish foraging and defense relationships and, so, would have been less likely to survive for long. They also would have appeared as less appealing to potential mates, so sexual selection also played a role. Moreover, non-cooperators would have been less able to provide for their young making their progeny less likely to survive. Through such mechanisms, relatively stable, loosely organized groups of early collaborative hunters emerged.

These cognitive changes were attended by modifications in individual psychology that were precursors to the moral psychology of modern humans. Along with broadening spheres of sympathetic concern, the structure of joint intentionality holds the seeds to a sense of fairness grounded in self-other equivalence. Given the structure of joint intentional action, early humans came to grasp that, in the context of collaboration, they were essentially interchangeable. Each individual was subject to the same normative demands of joint agency, and fulfilling those demands was necessary for collaborative success, but any individual could, in principle, fill any role (Tomasello, 2016, p. 56). Moreover, given that they understood that they were mutually interdependent on their actual and potential foraging partners each of which could fulfill the normative demands of collaboration, early humans were motivated to treat each other fairly. To do otherwise would threaten future collaboration either by alienating oneself from the group or by harming one's future partners, and this is to risk one's own livelihood (Tomasello, 2016, p. 60).

5. From Cooperative Dyads to Cooperative Communities

The dyadic cooperation of our early hominid ancestors seems to have remained a relatively stable strategy for quite some time. Groups of early hominids managed to survive and spread. However, by about 150,000-100,000 years ago, according to Tomasello, this way of life began to fracture under the weight of its own success. Around the time of the emergence of behaviorally modern human beings in the late Pleistocene (Boehm, 2012b, Chapter 4), the population size and geographic range of groups had grown large enough that the groups began regularly interacting and competing for the same resources. With groups in competition over resources, there was now a need for

greater cooperation within groups in order to compete with and defend from outsiders. The result was a scaling up of dyadic cooperative machinery to group level cooperation (Tomasello, 2016, p. 88; Tomasello et al., 2012, p. 681).

The group structure emerging in this period seems to have been rather complex. As groups expanded with a growing population, they began both to interact with outside groups more frequently and to fracture internally. There was also likely regular movement by individuals between groups. Eventually, larger groups splintered off into smaller “bands” that were still bound together as “tribes” by a shared history and culture. Such bands would join together for ritual purposes and to engage in intergroup conflict and joint defense (Tomasello, 2016, p. 88). Within this tribal organization, individuals faced two new challenges. First, they had to be able to recognize members of their cultural group—often strangers to them, since they may have intimate knowledge of the 20-50 or so members of their band but not of all of the members of the much bigger tribe—and to be recognized in return. This made conformity with cultural customs a necessity and helps to explain why human children are much more motivated to conform and enforce conformity than are other primates (Tomasello, 2016, p. 89). The second challenge they faced was that of coordinating to help and protect—and be helped and protected by—all of their actual and potential collaborators. This was, in essence, the entirety of the cultural group. As human cultures became more complex and as task-specific knowledge became more specialized with increasing division of labor, our modern human ancestors became even more interdependent for their survival. Moreover, in times of conflict, the whole of one’s tribe might be called on for protection, as such, it was in one’s best interest to ensure their wellbeing (Tomasello, 2016, p. 88).

The response to the first challenge was the expansion of ways of marking one's membership in the group both physically and behaviorally. Modern humans developed modes of dress and adornment to signal group membership as well as shared language and customs that are passed down to each generation. They also evolved strong propensities for loyalty to their cultural groups and for in-group favoritism, which was likely heightened in times of conflict (Buchanan & Powell, 2018; Tomasello, 2016, pp. 88–92).

In response to the second challenge, modern humans scaled up the cognitive machinery that evolved to manage dyadic cooperation. Joint intentionality evolved into cognitive processes of *collective intentionality* that placed the cooperative individual in the context of a group of collaborators with all of whom she was now deeply interdependent. Membership in a cultural group required further development of senses of agency and identity as modes for managing relationships between individuals who were not nearly as intimately known to one another as earlier dyadic cooperators.

In the context of larger cultural groups, social organization became more complex. Individuals were no longer engaging in one-off cooperative endeavors like going on a hunt but in myriad cooperative groups of a variety of sizes, working toward a variety of ends sometimes nested within one another, sometimes overlapping, and sometimes in conflict. One might have responsibilities as a member of a coalition of hunters, defensive tasks when the community was threatened by outsiders, and be occasionally engaged with coalitions targeting would-be internal aggressors. As in dyadic cooperative hunting, each of these tasks involves individuals inhabiting various roles, but, in this case, the collaborating groups are larger, and the roles are defined not simply by

the common meat-seeking interests of the collaborators but by something that is yet another step removed from them.

Tomasello argues that cognitive processes of collective intentionality evolved to manage this more complex form of collaboration. The structure of these processes is akin to the dual-level structure of joint intentionality except that the joint agent "we", in this context, becomes the cultural "we," as the group as a whole was now engaged in a collective collaborative endeavor to out-compete other cultural groups for scarce resources, and the "you" becomes plural, as one sees oneself in collaboration with multiple others instead of just a partner in a dyad.

Within this structure, cooperative endeavors need not start from scratch with each new project. They are now tokens of the various types of cooperation in which individuals could engage as members of the cultural group. Rather than generating new cooperative roles based on the joint intentions of the cooperators, endeavors begin with *conventional roles* that compose these various cooperative endeavor types already in place. Given that we are all Waziri—to borrow Tomasello's favorite fictional tribe—we all know and know that each other knows how to net fish and what is required of the net weavers, the chasers, and the netters in order to succeed. These roles are defined by the shared skills, knowledge, beliefs, and practices that are collectively shaped and passed down to each new generation as part of the cultural common ground (Tomasello, 2016, p. 93).⁶ They are, in this sense, "conventionalizations (standardizations) of the small-scale collaborative (and other) activities" of our earlier humans (Tomasello et al., 2012, p. 683). As part of the cultural common ground—that is, as part of what each of us knows

⁶ Cf. (Sterelny, 2014) on the unique coapcity of humans to teach and learn this trove of cultural knowledge.

that we all know—these conventional roles allow individuals to navigate into and out of various collaborative relationships over the course of a day, a week, and a life. Since “everyone knew that everyone knew how to perform in the ideal way,” even in engaging in new collaborative contexts with collaborators previously unknown to you, you could be fairly certain that they would know what was required of them for your collective success (Tomasello, 2016, p. 123). Thus, even across a large cultural group, one could be sure that those who were marked as members would be able to fulfill the necessary collaborative roles.

The kinds of cooperative situations in which behaviorally modern humans might have found themselves are vastly more varied, complex, and nuanced than the dyadic cooperation of early humans. Not all cooperative situations are aimed at attaining an immediately available goal. Sometimes acting cooperatively in the present is merely one not clearly connected step toward some long-term payoff, as may be, for example, getting the projector working for a colleague’s presentation when no one really sees the value of the meeting. In other instances, cooperation is just doing one’s part to maintain smooth functioning of the broader social organism. This might mean something as simple as walking on the right side of the path, cleaning up one’s own mess after cooking, or offering the better sleeping spot or something more complex like taking turns using a valuable resource or begging off a proposition from a potential mate because it would upset a compatriot.

Given that the unit of cooperation for modern humans is the cultural group as a whole, all of these are potentially examples of cooperative situations. As such, they are all shot through with “ought,” yet they don’t seem to be captured by the kinds of role

ideals relevant to the context of dyadic cooperation. Knowing when to turn down a proposition is a very different kind of thing to grasp than is knowing how to be good hunting partner. Enter social norms. Social norms are “a set of expectations that everyone in the group shared in cultural common ground about how individuals must behave in various situations to be cooperative...[and to ensure] the smooth functioning of social interactions in the group” (Tomasello, 2016, p. 98). Such expectations set the ground rules not just for engaging in all variety goal-directed cooperative endeavors but also for appropriate behavior as *one of us*, a member of the cultural group, a person. We Waziri turn down propositions from potential mates that we know to be sought after by our kin, and anyone who doesn’t know this norm or doesn’t abide by it is either deserving of sanction or must not be Waziri.

Since social norms are part of the cultural common ground and all members of the tribe are part of the cultural “we,” the way in which social norms are enforced is distinctive. Whereas joint agency involved being prepared to accept sanction from one’s collaborative partner for failure to fulfill collaborative role ideals, cultural agency involves potentially being subject to such sanction from any member of one’s cultural group. Being a cultural agent and managing one’s social identity involves subjecting oneself to the authority of the cultural “we.” With the advent of collective intentionality, then, we also have the development of group level social control: everyone, in principle, can hold everyone else to the oughts that bind them. All are responsible for enforcing the norms that ensure the wellbeing of the group.

This new normative environment, moreover, gave rise to the need for each individual to carefully manage her social identity. In the context of dyadic cooperation,

managing a cooperative identity was a matter of managing how others with whom you regularly interacted viewed you as a potential collaborator, but since modern humans were collaborating with the entire tribe, one's reputation as a cooperator became much more important. One's cultural identity is a matter not just of how one presents oneself to those with whom one is closest but of the gossip that spreads about one's past deeds and misdeeds. Since every member of the group knows the shared social norms and knows that you, as a member of the group, know those norms, any deviant behavior, if not directly sanctioned, is likely to be broadcast to other members damaging one's reputation and one's standing as *one of us*. Even if one of our modern human ancestors did not consider their current collaborators fit for future collaboration, they needed to be on their best behavior, as any misconduct could threaten their future prospects with others and, perhaps, even their membership in the tribe.

Of course, not all violations of social norms are likely to elicit the same severity of response from one's compatriots. If I absentmindedly walk on the wrong side of the path, it's unlikely I receive a harsh rebuke. If, however, we are on a net hunt, and I surreptitiously move my net ahead of the other hunters in an effort to capture more meat for myself, I am likely to be put down rather harshly.⁷ The difference is that in the latter instance my social norm violation also violates the sense of fairness grounded in self-other equivalence and threatens to do real harm to members of my band. As such, it is read not merely as doing something that we Waziri don't do but as a failure to show due respect and concern for others. The social norm against cheating on the hunt is, in this

⁷ Boehm relates the story of Cephu, a mature Mbuti Pygmy hunter, as reported by anthropologist Colin Turnbull. Cephu committed just this crime on a net hunting expedition and, having been caught, was publicly shamed and temporarily shunned by his band. Hunter-gatherer bands will sometimes even turn to capital punishment for such violations of moralized social norms (Boehm, 2012b, pp. 37–42).

sense, *moralized* in a way the other is not. When social norms have to do with matters of fairness and respect, or come to be understood as such, their enforcement becomes more stringent and their violations come to be understood as moral failures.

6. The Free-Rider Objection

We need now to consider an important objection to the evolutionary story we've been telling up to this point. Altruists and profligate cooperators are easy targets for free-riders. If you're generally willing to help, then I can maximize my returns by appearing as if I am too but, in the end, keeping the spoils for myself. Indeed, mathematical modeling has demonstrated that well-designed free-riders can take advantage of and undermine even a population in which a majority are altruists. If this is right, then even if modern humans devised methods for enhanced coordination and cooperation, they would have been beset by free-riders who continually took advantage of them and undermined that cooperation. Tomesello seems to try to skirt by this important problem by arguing that reciprocal altruism became a fitness conducive strategy once cooperation became obligate. If I recognize that I will continue to need your help, then I'm less likely to cheat you out of your share of the spoils. By the time of modern humans, selection by reputation could also play a role in weeding out free-riders. If you're a cheat, word will get around, no one will be willing to help you out, and this will be detrimental to the survival of your genetic line.

There are two weaknesses to this approach to the free-rider problem. First, it seems that on this account, we should expect free-riders to have disappeared by now. If free-rider genes have faced significant selective pressure for thousands of generations, they should have been eradicated. But, of course, this isn't what we see. Free-riders

persist even in our contemporary world. Indeed, the worry that one is being cheated is ever-present. We warn our children not to be too trusting, and we have the sneaking suspicion that, even though we are taught to be good, nice folks finish last. So, why are there still so many free-riders?

The second weakness in Tomasello's approach is its focus on a particular type of free-rider. Following a nearly universal trend among evolutionary theorists, he concerns himself only with would-be cheaters and thieves. These classic free-riders play along with a cooperative endeavor just to the point where they can insure a win for themselves either by cheating so that they end up with more of the spoils, stealing them outright, or deceiving cooperators into giving them up. Selfish free-riders, though, need not be deceptive cheats and thieves. Any individual who takes more than he gives is free-riding on those who do the lion's share of the work, and *selfish intimidators* are masters at this. Alpha-type bullies do not engage in any chicanery, rather, they simply use force (or threats of force) to take what they want from altruistic cooperators who can't fend them off. In some cases, of course, these bullies might play an important role in social coordination, suppressing violent tendencies of others and, in general, keeping order, but they can also be massively exploitative. Any species with a hierarchal tendency can fall victim to this kind of free-riding, as sitting atop a dominance hierarchy in which selfish aggression can be freely expressed can pay dividends in terms of fitness (Boehm, 2012b, pp. 64–66). When chimpanzees go monkey hunting, the alpha male will hang back and watch, and, when the hunt concludes, he will step in and claim the spoils from whoever made the kill. Bullies, though, cannot be dealt with in the same way as cheaters, so we'd better have something to say about how human-level social coordination deals with these

free-riders.

Let's take these two problems in order.

Given the existence of contemporary free-riders, the evolutionary processes at work in selecting against free-riding must have suffered from a bit of looseness.⁸ A nearly universal feature of accounts of the evolution of morality is that they posit that it was adaptive because of the benefits of cooperation that accrued to bands of altruistic individuals, but these benefits will be fragile if free-riders are not controlled in some way. How, then, are we to explain that free-riders continue to exist? What gaps in the processes of natural selection did they manage to slip through? The answer is to be found in the evolutionary conscience, the “still small voice that tells us how far we can go in serving our own interests without incurring intolerable risks” (Alexander, 1982, p. 102; quoted in Boehm, 2012b, p. 30).

Early humans came to have capacities for self-regulation that differed in kind and degree from chimpanzees and bonobos. They learned to control their behaviors not only out of fear of reprimand or rebuke—as chimpanzees clearly do in maintenance of their strict dominance hierarchies—but also out of recognition that they needed to appear cooperative to others in order to partake in future cooperation. This required that they develop methods of self-regulation that were responsive to a broader array of social and emotional cues. Importantly, the individuals who were successful at this did not need to

⁸ Looseness is usually characterized as a kind of inefficiency in natural selection that accounts for some behaviors as a side-effect of the substantial benefits of some other behavior. Boehm, for example, cites George Williams' example of looseness in the realm of reproductive behavior: “reproductive functions, perhaps to a greater extent than any other adaptations, are characterized by a considerable degree of looseness in timing and execution,” which may be advantageous in pairing up with mates but also accounts for widespread homosexual behavior among many animals (Boehm, 2012b, p. 56 quoting; G. C. Williams, 1966, p. 205). Looseness, however, might also account for unexpected traits that are not mere spandrels of highly advantageous ones but, rather, that seem to have sneaked through under the radar. That is what we find in the case of free-riding.

be free of anti-altruistic impulses like the inclination to cheat, dominate, or be lazy. They only needed to be able to suppress those impulses with enough regularity to *appear* altruistic. If they could conform, they could survive and pass on their genes, even those responsible for their anti-social tendencies.

This protoconscience was a preadaptation that was then leveraged as behaviorally modern human beings emerged with growing group size and interdependence. The evolutionary conscience then became a powerful tool for managing one's social identity. Those who could self-regulate in order to conform to social norms had a significant advantage, for they could maintain reputations as reliable cooperative partners, allowing them to accrue nutritional and reproductive advantages. They could also be shaped by their culture in an important way by being subject to the “preaching” of golden rules, those universal exhortations to good behavior, generosity, and fair-play. Golden rules are important in any contingent system of indirect reciprocity, where the good of the group is maintained by individuals giving generously and taking only what they need. They are a form of “prosocial ‘propaganda’,” and they work remarkably well in shaping behavior among human beings because our evolutionary conscience functions to internalize rules and motivate conformity to them (Boehm, 2012b, pp. 52–53). We are eminently indoctrinable, and our altruistic ancestors used this to their advantage. As such, we find the promulgation of golden rules among every extant human culture, including contemporary hunter-gatherers (Boehm, 2012b, p. 52).

The answer to the first challenge, then, is that free-rider suppression often functioned at the level of phenotype rather than genotype. If carriers of genes linked to a propensity to free-ride could manage to suppress that behavior by way of an evolutionary

conscience, then their genes wouldn't suffer the consequences of the various punishments that those of visible free-riders suffered. Social selection can only work on actually expressed traits, but if individuals got good enough at managing their reputations then, for all intents and purposes, they counted as altruists. This is why free-riders survived.

The second weakness, recall, consists in ignoring a particularly successful kind of free-rider: the bully. Bullies are important free-riders precisely because they have been so successful. In creatures with hierarchical tendencies such as humans, chimpanzees, and, with all likelihood, Ancestral *Pan*, being an alpha-male has been an extraordinarily fruitful strategy for securing the best nutrition and mates. So long as an individual can use force or its threat to keep subordinates in line, he can take from them what he wishes. Bullying, though, is also a kind of free-riding quite susceptible to suppression, for, unlike sneakier free-riders, bullies can't succeed without being recognized as bullies. Efforts at alpha-type suppression can be seen among our close evolutionary relatives. Subordinate chimpanzees sometimes form counter-dominant coalitions to attack, drive off, and even kill dominant males in order to secure access to females for themselves. Among bonobos, females sometimes join together to control the behavior of particularly aggressive males (Boehm, 2012b, pp. 95–96). These efforts at checking dominance are fragile. Chimpanzees and bonobos lack the cognitive processes for social coordination that early humans developed, and, perhaps more importantly, such counter-dominant enforcement is costly. Enforcing altruism puts an individual in the unenviable position of being easily singled out for retribution by the dominant male or others who are loyal to him, and it is far from clear why unaffected third-parties would ever be motivated to take this risk. This higher-order problem of altruism existed also for our human ancestors, so what is to

account for their willingness to more aggressively police anti-social behavior and enforce social norms?

Part of the answer is that as modern humans were forced to begin hunting larger undulates, meat-sharing became essential to their lifeways. These ancestors recognized the value of having able hunters in their groups, but in order to capitalize on the advantage this provides, those hunters need to be well-fed enough to remain energized and healthy for the hunt. Other dedicated social carnivores—wolves, for example—rely on evolved hierarchies to achieve this. Dominant males get the most meat, but they tolerate younger subordinates, who are essential to the success of the hunt, taking some for themselves. Modern humans, though, found ways to nearly equalize meat intake by leveraging their evolved modes of social control to construct egalitarian systems for the distribution of meat (Boehm, 2012b, pp. 143–144). These often involve dedicated butchers who distribute meat to the band and elaborate rituals of self-effacement by successful hunters (Boehm, 2012b, p. 43). Alpha-type bullies, however, threaten this equal distribution, so it becomes ever more urgent to keep them in check. Given the human propensities for social learning, rule internalization, and an evolutionary conscience, if counter-dominant coalitions could make an example of would-be aggressors, over time, others would be less and less likely to exhibit similar behavior. Indeed, they did and do make an example of them in just this way. Among contemporary hunter-gatherers, practices of public shaming, shunning, banishment, and even capital punishment are regularly used to enforce egalitarian norms (Boehm, 2012b, pp. 195–199 & chap. 9).

Another part of the answer to why our modern human ancestors began more

aggressively policing anti-social behavior lies in the theory of cultural group selection.⁹ Though it is contentious, there is a good deal of circumstantial evidence that suggests that there was substantial intergroup competition and conflict—even warfare—by the late Pleistocene (Buchanan & Powell, 2018, pp. 126–127).¹⁰ Bands of hunter-gatherers that evolved institutionalized modes of punishing norm violations—especially anti-egalitarian behavior—had a significant survival advantage in this context. Groups that had a greater number of altruists and third-party norm enforcers developed more cooperative social structures and coordinative practices. This allowed them to outperform and, eventually, replace groups engaged in less effective cooperation. The initial benefits of cooperation were economic. Increases in foraging yields, greater success in large game hunting, and better health and survival for members of the band led to increasing group sizes. Increased group size and better modes of coordination gave these cooperators an “advantage in raiding, border skirmishes, and full-scale military conflicts, with victorious groups populating the territories and commandeering the resources of vanquished groups” (Buchanan & Powell, 2018, p. 126). Moreover, the same psychological mechanisms that account for the human ability to internalize and abide by norms also account for our cultural learning abilities. Moralized norms stabilized institutions for the

⁹ Group selection fell into disfavor in the 1960s in the wake of Williams’ withering critiques (G. C. Williams, 1966), but it has experienced a renaissance over the past two decades as a component of theories of multilevel selection. A growing number of theorists argue that robust cooperation and coordinative processes are only likely to evolve via selection at the group level because of the significant costs of altruism and enforcement in terms of individual fitness (Buchanan & Powell, 2018, p. 123n9; Godfrey-Smith, 2011; Sober, Wilson, & Wilson, 1999; Wilson & Wilson, 2007).

¹⁰ The record of conflict is well established in the very late Pleistocene and early Holocene, but it is more contentious during the time that human modes of social control likely evolved. Kim Sterelny has argued that intergroup conflict was less likely to have occurred between bands of persistence predators like *Homo* and, so, is not likely to account for our evolved moral thinking and behavior (Sterelny, 2014). Buchanan and Powell, in response, point out that violent intergroup conflict does occur among other persistence hunters—wolves, for example—so this specialized predation does not rule out such conflict (Buchanan & Powell, 2018, p. 128).

transmission of various technologies and information about local ecologies to new generations, which also represents a significant advantage for groups that could find ways to enhance cooperation (Sterelny, 2014).

The upshot of these various processes was a *militant egalitarianism* among modern humans. Alpha-type bullies were a visible and powerful threat to the kind of meat-sharing practices that were required for harmony within the group and competition outside of it. As such, they were likely the initial targets of social control. They would have faced ridicule, shaming, shunning, banishment, and, in some cases, capital punishment, and only those who had enhanced capacities to suppress dominant proclivities could successfully reproduce. With time, egalitarian social norms could be inculcated into new generations, and those bands that succeeded in becoming more cooperative out-competed others. With social norms commonplace and a stronger evolutionary conscience taking hold, suppression of dominant alpha-types could spread to other kinds of free-riders further enhancing cooperation. By about 45,000 years ago, by which time cultural modernity had been fully phased in, modern humans had come to resent all free-riders and had developed both direct—physical and social sanction—and indirect—gossip and reputation assassination—institutionalized modes for dealing with these deviants. Modern humans had, by this point, made the transition from the hierarchically structured group life of Ancestral *Pan* to the militantly egalitarian lifeways we find in all contemporary hunter-gatherers (Boehm, 2012b, pp. 149–161).¹¹

¹¹ This militant egalitarianism was, of course, contingent upon the environmental and social conditions in which it arose. As such, it was incredibly fragile. Since free-riders (especially bullies) continued to reproduce, the tendency toward hierarchy and domination never vanished among humans, and it powerfully re-asserted itself once conditions were favorable. With the evolution of agriculture there came a need to protect, control, and distribute surpluses. This set the stage for alpha-type free-riders to re-establish hierarchical lifeways, and we, as such, are left with both legacies. For a story about how these changes might have come to pass, see (Gellner, 1990).

7. The Function of Moral Discourse

Bands of militantly egalitarian modern humans with highly institutionalized modes of social control, fully developed evolutionary consciences, and robust systems of social norms would also have had the capacity to talk about their normative circumstances. It is quite contentious exactly when the *Homo* line developed a capacity for symbolic language. *Homo heidelbergensis* probably used non-compositional, non-descriptive vocalizations to manipulate one another's behavior, and *Homo neanderthalensis* might have had the capacity to make many of the sounds modern humans make. It is broadly thought, however, that the use of symbolic language emerged in *Homo sapiens* sometime between 200,000 years ago—when anatomically modern humans begin to appear in the fossil record—and 50,000-60,000 years ago when they began emerging from Africa carrying language and culture with them.¹² Whatever the exact timing of the evolution of symbolic communication, it is unquestionable that modern humans by about 45,000 years ago had both the capacity for symbolic communication and quite robust modes of social control.

Modern humans, presumably, would have shared information about social norms in a variety of forms. We've already seen that they would have gossiped about one another quite extensively, as information about the cooperative reputations of others is important information about one's environment. They also likely expressed social norms in language for pedagogical and enforcement purposes, making clear what is and is not appropriate behavior for those who would be one of us. Preschool children use generic normative language—"It's wrong to do that!" or "No one should do that!"—to enforce

¹² See (Tomasello, 2010) (Pinker, 2010) (Hauser, Chomsky, & Fitch, 2002).

both moral and conventional norms on third parties. They will also try to explain to newcomers the “proper” way to play with a particular toy—even if they’ve just made up the rules—using such generic language (Tomasello, 2016, pp. 102 & 120). Contemporary hunter-gatherers also regularly engage in this kind of verbal pedagogy and enforcement (Boehm, 2012b). Both of these are strong pieces of circumstantial evidence that our ancestors did this, as well.

It is also quite probable that they had disputes about what they ought to do. It is inconceivable that any group of self-interested individuals experiencing the world from their own embodied perspectives and having necessarily limited knowledge of their physical and social environments could have engaged in complex, on-going cooperation without some squabbles. Given the capacity for linguistic communication, they likely had discursive methods for sorting out such disputes, for determining what one ought to do in the relevant circumstances.

The generic normative language used in their pedagogical, enforcement, and deliberative speech acts, however, likely did not consist in moral discourse in the sense we examined in the previous chapter. The norms with which they were concerned were thoroughly parochial. In cases where there was some uncertainty about what to do, the aim was to determine what *we* Waziri do, i.e., to determine out what *our* social norms are. Finding this out may not have been a simple matter, but it was always an empirical matter. They would have needed to examine similar cases, discuss how they react to those cases (given their internalization of the relevant norms), and perhaps consult with certain experts in the community who—like Miss Manners—could draw on long experience to channel the norms of the community. Once it was determined what the norms of the

community were, any dispute was settled even if individuals had other reasons for thinking that the norms of the community did not prescribe the best way to proceed.

Given the militant egalitarianism of these modern humans, though, it is unlikely that any recognized expertise granted those experts ultimate authority over what ought to be done. There may have been elders, but their purview was to advise not to pass judgment. Granting authority to a particular individual or group would have undermined the egalitarianism they worked so diligently to maintain, for such power can be used surreptitiously to gain advantage for oneself and one's kin. Moral authority is a potent kind of social power, and our subjects knew that there were still alpha-type deviants in their midst, even if they were successful at suppressing their propensity for dominant behavior. As such, they would not have granted anyone such power.

The generic normative language of these groups would have been akin to what I have called institutional discourse. Agents could deliberate individually or together about what they ought to do. None of them would have had any ultimate authority to determine an answer to this question, but whatever the answer was in any given case was determined by what the extant community norms were. Though individuals could be in error about what was required of them in any given case, the community as a whole could not be wrong.

This is quite unlike how I have described moral discourse. While I argued that it does exhibit a distributed pattern of authority such that no individual is ultimately authoritative about what anyone ought to do, it also withholds this authority from the group as a whole. It is non-dogmatic. Entire communities, entire tribes, entire populations can (and often are) wrong about what is morally required of them. As such, whatever

function this generic normative discursive practice had for modern humans (and might still have for us today), it is not the function that explains why moral discourse exhibits the distinctive structural features that it does. This discursive practice is not moral discourse.

To get at the function of moral discourse, we need to recognize some instabilities that are inherent to the generic normative discourse of our forebears. The first of these is commonplace enough. We are all, from time to time, unhappy about the demands our culture places on us and uncertain about whether they really support the common good. As reflective and self-interested beings, we can sometimes come to see that we have strong reasons to act in ways contrary to the demands of our community. Moreover, those demands often distribute the burdens of compliance unevenly. What comes to be seen as a violation of particularly stringent—perhaps *moralized*—norms is not something decided by reason but, rather, is subject to all variety of local pressures. Dressing shabbily for a festal celebration may be taken as a minor offense or it may be read as disrespecting the group, a serious violation not just of the community's norms of dress but a threat to the smooth functioning of the community. The reasons it might show up as the latter for a community are myriad. Perhaps there is a history of others challenging the community in just this way, maybe in times of significant external threat showing conformity is of heightened necessity, or maybe some forms of dress have come to offend the religious sensibilities of the community.

Whatever the causes, reflective individuals may come to resent being compelled to conform to such moralized norms if they become convinced that they are not, in fact, necessary for the well-being of the community or that they do, in fact, disproportionately

burden particular individuals. They may, for example, come to believe that the moralized norm against shabby dress unfairly affects those who cannot afford finer attire or who dress a particular way because of the peculiar traditions of their band or family, and, furthermore, that dressing as one pleases does no one any harm. These judgments may be grounded in their sense of fairness (derived, you will recall, from self-other equivalence). As such, they may come to see the moralized norm of the community as, in fact, in violation of their sense of second-personal morality. Conflicts such as this have the potential to destabilize the group causing individuals to decamp or factions to form, but the generic normative discursive practices of our modern humans have no mechanism for addressing such conflicts. They simply enforce conformity with social norms.

A second source of instability derives from intergroup interactions. Besides competition and war, modern humans engaged across cultural groups in a variety of ways. There is ample evidence of a significant degree of trading even across long distances. Lapis lazuli mined in pre-historic Afghanistan has been found in Egypt, Mauritania, the Caucasus, and elsewhere. There is even evidence that some trade was taking place before the advent of symbolic language. Contemporary hunter-gatherer tribes—and their prehistoric counterparts—also form military alliances and engage in exogamy and other forms of cultural exchange. Given the right environmental conditions, some groups of modern humans seem to have been quite amenable to peaceful intergroup relations (Buchanan & Powell, 2018, p. 134). These interactions with other cultures, however, required some way of stabilizing cross-cultural coordinative practices and social norms. Exogamy, war, and trade are not simple, one-off endeavors. They require the development and extension of systems of norms without which they will likely lead to

discord and conflict within the newly formed meta-group. Claims about what we Waziri do, however, will be of little use when we are not all Waziri.

This is the evolutionary crucible within which moral discourse was forged. The complex social practices required for managing intragroup stability, cooperation, and coordination, are unstable under conditions of complex interaction with other groups and when individuals come to believe that conformity harms them at too little benefit to the group. The source of instability is the lack of some means for rationally evaluating, defending, or revising a group's social norms; they simply demand compliance. These practices evolved to manage cooperation by imposing normative constraints on individual behavior, but they lack the means for explicit revision of those norms in the face of instability.¹³ Given the significant benefits of larger group size, increased stability, and intergroup cooperation, any groups that could devise means for correcting for these instabilities would have a fitness advantage. Moreover, these bands of hunter-gatherers already had all of the cognitive tools to solve the problem they faced. They simply needed a new kind of discursive technology that leveraged those tools in novel ways.

To correct for the instabilities we've identified, a practice is needed that allows not just for clarification of how the social norms of a group apply in a given instance but for the rational evaluation, defense, and emendation of those very norms. It must allow individuals who are recognized as part of the relevant cooperative community to challenge the norms of that community, and it must resolve such challenges under two important constraints. Like the generic normative discourse of our modern humans, it

¹³ This should not be taken to imply that these norms do not evolve over time. Any system of norms will evolve with changing environmental conditions, technology, and interests. What is absent is some way of explicitly, rationally revising these norms on ways responsive to different groups or factions. Norms may drift, but this drift is not responsive to reasons.

must limit the possibilities for domination by alpha-type free-riders. Given that large group size is beneficial, it also must reduce the risk of desertion by those who have come to feel wronged by the group's social norms and the risk of defection by groups with which the group aims to cooperate. In short, it must reduce the chance that resolution of the instability will fracture the group.

Moral discourse as we have identified it is just such a piece of discursive technology. It stands as a meta-practice to the generic normative discourses of our modern humans and leverages their already extant cognitive capacities for norm internalization and self-governance, latching onto their motivational psychology in the same way as the social norms enforced by the community. This explains why moral discourse exhibits the distinctive pragmatic structure of the prescriptive with agent-neutral inputs and agent-relative outputs. If it is to be effective at enforcing compliance with the norms of the community—including those that are amended via moral discourse—then it must take on some special normative significance for the agents whose normative statuses moral claims are about. They must have been habituated to respond appropriately to them.

Moral discourse avoids the threat of domination by alpha-type free-riders in the same way that the generic normative discourse of our modern humans does. It distributes authority to all members of the community, vesting no special authority to determine what is the case in any of them. Any member of the community in good standing can issue a moral claim in any of the contexts discussed in the previous chapter, i.e., in either individual or joint deliberation or in an attempt at second- or third-personal holding or

assessment.¹⁴ No such claim, however, is immune from challenge. Any member of the community in good standing can raise one, demanding that reasons be giving in defense of the initial claim. Any such challenge, though, can itself be challenged and must, again, be defended with reasons. That any moral claim is open to challenge makes it much harder—though not impossible—for individuals or groups within a community to leverage moral discourse to subjugate others. The demands of militant egalitarianism, then, serve to explain, in part, the epistemic structural norms of moral discourse.

What distinguishes moral discourse from institutional normative discourse is that moral discourse does not vest authority even in the community as a whole, and we now have an explanation for why the practice is governed by norms of this shape. A practice that does vest authority in the community is eminently vulnerable to the instabilities we have just discussed. Only a practice that makes the social norms of the community answerable to individuals can avoid this kind of instability. Moral discourse does this by opening even the accepted norms of a community to challenge from the perspective of individuals in good standing in that community. These norms must be defended in terms of reasons that the challengers find acceptable from their own perspectives. We might call these “public reasons” (Darwall, 2006), but in the very strong sense that they must be reasons that anyone could find acceptable, not just those individuals who already accept the shared norms of the community. Making social norms answerable to individuals reduces the risk that they will desert the group, thus maintaining a beneficially large

¹⁴ I put off for now the task of saying just what is required to be in good standing. One of the aims of Chapter 6 is to venture part of the answer to this question by addressing one important way in which the agent-relative inputs of moral prescriptives are shaped. For now, I shall only note that not just anyone has the standing to make a moral claim; one must be recognized as a member of the relevant community. The tension between this requirement of standing and the fact that moral discourse does not vest authority even in the community as a whole is fundamental to the vexing issue of moral relativism.

group size. It also enables groups to interact with one another across difference in social norms by providing a means to settle on shared norms answerable to individuals in both groups when necessary and when conditions are right.¹⁵

Moral discourse is a meta-normative vocabulary with a meta-coordinative function. It sits atop already up-and-running social-coordinative practices and serves to correct instabilities that arise in them by providing a necessary means for challenging, evaluating, defending, sharpening, and amending social norms. In the process, it changes the very nature of these norms. It is common among evolutionary theorists and metaethicists to try to demarcate moral norms from mere social or cultural conventions. There is a difference in kind between “One ought not to rape” and “One ought to silence her phone in the library” that it seems important to be able to capture (Buchanan & Powell, 2018, pp. 135–136). Evolutionary theorists have tended to try to capture this in terms of some norms *becoming* moralized. Tomasello, for example, argues that those norms that are properly related to our senses of sympathy and fairness will tend to elicit moral responses (Tomasello, 2016). Metaethicists, on the other hand, have tended to think of moral norms in terms of their universality. Moral norms are categorical, applying to everyone and enforceable by anyone. I think that the story is more nuanced. We do tend to be more concerned about norms properly related to the moral sentiments, but this alone is not what makes them moral. Rather, what makes them moral is the reason-giving practice in which they’re caught up, which makes ordinary social or cultural conventions ultimately answerable to individuals even though individuals are subject to their

¹⁵ Of course, the presumption is always in favor of doing what we’ve always done. A handful of dissenters are easily dismissed. Even so, there are occasions on which enough pressure arises to force such re-evaluation of communal norms.

authority. Moral norms are those that we have found important enough to examine and that we can enforce only if *we* can give reasons in their defense. They are, in this sense, *our* norms in a way that even our cultural norms are not. They are made by all of us for our own governance.

8. Conclusion

In taking rational ownership of the norms with which we govern our communal lives, we dramatically change their nature. Norms are often simply given to us. We find ourselves thrown into a community, a way of life, and we inhabit it. We do what people of our ilk do, what our kith and kin have taught us, what our ancestors or our gods require of us. These norms change slowly as our circumstances change. With new technologies and changing environments, what was once impermissible may become required, what was once required simply tolerated, and so on. The discursive technology of a meta-normative vocabulary like moral discourse, however, alters all of this, for now we can demand reasons for the ways we are told we ought to behave. When the reasons are found wanting, we can demand change by giving reasons for the norms we propose as replacements. It is likely that this new technology was used only in a limited way at first, tweaking the existing social order here and there. It is likely that tradition continued to rule, and that few found reason to question it. But the possibility now existed for moral upheaval, for the overthrow of the old systems, for rapid advances to keep up with changing environments. With time, we remade our norms and, so, ourselves, and we did so with great variety.

As I noted at the outset, the evolutionary account sketched here does not attempt to explain the content of the various moral dicta human moral communities have

discovered, only the practice by which they are discovered. This practice is one that we adopted to further our cooperative enterprises, to ensure that our groups did not fracture or dissolve in the face of dissent over the social norms that we found ourselves with, and to leverage opportunities for cooperation with other groups. This function, as we have seen, explains why moral discourse is governed by pragmatic structural norms that allow claims with agent-neutral entitlement to target and motivate particular individuals and why it exhibits a distributed pattern of authority, yet limits that authority to those one recognizes as members of one's community. Once such a practice was in place, however, it could be used to shape our cultures in various ways. At times it has led to insularity; the agent-relative inputs of moral prescriptives became more harshly policed and our moral communities contracted. We relied on moral discourse to enforce our norms, and in times of external threat or internal fractures, moral authority has even been usurped by those who could wield it to control the community. At other times, it has led to expansion. The agent-neutral inputs of moral prescriptives allow for new perspectives to gain purchase and challenge existing social orders. When times were good and the advantages of cooperation apparent, our moral communities broadened, and new voices begin to shape our thinking about what was required of us.¹⁶ Both of these are our legacies, and the evolutionary perspective developed here and paired with PALM allows us to see how they spring from the same source.

In the next two chapters, I turn to elucidating this tension. In Chapter 5, I offer an account of objectivity in terms of the authority structure of a practice and demonstrate that moral discourse exhibits this kind of objectivity. In Chapter 6, I examine one of the

¹⁶ See, (Buchanan & Powell, 2018, Chapter 6).

primary sources of insularity in our moral communities in the way that a particular kind of trust mediates the recognition in others of the standing to hold accountable, which is fundamental to membership in the community. In examining how these dimensions of objectivity and relativity are both inherent to moral discourse, we can come to understand how the question of moral relativism is one that we will never exorcise.

Chapter 5: Moral Objectivity

1. Introduction

PALM (the pragmatic account of linguistic meaning) is an avowedly Neo-Pragmatist, anti-Representationalist framework for the explanation of meaning in terms of use. I have drawn on work from Brandom, Williams, Price, Kukla, and Lance to fill out the account, and I've tried to demonstrate some of the benefits of thinking about language in this way. In this chapter, I put PALM to use to address an internecine dispute about what pragmatists should say about objectivity.

The very idea of objectivity seems to be wed to the Representationalist framework. An objective claim is one that purports to represent or correspond to an object (or an arrangement of objects, a state of affairs). An objectively *true* claim is one that succeeds in corresponding to its object; it is a *description* of it (Putnam, 2005, pp. 52–53). If this is what objectivity consists in, then it is irredeemably bound up with the core Representationalist commitments to theorizing meaning in terms of truth and reference and truth in terms of correspondence, i.e., the commitments that are at the root of debates about realism and anti-realism, cognitivism and non-cognitivism, and naturalism and non-naturalism. Pragmatists have roundly rejected this framework, so it makes sense to ask whether the discourse of objectivity gets thrown out with it. Neo-pragmatists are of two minds. Some, like Rorty, argue that talk of objectivity is tainted,

and risks reinvigorating this problematic framework. Instead, we should opt for something more down to earth: *solidarity*, “unforced agreement” among an ever-growing discursive community (Rorty, 1991e). Others—Brandom, McDowell, Price, Putnam, and others who are sometimes labelled New Pragmatists, for example—think there’s something there to be rehabilitated.¹ Rorty is “a disappointed metaphysical realist,” who sees that the guarantee of objectivity that our words map onto reality is an impossibility (Putnam, 2002, p. 101).² The rehabilitators don’t want to conceive of objectivity as such a guarantee but as a useful concept for capturing the idea that some of our claims but not others are “answerable” to something that is not directly in our control for their correctness (McDowell, 2000), that the conditions of applicability for some concepts transcend what anyone or everyone takes to be the case (R. Brandom, 1994), or that some of our claims, by virtue of how we use them, aim at covariance with something in our environment (Price, 2013, Chapter 2).

On which side should we come down: rejection or rehabilitation? Pragmatists ought to answer this question by asking what good the concept is. What does it do for us? If it’s all ill and no good, as Rorty thinks, then so much the worse for the discourse of objectivity. If, however, the concept makes a difference in practice, if it’s useful for us, then maybe we ought to find a way to hold onto it. In this chapter, I’m going to put forward an account of what good I think the concept of objectivity is by developing a conception of objectivity that falls out of PALM. I’ll plump for rehabilitation over rejection. The rehabilitated notion of objectivity I defend is a procedural account that follows some paths laid down by Brandom and by feminist scholars like Helen Longino

¹ See, (R. Brandom, 1994, Chapter 8; McDowell, 2000; Price, 2010, 2013; Putnam, 2002, 2005).

² On the roots of Rorty’s rejection of objectivity, see (Brandom, 2013) and (Levine, 2010).

and Elizabeth Anderson but that also takes some important lessons from Rorty.³ The account locates objectivity in the norms of authority of a discursive practice. As such, the PALM of a vocabulary will capture whether that vocabulary is, to some degree or another, objective. The utility of this procedural account is found in its usefulness as a guide both to which discourses are structured in such a way that “unforced agreement” among practitioners is a reasonable goal and to ways of either reinforcing or amending the norms of such practices in order to make them more open, democratic, and inclusive, i.e., more objective.

In the final third of the chapter, I’ll put the account I’ve developed to work to argue that moral discourse counts as a fully objective discursive practice, while its neighbors—prudential discourse and taste discourse—fall somewhere closer to the subjective end of the spectrum.

2. The Traditional Notion of Objectivity

A good place to start is the very basic question: What is objectivity? This question, unfortunately, has no simple answer.⁴ As with all philosophical terms of art, there is no one way that philosophers (or, for that matter, non-philosophers) have deployed this concept. There are, however, some themes constitutive of the traditional

³ I will not argue that the account of objectivity developed here is the *only* account available. I think there are others that do a different kind of work, for example, in helping us to get a grip on notions of linguistic covariance with external objects—Price’s e-Representation (Price, 2013, Chapter 3)—as well as how our claims manage to be *about* the objects and properties to which we grant authority over them (R. Brandom, 1994). These other notions of objectivity, though, are parasitic on the more fundamental, procedural account developed here.

⁴ Indeed, countless books have been written on the topic. See, for example (Nagel, 1986), (B. Williams, 1978), (Boghossian, 2007, Chapter 5), and (Axtell, 2015). For a conceptual history of objectivity, see (L. Daston, 1992; L. Daston & Galison, 1992; L. J. Daston & Galison, 2010). There are many ways of troubling the idea of objectivity. From the feminist perspective, see, (Lloyd, 1995), (Kukla, 2006), (Douglas, 2004); and for pragmatist concerns about the concept, some of which I’ll return to in a moment, see, (Rorty, 1979, 1991b, 1991f).

concept of objectivity that are nicely captured by Elizabeth Anderson. Objectivity, she claims, consists of a package of ideas put forward in the 17th and 18th centuries as a philosophical account of the superiority of Newtonian science. The package consisted in the following:

(a) Subject/object dichotomy: what is really (“objectively”) real exists independently of knowers. (b) Aperspectivity: “objective” knowledge is ascertained through “the view from nowhere,” a view that transcends or abstracts from our particular locations. (c) Detachment: knowers have an “objective” stance toward what is known when they are emotionally detached from it. (d) Value-neutrality: knowers have an “objective” stance toward what is known when they adopt an evaluatively neutral attitude toward it. (e) Control: “objective” knowledge of an object (the way it “really” is) is attained by controlling it, especially by experimental manipulation, and observing the regularities it manifests under control. (f) External guidance: “objective” knowledge consists of representations whose content is dictated by the way things really are, not by the knower (Anderson, 2015).

These ideas coalesce into an account of science.

[The] aim [of science] is to know the way things are, independent of knowers, and...scientists achieve this aim through detachment and control, which enable them to achieve aperspectivity and external guidance (Anderson, 2015).

This conception of the aims and practices of objective science is then deployed to demonstrate the conflation by Scholastic scientists, who conceived of “objects as intrinsically possessing secondary qualities and ends,” between the way things are in themselves and the way they appear to “emotionally engaged human knowers.” Science governed by the ideals of objectivity “enabled the successor scientists to avoid these errors and achieve an ‘absolute’ conception of the universe” (Anderson, 2015 citing; B. Williams, 1978).

The traditional notion of objectivity weaves together strands of several different kinds. It is committed to *ontological objectivity* in the sense that the aim of inquiry is to veridically represent the objects of inquiry, which exist in a world that is mind-independent. How the world is not up to those who perceive it, and it exercises authority

with respect to our claims about it. Objective claims are those that are *answerable* to the objects and properties that they are about. Veridical representation, moreover, is achieved via *epistemological objectivity*, which ensures that the representation is not marred by individual or cultural bias but is instead elicited by the object itself. The aperspectivity of epistemological objectivity is itself ensured by a kind of *methodological objectivity* that is constituted by a collection of norms that counsel detachment and distance from the objects and results of inquiry or judgment.

Once in place, this traditional conception of objectivity is available to be pressed into service toward a variety of ends. It is put to explanatory use to shed light on the success of the physical sciences. They employ research methods and structure their intellectual communities in such a way that the biases and errors of individual researchers and labs are minimized and controlled for. As such, their claims are constrained by the objects they are about, and they achieve a kind of aperspectival view of the world as it is in itself. The knowledge that is thus made available is eminently useful in controlling the world we inhabit.

Scientific understanding, moreover, is getting better. We are making intellectual progress, closing in on the “one full, objective, true, account of nature” (Kuhn, 1962, p. 171), and this too is explained by this realist account of objectivity. Objective methods of inquiry allow us to “climb outside of our own minds” to see the relationship in which we stand to the world we are trying to represent and, so, to remove the distortions caused by our individual perspectives, presuppositions, or biases (Nagel, 1986, p. 9). The more removed our perspective becomes, the more felicitous will be our representation of what is really there. This is what constitutes progress.

As we make progress, we come to see that along the way we got some things wrong. If we hadn't, it wouldn't be progress. Not only do individuals get things wrong, but entire communities get things wrong. The best scientific minds of their times thought that the earth was the center of the universe, that motion was absolute, and that phlogiston was released in combustion, but no one with a descent education believes those things any longer. In science, at least, it seems to us that no matter how well-supported one's claims are, it's still possible that they are not true. There is a possible gap between what we have reason to believe and what really is the case, and this gap is explained by the traditional notion of objectivity. Our claims in this domain aim to represent how things *really* are with a world that is out there, distinct from our beliefs about it. As such, we could always come to find that what we had the best reason to believe is not, in fact, the case. Our discourse is answerable to the world, and the world can prove that we've made a mistake.

Not all discourse, though, is answerable to the world. Some discursive practices follow the model of science or empirical inquiry more generally. They aim at veridical representation of the world that is there for us to discover, and they involve methods of detachment that aim at a "view from nowhere." These discourses are lauded as objective. Others, though they might pretend to objectivity, are not, in fact answerable to objects that exist apart from us either because the discourses are poorly structured or because the objects they purport to represent simply don't exist. These are second-class discourses at best. They are the domain of mere opinion, not knowledge. Their truths, if there are any, are socially constructed and, so, it is thought, somehow within our control, not imposed

on us by the world. This is the differentiating work of objectivity: it sorts the robust, realist discourse from all the rest.

3. Objectivity Challenged

Pragmatists insists that we cannot make sense of the idea that some discursive practices are more in touch with the world as it exists in itself than are others, that they offer better representations of that world and, so, count as (more) objective in this traditional sense.⁵ The reason we cannot so much as make sense of this idea is that it is impossible for us to “climb outside of our own minds” (Nagel, 1986, p. 9). In order to judge that one practice or vocabulary is more in touch with reality than another we would need to be able to take a sideways-on, vocabulary neutral view on our discursive practices and the world, but no such neutral standpoint exists. When we detach ourselves from one practice to examine it sideways-on, we inevitably inhabit another, which imposes its own structure on our view of how the first vocabulary mapped onto the world. We cannot step outside of our own skins. Pragmatists of all stripes recognize that this is our lot. There are no skyhooks that can lift us out of the language-games we deploy to inhabit, navigate, and understand our physical and social worlds. There is no God’s-eye view for us to take up in order to see whether some vocabularies are more in touch

⁵ The discursive practices up for comparison may be practices of a single community—say its moral and scientific discourses—or they may be the practices of two or more communities, e.g., the moral practices of western liberal democracies versus those of Western Nepalese villagers who defend the practice of *chhuapadi*. In the first case, realists may claim, for example, that scientific claims are objective while moral claims are subjective or relativistic or merely expressions of non-cognitive attitudes. In the second, a moral realist might claim that moral claims of the denizens of a western democracy in condemnation of *chhuapadi* are objectively true while those of the Nepalese practitioners are objectively false. On the other hand, she might claim that there is no objective fact of the matter, i.e., that moral claims can be true only relative to a moral framework or grid of justification. In that case, it may be true for the westerner that *chhuapadi* is wrong but true for the Nepalese that it is not. Note, however, that this kind of relativism can only be posited by a theorist who is committed to taking up a kind of external view on her own practices in the way that the realist does but, as we shall see, the pragmatist rejects. See, (Rorty, 1991e, pp. 23 & 30). Cf. (Harman, 1977; Harman & Thomson, 1996).

with language-independent reality than others (McDowell, 2000; Putnam, 2002; Rorty, 1991e, 1997, 2000).

Pragmatists, as such, challenge the ideal of aperspectivity. Any representation of the world will reflect the perspective, interests, and biases of the observers. We cannot relinquish the set of concepts and biases our enculturation has bestowed on us, nor can we detach ourselves from our practical interests as we investigate the world. Indeed, our biases and interests are what get our investigations off the ground.⁶

Once aperspectivity is rejected, the rest of the traditional package of objectivity begins to unravel. If we are always inhabiting an embodied perspective on the world, though we can strive for emotional detachment and value-neutrality in inquiry, we must recognize these as regulative ideals, never to be achieved.⁷ If we don't have access to this neutral perspective, then we cannot be assured of external guidance exerted on our beliefs by the objects themselves. We can never be certain that our interests and biases are not affecting the judgments we make. Moreover, feminist theorists have argued, following Quine, that the evidence we have about the objects of inquiry always underdetermines our theories of those objects. We make innumerable contingent choices with regard to attention, framing, interpretation, and theoretical virtues like simplicity and beauty, yet those choices are, in most cases, taken for granted in the actual practice of science and in philosophical reflection on those practices. Rather than pretend these choices are not

⁶ See, (Antony, 1993) and (Anderson, 2015).

⁷ These ideals are thought necessary to avoid projective errors, wishful thinking, dogmatism, and motivated reasoning, but feminist theorists have argued that they often do just the opposite. Detachment can lead to epistemic defects of emotional distance and an inability to correctly perceive certain emotionally laden phenomena, and when we represent ourselves as value-neutral inquirers we cannot help but fail to recognize the ways in which our values shape our inquiries. As such, we will never bring those values themselves up for critical examination.

made, we would do better to reflect on them and understand how they shape our theorizing (Longino, 1990, 2001).

Given that we cannot escape our language-games, Rorty—representing the first kind of neo-pragmatist I introduced—argues for a wholesale rejection of the discourse of objectivity. He links the desire for objectivity to a yearning for an easy way out of our difficult cognitive circumstances, for an authority to settle matters for us. This, he claims, is a mark of cultural immaturity. On the Representationalist notion of objectivity inquirers submit themselves to standards of correctness constituted by the bit of reality that is to be under investigation. Rorty conceives of this as akin to authoritarian religion with the world itself playing “the role of the non-human Other before which we are to humble ourselves.” In recognizing the authority of the world in itself we abase ourselves before something non-human, and, seen in this light, engaging in the discourse of objectivity “merely prolongs a cultural and intellectual infantilism” that, for Rorty, is endemic to the tradition from Plato right through to contemporary analytic philosophy (McDowell, 2000, pp. 109–110).⁸

The reason he sees this as intellectual and cultural immaturity is that a belief in objectivity wishfully denies the discursive or deliberative predicament in which we find ourselves. Engaging in certain areas of inquiry, collecting evidence, constructing arguments, and so on, we can come to be fairly confident in the answers at which we arrive. We can justify to ourselves and our peers the conclusions we’ve reached using the standard resources we have in whatever discourse we’ve taken up and congratulate ourselves on the work well done. Yet, even if we’ve done our level best at honest,

⁸ Also see, (Rorty, 2000).

meticulous inquiry, we lack any guarantee that we will be able to convince just anyone with whom we might discuss the matter. This shakes our confidence in our answer, especially when we really do encounter someone we cannot persuade. The discourse of objectivity, as McDowell reads Rorty, is “a philosophical attempt to shore up the confidence so threatened.” If there is an objective reality, it will set us straight, and anyone who fails to see things the way we do must be failing to make use of the universal human capacities for perception and reason. The idea of objectivity is the idea that “reality itself fills this gap in our persuasive resources” (McDowell, 2000, p. 112), but if we are to take responsibility as fully mature inquirers we must do the work of filling this gap ourselves. We must restore our lost confidence by the hard work of justification and rational persuasion (and, sometimes, force) not by the easy move of dismissing dissenters as rational failures. Objectivity offers us undeserved consolation.

In the place of the notion of objectivity, Rorty says we should aim instead for *solidarity*. The hope of solidarity is that we can construct ever expanding discursive communities, an ever larger “we” and ever more diverse *ethnos* to which we must justify our judgments and against which they must pass muster (Rorty, 1991e, p. 23). As the discursive community is extended to an ever more diverse *ethnos*, the claims that the present community holds to be well-justified, those to which they would ascribe truth, come under new challenges from new discussants with new perspectives as well as from old discussants with new technology and new insights. The judgments that stand the test of this expanding space of reasons are the ones that we will count as better justified, those to which the newly constituted community will ascribe truth. The test, though, is one that relies on persuasion, on respect for the opinions of one’s discursive peers, on curiosity

and a thirst for new information and new tests. The hope of solidarity is the building of such a community, which provides us the opportunity to continue testing and refining our judgments, to become better and better at navigating the world we find ourselves in (Rorty, 1991d, p. 39).

Like objectivity, the idea of solidarity gives some sense to the gap between truth and justification, to the notion that even if we've done our level best, we could still somehow have missed the mark. Unlike objectivity, however, solidarity doesn't require that we have an idea of the mark we might miss, only the idea that we might do better as our discursive community matures and expands. Rather than coming up short from correctly representing the world as it is in itself, recognizing that justification may fall short of truth is just recognizing a "gap between the actual good and the possible better. From a pragmatist point of view, to say that what is rational for us now to believe may not be *true*, is simply to say that somebody may come up with a better idea" (Rorty, 1991e, p. 23).⁹

Rorty also thinks that his pragmatist notion of solidarity can help us to make sense of the intuition of intellectual progress. He writes, for example, that "the intuition that we are making intellectual progress is simply the intuition that, in respect to self-consciousness and intellectual responsibility, we are getting farther and farther away from the cavemen" (Rorty, 1997, p. 176). Rather than convergence on a final truth, the intuition of progress is about growth, about getting better and better and navigating our

⁹ It is important to see that Rorty's idea is quite distinct from the Peircean idea of an ideal speech community that converges on truth at the end of inquiry. The Rortian discursive community expands, but there is no ideal that it might reach, no final conclusion against which we can measure the judgments we now make. Our position is more precarious than this. Any conclusions we might reach might, by some future community, be overturned. His is not an epistemic notion of truth.

physical and social world. Such growth is the product of expanding our intellectual communities, providing more and better tests for our commitments. What is important, according to Rorty, is that we can tell ourselves a story about how we got from where we were to where we are that is consistent with the demands of our reason mongering practices.

Besides the notions of aperspectivity, detachment, value-neutrality, what we must give up from the traditional notion of objectivity in adopting solidarity is a commitment to the idea that some discursive practices are more in touch with the world out there than are others. No discursive practice in which we engage is privileged over any other in the sense of being more in touch with reality, for we would have no way of making such a judgment. No part of culture better represents nature or exhibits the humility to bring mind into direct confrontation with the world, rather every discursive practice can't help but put us in touch with the world in some way for we can't help but be causally in touch with it (Rorty, 1991d, p. 36). Different areas of culture simply answer to different human needs (Rorty, 1991a, p. 8). Some discursive practices like the sciences may deserve special praise, but this is only because they have developed institutions that "give concreteness and detail to the ideal of 'unforced agreement'...[and] flesh out the idea of 'a free and open encounter'" that constitute the ideal of solidarity (Rorty, 1991d, p. 39), it is not because they are more detached from us or more in touch with whatever's out there.

We can draw on Helen Longino to add some useful detail with her understanding of scientific discourse as democratic discourse. She argues that "[c]ritical discursive interactions are social processes of knowledge production...[which] determine what gets

to remain in the public pool of information that counts as knowledge” (Longino, 2001, p. 129). As such, the aim of good knowledge producing practices should be to produce more and better critical discursive interactions, just as Rorty envisions with his notion of solidarity. What Longino adds to the idea is an examination of the kinds of norms that make such practices epistemologically effective. She claims, for example, that we must ensure *venues*, “publicly recognized forums for the criticism of evidence, of methods, and of assumptions and reasoning” (Longino, 2001, p. 129). We must also enforce the *uptake* of criticism, not just the tolerance of dissent. We must hear out dissenters and respond to them when such dissent is reasonable, but dissenters must also give up on lines of criticism that have been reasonably answered (Longino, 2001, pp. 129–130). Further, we must adopt *public standards* to which we refer in the evaluation of evidence, theories, hypotheses, and research practices. And, finally, such practices must have norms of *tempered equality* according to which intellectual authority is distributed. This ensures a diversity of perspectives; everyone could potentially dissent or demand reasons, but such equality must be tempered with respect to diversity in intellectual capacities and expertise. You’ve got to demonstrate you can play the game before your dissent counts for much (Longino, 2001, p. 131). Paired with Rorty’s ideal of solidarity, this gives us a nice picture of an ever-expanding discursive community committed to testing its judgments against new perspectives, one that is free and open and aimed at fostering unforced agreement while giving appropriate uptake to dissent. Practices that exhibit these norms to a higher degree count as exhibiting greater solidarity.

4. Rehabilitating Objectivity?

Rorty's prescription to abandon the discourse of objectivity is too strong. He is right to warn of the danger of backsliding, of ending up unconsciously accepting some of the problematic assumptions, metaphysical baggage, and wishful thinking of the Representationalist version of this discourse, but the concept of objectivity can be shorn of these notions and reclaimed to do some work that pragmatists should want to do.

Whether chocolate ice cream tastes good, whether the shade of green on the middle swatch is more pleasing, whether I ought to go for a bike ride. The answers to these questions seem to be premised on something particular about my reactions to the world, my tastes, and my interests. They are, in this sense, up to me in a peculiar sort of way, a way in which the answer to whether the tree out my window is a maple or an oak is not. Each of these has to do with some part of my experience that we are apt to label 'subjective'. It is something specific to me that determines the answers to these questions, and it is something specific to me over which I have some modicum of control. Though my tastes seem to be something simply thrust upon me, it is clear on reflection that they can change and that I can change them. I can learn to appreciate flavors, sensations, or experiences that I previously found unpleasant. I can refine my palate by exposing myself to new flavors. I can expand my interests in fictional genres by exploring new authors. Moreover, since what is in my best interest is ultimately a matter of my preferences upon examination and reflection, my interests can vary, as well.

The answers to some other questions are not up to me in this way. Whatever my responses or my preferences are, certain four-letter words constitute insults, the smaller fork on the left is for salads, and I ought to drive on the right-hand side of the road. I might think salad forks silly and believe that no words are endowed with the special

power inflict harm, but what matters in these cases are what my community does, not what I believe. If it is a convention in my community to use the little fork for salads, to drive on the right, and to take umbrage at certain four-letter words, then it is the case that those are salad forks, I ought to drive on the right, and those words are insults. The community simply can't be wrong on such matters.

Is the speed of light constant? Are red cars involved in proportionally more accidents than cars of other colors? Is representative democracy a more just form of government than autocracy? The answers to these questions and many others are neither up to me nor a matter of what my community does. What I or my community take to be the case, how we conduct ourselves, or what our preferences are has no (immediate) bearing on the correct answers to these questions. Changing my or our *thinking* about the speed of light does not change the speed of light.¹⁰ We could all agree that red cars are not involved in more accidents, yet we could all be wrong. It could be common wisdom backed by centuries of research and argumentation that representative democracies are more just, yet we might discover that wisdom mistaken. These questions have answers that transcend what anyone or everyone takes to be the case. To put it another way, these questions call for answers that purport to be objective in the sense that their status as true or false transcends our deontic attitudes of taking one another (or ourselves) to be entitled to them and related claims.

¹⁰ Cf. (Haslanger, 1995, p. 96): "Bridge abutments and fists to the face are "independently real" at least in the sense that no individual or community of individuals can simply think them out of existence...a change in my thinking, *by itself*, cannot make my body, my friends, or my neighborhood go out of existence, nor thankfully can a change in anyone else's." Note, however, that a communal change in thinking that affects our patterns of behavior could, in fact, make my neighborhood go out of existence or change the number of accidents involving red cars. The point is not that these are facts that we cannot change but that the truth or falsity of our claims is determined by something beyond what we individually or as a community take to be the case.

Rorty's idea of solidarity promises to make some sense of the notion that what is the case potentially outruns what anyone or everyone takes to be the case. What it cannot do, however, is tell us why some discursive practices are like the latter examples and some like the former, i.e., why some exhibit solidarity in this sense while others do not. One thing that goes missing when we lack the capacity to sort discursive practices along this dimension is a theoretical resource for explaining *why* "unforced agreement" is sometimes to be expected in a practice and sometimes not. We might be able to point to a practice like empirical discourse noting that persuasive resources might bring about broad agreement and to a practice like taste discourse noting that they likely won't, but, having rejected Representationalism, we cannot explain this feature of these practices. Rorty might respond that the notion of objectivity I want to rehabilitate also can't explain this distinction, it merely makes us feel as if we have an explanation by intimating that in the one case there's something out there for our claims to be about, while in the other there is not. I think, however, that we can avoid this intimation if we are careful about what objectivity consists in.

I'll argue that objectivity must be understood as a feature of the norms of authority of a discursive practice, not as a matter of its relation to something outside of itself. Moreover, the notion of objectivity that I offer here does work that the notion of solidarity cannot. Once we can explain the distinction between objective and non-objective practices in functional terms, we will have at our disposal a framework for diagnosing failures to achieve varying degrees of "unforced agreement" in our target practices. We will be able to say why a practice is no longer fulfilling the function for which it was designed. In some cases, this may be because the practice aims to fulfill a

need its users no longer have, in others it may be because the practice has decayed or been mutated either intentionally or unintentionally, under its own weight or by external forces. Without the notion of objectivity, though, we have no grip on these issues. We have only practices that may or may not aim at agreement and that may undergo changes over time, but no way to rationally assess them.

The question, then, is whether we have the resources to capture this sense of objectivity in a way that is consistent with pragmatist critiques of the traditional notion of objectivity. The central worry is that we may not because our starting point has been one of methodological phenomenism about norms. Following Brandom, we have accepted that normative statuses of commitment and entitlement are the products of normative attitudes of *taking* one another to be committed and entitled. Attribution is prior to status. Whether one is entitled to attribute commitment or entitlement is itself a matter of the attitudes of others attributing entitlement, and so on. Given this starting point, it is difficult to see how we might achieve any kind of *attitude-transcendent* constraints on our discursive behavior. How could it be that what one ought to say in a given case—a normative status she has—can outrun what anyone or everyone does say—which is constitutive of the normative statuses they attribute to her—if what they do say—sometimes in very complex ways, not merely as a matter of simple agreement—always settles what they should say since they are the only judges in the neighborhood? Showing that and how this is possible is now our challenge.

5. Brandom on Objectivity

The place to begin addressing our challenge is with Brandom, since his account raises the same concerns. The idea he aims to defend is that objectivity is a “feature of the

structure of discursive intersubjectivity” (R. Brandom, 1994, p. 599) such that it is an ineliminable feature of every discursive perspective “that there is a difference between what is objectively correct in the way of concept application and what is merely taken to be so” (R. Brandom, 1994, p. 600). To get at what this means, we need to take a step back and briefly review his deontic scorekeeping account of conceptual contents.

On Brandom’s model, language users are engaged in practices of giving and asking for reasons. The fundamental moves in these practices are the attribution of deontic statuses of commitment and entitlement on the basis of performances by other practitioners that, by virtue of being caught up in these practices in the right ways, count as speech acts, primarily acts of asserting. The attributions themselves are understood as deontic attitudes, i.e., as attitudes of *taking* or *treating* a performance as appropriate or inappropriate, entitled or not, as the undertaking of a commitment or simply an emission of sound. Each individual acts as a scorekeeper tracking the commitments and entitlements of her interlocutors. When an interlocutor makes a move, e.g., by uttering “Plato (my dog) is a beagle,” the scorekeeper attributes to her a commitment to the content of her utterance by taking her to be committed to inferentially upstream antecedents—ways one could answer the question “Why do you say that?”, e.g., by uttering “By his size, color, and distinctive howl”—and inferentially downstream consequents—for example, “Plato is a mammal”—of her claim.

As we’ve already seen, this inferential network is holistically constitutive of semantic content for Brandom. Given the holistic nature of this account, a scorekeeper can only attribute inferential commitments to a speaker by embedding her commitment in a network of background commitments, but, here, she faces a choice. She can either score

the move with respect to the background commitments that she herself accepts *or* with respect to the background commitments she attributes to the speaker. Scoring it according to the former gives her the commitments she takes the speaker to have acknowledged, i.e., it represents how the scorekeeper takes things to be from the speaker's perspective. Brandom calls this the *de dicto* reading. Scoring it in the latter way, on the other hand, gives her the commitments she takes the speaker *actually* to have undertaken, i.e., it represents how she takes things to be from her perspective. This is the *de re* reading.

In Brandom's lightly regimented idiom, the distinction between what is attributed *de dicto* and what is attributed *de re* is marked by what is contained within the scope of the "of" clause of a *de re* attribution. Consider Brandom's example of the Constable reporting to the Inspector that he had spotted "the desperate fugitive, a stranger who is rumored to be passing through the village" in the darkened courtyard the previous evening. The Inspector, however, believes that the man the Constable describes as "the desperate fugitive" is, in fact, "the Croaker, a harmless village character whom no one, least of which the constable (who knows him well), would think could be the desperate stranger." The Inspector can make explicit the "objective" content of the Constable's belief by way of an ascription *de re*: "The Constable claims of the Croaker (a man who could not possibly be the fugitive) that he is the fugitive." With this ascription, the Inspector attributes a commitment to the Constable that he saw someone in the courtyard last evening and that it was the fugitive, but he *undertakes* commitment only to its being the case that the Constable saw someone and withholds commitment to its being the fugitive. Rather, from his perspective, he anaphorically picks up the demonstrative

commitment of the Constable but specifies the target of the demonstrative in terms of his background commitments as “the Croaker” (R. Brandom, 1994, p. 595).

Here’s Brandom’s next important move in thinking about objectivity: this isn’t a choice. Scorekeepers don’t just keep a single set of books, they keep two. For each interlocutor, they (we) track both *de dicto* and *de re* attributions.¹¹ The Inspector attributes *de dicto* to the Constable commitment to the claim that the desperate fugitive was in the courtyard last evening, and he attributes *de re* to the Constable commitment to the claim *of* (or about) the Croaker *that* he was in the courtyard last evening on the basis of *his own* background commitments regarding the whereabouts of the Croaker (R. Brandom, 1994, p. 597). Keeping these two sets of books opens up a gap across which information can be extracted. The scorekeeper can learn from the speaker by weaving the commitments to which she takes him to be entitled into her own web of beliefs while screening out those that are in conflict with her background commitments.¹²

More importantly for us, the gap between the two sets of books is also where we get our first glimpse of objectivity in Brandom’s account. What follows from a claim is

¹¹ This is also central to Brandom’s account of communication as *intra-linguistic* interpretation (R. Brandom, 1994, pp. 477–480 & 588). Also see, (R. Brandom, 2007, pp. 667–668).

¹² One might worry at this point that the commitment to semantic holism raises an insurmountable problem for this account. Since the meaning of a claim is determined by its place in an inferential network, meaning can only be determined in conjunction with a set of background commitments. If, however, the background commitments of the scorekeeper and the speaker differ—as they must—then they cannot possibly *mean the same thing* by their claims. The words in the scorekeeper’s mouth cannot mean the same thing as they do in the mouth of the speaker to whom they are ascribed. If this is right, then the gap between perspectives looks unbridgeable. Brandom argues, however, that even though claims must be expressed differently from different points of view “what is attributed explicitly in a *de re* ascription can be the very same claim that would be acknowledged, using different words, in an assertion by the one to whom it is ascribed.” Speakers can understand what one another say, but doing so requires “mastering the coordinated system of score-keeping perspectives, not passing something nonperspectival from hand to hand (or mouth to mouth).” The content they both grasp in communication is “*essentially expressively perspectival*.” It can only be specified from a particular point of view, but we can manage to navigate across those points of view by attributing background commitments (sometimes shared with us, sometimes divergent) to our interlocutors. “Conceptual content consists in the systematic relations among the various pragmatic significances” that claims have in the mouths of various speakers (R. Brandom, 1994, pp. 590–591).

determined from the perspective of the scorekeeper by conjoining it with claims to which she acknowledges commitment, i.e., with what she takes to be the case, but this may differ from the commitments she attributes to the speaker. So, from the perspective of each scorekeeper, there is always a distinction for every other interlocutor between what that player acknowledges and what she is actually committed to, for what she is *actually* committed to is determined by what the scorekeeper takes to follow from her claim. The scorekeeper, then, can judge that when the speaker's acknowledged commitments diverge from those she's consequentially undertaken, she is *objectively* mistaken. What she takes to be the case is not what is the case (from the scorekeeper's perspective), and, so, what she says is not what she should have said. The scorekeeping perspective embodies a seems/is distinction. The "permanent possibility of a distinction between how things are and how they are taken to be by some interlocutor is built into the social-inferential articulation of concepts" (R. Brandom, 1994, p. 597).

Each interlocutor in a discursive practice is herself a scorekeeper tracking her own commitments and those of her interlocutors, and each of these perspectives maintains this distinction between how things are and how they are merely taken to be. Since this is the case, Brandom claims that no perspective is "privileged in advance over any other" (R. Brandom, 1994, p. 600). Each is in the same circumstance in the game scoring the moves of others from her own perspective, and so each is susceptible to being judged objectively mistaken by others. No perspective can claim to have a privileged "view from nowhere," and we cannot know, prior to sorting things out in the space of reasons, which perspective is better informed, more trustworthy, or worthy of belief. Figuring out who ought to be

counted as correct “is a messy retail business of assessing the comparative authority of competing evidential and inferential claims” (R. Brandom, 1994, p. 601).

The sense of objectivity that I have set out to salvage consists in the persistent possibility that what one takes to be the case or what everyone takes to be the case is not, in fact, the case. The traditional conception of objective correctness takes it that this is a matter of one’s application of a concept corresponding to how things are or how they would be taken to be from a privileged perspective, some view from nowhere, God’s-eye view, or view from the end of inquiry. Brandom’s account, on the other hand, aims to “reconstrue objectivity as consisting in a kind of perspectival form, rather than in nonperspectival or cross-perspectival content” (R. Brandom, 1994, p. 601). What ensures that there is a possible gap between what anyone or everyone takes to be the case and what is the case is that there is a fundamental symmetry between the perspectives of discursive practitioners. Every practitioner is a scorekeeper, and each could judge any other to have made an objective error and, so, could demand reasons for the claim the other has made. Sorting out which of the parties has the best reasons determines what commitments both should undertake, but this is always a provisional matter for both are subject to further assessment from others and in light of new evidence.

It’s easy to see how each scorekeeper can maintain this distinction between what one takes to be the case and what really is the case with respect to other practitioners. Where we have two sets of books to check, it’s just a matter of comparing them. It’s also fairly easy to see what Brandom’s tactic is for preventing what is the case from collapsing into what the community takes to be the case, i.e., from privileging the perspective of the community as a whole. His notion of *I-Thou* sociality does this work. *I-We* sociality takes

each individual to stand in relation to a community of speakers. What the individual speaker means by her words and whether her utterances are appropriate is determined by reference to what the community takes to be appropriate. Deviating from the privileged communal perspective is getting things wrong, but it is impossible for the community as a whole to get things wrong. According to *I-Thou* sociality, each individual stands only in relation to other individuals, not to the community as a whole. Each individual judges the speech of others not against the standard of what the community takes to be the case but of what she takes to be the case. Of course, the community plays a role, for it is via enculturation into the community that shared norms are propagated, but the community's perspective is not privileged. Each individual is encumbered with the task of sorting things out for herself, and each speaker only needs to be concerned with making claims that could be taken up by her immediate audience. In this way, what the community accepts comes into discursive contact with the claims made by individuals in a way that does not grant any authority to the community perspective. The claims an individual accepts that are consistent with her community's beliefs may still be challenged and will still require reasons in their defense. In most cases, discursive practitioners will not be content to accept, "That's just what *we* believe/do," as a good reason.¹³ No, the challenge for an account like Brandom's—and ours—is to show that the distinction between what one takes to be the case and what really is the case can be maintained when one considers one's own beliefs. The worry is that in this case, one only has one scorebook from which to work. The *de dicto/de re* distinction collapses in on itself, so how can one make sense of getting things wrong oneself?

¹³ Though, in some important cases we'll discuss below, they will accept precisely this sort of reason.

If each scorekeeper cannot maintain this distinction with respect to her own commitments, we risk the first-person perspective of the scorekeeper becoming, for her, “the infallible arbiter of the way things are” (Nyhof-DeMoor, 2011, p. 88). The result would be a subjectivist, problematically dogmatic understanding of discursivity. We couldn’t make sense of individuals reassessing or revising their commitments, accepting challenges from others, or feeling the need to provide reasons for her claims. Rather, we’d have a model in which individuals treat one another as fallible instruments from which additional information about the environment could be extracted but that could never undermine the commitments one already has. Anytime a conflict arises on this model, each individual could do nothing other than take herself to be right and her interlocutor mistaken. This would clearly be a deeply problematic result. If, on the other hand, Brandom can show that this distinction can be maintained even when the individual examines her own commitments, then he will have shown, by his lights, that nothing in his model entails the privileging of any individual or communal perspective, i.e., that from all perspectives there is a possible gap between how things seem and how they really are.

Brandom argues that any account that maintains the possibility that an individual could see her own commitments as possibly mistaken should be able to show that two unacceptable conditionals do not hold. The first he dubs the “No First-Person Ignorance Condition” (NFPI): $(p) [p \rightarrow (I \text{ claim that } p)]$. This says that if something is the case, then I take it to be the case. The second is the “No First-Person Error Condition” (NFPE): $(p) [(I \text{ claim that } p) \rightarrow p]$. This conditional says that everything that I claim is true. If the former holds, then each individual would be committed to the claim that she cannot

fail to know something that is the case. If the latter holds, then, for each individual, whatever she takes to be the case must be the case. These conditionals define the individual's perspective as privileged from her point of view. She cannot make sense of herself getting things wrong, and, so, for her there's no possibility of a gap between what she takes to be the case and what is the case. There's no objectivity of the kind we're after if these conditionals hold.

Brandom argues that neither NFPI nor NFPE are consequences of his scorekeeping account of discursivity. To show that NFPI does not hold, it will suffice to discover a claim that is incompatible with the consequent yet compatible with the antecedent. "I do not claim that p" is just such a claim. It is clearly incompatible with "I claim that p;" it is its denial. The key to seeing that it is compatible with "p" is to see that nothing precludes one from being entitled to both "p" and "I do not claim that p," even though it would be strange to undertake a commitment to both. This becomes evident if one considers my endorsement of the ascriptions "S claims that I do not claim that p, and S claims that p." There is nothing suspect in taking S to be entitled to both of these commitments, "for one involves what commitments S attributes to me, and the other involves what commitments S undertakes, and these do not collide." This ascription makes evident how I can take up a *third-person perspective* on my own present discursive statuses by attributing to another attitudes of undertaking commitments with regard to my own statuses. In doing so, I can avoid the seeming problem that, for me, "p" and "I do not claim that p" seem to be in conflict. It becomes clear that the conflict is not between the semantic contents of the two claims, but between their pragmatic import. "My denial that I claim that p collides with what I am doing (claiming that p), not with

what I am saying (that p)” (R. Brandom, 1994, p. 605).¹⁴ The upshot of being able to take up this third-person perspective on my own statuses is that, just as I can take others to have acknowledged commitments to which they are not entitled, I can take myself to have acknowledged commitments to which I might not be entitled.

A counterexample to NFPE is forthcoming along the same lines. Any claim *q* such that *p* and *q* are incompatible will serve to generate such a counterexample, for I could ascribe to myself both “I claim that *p*” and “*q* (which is incompatible with *p*)” without explicitly undertaking incompatible commitments. I could be entitled to the ascriptional commitment “I claim that *p*” and to “*q*” so long as I am committed to both *p* and *q*, yet, since *p* and *q* are incompatible, I would not be entitled to either. I could be entitled to the ascription of incompatible commitments to myself, and this is just what it is for me to have incompatible commitments and, so, to be in error about something (R. Brandom, 1994, p. 606).

These “objectivity proofs,” according to Brandom, show that nothing in his “deontic scorekeeping account of inferentially articulated conceptual contents” entails a collapse into the “mere privileging of one’s own perspective” (R. Brandom, 1994, pp. 601–604). Each scorekeeping perspective maintains the omnipresent possibility that things are not as one takes them to be, whether that one is an interlocutor, oneself, or one’s whole community. This seems, then, to fit the bill. We now have an account of objectivity that holds that no perspective is privileged in advance over any other. On this account we can make sense of the intuition that it’s possible that what anyone or even everyone takes to be the case is not, in fact, the case, but we haven’t succumb to the idea

¹⁴ This, of course, is just a version of Moore’s paradox.

that what makes it so is some non-human other to which our discourse is answerable. Rather, what makes it so is that each of us is always in a position from our own perspectives to call others to task for what they take to be the case. Each perspective is answerable to all others, and authority, in the end, resides in the reason-giving practice itself and the means it provides for sorting competing claims. That authority is just ours, for we are the ones who enforce the norms of such practices. We do not subjugate ourselves to some non-human Other.

6. Sorting Discursive Practices

This Brandomian notion objectivity does some of the work we hoped our rehabilitated pragmatist notion would be able to do but not all of it. Brandom demonstrates that his scorekeeping account of conceptual contents maintains the possibility of anyone or everyone getting things wrong, but he fails to provide any resources for making sense of our judgment that some discourses are more objective than others nor does it explain why we find this structural objectivity in the practices we do. It gives us nothing beyond what Rorty's notion of solidarity could already provide. We need to do better. If we can't sort practices in terms of objectivity, then the notion we've developed is just an idle cog. It won't allow us to diagnose failures in those practices or work toward bettering our chances at "unforced agreement."

PALM points us in the right direction, drawing our attention away from inferential norms and toward the norms of authority that give discursive practices their distinctive shapes. As we've seen, authority norms differ from practice to practice, so they may provide just the resources we need to distinguish objective from non-objective practices and to understand what makes objective practices useful for the purposes to which we put

them. By attending to these norms, we can build on Brandom's notion of objectivity to answer the problems we've raised.

To see the role that authority norms play in an account of objectivity, it will be useful to return to Brandom's account to see just where it goes wrong. Brandom's "objectivity proofs" show that nothing in his account forces a privileging of one perspective over others, but this leaves open the possibility that there are some instances in which one perspective is, in fact, privileged. Brandom notes this only in a footnote to the objectivity proofs in which he mentions that the recipe he provides will produce counterexamples to the universally quantified conditionals only if *nonascriptional* claims are used as examples (R. Brandom, 1994, p. 602n99), but why? What goes awry when we turn to ascriptions and why does this happen?

Can the distinction between "p" and "I claim that p" be maintained when we plug an ascriptional claim into NFPE?? Let "p" be "I believe that q."¹⁵ In this case NFPE gives us "If I claim that I believe that q, then I believe that q." Now, let's consider a claim that is incompatible with the consequent. "I do not believe that q," is clearly such a claim. Is this incompatible with the antecedent of the conditional? That is, could I be entitled to "I claim that I believe that q" even if I am committed to "I do not believe that q"? It should be clear that the answer is no. Commitment to "I do not believe that q," a claim that ascribes a commitment to me such that q is not a commitment of mine, precludes entitlement to ascribe to myself commitment to q. When I ascribe ascriptional commitments to myself, my ascription seems to settle the matter. In this case, NFPE does,

¹⁵ This is equivalent, in Brandom's idiom, to "I acknowledge (attribute to myself) commitment to q," but putting in terms of "believing that q" allows us to avoid some rather convoluted sentence structure in what follows.

in fact, hold. There's no daylight between what I take to be the case and what is the case regarding my espoused beliefs, so there's no way for me to be wrong about what they are.¹⁶

Now consider NFPI: If I believe that *q*, then I claim (or believe) that I believe that *q*. This time consider my ascribing the following to *S*: “*S* claims that I claim that I believe that *q*, and *S* claims that I do *not* believe that *q*.” Here I am attributing incompatible commitments to *S*, for it seems that *S* could not be entitled to both “Tom does *not* believe that *q*” and “Tom claims that he believes that *q*.” My claim that I believe that *q* settles the matter of whether or not I believe that *q*. It's just not clear what evidence *S* could have to the contrary that would be robust enough to override my own self-ascription.¹⁷ From my first-person perspective there just is no gap between “I believe that *q*” and “I believe (or claim) that I believe that *q*.” I cannot be mistaken about my own thoughts and perceptions, nor can I be ignorant of what they are. So long as *S* is committed to attributing to me acknowledgment of the claim that I am thinking that *q*, she could not be entitled to her commitment to the claim that I am not thinking that *q*.

¹⁶ This is consonant with the tradition in analytic epistemology of treating percepts, observation reports, occasion sentences, and the like as being indubitable. One just can't be wrong about how things look or appear to one and one can't be mistaken about what one takes to be the case, even if one can clearly be mistaken about what is the case. As we shall see, this is a product of the structure of the authority norms that govern the practice of making such reports.

¹⁷ I am making some pretty significant simplifying assumptions here. The first is that interlocutors are being *sincere* in their avowals, for, of course, *S* might be able to present evidence that I don't believe what I claim to believe and that I am trying to deceive by lying. The point is that we don't know what evidence could possibly stand to challenge what we take to be a sincere avowal of belief, whereas we do usually know what kind of evidence could stand to challenge an empirical claim, *especially* if we take it to be sincere. Behavioral evidence is sometimes relied upon to claim that someone doesn't *really* believe what they avow, but if the claimant continues to insist, even upon careful introspection and examination of her behavior, that it *is* what she *really* believes, it seems the avowal is authoritative. The only other response available seems to be to fall back to an argument about just what we mean by “believes.”

7. Rorty on the Mental

Recall that according to PALM, the authority norms of a practice divide into two species. The first of these determine the propriety of a speech act in terms of the relative social-normative standing of the speaker, e.g., whether she is in a position to give orders or make requests. The second determine the propriety of a speech act in terms of the epistemic authority required for it. The sense of epistemic authority intended here is not that which might be attributed to an expert in a field but rather the kind that can be secured by giving reasons in response to challenges or queries. The epistemic norms of a discursive practice are just those norms that govern the distribution of this kind of authority among practitioners. Since this authority is distributed by way of the practice of making claims, challenging claims, and responding to challenges, the relevant norms are those that determine which claims count as entitled, what might count as an appropriate challenge to a claim, and what is called for in response to a challenge.

We just saw that Brandom's recipe for constructing counterexamples to NFPE and NFPI fails when we turn to claims that ascribe beliefs, thoughts, perceptions, or other mental states to an individual. The reason it fails is because in these cases the individual's sincere avowal strikes us as immune from challenge, and now we have a framework within which we can capture this fact. The epistemic norms of the discursive practice of self-ascription of mental states like occurrent thoughts, beliefs, perceptions, and, as we'll see soon, preferences don't allow for (or at least severely limit the scope of) challenges to sincere avowals that the speaker is in one of these states. In this way, the epistemic norms of these practices vest authority over claims about one's mental states in the individual whose mental states they are. They privilege the individual's perspective.

It's important to grasp just what this privileging of perspective amounts to when conceived in terms of PALM. In his famous early paper, "Incorrigibility as the Mark of the Mental," Rorty used Sellars's "myth of Jones" to argue that what makes mental states what they are is the special way that we treat them in discourse, just as I claim here (Rorty, 1970, p. 417; Sellars, 1956). There is a "convention, [or]...linguistic practice, which dictates that first-person contemporaneous reports of such states are the last word on their existence and features" (Rorty, 1970, p. 414). As such, our reports on our own mental states are incorrigible insofar as "there is no assured way to go about correcting them if they should be in error" (Rorty, 1970, p. 417). This incorrigibility, though, is not explained by some metaphysical facts about mental states or our epistemic access to them. Rather, it is explained "'naturalistically'...in terms of the linguistic practices" we have adopted (Rorty, 1970, p. 416).

According to Sellars's myth, the genius Jones introduced his Rylean brothers and sisters to "thoughts" conceived of as "inner" states of an individual that help to explain certain kinds of behavior, and he introduced "sensations" as "'inner' states postulated to explain the occurrence of certain thoughts (e.g., the thought that there is a red triangle before me, when there isn't)" (Rorty, 1970, p. 411). These postulated theoretical entities, according to the myth, had "certain intrinsic features." Thoughts "were true or false, and were *about* things...They shared...the 'semantical features of sentences'" (Rorty, 1970, p. 411). Sensations, for their part, had intrinsic features like "being 'of red' and 'of a triangle'" (Rorty, 1970, p. 411). In this way, they shared logical space with their physical counterparts, but their pragmatics are decidedly different. For our purposes, what is important to recognize is that sensations and thoughts were introduced as theoretical

entities to serve an explanatory purpose. Their existence was inferred from the behavior of individuals to fill the explanatory needs of those with whom those individuals interacted. They were not objects of introspection and the owners of these thoughts and sensations were not taken to be authoritative reporters on their existence or content.

In Rorty's extension of the Sellarsian myth, Jones's successors discovered that the totality of Smith's behaviors could best be explained if they relied upon Smith's own reports when they conflicted with the behavioral evidence available to them. Over time and with further confirmation, this discovery gave rise to a convention that, when it came to Smith's thoughts and sensations, what Smith said went. This is Rorty's "naturalistic" explanation of incorrigibility: Jones's successors found it useful to adopt norms of discourse such that there is no accepted procedure for rationally coming to believe that Smith thinks that *not-p* when Smith avows that *p* (Rorty, 1970, p. 416).

From the perspective of PALM, we tell a similar story. Self-ascriptions of mental states are recognitive speech acts with agent-relative inputs and agent-neutral outputs. They lay claim to authority to commitments available only from the individual perspective and enter their contents into the public space of reasons. What is distinctive of discourse about mental states is that it privileges these recognitives, granting them unimpeachable epistemic authority. When someone makes a sincere avowal about what they are presently thinking, what they believe, how things look to them, what their preferences are, and so on, we take those avowals as the last word on the existence and features of these states by accepting that there is no way to issue legitimate challenges to these claims. They simply settle the matter. PALM explains the acceptance of norms that vest authority in the individual in this way by appeal to the function of this practice.

Much like Rorty, we might construct an account such that the greatest predictive success was to be had by taking individual speakers to be authoritative on such matters.

We can treat this as a model for looking at other discursive practices. Looking back to Brandom's original recipes for counterexamples to NFPE and NFPI, for example, we now see that he clearly had empirical claims in mind. Brandom's project is, centrally, an attempt to give an inferentialist account of empirical discourse, so when he turns to the "objectivity proofs" he is tacitly relying on the epistemic norms of empirical discourse. These norms are responsible for our judging that taking another to be committed to attributing to oneself commitment to "p" and to "I claim that not 'p'" we are not attributing incompatible commitments to them, for they make it the case that my individual perspective is not authoritative with regard to what is empirically the case.

Empirical discourse exhibits a default and query epistemic structure:

[C]laimants possess and are properly granted default epistemic assertional (doxastic) entitlement: entitlement that does not depend on the claimant's doing or having done any *specific evidential work* or possessing any *citable reasons* for his commitment. Default entitlement is *situationally* unearned and imposes no *standing* deference requirement. However, it is open to defeat by failure to answer a *contextually appropriate query* (M. Williams, 2015, pp. 263–264).

In principle, anyone can lay claim to entitlement to an ordinary empirical declarative; its normative inputs are agent-neutral. Such a claim to entitlement is granted by default. The agent need not have done any particular evidential work or drawn out any particular set of inference prior to having made the claim. She need only be counted as a responsible, reliable epistemic agent by her community, i.e., as a member of the discursive community. Such agents face a standing commitment to respond to contextually appropriate queries regarding the claims they make. Ordinary empirical declaratives, in this way, are not *free moves* even if entitlement to them is had by default. One must earn

her right to maintain the epistemic positions she stakes out by defending them against challenges and answering to queries.

We find a similar normative structure when we turn our attention to empirical observatives. Observatives, recall, are *recognitive* speech acts with agent-*relative* normative inputs and agent-neutral normative outputs that serve as language-entry transitions from first-personal perception of the passing scene, i.e., they register an agent's recognition of objects and properties in the world. Observatives are governed by the same default and query structure as other empirical claims. The pragmatic structure of observatives, though, introduces a complication. Observatives are, by their very nature, personal. No one else can be entitled to the observative that I utter because it gives expression to *my* first-personal uptake of features of the passing scene. The attention of others can be called to those features. They can rely on my uptake of those features in their own reasoning if they take me to be a reliable reporter. They cannot, however, have *my* experiences. When observatives are entered into the space of reasons, there are, as such, two ways they can be taken up, governed by two sets of epistemic norms. The first, and most common, is that they are treated as claims about what is the case. They are default entitled, but open to well-motivated challenge or query in just the way that any other empirical claim is. Observatives play a distinctive role in empirical discourse—they prevent frictionless spinning in the void—but they do not play a determinative role. They do not alone settle any empirical matter, for they can always be challenged and defeated. No observative is ultimately authoritative or immune for challenge.¹⁸

¹⁸ One might object that some observatives are, in fact, incorrigible, i.e., that we could not make sense of a challenge to some observatives when they are sincerely uttered. Consider the Moorean observative: "Here is one hand." Or the claim, by a known reliable reporter: "This is red." Embedded in appropriate contexts, these claims seem beyond doubt. There just is no possibility of error here. Does this threaten the objectivity

The second way these recognitives may be taken up does treat them as both distinctive and determinative, but *only* as reports of how things seem to the speaker rather than of how things are. As with other recognitives that give expression to my first-personal uptake of my own mental states, there is no legitimate way to challenge me with regard to what I take my perceptual experiences to be like. I not only have default entitlement to my reports on my own experiences, this default entitlement cannot be undermined. If I claim that x looks to be red, then, for me, it does, and, so long as I am counted as a competent user of the relevant concepts, there just is no way to challenge me regarding the qualities I claim my experiences to have. I am the ultimate authority about how things appear or seem to me.¹⁹

of these claims or the empirical discursive practices in which they are advanced? No, it does not follow that they lack objective purport or that the practices in which they are embedded are somehow non-objective. To see that this is the case, notice the role that context is playing here. The context in which the speech act is delivered is what is doing the work of ruling out the possibility of error in the Moorean and “red” cases, for there is no way that the challenger could justifiably raise a *real* possibility of error. There is nothing, that is, to which she could point that makes it seem reasonable to reassess the grounds of one’s belief. We can imagine these speech acts, however, in strange contexts in which the real possibility of error does exist. For example, if one has just ingested some hallucinogenic substance then a real possibility of error has been introduced. In such contexts, a real possibility of error does exist, and challenges based on that possibility are warranted. In most contexts we do not even know what it would mean to doubt “Here is one hand,” but this does not entail that the speaker’s authority is unimpeachable, only that the epistemic context is such that no real possibilities of error are to be found. What we find, then, are two ways in which a recognitive speech act might be immune from challenge or, perhaps more perspicuously, two different explanations for their immunity. The first explanation is that some recognitive speech acts are immune from challenge in virtue of the authority norms of the discursive practices in which they are embedded. We have found it useful for one reason or another to vest authority in an individual’s perspective in such practices. The second explanation is that some recognitive speech acts—some observatives, in particular—are immune from challenge because nothing in the epistemic, discursive context could warrant a challenge, even if we can ‘imagine’ some contexts in which challenges may be warranted. Of course, exactly how we should understand the epistemic, discursive context of a speech act is a tricky business. I am inclined to think that the discursive frame that determines such context rules out certain possibilities of error from the start and that imaginative scenarios that are often used by the skeptic to raise error possibilities play on modal ambiguities. I am indebted to Michael Williams for pushing me to think more carefully about this point. Cf. Wittgenstein’s discussion of Moore’s argument, (Wittgenstein, 1972).

¹⁹ When we seek to have our observatives taken up in this way or when they have been challenged and we want to retreat to safer ground, we use words like “looks,” “seems,” or “appears” to hedge our commitments (R. Brandom, 2015, pp. 105–109; Devries & Triplett, 2000, Chapter 3; Sellars, 1956, sec. 6). In claiming that x *looks* red, I am withholding my full endorsement of the claim “x is red.” The latter is open to challenge on a variety of grounds, but regarding the former, I am incorrigible. This is not because the former has some special metaphysical status or gives me some special access to the passing scene. Rather,

The flip side of these epistemic norms imposes symmetrical demands on one's interlocutors. Those who challenge ordinary empirical declaratives also must earn the right to maintain the positions they stake out. Challenges, like claims, are default entitled, for they are themselves just ordinary empirical claims in the standard cases. Also, like ordinary empirical declaratives, challenges are not *free moves*, for the challenger herself can face a challenge to demonstrate that the challenge she has issued really is contextually appropriate. And on this pattern goes, for responses to challenges can themselves face contextually appropriate challenges and queries, and no one is in a position to settle a matter once and for all. Even the most entrenched empirical position can, in the right contexts and with the right evidence, face a well-motivated challenge.

Empirical discourse does not vest authority in any individual or group of individuals but relies on a process of the sorting out of evidential and inferential claims to settle disagreements. No settled position is immune from further examination. This is the pattern of reasoning generated by the epistemic normative structure that is constitutive of objectivity in the sense that I think pragmatists should adopt. This notion of objectivity is wholly anti-Representationalist. Rather than conceiving of objectivity in terms of the objects of our discourse (some non-human Other) imposing their authority on us and our claims being answerable to them for their correctness or the aim of objective discourse as getting closer and closer to what's really there, we understand objectivity in terms of a structure of epistemic norms that ensures that no perspective is privileged in advance over any other. Like the Representationalist conception, this notion of objectivity makes

this is an artefact of the fact that we have found it useful to adopt a practice of treating each other as authoritative with regards to how things seem to us even though we do not grant each other such authority with regards to how things are. Even so, how things seem to one is often taken by others as good evidence about how things are. This is the distinctive role of observatives in ordinary empirical discourse.

sense of the possibility that anyone or everyone could possibly be mistaken, but it does so by understanding this in terms of the possibility of new challenges arising from new perspectives as a result of an expanding discursive community, new experiences, or new technologies. We can always improve upon our epistemic circumstances, and, in doing so, we get better and better at navigating our physical and social environments.

This pragmatist sense of objectivity also does real work. It provides a framework for sorting between practices that are more-or-less objective and those that are more-or-less subjective in terms of the authority structures of those practices. When we identify a practice that is more-or-less objective, we can then ask why it exhibits this structure. Why is it the case that we have found it useful to adopt these epistemic norms? What function do they serve in the lives our communities? Having an answer to this question, we can then assess these practices: Are they worth maintaining or are they no longer serving the function for which they evolved? If they are worth hanging onto are they functioning as well as they could be? If not, why not? What could we do to revise the norms of these practices or the ways in which we engage in them to better serve these functions? One thing we might discover is that objectivity in this sense is in tension with relativism in the sense of ethnocentricity. Objectivity is served by increasing the scope of our discursive communities, by bringing to bear more and more new perspectives to challenge the claims we accept. But there are forces that push against this expansion of our *ethnos*. Fear of others, lack of safety and stability, scarcity of resources, external threats, or the perception of any of these things brought on by demagoguery all stand in the way of improving our epistemic circumstances, for example. The response should be to push back against these forces, to aim to expand our *ethnos* and, so, make our practices more

objective. In this way, this pragmatist notion of objectivity builds on Rorty's notion of solidarity, but it expands on it, as well, for it tells us in which practices we ought to seek such solidarity and in which it is not a worthwhile aim (Rorty, 1991e, p. 23).

Finally, this pragmatist notion of objectivity points toward the sorts of reforms that might make a malfunctioning objective discursive practice function more usefully. We ought to pay close attention to the epistemic norms of the practice, how they are enforced, when they are violated, when we are lax in their enforcement, and so on. As noted earlier, Longino provides a useful framework for thinking along these lines when she identifies the social structures of scientific discourse that account for its democratic openness. We ought to examine how social, political, or economic authority affects whether or not one's claims are taken to be open to challenge, whether forums for disseminating information—journals, for example—encourage close examination and evaluation of claims others have made or primarily reward novelty, and whether standards of assessment are publicly available and broadly shared.

8. Moral Objectivity

PALM contains the resources for a pragmatist friendly account of objectivity, but what of the PALM for moral discourse? Is moral discourse objective in the sense just developed? To get at this question, the first thing we need to do is home in on the speech acts where objective discourses differ most dramatically from more or less subjective discursive practices. In examining empirical observatives, we have just seen that they are subject to two dramatically different kinds of uptake. In the first instance, these recognitives are treated as unimpeachable reports on how things seem to the speaker from her own first-personal perspective. In the second, they are taken up as defeasible claims

about how things are that are subject to the default and query structure of empirical discourse. If taken up in the former way, these recognitives are treated as part of our subjective discourse about mental states. If the latter, they are treated as part of our objective discourse about how things are in our shared world. This distinction, I suggest, is central to understanding the objectivity of moral discourse.

To see how, let's consider two neighboring discursive practices. The first, which I'll call *taste* discourse, is the practice in which we decide what we should do given our tastes and preferences. Should I order the IPA or the pilsner? Should I buy you the red tie or the blue one for your birthday? The second is *prudential* discourse. Like taste discourse, this is a practical discursive practice in which we decide how to act, but rather than basing our decisions simply on what one expresses as their tastes or preferences, we consider what is in the individual's best interests. Both practical discourses are paradigmatically comprised of three types of speech acts. The first are *prescriptives*. These are speech acts with agent-neutral inputs and agent-relative outputs, i.e., their authority derives from the space of reasons, but they have particular practical significance for certain individuals. Prescriptives, in all practices in which they occur, exhibit a default and query structure. No one's prescriptives are authoritative; this is just what it is for them to have agent-neutral inputs. The second are *declaratives*. Practical discursive practices must contain claims to undertake commitments with regard to how things are both with normative matters and with empirical matters. Declaratives, with their agent-neutral inputs and outputs, are the speech acts we typically use to do this, and, as always, they exhibit a default and query structure. No one's declarative is ultimately authoritative. The third kind of speech act are *recognitives*. These comprise both ordinary

obeservatives, speech acts that express recognition of the tastes and preferences of others on the basis of observed behavior, and recognitives that express one's own tastes or preferences.

Let's consider an example of taste discourse to get a feel for how these different speech acts function in the practice. Suppose Jill and Jackson are trying to decide which beer to order for me before I arrive at the bar. They consult each other regarding their knowledge of my tastes issuing recognitives about my past behavior. They advance claims about what I like or dislike, and each of these claims is default entitled but open to contextually well-motivated challenges. Jackson claims that I particularly like bitter beers, as he has tasted the beers I've ordered in the past and that's the main thing he's gleaned from them. Jill, who is more of a beer connoisseur, demurs noting Jackson's unrefined palate and her knowledge that I prefer a brew with a citrusy nose. At some point, whatever conclusion they arrive at about my tastes is translated into a beer order by consulting the tasting notes on the menu.

Now reconsider this scenario in my presence. The first thing to notice is that I can immediately settle the question about my preferences. Once I give them expression, no one is in a position to challenge me with regard to what they are. My default entitlement to this recognitive is absolute. I could, however, make a mistake about which brew will best fulfill those preferences. If I express an interest in ordering beer A, Jill might argue that B would be a better choice, as it's better balanced with a more pronounced citrus nose. We might go on to debate the merits of the various beers on the list until we come to some agreement about which is most likely to fit my tastes. Even once we reach such an agreement, the server might interject to challenge our consensus from her perspective.

As in empirical discourse, no conclusion is ever final, as there is always the potential of challenge from a new perspective the authority of which is equal to the others. No perspective or set of perspectives is privileged in advance in this conversation. Jill or the server could turn out to be right and I could turn out to be wrong. Though what is in question is how to enact my preferences in a decision about what to order, which *prescriptive* I ought to endorse, i.e., which is binding on my actions, is not determined only by what I take to be the case.

This, however, does not mean that taste discourse is fully objective in the way that, for example, empirical discourse is. To see why it is helpful to consider an idealization of the above scenario. Suppose that I am an *ideally informed* agent, i.e., that I know all the empirical facts in any way relevant to the question at hand. I have a clear sense of the flavor profile of each of the beers on the menu. In this case, my preferences are fully determinative of what I ought to do, i.e., of what taste prescriptives are true or are binding on me, and my expression of those preferences in my decision about what to do—which beer to order—is authoritative. Given the context, the norms of the practice grant authority to me. If there is no challenge available on empirical matters, there just is no contextually well-motivated challenge possible to my entitlement to the prescriptive “I (Tom) ought to order beer A.” What this shows is that what lent the air of objectivity to this case is the way in which empirical matters bear on it not the norms of taste or preference discourse. Once we bracket empirical questions, the norms of taste discourse come to the fore. We find, then, that in matters of what one ought to do given one’s taste, the agent whose tastes are in question is granted a special authority with regard to which prescriptives are true. Their truth follows not from what reasons can be given in their

support but from her own cognitive expression of her tastes about which she is the ultimate authority. These cognitives serve both a distinctive and determinative role in taste discourse.

I think that the right thing to say about taste discourse is that it is neither wholly subjective nor wholly objective, but that, given the way that cognitives about tastes and preferences are treated in the practice, it falls pretty far toward the subjective end of a spectrum between the two. It is a practical discursive practice that vests a great deal of authority in the individual who is targeted by the prescriptives issued in the practice. Compare this to prudential discourse. Whether I ought to put \$200 per month or \$500 per month in my retirement account is a question that depends on many factors: my income, my family circumstances, my medical needs, my zip code, my degree of future discounting, and so on. Many of these require the settling of empirical matters, but some of them have to do with what it is in my best interest to do. This is something about which I *can* be mistaken according to the norms of prudential discourse. Some people have a very hard time seeing what is in their own best interest; we think that people act against their interests quite often; and we tend to think that it requires a good deal of careful reflection to grasp what one's long-term interests are. This reflection is often clouded by various cognitive biases.²⁰ The cognitives that give expression to one's own interests are not treated as authoritative. One can be challenged on what one takes to be in her interests. Such challenges and responses to them, however, must eventually be premised upon what one's preferences are, and this is something about which one is authoritative. If, upon reflection, I continue to have a strong preference for present

²⁰ See, for example, (Kahneman, 2013; Thaler & Sunstein, 2009).

consumption and, coherent with this, I prefer to live my golden years in squalor, then, once this is established, there's just no room left to challenge my judgment that I ought to put less into retirement no matter how ill-advised you take these preferences to be.

We can see this if we consider an ideally coherent agent. Though we might think that what they take to be a prudent course of action is incredibly immoral, distasteful, hurtful, or just plain weird, we cannot challenge such a judgment without violating the assumption of ideal coherence. If their preferences are well-ordered and coherent²¹ and they are fully empirically informed, they just couldn't be wrong about what they prudentially ought to do.²² Prudential discourse, then, falls a bit further toward the objective end of the spectrum, but is far from being a fully objective discursive practice. An individual's judgments about her interests are not determinative of what she ought to do, but her recognitives about her tastes turn out to be what matters hinge on in the final analysis.

To place moral discourse on this spectrum, it's now clear that we need to determine how recognitives are taken up in the practice. Moral discourse trades in at least two important varieties of recognitives: *moral observatives* and what I'll call *value recognitives* ("valuings," for short). Let's take each in turn.

Moral observatives are expressions of one's first-personal uptake of the moral features of the passing the scene. They register what is essentially a *prima facie* judgment of what is morally salient *from an individual perspective* (agent-relative inputs). They are claims like, "The cost of that life-saving drug is unjust," "Clearly, the thing to do here is

²¹ In the sense canonically defined by von Neumann and Morgenstern, (Neumann & Morgenstern, 1944).

²² For a very informative discussion of such agents and their role in our reflection on these matters, see (Street, 2009).

stop the drug trial,” or “That crime was grossly misogynistic.” Presumably, a properly trained-up moral agent learns to see moral states of affairs in much the same way that a properly trained-up physicist can see *mu mesons* in a cloud chamber or an experienced birder can tell that that’s a house finch rather than a purple finch at a glance. One can see instances of injustice, perceive that someone has been harmed, and see when a state of affairs is unfair or some outcome is racially discriminatory.²³ This just means that one has been trained into a reliable disposition to respond differentially to observational stimuli with the deployment of moral concepts. These moral observatives, like their empirical counterparts, are taken up as claims about what is morally the case. Their outputs are agent-neutral, seeking uptake from whomever might be in a position to give it without respect for their social-normative standing. As such, they are not immune to challenges. An interlocutor might query the claimant with regard to the concepts she has deployed, with regard to whether her partiality for one of the involved parties has affected her read on the moral features, or whether she is really a reliable reporter on such matters. “Why do you think it was misogynistic? The shooter seems like a troubled individual suffering a psychotic break, but I don’t think he hates women.” The claimant can, of course, retreat to a claim about how things seem to her—“Well, it seemed misogynistic *to me*”—but, as in the empirical case, this retreat to ‘looks’ talk constitutes a failure to take a position about how things are (morally) in the world. If her claim is to be treated as a claim about how things morally are—and, so, relevant to what one ought morally to do—then she must be prepared to defend it against challenges. “He may have been troubled, but his

²³ This is obvious in the widely accepted practice of testing proposed theories of morality and moral principles against intuitions about real and imagined cases, i.e., the process of reflective equilibrium (Rawls, 1971).

action is part of a larger pattern of behavior that serves to reinforce patriarchal social norms by using violence to police the behavior of women. Given his targets and the manifesto he wrote, this is clearly an act of misogyny.”²⁴ Moral observatives, like their empirical counterparts, are granted default entitlement but are open to query and challenge. They are never decisive on their own about what is morally the case, but they can provide strong defeasible reasons for thinking that things are thus-and-so. So far, moral discourse is tending toward the fully objective end of the spectrum, as no individual perspective is privileged in the case of moral observatives.

Let’s turn next to valuings. Valuings give expression to an individual’s first-personal take on what is good, valuable, worthwhile, desirable, preferable, and so on. Rather than expressing uptake of some aspect of the passing scene, they express an individual’s commitments with regard to what is good, preferable, and so on. They range widely from “Chocolate ice cream is good” to “Genocide is horrific,” and they play various roles in moral discourse. Expressions of value, desire, etc., are relevant, for example, in determining the pain or pleasure that a proposed action might cause for others and so may be relevant to determining the moral worth of the action. Such expressions might also be authoritative with respect to the value of special relationships—say between partners or between parent and child—and, insofar as we want to leave room for the possible moral relevance of such relationships, we may recognize that an individual’s own estimation of the value of that relationship is the best gauge of its actual value. Mostly, though, valuings give expression to our judgments about what’s valuable. We might say, for example, that autonomy in health care choices is

²⁴ For an account of misogyny as the enforcement arm of patriarchy, see (Manne, 2018).

valuable, that rent seeking is bad, or that racial discrimination is detestable. In one way, these expressions are just register an individual's preferences. If I say, for example, "I value autonomy," this might be taken up by others in the same way as "That looks red." They might respond by saying, "Yes, Tom, we know you do. That's very nice" and then move on to more serious matters. In moral discourse, however, these expressions are subject to being taken up not simply as expressions of the speaker's mental states but as claims about how things are in a shared world. Such claims can be disagreed with, argued over, challenged, examined, revised, and so on, and moral discourse is the space in which we do that. Something is *valued* if someone values it, but something is *morally valuable* only if its value can be defended to others by giving reasons that could be taken up by anyone, that are agent-neutrally acceptable.

One place that this epistemic structure of moral discourse comes out quite clearly is in the consideration of excuses, exculpations, and justifications for action (or inaction).²⁵ Suppose I direct a second-personal assessment at Sam: "It was wrong for you to take those blocks from your friend. You already had more than he did, and it hurt his feelings." Sam, being a five-year-old, is apt to respond by saying earnestly that he needed them to finish building his bridge, i.e., that he valued them as necessary for the completion of his project. Of course, none of us will accept this attempt by Sam to evade the moral appraisal, for the fact that they were necessary for his bridge does not bear on whether he was wrong to take the blocks. If we are in a pedagogical frame of mind, we might remind Sam that what is right and what he wants are two different things and that just doing what he wants can lead him to do things that are not right. We might also ask

²⁵ Excuses and exculpations are often fertile ground for philosophical reflection, as Austin taught us long ago (Austin, 1957).

him to reflect on how he would feel if the blocks were taken from him and on whether it would be ok for his playmate to take them for the same reason that he gave. If we are not feeling so pedagogical, we might just reprimand Sam and give the blocks back to the other child. Either way, reasons that boil down to “But I want it!” are roundly rejected as good moral reasons.

Similarly, if we consider again an ideally coherent agent—let’s take Ideally Coherent Caligula²⁶—we can elicit the same intuitions. Caligula takes special pleasure in imposing maximal suffering on others. So much so, in fact, that it is his primary aim in life.²⁷ Being ideally coherent, Caligula’s valuing of maximizing the suffering of others is entirely logically and instrumentally consistent with all of his other values, and he is not mistaken about any of the non-normative facts, including, for example, how others will treat him in response to his seeking to act on this value. Nonetheless, Caligula’s claim that it is what he has most reason to do, what is in his best interest, or what is prudent for him have no bearing at all on the truth of the moral prescriptive, “Caligula ought not to cause others to suffer.” We don’t—and shouldn’t—care if that’s what he values doing, only whether there could possibly be good reasons “grounded in something independent of [Caligula’s] stance” that could be given in defense of his intended actions (Darwall, 2006, pp. 56–57). If he were to say, “Causing suffering is valuable,” we would demand an argument to the effect that it ought to be valued.

There is, then, a symmetry between first-personal cognitives and empirical observatives. In both cases, one is taken to be authoritative with regard to the content that

²⁶ See, for example, (Gibbard, 1999; Street, 2009).

²⁷ I hope it’s clear that this is a caricature. The real Caligula, though he was certainly brutal, was also subject to some of history’s most damning propagandizing (Barrett, 1989).

is given uptake: how things seem in the empirical case and what one values in the evaluative case. In neither case, though, is the recognitive expression of this uptake taken up as authoritative with respect to what is the case, i.e., what is empirically or evaluatively true. Instead, these recognitives are taken up as inputs to the discursive practice as defeasible evidence for something's being the case. "I value the freedom to choose my own path in life" is defeasible evidence of the ability to choose for oneself being valuable just as "Lo! A rabbit!" is defeasible evidence that there is a rabbit in the vicinity.

One important upshot of this analysis is that we now see the entry point for new perspectives to come to bear in moral discourse. The valuing of individual practitioners are not inert in the practice. Someone who has had experiences that others have not, who has lived a life subjugated by the accepted social norms, or who comes from a different culture may very well perceive the moral landscape differently than those who were enculturated into our moral discursive practices and who share our perspectives. Moral discourse provides for the inclusion of these varied perspectives by giving uptake to speech acts of valuing, but it requires reasons to be given that any practitioner might be able to accept. This, of course, is not to say that such reasons are likely to be accepted by others. This is the topic of the next chapter.

On the account of objectivity developed in this chapter, moral discourse counts as a fully objective discursive practice. It makes sense of the idea that anyone or everyone could possibly be wrong, and it does not privilege any perspective over any other prior to engaging in the practice. All claims are open to challenge and must be backed by reasons, and no individual, nor the community as a whole, is counted as finally authoritative. Any

moral claim, no matter how broadly accepted or how long held, could always be challenged and, possibly, revised.

9. Religion and Morality

One might object that I have sketched an overly secular picture of moral discourse. It sounds vaguely like the discursive practice in which we engage, but our moral practices are deeply informed by tradition and religious worldviews. I've left no room for these in my account, yet, if we look around, we'd be hard-pressed to find a moral debate that does not have some religious dimensions, we'd be equally hard pressed to explain why we find various issues morally important if it weren't for the influence of religion, and we'd be ignoring what most people considered morality to be for a large swath of human history. The problem is that religion doesn't play the game I've been examining. The speech acts that play the role of cognitives in religious-moral discourse are not cognitives at all but commands. They have agent relative inputs and agent-neutral outputs like cognitives, but the authority they claim is social-normative rather than epistemic. The moral edicts of the deity are final, absolute, and unchallengeable, for the deity is infallible. They settle moral questions once and for all and must be obeyed. If these edicts are really part of moral discourse, then moral discourse would not be objective in the way I have been claiming, for the deity's perspective would be privileged in advance over every other.

It's controversial exactly how we should understand such edicts. On the one hand, they may be the divine distillation of an absolute moral truth that is somehow written into the very fabric of the universe. The idea would be that god— meant in a sort of all-inclusive way to capture whatever deity or deities are relevant—is constrained in such a

way that god can only issue edicts that comport with the way things must be. This is analogous to the kind of understanding of god present, for example, in Leibniz's *Theodicy* according to which god's will is constrained from preventing some evils by certain requirements derived from god's ends.²⁸ On this line, human beings could hope to fallibly interpret the moral law, but god (and perhaps god's earthly representatives) is (are) its only infallible expositor. As such, when god issues moral edicts—in this case, god's infallible uptake of the moral law—god cannot be challenged. On the other hand, and perhaps more popularly, we might understand god's moral edicts as an expression of god's will. On this version, they are not recognitives at all but simply *divine moral commands*. Such commands settle the question of what one is morally required to do. They are not challengeable, and god's entitlement cannot be impugned. According to either of these possibilities, though, there is an authoritative moral perspective.

Though there is not space for a careful and complete treatment here, I want to suggest that religious-moral discourse, i.e., the practice in which we determine what one ought to do with reference to a particular religious system or worldview, is deeply intertwined with yet distinct from moral discourse as it has been examined herein.

Moral practice and moral discourse—like all social practices and, indeed, all technologies—are continuously evolving in response to changing physical and social environments. In its earliest stages of development, moral practice and discourse was likely a purely public affair. Individuals conformed to moral rules when others were watching but were perfectly willing to ignore them when they thought no one was looking (Kitcher, 2011, p. 111). This haphazard moral practice likely conferred

²⁸ G. W. Leibniz, *Leibniz: Philosophical Essays*, trans. Roger Ariew and Daniel Garber, 1st edition (Indianapolis: Hackett Publishing Company, 1989).

competitive advantages on cultural groups in which it developed by helping to facilitate greater social coordination, yet “[a]n ability to achieve conformity across a *broader* range of contexts would yield an *extra* edge in cultural competition” (Kitcher, 2011, p. 112). If individuals only refrain from defecting from cooperative foraging, for example, when they take themselves to be under observation, then there is the threat of a large free-rider problem and a loss of productivity for the group. It would be best if this could be remedied in some way. Furthermore, as groups grow in numbers, it becomes more difficult to keep all members under observation. This creates even greater pressure toward developing some means for enforcing public moral rules even when no one is watching.

One possible technique for enforcement of public rules in private moments is the instillation of a form of conscience. The guilt one feels for cheating one’s compatriots creates pressure toward conformity even if the likelihood of external sanctions is minimal. This, however, is a rather complex psychological mechanism that took a long time to evolve. It seems that instead many early human cultures hit on a technique that merely extends the already extant moral psychology—in particular, the tendency to follow the rules when being monitored by others—by appeal to unseen (and unseeable) enforcers. To the extent that individuals convince each other and, more importantly, their offspring that an unseen someone is always watching, fear of punishment (in the form of bad luck or an unpleasant hereafter) can serve to increase their tendency toward conformity to the moral rules with the effect of increasing their evolutionary fitness. This is a pattern that, as Philip Kitcher notes, repeats across a variety of extant cultural descendants of these earlier groups. “Western monotheisms [leverage]...an omniscient

deity who observes all, who judges, and who punishes lapses from commandments” (Kitcher, 2011, p. 112). In other religious traditions, “ancestors continue to observe the actions of the descendants and to retract their favors if the commands are broken” and “an ‘all-father’ ... ‘from his residence in the sky watches the actions of men [and] ‘is very angry when they do things they ought not to do, as when they eat forbidden food’” (Kitcher, 2011, pp. 112–113).²⁹

Once this idea of an unseen enforcer is firmly in place within a social group, it is easy to see how its powers expand. Rather than moral demands emanating from other individuals and the needs of the social group, commands can be attributed to the “all-father” and identified with its wishes. Groups with privileged access to the unseen enforcer or whose interpretations of its wishes are privileged—elders, priests, or other members of the clergy, for example—are now in a position to leverage their new found social capital to enforce existing moral rules as well as new ones they might devise (and which may be to their own benefit). Of course, such privileged groups would have existed prior to this development in social technology, but their ability to shape the behavior of others is increased with their access to these new tools, and with the growth of surpluses and accompanying needs for social control and communal defense that arose with the advent of agriculture, there is an evolutionary opportunity for social groups that can instill these beliefs in their members.³⁰

I am not claiming that this is a negative development in the history of moral practice and discourse. On the contrary, the evolved entanglement of religious and moral

²⁹ Kitcher is citing (Westermarck, 1908, p. 671).

³⁰ Compare, for example, Gellner’s account of the cognitive and technological transitions that accompanied the agricultural revolution, (Gellner, 1990). Also see, (Henrich, 2017).

practice is likely responsible for much of the successful social coordination that has allowed humankind to expand rapidly across the globe, develop complex civilizations, suppress many forms of interpersonal violence, extend our lifespans, and so on. It is important, though, to notice two things. First, religious-moral practice and discourse likely evolved in a context of already extant moral practices by leveraging the psychological mechanisms for greater social control in new contexts brought about by the development of agriculture. Second, religious-moral discourse represents a break from what began as a practice with broadly distributed authority. Rebounding from the egalitarianism that represents the split of early human meat-sharers from the evolutionary ancestors we share with chimpanzees, we see a return of some degree of dominance and authoritarianism with the advent of organized, geographically dispersed salvation religions.

The privileged position of the unseen enforcer allows it to “speak” (through the proxies of sacred texts, elders, and priests) with an authority that others do not have. Its moral commands settle once and for all what is required, and its claims are not open to challenge. Yet, its claims (and claimed authority) have *in fact* been challenged, and such challenges have, at least since the Enlightenment, been couched in moral terms. Moral language is used to enforce extant moral codes but also to challenge those very codes and practices in light of new perspectives. It is this dynamic that provides evidence that the original egalitarianism of moral discourse remains one of its core features, for it shows that though the mechanisms of morality have been leveraged toward authoritarian ends through religious-moral discourse, the underlying critical practice has remained an important part of our moral repertoire. This becomes especially clear in cases in which

the elders and priests (the powers that be) have tried to use morality to enforce practices that are exclusionary and oppressive. The oppressed, in such cases, have often found it useful to claim a moral voice of their own in order to challenge the claims of the oppressors, i.e., they have used moral discourse as part of a package of tools to challenge their social mistreatment. Such cases are instances in which we see moral discourse fulfilling its telos of facilitating ongoing social coordination while mitigating the instabilities inherent in the interaction between our social needs and our individual psychologies. The oppressed are holding the moral claims of the oppressors to their own discursive standards, demanding that they be justified *to them* and arguing that they cannot be so justified. In such cases, a bare claim to moral authority is useless. The moral commands of the deity are brought into the space of reasons and treated as challengeable in the very same ways as are the moral claims of others.

This dynamic demonstrates that the central practices of moral discourse continue to operate by their evolved structure of epistemic norms according to which authority is not vested in any individual or group perspective. Moral assertions, observatives, and recognitives are all treated according to this default and query structure, and no individual's perspective is treated as privileged in advance of inquiry. Moral discourse, in this way, mirrors the authority structure of ordinary empirical discourse and is objective in at least one central sense in which it is.

10. Conclusion

The sense of objectivity in which both moral and empirical discourse partake is not merely opposed to subjectivity but to dogmatism, for no claim stands absolute. No matter its pedigree or the reasons that have been mustered in its support in the past, every

claim is open to new challenges from new perspectives. This is just what it means for each perspective to have authority but none to be privileged. In empirical discourse this plays out as we devise new experiments and develop new technologies that raise problems for long held beliefs. The experimenters, the inventors, the users of these new methods and devices enter claims into the space of reasons from perspectives that were previously unavailable to us. These claims challenge old beliefs, pushing us to forever revise our understanding of the world that they purport to be about. In moral discourse, new perspectives that challenge old ways of thinking also arise in a variety of ways. Our values shift as our physical and social environments change. Voices that were once silenced—those of peoples who have been oppressed or marginalized by systems of patriarchy, racism, and capitalism, for example—find a way to be heard. Cultures intertwine bringing new patterns of thought and value to bear on one another. In these and many other ways, new perspectives enter into the space of reasons challenging us to forever revisit and re-examine our moral beliefs, to justify them in light of new challenges, or to revise or jettison them. The sense of objectivity I have outlined here is a sense tied to a project of inquiry, not an end. Empirical discourse evolves as we discover new methods and new technology. Moral discourse evolves as we face new challenges and hear once silenced voices. Neither of these projects is one that is likely to come to anything resembling completion.

It is important not to be too Pollyannaish about the prospects of moral objectivity. I have argued that moral discourse aims to incorporate new perspectives, new valuing, and new voices in order to bring new tests to bear on the moral commitments of a community, but there are two problems with this blanket characterization. The first is that

most of us are unlikely to be moved by mere reasons. We will ignore them, fail to fully integrate them, or twist them beyond recognition in order to hold our core commitments unscathed. Moral discourse can only function when we are engaged in ways that make us responsive to moral reasons, when we are working toward shared goals and have some shared commitments. Moral discourse, however, is not the only tool of persuasion we have. Moral change occurs when individuals come to see that they are not actually living some of their core commitments, when their identity is challenged by this recognition, or when they come to take up the values of someone close to them. In these cases, moral discourse is along for the ride; it's there to help us make sense of and justify the shift in values as it happens or, sometimes, after the fact. It is the rare case where it is actually in the driver's seat.

The second problem with this characterization is that our moral discursive communities come with built in immune systems. We are reluctant to let outsiders have a voice. We exclude perspectives, even when we are engaged in shared projects with the individuals whose perspectives they are. If we can manage coordination without moral alignment, we will, and, under many conditions, we will shrink the moral discursive community. There's an element of moral relativism in our practices with which the account of this chapter has not grappled. That's the task of Chapter 6.

Chapter 6: Discursive Communities, Trust, and the Norms of Standing¹

1. Introduction

There is a centuries old practice in western Nepal called chhaupadi. During menstruation and after giving birth women are excluded from communal life. They are thought to be “impure,” bearing a curse, and, as such, a danger to their family and their community. If they touch a tree, it will no longer bear fruit. If they touch a man, he will become ill. If they consume dairy products, the cow will no longer produce milk. To prevent the spread of this curse, menstruating women are sequestered in tool sheds, cow sheds, or specially constructed menstruation huts (*chhau goth*) for the duration of their periods, and new mothers are confined with their newborns for ten to fourteen days after birth. They are provided with food and water, but their accommodations are often poorly constructed, leaving women and girls exposed to the elements and to dangers like snakes and scorpions. This practice is enforced on around nineteen per cent of women age fifteen to forty-nine in Nepal even though it has been illegal for over a decade, and, each year, a number of women and girls die from exposure, asphyxiation, dehydration, snake bites, and other causes (Amatya, Ghimire, Callahan, Baral, & Poudel, 2018; Mijar, 2017; Sharma & Schultz, 2019).

¹ Sections 2-7 draw extensively from my “Trust, Communities, and the Standing to Hold Accountable,” Copyright © 2017 Johns Hopkins University Press. This article first appeared in *The Kennedy Institute of Ethics Journal*, Volume 27, Issue 2 Supplement, June, 2017, pages E1-E22.

Most readers, I assume, will agree with my judgment that the practice of chhaupadi is a moral abomination. It is clearly wrong to treat women whose bodies are undergoing an entirely normal and necessary biological process as “impure,” and it is horrific that a number of women die each year secluded in their huts, sometimes suffocating in the fumes of fires they build to stay warm on cold winter nights. Women are and ought to be full members of our moral community with all of the rights, privileges, and duties that this entails. They ought not to be sequestered, shunned, or ill-treated in other ways simply because their bodies differ from those of men. Moreover, I think that this position is fully justifiable and, indeed, true and that it applies universally. No matter the contingencies of environment, religion, and culture, no woman ought to be treated in this way. To enforce this practice and to turn a blind eye to it are both moral wrongs.

But, of course, I would think this from my perspective as a denizen of a Global North, post-industrial society, where women’s rights have been the object of struggle for well over a century. It’s deeply ingrained in my culture and my own moral sensibilities that differences in gender identification or biological sex do not justify differential moral treatment. It’s a judgment shared nearly universally by my peers, and, of course, I can justify this claim to them since we share so much in the way of moral judgment. If, as the too often repeated quotation from Rorty has it, “truth is what your contemporaries let you get away with saying,” then what I say is even true, at least for us (Rorty, 1979, p. 176). But is it true for the Nepalese Hindus who enforce chhaupadi on women and girls? Would they let one another get away with saying it? If they would not, then what are we to make of our judgment of them? Is it in some way illegitimate? Does it apply standards

of conduct to them that are only meant for us? Perhaps our moral judgments must be relativized to this or that moral system or to this or that culture or individual or justificatory framework.

If that's the case, then maybe my judgment of chhaupadi is not legitimate. Perhaps they are beyond the scope of my moral assessment. If they inhabit a different framework, a different "grid" of justification, a different form of life, then perhaps it's not my place to judge them. After all, if they really do inhabit a different justificatory framework, then nothing that I say would be likely to convince them that this practice is wrong. My moral concepts, it seems, just wouldn't get a grip in their framework, and theirs would be alien in mine. They would continue thinking that what they do is necessary to protect their community while I would argue endlessly that the real harm being done is to the women and girls of their community. We'd face an unbridgeable gulf in our moral outlooks, but we have no external point of view from which to judge one or the other correct. Thus, any attempt to impose my view of things on them must be mere hubris, an act of moral hegemony.

Most of us, I think, feel the pull of intuitions in both directions. Most of us are what Rorty calls "wet" liberals wanting to treat moral claims as objective and universal yet finding something compelling in the thought that our cultures, none of which is more valid than the next, inescapably shape our moral perspectives (Rorty, 1991c, p. 203). Since none of us can extricate ourselves from the grip of culture, no one can help but make judgments about inhabitants of other cultures from one's own limited perspective. But why should those judgments count for anything? Why isn't the appropriate response from those being judged to tell those doing the judging to buzz off?

PALM has the resources to make sense of these competing intuitions. It shows us very clearly why and how moral discourse is both objective yet, in a particular way, relativistic. The key insight is the recognition that there are two kinds of normative inputs for moral claims, whether they are prescriptives, declaratives, or recognitives of one of the sorts we examined in the last chapter. The first kind consist in the agent-neutral epistemic authority one can claim from the space of reasons. We examined the epistemic norms of moral discourse in the last chapter and determined that it counts as an objective discursive practice by the standards we deployed. It is democratic. It does not vest authority in any individual or group perspective and, as a result, moral claims are neither subjective nor dogmatic. The second kind consists in agent-relative, social-normative authority that is necessary to felicitously pull off a moral speech act. These agent relative norms of authority—the central topic of this chapter—circumscribe the moral discursive community by limiting who will be recognized as having the standing to engage in it. By attending to these diverging branches of the PALM Tree, we capture the diverging intuitions of objectivity and relativism and come to see that the tension between them is a central feature rather than a bug of moral discourse.

The last chapter left us with the impression that *anyone* can enter into an objective discourse. The standing to do so is universally held since no perspective is privileged in advance over any other. This, however, is a significant idealization, for there are many reasons we might reject the standing of others to enter into various objective discursive practices. In order to enter into scientific discourse, one needs to demonstrate one's credentials as a *bona fide* scientist. If we have reason to think you lack the requisite background knowledge, disposition to reason carefully, or, dare I say, objectivity, we will

simply not take your claims seriously. In the commerce of daily life, we don't recognize those we take to be incompetent, ill-motivated, cognitively compromised, or very young children to have the standing to make claims or to issue challenges or queries.

I argued in the last chapter that empirical observatives have agent-relative inputs, that they are essentially individuating, giving expression to an individual's first-personal uptake of the passing scene. Only the individual speaker could be entitled to the observative, and default entitlement depends on her being in the right place and time to have made the observation, on her being a competent user of the language, and on her being taken to have a reliable disposition to respond differentially to stimuli. My present claim is that assessment of agent-relative entitlement is not limited to observatives. Though the particular norms differ, we also assess the standing of interlocutors to issue declaratives, initiate queries, and level challenges. If we decide that a speaker is not, in the relevant sense, *one of us*, e.g., if we take them to lack the battery of conversationally relevant concepts or the ability to draw contextually appropriate inferences, then we will simply ignore what they have to say. The two branches of the PALM Tree are both necessary to chart the normative structure of any discourse.

Moral discourse also has agent-relative norms of authority, and they are even more stringent than those for ordinary empirical discourse. They limit the moral discursive community to a subset of those with whom we readily engage in other kinds of discourse and inquiry. It is this circumscription of the moral discursive community that gives us the sense that moral discourse is in some way less objective than ordinary empirical discourse. At the end of this chapter, I will suggest that relativism is the wrong conclusion to draw from this, but first we need to understand why the moral discursive

community is circumscribed in this way. We need an account of the norms of authority operative in moral discourse, those that determine who has the standing to make and challenge claims and to use moral prescriptives to hold others to the oughts that bind them. My account advance in two steps. The first is an investigation of the *norms of holding* for moral discourse, i.e., those norms that determine who has the standing to hold others to the moral oughts that bind them. I argue that, unlike the norms of holding of institutionalized hierarchies, whether someone has the standing to hold one morally to account is a matter of whether they are *recognized* as having such standing. As such, we need to examine when such recognition might be rightly withheld. I conclude that a fundamental condition for recognizing the authority of another to hold you to the oughts that bind you is a kind of mutual trust that the other will recognize you in return. In the second step, I argue that all moves within moral discourse constitute acts of holdings, so at least some norms of holding are operative for all moral claims. Only those who we recognize as having some standing to hold us to account have the standing to engage in moral discourse.

2. Norms of Holding

Beginning with Strawson's influential paper "Freedom and Resentment," there is a rich literature that examines what it is to *take* someone to be responsible. Much of this inquiry is endeavored on the way to answering questions about what it is to *be* responsible or as part of some other broader project.² More recently, however, some philosophers

² See, for example, (Darwall, 2006; Korsgaard, 1992; Maher, 2010; Oakley, 1991; Smith, 2007; Strawson, 1962; Wallace, 1994; Watson, 1996).

have turned their attention to what it is to *hold* someone responsible for some action and to the conditions under which such holdings are appropriate or felicitous.

Considering conditions that might undermine one's standing to hold another to account, for example, G. A. Cohen has examined the challenges posed by hypocrisy and complicity. Reflecting on the Israeli-Palestinian conflict, he argues that both parties face a "powerful *tu quoque* challenge" that undermines each side's standing "to point the finger at the other with no comment on his own glass house" (2006, pp. 110–111). Cohen argues that Israel, for example, lacks the standing to condemn Palestinian acts of terror without first examining its own role in causing the grievances to which terrorist acts are a response and in creating the conditions under which the terrorist response is the only one available (2006, pp. 114–115). By Cohen's lights, hypocrisy and complicity undermine the standing to hold the other to account.

R. A. Duff has also joined the fray, arguing that one's standing to hold accountable can be undermined if one has previously wronged the person one is trying to blame or if one incited the wrongdoing for which one is leveling blame (2010, p. 129).³ Linda Radzik has advanced this line of inquiry proposing three principles that she claims underwrite the norms of standing to hold responsible. She has argued that "the importance of liberty in self-regarding behavior, the moral significance of special interpersonal relationships, and the interests victims have in asserting their own authority" can each give rise to reasons that undermine one's standing to hold in particular cases (2011, p. 597).

³ Duff is not alone in examining blame in particular. See the articles in, (Coates & Tognazzini, 2013), especially (Bell, 2013).

It appears, then, that there is growing consensus that certain circumstances can undermine one's standing to hold another to account but little in terms of a unified account of the standing to hold. In what follows, I attempt to advance such an account first by developing a sketch of what constitutes a speech act of holding and then extending Radzik's framework to identify a set of conditions under which it would be reasonable for one to withhold recognition of another's entitlement to engage in such acts.⁴

3. The Topography of Holding Responsible

A good place to start thinking about what constitutes an act of holding is with Colleen Macnamara's account. She offers a picture that aims to capture the ways in which the complex attitudes and activities of taking others to be responsible fit together in what she dubs *the participant stance*, i.e., the orientation we take toward other normatively bound beings.⁵ The picture has three concentric circles that together represent the participant stance with all the attitudes and activities that are constitutive of it. In the outermost circle but excluded from the other two are attitudes and activities involved in treating another *as* a person other than those involved in appraising her conduct or hold-

⁴ I focus on the standing to engage in *speech acts* of holding others to the oughts that bind them because my central interest here is in the practice of moral *discourse*. The sphere of moral discourse, however, does not exhaust the space of moral practices. The space of morality is also shaped by our physical actions, by habits of mind and reactive attitudes, by the ways in which we learn to react subtly to one another in social space, and by our implicit attitudes. One can engage in non-discursive acts of holding, for example, by engaging in boycotts, marches, and sit-ins, by the use of physical force, or by subtly reacting to behaviors in all sorts of ways.

⁵ The "participant stance" is Macnamara's label for a concept that she finds implicit in Strawson's "Freedom and Resentment." It is, she says, "the complex mental orientation we take toward another which modulates our patterns of salience, presumptive interpretations, and leaves us susceptible to certain emotions and types of interactions" (Macnamara, 2011, p. 83n1). The "participant stance" is defined in opposition to the "objective stance," which is the way we comport ourselves toward inanimate objects and those we deem "incapacitated in some or all respects for ordinary interpersonal relationships" ((Macnamara, 2011, p. 84) citing (Strawson, 1962)), as well as in relation to "participant *reactive attitudes*," which "are those distinct emotional states we are susceptible to when we adopt the participant stance" (2011, p. 83n1)).

ing her responsible. The examples are multitudinous and varied: making pacts with another, telling her secrets, expecting her to avoid emotional pain, regarding her with suspicion, falling head over heels for her. In the middle ring but excluded from the center circle we find attitudes and activities that are not yet appraisals or judgments of conduct but that engage another regarding her actual or potential conduct. Offering advice to another on how she should proceed, consulting with her about her potential actions, asking her for her reasons, and engaging in soul-searching moral inquiry all fall into this space. Finally, the innermost circle contains only those attitudes and activities constitutive of holding others responsible for their conduct. Here we can include both participant reactive attitudes involved in holding oneself and others accountable—shame, disgust, disappointment, etc.—as well as acts of rebuke, condemnation, cajoling, and demanding (Macnamara, 2011, pp. 97–98; also see Wanderer, 2014, pp. 64–65).⁶ These communicative acts often take the form of speech, but we also hold others accountable through things like protest, direct action, boycott, and sanctions. These all would fall within Macnamara’s inner circle, but our focus is on second-personal speech acts, i.e., attempts to directly address another through speech regarding her actual or potential behavior.⁷

⁶ It is disputed whether this inner circle is comprised only of acts of negative judgment and sanction or whether praise also has a place here. For discussion, see (Macnamara, 2011, pp. 92–93; cf. Smith, 2008, p. 381).

⁷ Since my present focus is on second-personal speech acts of holding others to deontic moral norms and my final aim is an account of the agent-relative authority norms of moral discourse, my arguments are intended only to identify potential challenges to one’s standing to felicitously pull off such speech acts. The norms governing other ways of holding accountable will inevitably differ from those identified here. For instance, the norms of holding associated with acts of protest or boycott do not necessarily require that their targets recognize the standing of their originators in order to be effective. I thank Mark Lance for helping me think more clearly about this distinction.

Within the inner circle of the participant stance, Macnamara identifies “two faces of holding others responsible” (2011, p. 89). The first of these faces encompasses attitudes and activities of *appraisal*. These are “forms of emotional reaction that mark the moral meaning of others’ morally significant actions” (2011, p. 89). Such responses are *evaluative*, but not necessarily within the deontic realm. Our emotional reactions to others may but need not necessarily indicate their adherence to or violation of deontic norms but often mark their exhibiting some virtues or vices or causing some pleasures or pains for themselves or others. Activities of the *accountability* face, on the other hand, are always responses to violations of deontic norms. This is the face of holding others to their obligations or, as Macnamara puts it, “of *holding* someone to the oughts that bind them” (2011, p. 90). There is a second distinction between the two faces that is more important for our inquiry. The accountability face but not the appraisal face *always* involves some kind of *communicative expression*. In order to hold someone to the oughts that bind her, i.e., to enforce the norms that are in place, one must by some means communicatively engage her. A judgment or appraisal that remains unexpressed in word or conduct is normatively inert. It necessarily fails to *hold* its target to *anything* at all.⁸

What we see in the two faces are two distinct senses of “holding responsible.” In the appraisal face sense, one *takes* one to be responsible. Suppose you happen upon two children. Fatima is crying, while Anika is sitting happily playing with a toy truck. Having just seen Fatima playing with the truck, one surmises that Anika has taken it from Fatima

⁸ Within the accountability face, we can make a further distinction between those holdings that are forward-looking and those that are backward-looking, i.e., between what I have called holdings and assessments. Holding stands to assessing (paradigmatically blaming) “as preventative medicine stands to curative medicine” (Wanderer, 2014, p. 66). While this is an important distinction in some settings, I will not differentiate between the two in what follows. The sorts of things that undermine one’s standing to hold are also things that undermine one’s standing to assess.

eliciting her tears. One takes Anika to be responsible for upsetting Fatima. This *taking* to be responsible might involve feeling that Anika has evinced some disregard for Fatima, harboring some resentment toward Anika, and, perhaps, judging that Anika ought not to have taken Fatima's toy without permission. Being merely a bystander, however, one might feel it is not one's place to correct Anika or to return the toy to Fatima. In fact, one might go along in one's business without any sort of behavioral expression of one's reactive attitudes and judgments regarding Anika. This is holding responsible as merely *taking* to be responsible. Now contrast Anika's father, who has also happened upon the scene. Judging Anika to be responsible for Fatima's tears, her father goes a step further. He reproaches Anika for taking the truck, tells her to apologize, and returns the truck to Fatima. Anika's father has held her responsible in a sense stronger than merely *taking* her to be responsible. He has enforced a norm by holding her to account for her actions, and he has done so by way of a set of communicative acts. Rather than merely taking her to be responsible, Anika's father has *held her accountable* for her actions.

The accountability face differs from the appraisal face, then, in that it encompasses only those takings to be responsible that are communicatively expressed with the aim of holding others to deontic norms. Jeremy Wanderer has pointed out that the requirement of communicative expression makes the accountability face "voluntary" (2014, p. 65). His idea seems to be that while we cannot control our reactive attitudes, we are in control when we communicatively express those attitudes "with the intent of rebuking" (Wanderer, 2014, p. 65). In forming the intention to rebuke another, we judge that the communicative act will at least potentially be effective, that it "is a worthwhile undertaking" (Wanderer, 2014, p. 65). A nice result of this would be that the case for norms of

holding would be bolstered by the idea that we are voluntarily responsive to such norms, but I want to caution against moving too quickly in this direction. My reason for hesitation is some reticence over the role of intention in the communicative act. Consider an example. Anika's father happens upon the scene described above and immediately and without intending to rebuke, takes the toy from Anika and returns it to Fatima. This act communicates something to Anika, all the same. She has been corrected. Given the context, she comes to understand that she ought not have taken the toy from Fatima even if it is not her father's intention to communicate this. It is not clear to me that this differs in any significant way from our earlier case in which Anika's father did intend to rebuke her. In fact, I think that much of our behavior around holding others to deontic norms takes this form: our reactive attitudes *seep out* whether or not we intend to express them. This doesn't make their expression any less communicative, for, as I see things, it is not the intent but the pragmatic structure of the act that defines it. Even when "unintentional," communicative acts seek certain kinds of uptake from their targets and function *as communicative acts* when they achieve such uptake, regardless of intention.

Communicative acts have "a distinct internal aim, mode of achieving it, and success conditions"(Macnamara, 2011, p. 90). Sanctioning behaviors like rebukes, for example, have as their aim what Macnamara calls "first-personal practical uptake of the ought-violation" by the individuals toward whom they are directed. They aim "to get the wrongdoer to acknowledge her wrongdoing [as a violation of the relevant norm], feel remorse, apologize, make amends, and commit to doing right in the future" (Macnamara, 2011, p. 90). This is achieved through the imposition of burdens on the one being rebuked. The burdens, in this case, are emotional. She feels "the sting of reproof" (Macnamara, 2011,

p. 90). Finally, the rebuke is successful when “it is met with full first-personal practical uptake of the ought-violation” (Macnamara, 2011, p. 90). It is this constellation of internal aim, mode, and success conditions that makes a communicative act of rebuke the act that it is. It does not matter whether the person doing the rebuking does so with intention or merely as a result of reactive attitudes that have seeped out in behavior. It does matter, though, whether expression is given to the attitudes at all. If it is not, no rebuke has occurred.⁹

This result does not undermine the case for standing conditions for holding accountable. I have argued that holdings need not be intentional, but this does not place them beyond one’s control. So long as it is possible for one to suppress an unintentional seeping out of reactive attitudes it seems reasonable to think there may be norms governing when one ought to do that. Compare, for example, the ability to suppress or contain unintentional outbursts of joy or anger when they would be contextually inappropriate. Moreover, as I have urged throughout, we must understand structural discursive norms as norms of assessment, i.e., as the norms interlocutors use in scoring one’s moves. As such, even unintentional acts of holding could be norm bound.

4. Who Has the Standing to Hold?

Who, if anyone, has the standing to hold another to moral norms? There are clear cut cases of standing conditions for other kinds of holdings, but the deontic moral case seems at least a little more complicated. Consider a prudential case. It would, we can assume, be better for you to skip dessert. All things considered, eating dessert amounts to the consumption of calories you don’t need that would be stored as fat and, in the long

⁹ Cf. (Macnamara, 2015).

run, be a detriment to your health. Yes, you'd enjoy dessert, but the long-term ill effects are a high price to pay for that small pleasure. Suppose this is your own assessment, and, given your interests, it is a correct assessment. You *really* ought not to eat dessert. Even so, it would be out of line for me, a stranger sitting at the next table, to lean over and whisper "You shouldn't have dessert tonight," when the waiter comes calling. Even if I'm a mind reader and somehow know with full certainty your interests and what they dictate, I would be terribly out of line if I chastised you as you considered the dessert menu. It's just not my place. Someone else—your best friend, for example—might be able to hold you to your prudential obligation to refrain from dessert, but a stranger at the next table simply lacks the standing to do so.

The judgment that it's not my place could originate from my own laziness or the fear that my own judgment about the situation is mistaken. It might be that I worry that interjecting will require me to follow through in ways for which I am not prepared. Sometimes, though, the judgment that it's none of my business, that I ought to refrain from holding you to account, is a judgment "that it would be *wrong* to" do so (Radzik, 2011, p. 582), that it would violate some salient norms. Now, one might think that what's salient are social norms delineating the private sphere and governing interpersonal conduct. I want to suggest that these same norms function as norms of standing within the practice of prudential discourse. Prudential discourse is aimed at determining what one ought to do to fulfill one's interests. This often requires consultation with others, as our interests are often opaque to us and our values are very rarely stable. We need others to challenge what we take our values to be, to aid us in their examination so that we might determine whether they are values we really ought to endorse and on which we ought to

act. But too many voices would needlessly complicate our lives. To have strangers assuming they know our preferences or challenging us to defend them would foster anxiety, interfere with all sorts of everyday social interactions, and possibly lead to a kind of paralysis in which we are unable to finally decide what we ought to do. It would undermine the very purpose of the practice in which we're engaged, as such, we have adopted norms of standing within that practice that circumscribe the community of individuals who can engage us in this way to those with whom we are engaged in shared projects or whose interests are deeply intertwined with our own in other ways.

Still, one might think that the moral case and the prudential case differ significantly. Darwall seems to defend the universal standing to hold in arguing that "the moral perspective [is] an impartially disciplined version of the second-person standpoint" (2006, p. 102). His idea is that when we take up the mantle of morality, we speak not as a particular individual but "as an equal participant in the first-person plural ("we") of the moral community" (2006, p. 102). We address another as an equal member of this community and on its behalf. In doing so, we are not claiming any special authority for ourselves, we are merely claiming standing as part of the "we" (Radzik, 2011, pp. 584–588). I think that this follows from a confusion between the standing to issue moral judgments and the standing to hold, but, in the end, I will argue that even the standing to issue moral judgments is circumscribed by norms of holding. To make this case, I want to examine a parallel between deontic moral holdings and another speech act with a similar pragmatic structure.

Orders come with clear cut standing conditions. As your professor, I can order you to stow your laptop away during class, but, should a passerby in the hall poke her

head in the door and token the same type of utterance, her speech act will be infelicitous. She lacks the standing—the authority—requisite for carrying out this speech act. Similarly, to use a well-trod example, if I am walking by a parade ground and overhear the drill instructor's order to her cadets to drop and give her twenty, I am under no obligation to begin doing pushups. She has the authority to order her cadets to do so, but I, as a civilian, am beyond the reach of that authority.

In these cases, there is a defined structure of authority—a hierarchy—that is known and recognized by the relevant parties and that is the product of broadly institutional relationships between them. This structure defines the standing conditions for issuing the speech acts in question. As a professor, I have the standing to issue an order regarding the use of laptops in my classroom, but this authority does not extend to orders about your time spent on Snapchat outside of class. These examples differ from the cases under investigation in at least two key respects. First, in the case of holding to deontic norms there is not, in general, a well-defined, institutional, and widely recognized structure of authority. There may be such a structure in special cases such as a parent issuing a holding to a child, a teacher to a student, or clergy to a parishioner, but, in general, our cases are messier and, often, negotiable. One dear friend might have the standing to hold me to quitting smoking, for example, while I may count it is an insult if another reminds me of my commitment when she sees me at the corner store. Second, orders are what Kukla and Lance call *constative holdings* (2009, pp. 111–112). These speech acts create new normative statuses rather than call attention to and enforce already existing ones. Before the drill instructor gives the order to do push-ups, there is no sense in which the cadets were already obligated to do them. The status of *one required to do push-ups* is

newly created in the act of ordering. Deontic moral holdings do not, by their very nature, create new normative statuses but rather have as their aim enforcing already existing commitments. They are *alethic holdings*. I ought to quit smoking, and my friend telling me so only aims to enforce this *true* prescriptive (Kukla & Lance, 2009, Chapter 5; Wanderer, 2014).¹⁰

Recall that, following Kukla and Lance, we are thinking of speech acts as functions on normative statuses. They take normative statuses as inputs, and their outputs are alterations of normative statuses. On either end these can be either agent-relative or agent-neutral. The former are indexed to individuals on the basis of their position in some social-normative space. The latter are, in principle, universal (2009, Chapter 1). Orders and alethic holdings exhibit a structural similarity in that both have agent-relative normative inputs; it's not the case that just anyone has the requisite authority, i.e., normative status or standing, to successfully pull off the speech act. They are distinguished, however, in that alethic holdings *also* have agent-neutral inputs; they are premised on the already existing commitments of their targets while orders create new commitments.

Darwall's universal standing position, I believe, arises in part out of a conflation of standing to hold and standing to issue a moral prescriptive. The prescriptive, (1) "Tom ought not to belittle his partner" is one to which any member of the relevant moral community could be entitled if they were prepared to give reasons in its defense. It is, of

¹⁰ There's an interesting and important question about how exactly alethic holdings function to normatively hold their targets to already existing commitments. Kukla and Lance (2009) and Wanderer (2014) each offer accounts of how they achieve their function. In general, we can say that alethic holdings both elicit first-personal practical uptake of the relevant norm and render that norm more salient by making it the case that the target of the holding is beholden not just to the norm but also to the originator of the holding. If my friend signs on the dotted line even after hearing my "Don't do it!", she has violated a norm, but she has also, in some sense, violated our relationship. That there is such a relationship to be violated is a necessary condition of the standing to hold.

course, true that I ought not to belittle my partner, and anyone in the community who can produce the reasons why secures entitlement to this utterance. Such reasons as she may have, however, do not necessarily entitle her to tell *me* (2) “Tom! Don’t belittle Kate!”

(2) is not a prescriptive, but an imperative. It has agent-relative inputs as well as agent-relative outputs. Like an order, this imperative cannot be uttered felicitously by just anyone *even if* she is entitled to (1), the prescriptive that underwrites it. Entitlement to (1) is necessary but not sufficient for entitlement to (2). We saw this already in the case of counsels of prudence discussed above. The other diner might have been entitled to utter the prescriptive “S should not order the cake” to his dining partner, but it would still be infelicitous to issue a second-personal holding directed at S. The reason, I claimed, has to do with a presumption of privacy. It would be treacherous to navigate a world in which just anyone would be entitled to render second-personal just any prudential obligation that one might have. As such, we have adopted norms of standing to hold that recognize only those to whom we’ve entrusted our interests or with whose interests our own are intertwined. Unlike the case of orders within well-defined, institutionalized structures of authority, exactly who has this standing is always negotiable, but one can see one’s way to the sorts of reasons that might be relevant to such negotiation. I see no good reason for thinking that the moral case should be treated differently than the prudential case. The standing to issue an alethic holding to a deontic moral norm accrues to particular agents on the basis of their position in normative space. What we need now is some grasp on the relevant position.

5. Agency Threats and the Standing to Hold

Linda Radzik's work on differentiating the standing to sanction provides a promising point of departure for identifying the sorts of reasons that might undermine one's standing to hold another to the moral oughts that bind her. Following Darwall, Radzik begins from the assumption that the standing to hold to deontic moral norms is universal but then argues that this entitlement is defeasible such that the target of an attempted holding might have a justifiable claim against one trying to hold her to a moral obligation (2011, p. 592). If this is right, then one might find oneself with "an obligation to refrain from sanctioning a particular kind of wrong" even when one is entitled to the concomitant moral prescriptive (2011, p. 590). The target would have a justifiable, second-personal claim *against* the prospective sanctioner that undermines the sanctioner's standing to hold the target accountable, and such a claim would implicitly appeal to the norms of standing we're trying to identify. Let's examine three cases that Radzik develops in which one's standing to hold accountable is defeated by considerations having to do with respect for agency (2011, pp. 592–593).

Radzik's first case is self-regarding behavior. She argues that holding a person with respect to her purely self-regarding behavior interferes with "the agent's ability to develop trust in her own judgment," which has the effect of undermining her agency (2011, p. 593). Agents need space to develop their decision-making capacities and shape their identities. Introducing too much noise in the form of external voices aiming to guide behavior threatens to derail the process and, as such, is an affront to agency itself. If this is right, then the agent has a claim against the would-be holder or sanctioner that she refrain from holding or sanctioning. Such a claim is defeasible, of course, as there may be

reasons that do, in some cases, justify a degree of paternalism, but the presumption is against such interference.

A second case Radzik explores is that of wrongs within what she calls “special relationships.” Within the bounds of romantic relationships, friendships, family, and activist groups, for example, outsiders lack standing to hold or sanction insiders with regard to behavior internal to the relationship. Such relationships are central to our self-conceptions and vital for our well-being, but they only function well when they are afforded “degrees of privacy, intimacy, and trust” (2011, p. 593). Outside interference can undermine their constitutive bonds. As such, the parties to such relationships have a claim against would-be interveners that they refrain from holding or sanctioning with respect their behavior vis-à-vis one another. As with the previous case, such a claim is defeasible. Intervention might be justified in the protection of the physical well-being of the parties or in cases that involve children, for example (2011, p. 594).

Finally, Radzik’s third case is that in which bystander sanction interferes with “the victim’s ability to find vindication in the aftermath of wrongdoing” (2011, p. 597). Here again we find a reason that has to do with respect for agency, but in this case it’s that of the victim rather than the target of the holding. Were a third-party to come to the rescue, the victim of the wrong may find herself further marginalized. She interpellated as one who cannot stand up for herself. Her standing as a moral agent is weakened, and the likelihood that the pattern of wrongdoing that has occurred will be replicated is heightened because of her diminished standing.

Radzik’s cases point to a variety of ways one’s standing to hold might be challenged. What unifies them is that in each case recognition of someone’s standing to hold

would in some way threaten the agency of some relevant party. This, I think, is central for grasping agent-relative standing required to felicitously engage in speech acts of moral holding. I argued in Chapter 4 that moral discourse is a meta-coordinative, meta-normative discursive practice that is constrained by a need to maintain a stable cooperative community in which domination is kept in check and in which there is minimal motivation for defection. If the norms of the practice allowed just any member of the community to engage in acts of holding toward any other, the non-domination constraint would be threatened precisely because such profligate acts of holding would threaten the agency of their targets, subjugating them to the judgments of others and leaving no space for individuals to judge for themselves what they ought to do. This claim, of course, needs further support, so let's turn now to examining some further ways in which standing might be undermined because it threatens someone's agency. In particular, I want to look at the case of hypocrisy.

6. Trust and the Standing to Hold

When one party tries to blame another for a violation of moral norms of which the first party herself is also guilty we hear challenges like "look who's talking" or "that's the pot calling the kettle black" (Cohen, 2006, p. 108). "Judge not, that ye be not judged," we are reminded in scripture. It is argued sometimes that the root of this *tu quoque* challenge is that in practicing hypocrisy, one is making an exception for oneself. Moral rules apply equally to all, but if I am guilty of a moral wrong and then blame you for a wrong of the same kind, I have tried to claim for myself some elevated status, to carve out an exception such that my action was acceptable while yours was condemnable. This challenges a deep Kantian commitment that the same rules must apply universally.

This explanation misses the mark. It explains what makes hypocrisy wrong and may justify one in condemning it. It sheds no light, however, on why committing this particular wrong undermines one's standing to hold others accountable for their wrongs. One need not be a moral saint to hold others to account. On the contrary, a sinner may herself be better placed to see just which acts are going to have terrible consequences. Why, then, does hypocrisy sometimes undermine the standing hold another to account for the same wrong of which the holder is guilty? Why can't the pot call the kettle black?

Recognizing the standing of a hypocrite to hold you to the moral oughts that bind you, like the cases Radzik canvassed, threatens your agency. An agent's willingness to judge hypocritically indicates that she is unlikely to recognize the standing of others to hold her to the moral obligations that she has. The charge of hypocrisy sticks when the individual facing the charge refuses to recognize her past wrongs. If, in response to the charge of hypocrisy, one takes responsibility for the wrong one has done and attempts, if possible, to correct for that wrong, one is likely to reclaim the standing to hold. It is only the unreformed hypocrite, the one who denies she ever did wrong, who is refused recognition. In accepting the charge of hypocrisy and trying to make amends, one demonstrates one's willingness to allow others to hold her to the moral oughts that bind her. She demonstrates that she will recognize their standing in the hopes that they will, in turn, recognize hers.

The unreformed hypocrite, on the other hand, wants her standing to hold recognized but refuses to recognize the standing of the target of her holding. If her standing were recognized, this would institute an asymmetrical relationship in which the hypocrite can exercise her authority to shape the other's behavior by adding second-personal force

to her moral obligations but refuses to allow the other to do the same in return. One way to characterize such a relationship is in terms of the kind of hierarchies from which the standing to issue orders derives. Recognizing the standing of the hypocrite is akin to accepting that she inhabits a role of hierarchical authority over the obligations that are binding on one's conduct. It is as if she is the drill sergeant or the parent who gets to issue orders about what you ought to do, and this is a threat to one's agency.

One might object, though, that moral holdings have an epistemic dimension. Recognizing the hypocrite's standing does not grant her so much authority, for one can always challenge a holding by challenging the prescriptive on which it is premised. Indeed, even in felicitous cases of holding others to account, it is often necessary not only to draw their attention to the oughts that bind them but also to persuade them that *this action* is really what the norms they endorse requires of them. It requires an argument. The hypocrite, then, might be able to draw your attention to the norms that bind you without exercising undue influence since the prescriptives she endorses can be challenged. If, however, the hypocrite already refuses to recognize your standing to hold her to the moral norms that bind her, what assurance do you have that she will recognize your standing to hold her to the epistemic norms that bind her? I think that we take hypocrisy as an indication that one refuses to recognize the standing of others to hold her to account to any of the oughts that bind her, and this is why we vehemently refuse to recognize the standing of the unreformed hypocrite to hold us to account.

More succinctly, we might say that hypocrisy undermines trust. I will not recognize your authority to hold me if I have no reason to trust that you will recognize my authority to hold you to your commitments or to challenge your attempts to hold me. I am

justified in withholding my recognition of standing insofar as recognizing your standing would undermine my agency. Your refusal to allow me to hold you accountable for your past violation is evidence that you do not recognize my standing to hold. As such, I have reason not to trust that you will recognize my standing to hold you accountable to any norms at all.

7. Trust and Divided Communities

During the 2016 US Presidential campaign and since the election, many of us have thought about how to hold accountable those who voted for Donald Trump. We tried, before the election, to hold prospective Trump voters to their own avowed anti-racist, anti-xenophobic, anti-sexist, anti-ableist, etc., moral commitments, and since the election, many of us have tried to hold to account those who did vote for him for enabling and emboldening his deeply troubling agenda. Such attempts have often been met with challenges to our standing to hold these folks accountable, even when they are old friends or family. We are often met with refrains of “Who are you to tell me...” or “What gives you the right...”

Like the case of hypocrisy, these refusals to recognize the would-be holder’s standing also issue from a lack of trust that the holder will reciprocally recognize that of the target. Recognizing the standing of the holder when reciprocal recognition is unlikely subsumes the target in a hierarchy, just as recognizing the unreformed hypocrite does. In this sort of case, the target is being held or blamed by someone who is nominally part of her community. She is a fellow citizen. She accepts many of the same moral principles that the target accepts and aims to live by them. The holder may be a family member, an old friend, a college roommate, or a co-worker. Many of these relationships, though, are

rather thin or one-dimensional. In my own experiences, I have met such responses when trying to talk to family and friends who I only see a few times a year or interact with only on social media. These are relationships that were once more robust, but as I've moved from my rural hometown to a metropolitan area, and visits have become fewer and farther between, these relationships have become much less normatively rich than they once were.

These attempts to hold others to the moral oughts that bind them are more likely to be successful where there are deep, trusting relationships, built on varied interactions that present numerous low-stakes opportunities to hold one another accountable or to witness one another acting out our moral commitments. Successful acts of holding are usually embedded in deep and long-tenured relationships between individuals as well as within communities that are bound together in varied and complex ways. We are not merely strangers passing in the night or even *merely* classmates or colleagues. In our communities, we may be related to some one individual as colleague, cycling partner, bowling buddy, and fellow Rotarian. Our relationships, that is to say, are usually multi-dimensional and normatively saturated. We engage across diverse normative environments that present many opportunities for normative holding in different settings. I might, for example, hold my cycling partner to account for slacking on his training and my fellow Rotarian to account for failing to pay dues. The kind of trust that assures one of reciprocal recognition of standing is built in these varied, low-stakes incidences of holding. It is this trust that is called on when one claims the standing to hold in the moral context. Where it is absent, we are likely to withhold recognition in defense of our own agency,

and we are not wrong to do so. The agent-relative norms of moral holding do not grant standing where these thick, multi-dimensional relationships are missing.

One might accept that lack of reason to trust undermines standing in cases in which the parties are moral equals, but object that this is not one of those cases. There are many instances of legitimate asymmetric holding relationships: child/parent, teacher/student, expert/novice. By continuing to support Trump in the face of all that he has said and done, one might think our friends and family have shown themselves to lack moral expertise. Perhaps one would be aiding the development of their agency by holding them to their moral commitments, even though it would clearly be a mistake for one to recognize their standing in return. This would be a kind of training relationship in which one helps them to see how their moral commitments ought to manifest in their actions. This objection cannot be lightly dismissed, though it seems to me that the evidence that someone is a moral neophyte would have to be rather substantial before it could override the presumption that we ought to treat other adult human beings as fellow, full-fledged moral agents. It could not be enough that we disagree only on the candidate they ought to have supported. It would not even be enough if we disagreed on a wide variety of moral judgments so long as they manage to make reasonable judgments and live out their moral commitments in much of their conduct, i.e., they acted like moral agents. What is important to see is that mutual recognition of the standing to hold is not predicated on agreement. We can morally disagree on a great many particular cases yet still recognize one another's standing to push each other to defend or reconsider our judgments and correct

each other when we've gone astray. Mutual recognition is where moral debate begins, not its culmination.¹¹

8. From Standing to Hold to Standing to Assert

We've now established that the standing to hold accountable is differentiated. Whether one has the standing to hold another accountable is a matter of whether the target can justifiably withhold recognition of such standing, i.e., whether she might be in a position to claim that recognizing the standing of the speaker to hold her accountable would in some way jeopardize her agency. We have also established that, in a central class of cases, her agency might be jeopardized if she lacks reason to trust (or has reason to distrust) that the speaker would reciprocally recognize her standing to hold the speaker accountable to the oughts that bind her.

The next step of the argument is to show that the norms of holding apply to a broader class of speech acts than is readily apparent. Insofar as they attend to them, most theorists treat norms of holding as applying only to explicit speech acts of holding. These are themselves varied. One might hold another to the oughts that bind her using an imperative, e.g., when a father tells his child to return the stolen toy, or a prescriptive used in the context of second-personal holding/assessment, e.g., when I remind you that you ought to finish that referee report tonight. Alternatively, you can hold another to her already existing obligations by cajoling her, guilt-tripping her, requesting that she do the thing ought to do, or issuing a third-personal holding/assessment within earshot. Moreover, each of these ways of holding comes with subtly different entitlement conditions. I,

¹¹ I thank an anonymous reviewer for the Kennedy Institute of Ethics Journal for bringing this objection to my attention.

for example, might be entitled only to request that you keep your promise to our mutual friend, while she might be entitled to demand it, and your partner might be entitled to guilt you into it (Kukla & Lance, 2009, pp. 105–113). There are, however, more indirect ways of holding others to the oughts that bind them. I might lack the standing to tell Mandy, a mere acquaintance, that she ought to stop taking advantage of the kindness of her friends in the way that she does, but I might be able to raise concern among her closer friends by telling them what I think Mandy ought to do. In some cases, the most effective way to get someone to act is not to confront them but to get others to do your bidding for you. If you can get enough others on board, they might influence the behavior of the target of your holding. This happens often among groups of friends and acquaintances and in small organizations, but it also happens on larger scales. I lack standing to tell many politicians directly how they ought to behave, but I certainly have standing to make it known to others so that we might, together, have greater influence.

Once we've got these indirect ways of holding in view, we can start to notice even more circuitous ways to hold others to the oughts that bind them. Consider, for example, the moral revolutionary. Dr. King used many different tactics to call the nation, its citizens, and its leaders to account for the wrongs of slavery, Jim Crow segregation, disenfranchisement, economic oppression, and unjust wars. One of his tactics was to take on the mantle of morality, to claim the moral high ground, and to argue the case that black men, women, and children had been horribly mistreated for far too long. Now, King, in some ways, was issuing a direct holding to those responsible for this mistreatment and to those who were in a position to do something about it, but he was also aiming to change the national conversation around race. He was trying to make the moral argument that

everyone deserved equal treatment no matter the color of their skin. If effective, this argument might persuade at least some to realize their own wrongs and those of their friends and neighbors, i.e., those of the people they might hold to account. Even if King lacked the standing to knock on a white man's door and take him to task for his behavior, he might convince some of those closest to him that they ought to do just that. The effect of his shifting the moral landscape was that others would hold accountable those with whom they stood in appropriate relations, i.e., those who would recognize their standing.

Every moral utterance is an act of holding along with whatever other speech acts it performs. This is a corollary of understanding moral discourse as the practice within which moral obligations, permissions, and values are forged. In engaging in the practice, one has the opportunity to nudge the thinking of other practitioners, to shift their moral perspectives and realign their commitments. If that happens, then they will come to hold themselves and one another to the newly shaped norms of their moral community.¹² As such, one's act of uttering a moral prescriptive or issuing a valuing in the context of moral debate or deliberation can have the effect of holding others to the oughts by which the speaker takes them to be bound.

If each moral utterance is an act of holding—perhaps in some attenuated sense—then it will be governed by some norms of holding. Holdings are interpersonal acts to which not just anyone is entitled. There are shades of intimacy, trust, and power that shape the entitlement conditions for holding, and some of these will be relevant even to holding in the attenuated sense of engaging in moral discourse. They will not be precisely

¹² I certainly do not mean to imply that this is ever a quick or easy process, the point is simply that every moral utterance could potentially shift the discourse and, so, lead to the recognition of and holding one another to new commitments.

the same norms of holding relevant to the sorts of cases we discussed earlier. Someone might have the standing to engage in moral discourse on a topic relevant to your behavior but still lack the standing to direct a second-personal holding at you. They might, for example, be part of the moral discursive community that debates the permissibility of abortion but not be entitled even to gently hold you to your duty to gestate if they conclude that, in your circumstances, abortion is wrong (Kukla & Lance, 2009, p. 110). So, what are the relevant norms of holding? To put it another way, what might be legitimate reasons for withholding recognition of someone's standing to engage in moral discourse?

Though I do not think that it is exhaustive, I suggest that one central ingredient in the standing to engage with a moral discursive community was already revealed above. A participant in a moral discursive community can rightly withhold recognition of another's standing to engage in the practice when she has reason to expect that the other will not reciprocally recognize her own standing (or, perhaps, when she lacks reason to expect otherwise). If she does not trust that the petitionary participant will give uptake to the moral reasons she musters, make herself subject to holdings by members of the community, or engage in moral dialogue and deliberation in good faith, then she is justified in rejecting their petition to enter into the discourse.¹³

Doing otherwise—recognizing their standing—would threaten the agency of all discursive participants. Return to the example with which I began the chapter, my moral judgment that *chhaupadi* is wrong. A western Nepalese man who expects his daughters to

¹³ In reflecting on this point, it is useful to keep in mind the I-Thou structure of sociality that we have followed Brandom in adopting. The point of contact between a petitionary participant and a moral discursive community is always some individual (or group of individuals) with whom the petitionary participant hopes to enter into dialogue. That individual (or group) is the one who must recognize or withhold recognition from her on the basis of the agent-relative authority norms of the practice.

practice chhaupadi has legitimate reason to be suspect of any outsider who tries to convince him that the practice is morally wrong. Outsiders tend not to understand or value local customs or religious beliefs, but this is only part of the reason. More importantly, he has no reason to think that, should the tables be turned, the outsider would be willing to hear the moral reasons he can muster in critique of some central practices of the outsider. The Nepalese man cannot trust that, should he recognize the standing of the outsider as a participant in his moral discourse, the outsider would reciprocally recognize him. In fact, if he were to take a quick glance at history, assuming this outsider is from Western Europe or North America, he might find ample reason to believe that the outsider positively would not recognize his standing but would rather treat him as morally inferior, in need of reform and education. Recognizing the standing of the outsider, then, grants him undue authority over the moral matters of the Nepalese man's moral discursive community. It would instantiate a hierarchy where none previously existed, and this would threaten the agency of all of its members. They would be placing themselves at the whim of one over whose action they have no sway and who, moreover, would not recognize the reasons they present to challenge his claims as reasons worth consideration.

As I argued in the case of holding, the constraint of non-domination that has shaped the structural discursive norms of moral discourse is responsible for our adopting a norm of standing that only grants standing within the practice to those who are likely to reciprocally recognize the standing of other practitioners. The practice simply could not fulfill its function if it allowed for individuals to exercise authority within the practice who would not make themselves subject to the judgment of others. It would lead to the

kind of fracturing and defection of a cooperative community against which I have argued moral discourse evolved as a prophylactic.

9. Earning Trust

The standing to hold others to the oughts that bind them is a product of the recognition of those others, and they can have both legitimate and illegitimate reasons for withholding such recognition. The legitimate reason we've been examining has to do with a lack of reason to trust that the other will reciprocally recognize one's own standing, will allow herself to be held to the oughts that bind her, will hear and respond to challenges, and so on. A natural question, then, is what could ground such trust?

Trust is never automatic. It does not result from familial or tribal bonds. It does not inhere in one's genes nor does it automatically accrue to those with whom one interacts regularly. We can all think of examples of acquaintances—perhaps even very close kin—who we do not recognize as having the standing to hold us morally accountable, perhaps because it would be hypocritical of them, perhaps because we differ so drastically in world views, or maybe just because we think that they are morally inept. Moreover, we interact daily with folks whose moral judgment we'd take to be suspect and whose standing we would not recognize. Commerce may require a degree of trust engendered by law, but it does not require morality. We manage to interact with many to whom we would not grant the standing to hold.

Moral saints don't necessarily deserve this trust, and sinners sometimes do. Moral perfection is impossible. The complexity of our world makes it such that even with careful planning and good intentions, one's actions can lead to morally disastrous outcomes, and no one acts solely on good intentions. But there are folks who seem to get close.

These moral saints, being human, are liable to be prideful, to think that they can judge others without opening themselves to judgment from their moral inferiors. They might look down on others, not with disdain, but with pity. If they do, then they do not deserve to have their standing to hold recognized, for they are unlikely to recognize the other's in return. If, however, they demonstrate not only moral aptitude but also a degree of humility, they might engender the trust from which recognition follows. On the other hand, moral sinners, though they've failed time and again to live up to the oughts that bind them, may deserve to have their standing recognized. If they recognize their wrongs, if they are repentant for them, if they hear the criticisms of others and take others to have the standing to hold them to account, then there is no legitimate reason to withhold recognition. We might even be more apt to accept sinners into our moral discursive communities than we are saints.

One does not get to be a member of the moral discursive community by being a good moral agent, by following the of rules morality, or simply by being a member of the relevant linguistic community. Rather, one's standing as a moral discursive practitioner is recognized in virtue of the other varied relationships in which one stands to other practitioners. This standing is the product of living in thick, multi-dimensional, and normatively saturated relationships with one another. When we have many opportunities to hold one another accountable to the *low-stakes* oughts that bind us, for example, to show up to a meeting, get in a workout, pay our dues, volunteer our time, or post our score for the golf league, we build trust. We come to see that the other will recognize our own standing to hold them accountable, and so we have no reason not to recognize theirs. These opportunities are the product of engaging in projects with one another, being involved in our

communities, and having shared goals and commitments across a variety of lived contexts. This suggests that earning the standing to hold others accountable to the moral oughts that bind them requires building a shared community with them. Only once we've had ample opportunity to engage in a variety of low-stakes, trust building ways will we have enough trust in reciprocal recognition to recognize another's standing. Morality doesn't build communities, communities build morality, which then holds them together. Moral discourse is, in this sense, a *normatively open* practice in which the standing to participate is earned outside of the practice itself (Lance, 2008).

10. Conclusion

The “wet” liberal’s dilemma is really a version of a more local problem. Kukla and Lance claim that the “distinction between prescriptive truth-claims and holdings with agent-relative entitlements is inchoately reflected in the common (and nearly incoherent) opinion that ‘abortion is wrong’ but ‘it’s nobody’s business to judge’ women who have abortions.” What folks who say this are really trying to express is the idea that though they judge abortion to be wrong, no one (least of all a stranger or government) has standing to hold a woman to her obligation to gestate (Kukla & Lance, 2009, p. 110). Kukla and Lance are right to note that the entitlement conditions for the prescriptive and the second-personal holding come apart, but they are wrong when they say that “it is *everybody’s* ‘business’ to judge” what is morally right or wrong, permissible or impermissible. The “wet” liberal’s dilemma is the result of the fact that moral prescriptives also have some agent-relative inputs, i.e., that the standing to issue prescriptives—to engage in moral judgement and discourse—is also differentiated. Not just anyone can hold another

to the oughts that bind her *nor* can just anyone enter into the space of reasons a claim about what those oughts are.

This does not mean that we cannot or should not morally judge the actions of members of other moral communities. Chhaupadi is morally abhorrent, and there's nothing wrong with my making that judgment. Moral claims are objective; nothing about the particularities of an individual agent or community can count as a legitimate challenge to a moral judgment, so the fact that someone doesn't recognize my standing or accepts certain mythical beliefs that I reject does not stand as a legitimate defeater for my judgment about the moral oughts that bind them. Yet, just as there is nothing wrong with my making the judgment, there is nothing wrong with a western Nepalese man rejecting my standing to do so. He *should* tell me to buzz off precisely because he has no reason to trust that I will reciprocally recognize *his* standing to engage in moral discourse and to challenge *my* claims. What makes the "wet" liberal queasy about making moral judgments about moral communities of which he is not a member is precisely that he recognizes that such rejections of recognition are legitimate, but he draws the wrong lesson when he thinks that this means he cannot or should not make the judgment at all. That his standing to do so within the moral discursive community of the western Nepalese will not be recognized means only that there is further moral work to do; the judgment itself is just the beginning.

This isn't hopelessly relativistic. It doesn't entail that we should all just bite our tongues and stay out of each other's business. Rather, it implies that the objectivity of moral discourse is a fragile work in progress. Objectivity as we have conceived of it is a

matter of every claim being open to challenge and no perspective being privileged in advance over any other, but our moral communities come with built in immune systems. We are not going to let just anyone challenge our long-held moral commitments; we are going to exclude others who threaten to undo the status quo. Moral discursive communities sometimes fragment, as a result. We begin to see even others in our own communities as dangerous outsiders, and so refuse to engage in moral discourse with them, opting instead to stand back and judge them (often vociferously) from afar. But moral communities can also grow, and we've now seen one way this can happen. As we engage in thick, multi-dimensional, and normatively saturated projects with one another and develop a mutual trust that the other will reciprocally recognize our standing to hold them accountable, we let one another into our respective moral communities. These bridges of trust provide the means by which we can claim the standing necessary to engage in moral discourse and, eventually, to hold others to the oughts that bind them. Building a bridge like this is one way in which one of us "wet" liberals might one day earn the standing to call the Nepalese practitioners of chhaupadi to account, but until we've built such bridges, we'll have to rely on those within that moral community to do this work and to fall back on other, more tenuous modes of holding that don't require such mutual recognition.¹⁴

Lest this all sound too easy, however, we must recognize that we always face the challenge of a kind of ethnocentricity that leads to a contracting rather than expanding moral discursive community. Mutual trust is a luxury that is available to communities with resources to spare, living in relative prosperity and safety (Rorty, 1998). The threat

¹⁴ I do not want to insinuate that building trust is the only way to expand our moral communities. There are others. Rorty, for example, identifies practices of sentimental education in which the moral sentiments of successive generations are brought into alignment with the western liberal tradition he identifies as valuable (Rorty, 1998). On this point, also see, (Baier, 1991).

of scarcity or danger from outsiders in the form of competition for resources, disease, or invasion or even the mere perception of such a threat can cause our communities close in on themselves, retreating into the safety of relative isolation. We begin thinking in exclusionary, tribal terms, and rather than looking for ways to expand our moral discursive community, we simply judge the outgroup as both morally wrong and dangerous and refuse to engage them in trust-building ways. Moral practices are *adaptively plastic* insofar as “*evolved human nature is both an obstacle to moral progress and an enabler of it, depending on the environment and the degree to which it resembles certain conditions that were prevalent in the environment of evolutionary adaptation*” (Buchanan & Powell, 2018, p. 187, emphasis in original). When we experience abundance and relative safety, we are open to engaging in moral practices with others, expanding our moral community, and holding up our commitments to more and varied challenges. When we experience scarcity and threat, we do just the opposite. Without mutual trust in reciprocal recognition on which to build, we withhold recognition of the standing of the outsider to hold us to account, and we begin identifying more and more groups and individuals as outsiders.

This is our lot. Moral discourse is a fragile, adaptively plastic practice; one that, under the right conditions, can hold communities together fostering greater coordination and cooperation and leading us to treat one another in ways that all of us find justifiable, i.e., to something that we might be tempted to call moral progress, but also one that can cause our communities to balkanize in the face of threats and scarcity. The objectivity of moral discourse, the ever-expanding source of tests for our moral commitments, is a significant achievement, but the relativism of moral discursive communities—their tendency to be stingy in granting standing to petitionary practitioners—is a feature, not a bug, of

this evolved tool of ours. It will always be with us, and we will always be pulled in these two directions so long as we continue to find moral discourse to have some utility for us or for as long as its institutional inertia persists.

Conclusion

Moral discourse is one of the many practices human beings have evolved to better coordinate our complex social, normative lives so that we might take advantage of the benefits of ongoing cooperation. We have come to understand this practice to be structured by a variety of norms of authority, pragmatic norms, and inferential norms that give it the distinctive features it exhibits. These norms are responsible for its tight but defeasible connection to motivation, its relationship to our reactive attitudes, its distributed structure of authority, i.e., its objectivity, and the fact that only a circumscribed group of individuals has standing to engage in any particular instantiation of a moral discursive practice, i.e., its relativity. Moreover, we have explained these structural discursive norms in terms of their making the practice suitable to answer to the specific needs and purposes of the beings who use it, in particular, we have argued that it is useful as a meta-coordinative, meta-normative vocabulary that provides means for resolving tensions in underlying coordinative practices in ways that do not fracture the group or allow for domination by alpha-type free-riders.

I would like to close by expressing a few of my hopes in putting forward this *coping* picture of moral discourse. The first is that this picture brings morality down to earth. The picture is one of moral practices that *belong* to us, of moral obligations that are *answerable* to us, and of moral authority that *rests* with us. It does make us answerable to

any non-human others; it does not impose itself upon us. Rather, we create it and revise it as we work together to build and navigate our physical and social realities. Yet, because what the practice demands of us is not in the control of any individual or group of individuals, it still presents to use as *objective*. Moral truths outrun what anyone or everyone takes them to be, and so they seem as though they are timeless or transcendent, even though they are not. This is one of the features that gives morality its distinctive force. My second hope, then, is that this coping picture of moral discourse takes some of the mystery out of how moral practices that are *only* ours can still impose such obligations on us or, perhaps more perspicuously, how we can use it to impose such obligations on ourselves.

My final hope is that the resources developed in this work can be leveraged to better navigate moral discursive spaces. Now that we can plainly see the function of moral discourse and the conditions under which it evolved to fulfill those functions, I hope that we can make better predictions about when moral discourse is likely to be useful to us and when it is likely to fail, and, so, when we should rely on it and when we should not. More importantly, perhaps, I hope that the account herein also allows us to see that we can shape our social lives in ways that allow moral discourse to better fulfill its telos as well as in ways that inhibit it. We have a tool at our disposal that can facilitate smoother social coordination in a way that minimizes risks of domination and fracture, but in order to avail ourselves of this tool, we need to minimize perceptions of scarcity and threat both from within and without. We cannot—as some are presently wont to do—play to fear and hope to still have a well-functioning moral life. We must, as Rorty suggests, continue the long project of sentimental education. And, finally, we must create

many more opportunities to build trust in mutual recognition through the fostering of thick, multi-dimensional, normatively saturated relationships with members of broad and diverse communities. Only in taking these steps can the tool of moral discourse continue to be of use to us; if we refuse them, it's likely to become obsolete.

[Page intentionally left blank]

Bibliography

- Alexander, R. D. (1982). *Darwinism and Human Affairs*. Seattle, WA: University of Washington Press.
- Altham, J. E. J. (1987). The Legacy of Emotivism. In G. MacDonald & C. Wright (Eds.), *Fact, Science, and Morality*. Oxford: Blackwell.
- Amatya, P., Ghimire, S., Callahan, K. E., Baral, B. K., & Poudel, K. C. (2018). Practice and lived experience of menstrual exiles (Chhaupadi) among adolescent girls in far-western Nepal. *PLoS ONE*, *13*(12). Retrieved from <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6287853/>
- Anderson, E. (2015). *Feminist Epistemology and Philosophy of Science*. Retrieved from <https://plato.stanford.edu/archives/fall2015/entries/feminism-epistemology/>
- Antony, L. M. (1993). Quine as Feminist: The Radical Import of Naturalized Epistemology. In L. M. Antony & C. Witt (Eds.), *A Mind of One's Own*. Retrieved from <https://www.taylorfrancis.com/>
- Austin, J. L. (1946). Other Minds, Part II. *Proceedings of the Aristotelian Society*, *148*, 148–187.
- Austin, J. L. (1957). A Plea for Excuses. *Proceedings of the Aristotelian Society*, *57*(1), 1–30.
- Austin, J. L. (1962). *How to Do Things with Words*. Cambridge: Harvard University Press.
- Austin, J. L., Strawson, P. F., & Cousin, D. R. (1950). Symposium: Truth. *Proceedings of the Aristotelian Society, Supplementary Volumes*, *24*, 111–172.

- Axtell, G. (2015). *Objectivity*. Malden, MA: Polity.
- Ayer, A. J. (1952). *Language, truth, and logic*. New York: Dover Publications.
- Baier, A. C. (1991). *A Progress of Sentiments: Reflections on Hume's Treatise*.
Cambridge, MA: Harvard University Press.
- Barnes, W. H. F. (1934). A Suggestion About Value. *Analysis*, 1(3), 45–46.
- Barrett, A. A. (1989). *Caligula: The Corruption of Power*. Yale University Press.
- Båve, A. (2009). A Deflationary Theory of Reference. *Synthese*, 169(1), 51–73.
- Bell, M. (2013). The Standing to Blame: A Critique. In D. J. Coates & N. Tognazzini
(Eds.), *Blame: It's Nature and Norms* (pp. 262–281). Oxford ; New York: Oxford
University Press.
- Belnap, N. (1990). Declaratives Are Not Enough. *Philosophical Studies*, 59, 1–30.
- Blackburn, S. (1984). *Spreading the word: Groundings in the philosophy of language*.
Oxford; New York: Clarendon Press ; Oxford University Press.
- Blackburn, S. (1988). Attitudes and Contents. *Ethics*, 98(3), 501–517.
- Blackburn, S. (1993). *Essays in Quasi-Realism* (First Edition). New York: Oxford
University Press.
- Blackburn, S. (2001). *Ruling Passions: A Theory of Practical Reasoning*. Oxford:
Clarendon Press.
- Blackburn, S. (2006). Antirealist Expressivism and Quasi-Realism. In D. Copp (Ed.), *The
Oxford Handbook of Ethical Theory* (pp. 146–162). Oxford: Oxford University
Press.
- Block, N. (1986). Advertisement for a Semantics for Psychology. *Midwest Studies in
Philosophy*, 10(1), 615–78.

- Boehm, C. (2012a). Ancestral Hierarchy and Conflict. *Science*, 336(6083), 844–847.
- Boehm, C. (2012b). *Moral Origins: The Evolution of Virtue, Altruism, and Shame* (1 edition). New York: Basic Books.
- Boghossian, P. (2007). *Fear of Knowledge: Against Relativism and Constructivism*. Oxford University Press, USA.
- Boyd, R. (1988). How to Be a Moral Realist. In G. Sayre-McCord (Ed.), *Essays on Moral Realism* (pp. 181–228). Ithaca, NY: Cornell University Press.
- Brandom, R. (1994). *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Cambridge, MA: Harvard University Press.
- Brandom, R. (1998). Insights and Blindspots of Reliabilism. *The Monist*, 81(3), 371–392.
- Brandom, R. (2000). *Articulating Reasons: An Introduction to Inferentialism*. Cambridge: Harvard University Press.
- Brandom, R. (2002). The Centrality of Sellars’s Two-Ply Account of Observation to the Arguments of “Empiricism and the Philosophy of Mind.” In *Tales of the Mighty Dead: Historical Essays in the Metaphysics of Intentionality*. Cambridge, Mass: Harvard University Press.
- Brandom, R. (2007). Inferentialism and Some of Its Challenges. *Philosophy and Phenomenological Research*, 74(3), 651–677.
- Brandom, R. (2008a). *Between Saying and Doing: Towards an Analytic Pragmatism*. Oxford University Press, USA.
- Brandom, R. (2008b). Replies. In B. Prien & D. Schweikard (Eds.), *Robert Brandom: Analytic Pragmatist*. Heusenstamm bei Frankfurt: Ontos Verlag. phl.

- Brandom, R. (2010). Brandom's Responses. In B. Weiss & J. Wanderer (Eds.), *Reading Brandom: On Making It Explicit* (1st ed., pp. 295–365). New York: Routledge.
- Brandom, R. (2011). *Perspectives on Pragmatism: Classical, Recent, and Contemporary*. Cambridge, Mass: Harvard University Press.
- Brandom, R. (2013). An Arc of Thought: From Rorty's Eliminative Materialism to His Pragmatism. In A. Groeschner, C. Koopman, & M. Sandbothe (Eds.), *Richard Rorty: From Pragmatist Philosophy to Cultural Politics* (pp. 23–30). London ; New York: Bloomsbury Academic.
- Brandom, R. (2015). *From Empiricism to Expressivism: Brandom Reads Sellars*. Cambridge, Massachusetts: Harvard University Press.
- Brink, D. O. (1989). *Moral Realism and the Foundations of Ethics* (1st edition). Cambridge ; New York: Cambridge University Press.
- Brink, D. O. (2001). Realism, Naturalism, and Moral Semantics. *Social Philosophy and Policy*, 18(2), 154.
- Buchanan, A., & Powell, R. (2018). *The Evolution of Moral Progress: A Biocultural Theory*. New York: Oxford University Press.
- Carnap, R. (1935). *Philosophy and Logical Syntax*. Providence, RI: American Mathematical Society.
- Chrisman, M. (2008). Expressivism, Inferentialism, and Saving the Debate. *Philosophy and Phenomenological Research*, 77(2), 334–358.
- Chrisman, M. (2015). *The Meaning of "Ought": Beyond Descriptivism and Expressivism in Metaethics* (1 edition). New York, NY: Oxford University Press.

- Coates, D. J., & Tognazzini, N. A. (Eds.). (2013). *Blame: Its nature and norms*. Oxford ; New York: Oxford University Press.
- Cohen, G. (2006). Casting the First Stone: Who Can, and Who Can't, Condemn the Terrorists? *Royal Institute of Philosophy Supplement*, 81(58), 113-.
- Craig, E. (1991). *Knowledge and the State of Nature: An Essay in Conceptual Synthesis* (1 edition). Oxford : New York: Clarendon Press.
- Darwall, S. (2006). *The Second-Person Standpoint: Morality, Respect, and Accountability*. Cambridge, Mass: Harvard University Press.
- Darwall, S., Gibbard, A., & Railton, P. (1992). Toward Fin de Siècle Ethics: Some Trends. *Philosophical Review*, 101(1), 115–189.
- Daston, L. (1992). Objectivity and the Escape from Perspective. *Social Studies of Science*, 22(4), 597–618.
- Daston, L., & Galison, P. (1992). The Image of Objectivity. *Representations*, (40), 81–128.
- Daston, L. J., & Galison, P. (2010). *Objectivity*. New York; Cambridge, Mass.: Zone Books.
- Davidson, D. (1963). Actions, Reasons, and Causes. *Journal of Philosophy*, 60(23), 685–700.
- Davidson, D. (1986). A Coherence Theory of Truth and Knowledge. In E. LePore (Ed.), *Truth and Interpretation. Perspectives on the Philosophy of Donald Davidson* (pp. 307–319). New Jersey: Blackwell.

- Devries, W. A., & Triplett, T. (2000). *Knowledge, Mind, and the Given: Reading Wilfrid Sellars's "Empiricism and the Philosophy of Mind."* Indianapolis: Hackett Publishing Company.
- Douglas, H. (2004). The Irreducible Complexity of Objectivity. *Synthese*, 138(3), 453–473.
- Dreier, J. (2004). Meta-Ethics and The Problem of Creeping Minimalism. *Philosophical Perspectives*, 18(1), 23–44.
- Duff, R. A. (2010). Blame, Moral Standing and the Legitimacy of the Criminal Trial. *Ratio*, 23(2), 123–140.
- Eklund, M. (2009). The Frege-Geach Problem and Kalderon's Moral Fictionalism. *The Philosophical Quarterly (1950-)*, 59(237), 705–712.
- Faerna, Á. M. (2014). Rorty and Dewey on Warrant. *Contemporary Pragmatism*, 11(1), 15–26.
- Field, H. (1994). Deflationist Views of Meaning and Content. *Mind*, 103(411), 249–285.
- Foot, P. (1972). Morality as a System of Hypothetical Imperatives. *The Philosophical Review*, 81(3), 305–316. Retrieved from JSTOR.
- Fricke, M. (2009). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press, USA.
- Gauthier, D. (1998). Why Contractarianism? In J. Rachels (Ed.), *Ethical Theory 2: Theories About How We Should Live*. Oxford University Press.
- Geach, P. T. (1960). Ascriptivism. *Philosophical Review*, 69, 221–225.
- Geach, P. T. (1965). Assertion. *Philosophical Review*, 74(4), 449–465.

- Gellner, E. (1990). *Plough, Sword and Book: The Structure of Human History*. Chicago: University of Chicago Press.
- Gibbard, A. (1990). *Wise choices, apt feelings: A theory of normative judgment*. Cambridge, MA: Harvard University Press.
- Gibbard, A. (1999). Morality as Consistency in Living: Korsgaard's Kantian Lectures. *Ethics*, 110(1), 140–164.
- Gibbard, A. (2003). *Thinking how to live*. Cambridge, Mass.: Harvard University Press.
- Godfrey-Smith, P. (2011). *Darwinian Populations and Natural Selection* (Reprint edition). Oxford University Press, U.S.A.
- Grice, P. (1991). *Studies in the Way of Words*. Cambridge, Mass.: Harvard University Press.
- Grover, Camp, Jr., J., & Belnap, N. (1975). A Prosentential Theory of Truth. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 27(2), 73–125.
- Hannon, M. (2015). The Importance of Knowledge Ascriptions. *Philosophy Compass*, 10(12), 856–866.
- Hannon, M. (2018). *What's the Point of Understanding?* Retrieved from <https://static1.squarespace.com/static/551587e0e4b0ce927f09707f/t/5a8066490d92973fece0759e/1518364236287/Mind+-+clean+version.pdf>
- Hannon, M. (2019). *What's the Point of Knowledge?: A Function-First Epistemology*. Oxford, New York: Oxford University Press.
- Hare, R. M. (1952). *The Language of Morals*. New York: Oxford University Press.

- Harman, G. (1977). *The Nature of Morality: An Introduction to Ethics*. New York: Oxford University Press.
- Harman, G. (1986). *Change in View*. Cambridge, MA: MIT Press.
- Harman, G., & Thomson, J. (1996). *Moral Relativism and Moral Objectivity* (1 edition). Cambridge, Mass., USA: Wiley-Blackwell.
- Hart, H. L. A. (1951). The Ascription of Responsibility and Rights. *Proceedings of the Aristotelian Society*, 49 (1948-1949)(New Series), 171–194.
- Haslanger, S. (1995). Ontology and Social Construction. *Philosophical Topics*, 23, 95–125.
- Haugeland, J. (1982). Heidegger on Being a Person. *Noûs*, 16(1), 15–26.
- Hauser, M. D., Chomsky, N., & Fitch, W. T. (2002). The Faculty of Language: What Is It, Who Has It, and How Did It Evolve? *Science*, 298(5598), 1569–1579.
<https://doi.org/10.1126/science.298.5598.1569>
- Henrich, J. (2017). *The Secret of Our Success: How Culture Is Driving Human Evolution, Domesticating Our Species, and Making Us Smarter*. Princeton Oxford: Princeton University Press.
- Horgan, Terence, & Timmons, M. (1991). New Wave Moral Realism Meets Moral Twin Earth. *Journal of Philosophical Research*, 16, 447–465.
- Horgan, Terence, & Timmons, M. (1992). Troubles for New Wave Moral Semantics: The “Open Question Argument” Revived. *Philosophical Papers*, 21(3), 153–175.
- Horgan, Terry, & Timmons, M. (2006). Cognitivist Expressivism. In Terry Horgan & M. Timmons (Eds.), *Metaethics After Moore* (pp. 255–298). Oxford: Oxford University Press.

- Horwich, P. (1998a). *Meaning*. Oxford: Clarendon Press.
- Horwich, P. (1998b). *Truth* (2nd ed.). Oxford: Clarendon Press.
- Horwich, P. (2001). A Defense of Minimalism. *Synthese*, 126(1–2), 149–165.
- Hume, D. (1751). *An Enquiry Concerning the Principles of Morals*. Hackett Publishing.
- Jackson, F. (1997). Naturalism and the Fate of the M-Worlds II. *Aristotelian Society Supplementary Volume*, 71(1), 247–268.
- Jackson, F. (1998). *From metaphysics to ethics a defence of conceptual analysis*. Oxford; New York: Clarendon Press ; Oxford University Press.
- Joyce, R. (2007). *The myth of morality*. Cambridge: Cambridge University Press.
- Joyce, R. (2015a). Mackie’s Arguments for the Moral Error Theory. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2012). Retrieved from <https://plato.stanford.edu/entries/moral-anti-realism/moral-error-theory.html>
- Joyce, R. (2015b). Moral Anti-Realism > Projectivism and Quasi-realism (Stanford Encyclopedia of Philosophy). Retrieved June 18, 2019, from <https://plato.stanford.edu/entries/moral-anti-realism/projectivism-quasi-realism.html>
- Kahneman, D. (2013). *Thinking Fast and Slow*. New York, NY: Farrar, Straus and Giroux.
- Kalderon, M. E. (2007). *Moral Fictionalism (Lines of Thought)*. Chicago: Oxford University Press, USA.
- Kalderon, M. E. (2008). Moral Fictionalism, the Frege-Geach Problem, and Reasonable Inference. *Analysis*, 68(2), 133–143.

- Kant, I. (1785). *Groundwork of the Metaphysics of Morals* (2nd edition; M. Gregor & J. Timmermann, Eds.). Cambridge: Cambridge University Press.
- Kitcher, P. (2011). *The Ethical Project*. Cambridge, MA: Harvard University Press.
- Korsgaard, C. M. (1992). Creating the Kingdom of Ends: Reciprocity and Responsibility in Personal Relations. *Philosophical Perspectives*, 6, 305–332.
- Korsgaard, C. M., & O’Neill, O. (1996). *The Sources of Normativity*. Cambridge ; New York: Cambridge University Press.
- Kraut, R. (1990). Varieties of Pragmatism. *Mind*, 99(394), 157–183.
- Kuhn, T. (1962). *The Structure of Scientific Revolutions* (2nd ed.). Chicago: University of Chicago Press.
- Kukla, R. (2000). Myth, Memory and Misrecognition in Sellars’ “Empiricism and the Philosophy of Mind.” *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 101(2–3), 161–211.
- Kukla, R. (2006). Objectivity and Perspective in Empirical Knowledge. *Episteme: A Journal of Social Epistemology*, 3(1), 80–95.
- Kukla, R. (2018). That’s What She Said: The Language of Sexual Negotiation. *Ethics*, 129(1), 70–97.
- Kukla, R. (2019, February 4). Sex Talks. Retrieved February 11, 2019, from Aeon website: <https://aeon.co/essays/consent-and-refusal-are-not-the-only-talking-points-in-sex>
- Kukla, R., & Lance, M. (2009). “Yo!” and “Lo!”: *The Pragmatic Topography of the Space of Reasons*. Boston, MA: Harvard University Press.

- Kukla, R., & Lance, M. (2010). Perception, Language, and the First Person. In B. Weiss & J. Wanderer (Eds.), *Reading Brandom: On Making It Explicit* (1st ed., pp. 115–128). New York: Routledge.
- Lance, M. (1997). The Significance of Anaphoric Theories of Truth and Reference. In *Truth*. Atascadero: Ridgeview.
- Lance, M. (2008, March). *Temporal Structures of Norms or A semi-precise characterization of the self-consciousness of language (2008 University of Chicago Lipkind Lecture)*.
- Lance, M. (2015). Life Is Not a Box-Score: Lived Normativity, Abstract Evaluation, and the Is/Ought Distinction. In M. Williams, S. Gross, & N. Tebben (Eds.), *Meaning Without Representation: Essays on Truth, Expression, Normativity, and Naturalism*. Oxford: Oxford University Press.
- Lance, M., & O’Leary-Hawthorne, J. (1997). *The grammar of meaning: normativity and semantic discourse*. Cambridge: Cambridge University Press.
- Laukötter, S., Prien, B., & Schepelmann, T. (& others). (2008). Are Fundamental Discursive Norms Objective? In *Robert Brandom: Analytic Pragmatist*. Heusenstamm bei Frankfurt: Ontos Verlag.
- Leeds, S. (1978). Theories of Reference and Truth. *Erkenntnis* (1975-), 13(1), 111–129.
- Lenman, J. (2014). Moral Naturalism. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2014). Retrieved from <https://plato.stanford.edu/archives/spr2014/entries/naturalism-moral/>
- Levine, S. (2010). Rehabilitating Objectivity: Rorty, Brandom, and the New Pragmatism. *Canadian Journal of Philosophy*, 40(4), 567–589.

- Lewis, D. (1969). *Convention: A Philosophical Study*. Oxford: Blackwell.
- Lewis, D. (1979). Scorekeeping in a Language Game. *Journal of Philosophical Logic*, 8(1), 339–359.
- Little, M. O., & Macnamara, C. (2008, September 20). *Between the Optional and the Obligatory: Toward a Theory of Deontic Pluralism*. Unpublished.
- Lloyd, E. A. (1995). Objectivity and the Double Standard for Feminist Epistemologies. *Synthese*, 104(3), 351–381.
- Longino, H. E. (1990). *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry*. Princeton University Press.
- Longino, H. E. (2001). *The Fate of Knowledge*. Princeton University Press.
- Mackie, J. (1977). *Ethics: Inventing Right and Wrong*. London: Penguin.
- Macnamara, C. (2011). Holding others responsible. *Philosophical Studies*, 152(1), 81–102.
- Macnamara, C. (2015). Reactive Attitudes as Communicative Entities. *Philosophy and Phenomenological Research*, 90(3), 546–569.
- Maher, C. (2010). On Being and Holding Responsible. *Philosophical Explorations*, 13(2), 129–140.
- Manne, K. (2018). *Down Girl: The Logic of Misogyny*. Oxford University Press.
- McDowell, J. (1979). Virtue and Reason. *The Monist*, 62(3), 331–350. Retrieved from JSTOR.
- McDowell, J. (1994). *Mind and World*. Cambridge, Mass: Harvard University Press.
- McDowell, J. (2000). Towards Rehabilitating Objectivity. In R. B. Brandom (Ed.), *Rorty and His Critics* (1 edition, pp. 109–122). Malden, Mass: Wiley-Blackwell.

- Melis, A. P., Hare, B., & Tomasello, M. (2008). Do chimpanzees reciprocate received favours? *Animal Behaviour*, 76(3), 951–962.
- Mijar, D. B. (2017, April 5). Perspectives: The Practice of Chhaupadi. Retrieved June 13, 2019, from She's the First website:
<https://medium.com/@shesthefirst/perspectives-the-practice-of-chhaupadi-8b398d4524dc>
- Moore, G. (1903). *Principia Ethica*. Cambridge: Cambridge University Press.
- Morris, C. W. (1938). *Foundations of the Theory of Signs*. Chicago: University of Chicago Press.
- Muller, M. N., & Mitani, J. C. (2005). Conflict and Cooperation in Wild Chimpanzees. In *Advances in the Study of Behavior* (Vol. 35, pp. 275–331).
[https://doi.org/10.1016/S0065-3454\(05\)35007-8](https://doi.org/10.1016/S0065-3454(05)35007-8)
- Nagel, T. (1986). *The View from Nowhere*. Oxford University Press, USA.
- Neumann, J. von, & Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.
- Nyhof-DeMoor, M. (2011). *Brandom and Hegel on Objectivity, Subjectivity and Sociality: A Tune Beyond Us, Yet Ourselves*. Retrieved from
https://www.academia.edu/15036039/Brandom_and_Hegel_on_Objectivity_Subjectivity_and_Sociality
- Oakley, J. (1991). *Morality and the Emotions*. London ; New York: Routledge.
- O'Leary-Hawthorne, J., & Price, H. (1996). How to stand up for non-cognitivists. *Australasian Journal of Philosophy*, 74(2), 275–292.
<https://doi.org/10.1080/00048409612347251>

- Paul, L. A. (2016). *Transformative Experience* (Reprint edition). Oxford: Oxford University Press.
- Penco, C. (1999). Interview with Robert Brandom. *Epistemologia*, 143–150.
- Peregrin, J. (2012). What is Inferentialism? In L. Gurova (Ed.), *Inference, Consequence, and Meaning: Perspectives on Inferentialism* (pp. 3–16). Retrieved from <http://jarda.peregrin.cz/mybibl/PDFTxt/580.pdf>
- Perry, J. (1979). The Problem of the Essential Indexical. *Noûs*, 13(1), 3–21.
- Pinker, S. (2010). *The Language Instinct: How the Mind Creates Language*. New York, NY: Harper Collins.
- Price, H. (1997). Naturalism and the Fate of the M-Worlds I. *Proceedings of the Aristotelian Society, Supplementary Volumes*, 71, 247–282.
- Price, H. (2003). Truth as Convenient Friction. *Journal of Philosophy*, 100(4), 167–190.
- Price, H. (2010). One Cheer for Representationalism? In R. E. Auxier & L. E. Hahn (Eds.), *The Philosophy of Richard Rorty* (pp. 269–289). Chicago: Open Court.
- Price, H. (2013). *Expressivism, pragmatism and representationalism*. Cambridge: Cambridge University Press.
- Price, H., & Jackson, F. (1997). Naturalism and the Fate of the M-Worlds. *Proceedings of the Aristotelian Society, Supplementary Volumes*, 71, 247–282.
- Putnam, H. (2002). *The collapse of the fact/value dichotomy and other essays*. Cambridge, MA: Harvard University Press.
- Putnam, H. (2005). *Ethics without ontology*. Cambridge, Mass.; London: Harvard University Press.

- Quine, W. V. O. (1986). *Philosophy of Logic: 2nd Edition* (2 edition). Cambridge, Mass: Harvard University Press.
- Radzik, L. (2011). On Minding Your Own Business: Differentiating Accountability Relations Within the Moral Community. *Social Theory and Practice*, 37(4), 574–598.
- Ramsey, F. P. (1927). Facts and Propositions. *Proceedings of the Aristotelian Society*, 7(1), 153–170.
- Rawls, J. (1971). *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- Recanati, F. (2004). *Literal Meaning*. Cambridge: Cambridge University Press.
- Rorty, R. (1970). Incommensurability as the Mark of the Mental. *The Journal of Philosophy*, 67(12), 399–424. <https://doi.org/10.2307/2024002>
- Rorty, R. (1979). *Philosophy and the Mirror of Nature*. Princeton University Press.
- Rorty, R. (1991a). Introduction. In *Objectivity, Relativism, and Truth: Philosophical Papers* (Vol. 1, pp. 1–17). Cambridge: Cambridge University Press.
- Rorty, R. (1991b). *Objectivity, Relativism, and Truth: Philosophical Papers* (Vol. 1). Cambridge: Cambridge University Press.
- Rorty, R. (1991c). On ethnocentrism: A reply to Clifford Geertz. In *Objectivity, Relativism, and Truth: Philosophical Papers* (Vol. 1, pp. 203–210). Cambridge: Cambridge University Press.
- Rorty, R. (1991d). Science as Solidarity. In *Objectivity, Relativism, and Truth: Philosophical Papers* (Vol. 1, pp. 35–45). Cambridge: Cambridge University Press.

- Rorty, R. (1991e). Solidarity or objectivity? In *Objectivity, Relativism, and Truth: Philosophical Papers* (Vol. 1, pp. 21–34). Cambridge: Cambridge University Press.
- Rorty, R. (1997). What Do You Do When They Call You a 'Relativist'? *Philosophy and Phenomenological Research*, 57(1), 173.
- Rorty, R. (1998a). Human Rights, Rationality, and Sentimentality. In *Truth and Progress: Philosophical Papers* (pp. 167–185). Cambridge: Cambridge University Press.
- Rorty, R. (1998b). *Truth and Progress: Philosophical papers* (Vol. 3). Cambridge: Cambridge University Press.
- Rorty, R. (2000). Response to McDowell. In R. B. Brandom (Ed.), *Rorty and His Critics* (1 edition, pp. 123–128). Malden, Mass: Wiley-Blackwell.
- Rorty, R., & Price, H. (2010). Exchange on “Truth as Convenient Friction.” In M. de Caro & D. Macarthur (Eds.), *Naturalism and Normativity*. Columbia University Press.
- Ryle, G. (1950). “If”, “So”, and “Because.” In M. Black (Ed.), *Philosophical Analysis: A Collection of Essays* (pp. 302–318). Abingdon-on-Thames: Routledge.
- Sayre-McCord, Geoffrey (Ed.). (1988). *Essays on moral realism*. Ithaca, N.Y: Cornell University Press.
- Schroeder, M. A. (2010). *Being for: Evaluating the semantic program of expressivism*. Oxford: Oxford University Press.
- Searle, J. R. (1962). Meaning and Speech Acts. *Philosophical Review*, 71(4), 423–432.

- Searle, J. R. (1969). *Speech Acts: An Essay in the Philosophy of Language*. London: Cambridge U.P.
- Sellars, W. (1954). Some Reflections on Language Games. In R. Brandom & K. Scharp (Eds.), *In the Space of Reasons: Selected Essays of Wilfrid Sellars* (pp. 28–56). Cambridge: Harvard University Press.
- Sellars, W. (1956). Empiricism and the Philosophy of Mind. In W. A. Devries & T. Triplett (Eds.), *Knowledge, Mind, and the Given: Reading Wilfrid Sellars's "Empiricism and the Philosophy of Mind"* (pp. 205–276). Indianapolis: Hackett Publishing Company.
- Sellars, W. (1957). Counterfactuals, Dispositions, and the Causal Modalities. In M. Scriven, G. Maxwell, & H. Feigl (Eds.), *Minnesota Studies in The Philosophy of Science, Vol. II*. University of Minnesota Press.
- Sellars, W. (1962). Philosophy and the Scientific Image of Man. In *In the Space of Reasons* (pp. 369–408). Cambridge: Harvard University Press.
- Sellars, W. (1989). *The Metaphysics of Epistemology: Lectures by Wilfrid Sellars* (First Edition edition; P. W. Amaral, Ed.). Atascadero, Calif: Ridgeview Pub Co.
- Sellars, W. (1993). *Science and Metaphysics: Variations on Kantian Themes*. Atascadero, Calif.: Ridgeview Pub Co.
- Sharma, B., & Schultz, K. (2019, January 10). Woman and 2 Children Die in Nepal Menstruation Hut. *The New York Times*. Retrieved from <https://www.nytimes.com/2019/01/09/world/asia/nepal-menstruation-taboo.html>
- Simpson, M. (2017). *Non-Representationalism and Metaphysics*. Cambridge, UK: Ph.D. dissertation, Robinson College, University of Cambridge.

- Smith, A. M. (2007). On Being Responsible and Holding Responsible. *The Journal of Ethics*, 11(4), 465–484.
- Smith, A. M. (2008). Control, responsibility, and moral assessment. *Philosophical Studies*, 138(3), 367–392.
- Smith, M. (1994). *The Moral Problem*. Cambridge, MA: Blackwell.
- Smyth, N. (2017). The Function of Morality. *Philosophical Studies*, 174(5), 1127–1144.
- Sober, P. E., Wilson, P. D. S., & Wilson, D. S. (1999). *Unto Others: The Evolution and Psychology of Unselfish Behavior* (New Ed edition). Cambridge, MA: Harvard University Press.
- Sterelny, K. (2014). *The Evolved Apprentice: How Evolution Made Humans Unique*. Cambridge: A Bradford Book.
- Stevenson, C. L. (1972). *Ethics and language*. New Haven: Yale University Press.
- Strawson, P. F. (1962). Freedom and Resentment. *Proceedings of the British Academy*, 48, 1–25.
- Street, S. (2006). A Darwinian Dilemma for Realist Theories of Value. *Philosophical Studies*, 127(1), 109–166.
- Street, S. (2009). In Defense of Future Tuesday Indifference: Ideally Coherent Eccentrics and the Contingency of What Matters. *Philosophical Issues*, 19(1), 273–298.
- Szabó, Z. G. (2005). Semantics versus Pragmatics. In Z. G. Szabó (Ed.), *Introduction*. Oxford University Press.
- Taylor, C. (2010). Language Not Mysterious? In B. Weiss & J. Wanderer (Eds.), *Reading Brandom: On Making It Explicit* (1st ed., pp. 32–47). New York: Routledge.

- Tebben, N. (2015). Anti-Representational Semantics: Four Themes. In S. Gross, N. Tebben, & M. Williams (Eds.), *Meaning Without Representation: Expression, Truth, Normativity, and Naturalism*. Oxford; New York: Oxford University Press.
- Thaler, R., & Sunstein, C. (2009). *Nudge*. Penguin Books.
- Tomasello, M. (2010). *Origins of Human Communication*. MIT Press.
- Tomasello, M. (2016). *A Natural History of Human Morality*. Cambridge, Massachusetts: Harvard University Press.
- Tomasello, M., Melis, A. P., Tennie, C., Wyman, E., & Herrmann, E. (2012). Two Key Steps in the Evolution of Human Cooperation: The Interdependence Hypothesis. *Current Anthropology*, 53(6), 673–692.
- Travis, C. (2005). A Sense of Occasion. *The Philosophical Quarterly*, 55(219), 286–314.
- van Roojen, M. (2016). Moral Cognitivism vs. Non-Cognitivism. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2016). Retrieved from <https://plato.stanford.edu/archives/win2016/entries/moral-cognitivism/>
- Wallace, R. J. (1994). *Responsibility and the Moral Sentiments*. Cambridge, Mass: Harvard University Press.
- Wanderer, J. (2008). *Robert Brandom*. McGill-Queen's University Press.
- Wanderer, J. (2010a). Brandom's Challenges. In B. Weiss & J. Wanderer (Eds.), *Reading Brandom: On Making It Explicit* (1st ed., pp. 96–114). New York: Routledge.
- Wanderer, J. (2010b). Inhabiting the Space of Reasoning. *Analysis*, 70(2), 367–378.
- Wanderer, J. (2014). Alethic Holdings. *Philosophical Topics*, 42(1), 63–84.
- Warren, M. D. (2015). Moral inferentialism and the Frege-Geach problem. *Philosophical Studies*, 172(11), 2859–2885. <https://doi.org/10.1007/s11098-015-0447-8>

- Warren, M. D. (2018). Building bridges with words: An inferential account of ethical univocity. *Canadian Journal of Philosophy*, 48(3–4), 468–488.
<https://doi.org/10.1080/00455091.2017.1422630>
- Watson, G. (1996). Two Faces of Responsibility. *Philosophical Topics*, 24(2), 227–248.
- Westermarck, E. (1908). *The Origin and Development of the Moral Ideas*. London: Macmillan and Company.
- Wilk, T. (2017). Trust, Communities, and the Standing to Hold Accountable. *Kennedy Institute of Ethics Journal*, 27(2 Supplement), E1–E22.
- Williams, B. (1978). *Descartes: The Project of Pure Enquiry*. New Jersey: Humanities Press.
- Williams, G. C. (1966). *Adaptation and Natural Selection: A Critique of Some Current Evolutionary Thought*. New Jersey: Princeton University Press.
- Williams, M. (2001). *Problems of Knowledge: A Critical Introduction to Epistemology*. Oxford: Oxford University Press.
- Williams, M. (2010). Pragmatism, Minimalism, Expressivism. *International Journal of Philosophical Studies*, 18(3), 317–330.
<https://doi.org/10.1080/09672559.2010.492116>
- Williams, M. (2013a). How Pragmatists Can Be Local Expressivists. In H. Price (Ed.), *Expressivism, pragmatism and representationalism* (pp. 128–144). Cambridge: Cambridge University Press.
- Williams, M. (2013b). Knowledge, Ascriptivism and Defeasible Concepts. *Grazer Philosophische Studien*, 87(1), 9–36.

- Williams, M. (2015). What's so Special About Human Knowledge? *Episteme*, 12(2), 249–268.
- Wilson, D. S., & Wilson, E. O. (2007). Rethinking the theoretical foundation of sociobiology. *The Quarterly Review of Biology*, 82(4), 327–348.
- Wittgenstein, L. (1953). *Philosophical Investigations* (3rd ed.; G. E. M. Anscombe, Trans.). Oxford: Blackwell Publishers.
- Wittgenstein, L. (1972). *On certainty* (G. E. M. Anscombe & G. H. von Wright, Eds.). New York: Harper & Row.
- Wright, C. (2006). Intuitionism, Realism, Relativism and Rhubarb. In P. Greenough & M. Lynch (Eds.), *Truth and Realism* (pp. 38–60). Oxford: Clarendon Press.

Curriculum Vitae

Thomas Michael Wilk was born on March 17, 1983, in Cairnbrook, Pennsylvania. He is a 2001 graduate of Bedford Area High School. In 2005, he earned his B.A. in philosophy and economics from Gettysburg College, where he was the recipient of the Chan L. Coulter Award for outstanding work in philosophy. He completed his M.A. in Philosophy at George Mason University in 2009, earning the Outstanding Graduate Student Award. He began his Ph.D. at Johns Hopkins in 2010, where he was twice the recipient Miller Essay Prize for an outstanding essay in philosophy. During the 2016-17 academic year, his work was funded by the David Sachs Dissertation Fellowship. He has also been funded by the Leon Lauer Fellowship and the Hodson Fellowship in the Humanities. In 2011-12, Tom served as the President of the Hammond Society at Johns Hopkins. His research interests lie in the philosophy of language, ethics (meta-ethics and applied), epistemology, and social and political philosophy, all approached from the perspective of neo-pragmatism. His publications include "Inferences, Experiences, and the Myth of the Given" in *Logos & Episteme*, "Trust, Communities, and the Standing to Hold Accountable" in the *Kennedy Institute of Ethics Journal*, and "The Right Way to Win Over Posterity" in *Hamilton and Philosophy*. Tom has taught a wide variety of courses in philosophy at Johns Hopkins University, George Mason University, George Washington University, Gettysburg College, Goucher College, and the United States Naval Academy. He lives in Alexandria, VA, with his wife, Kate, sons, Sam and Jack, and his beagle, Plato.