

**From the holdings of the Department of Special Collections , Princeton University Library,
Princeton, NJ**

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyright law. We do not charge any permission or usage fees for the publication of images of material in our collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff. For more information regarding copyright, credit, and citation please review the information on the Copyright, Credit and Citation Guidelines page of our website.

<https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines>. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. <https://library.princeton.edu/special-collections/ask-us>

Doctoral Dissertations, 2011-2026

Philosophy

2016

On the Varieties of Normativity

Wodak, Daniel

Princeton, NJ : Princeton University

<http://arks.princeton.edu/ark:/88435/dsp01z316q407c>

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. <https://www.copyright.gov/title17/title17.pdf>

ON THE VARIETIES OF NORMATIVITY

DANIEL WODAK

A DISSERTATION
PRESENTED TO THE FACULTY
OF PRINCETON UNIVERSITY
IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE
BY THE DEPARTMENT OF
PHILOSOPHY

ADVISER: MICHAEL SMITH, GIDEON ROSEN,
AND PHILIP PETTIT

SEPTEMBER 2016

© Copyright by Daniel Wodak, 2016.

All rights reserved.

Abstract

Some varieties of normativity, like morality, seem to have an authoritative status. Other varieties of normativity, like etiquette, do not seem to have this status. One way to put this idea is that when we know that the dictates of morality and etiquette conflict, we should prioritize the former. Another is that we can know that etiquette requires us to perform some action while knowing that there is no reason to do so, because in these circumstances the requirements of etiquette are trivial or terrible.

This work systematically explores the varieties of normativity that are non-authoritative, or “merely formal”. Chapter Two develops the position that we can know that the dictates of standards like morality and etiquette conflict, and make a mistake about what to do by prioritizing the wrong standard. This is a mistake about the status, rather than the contents, of standards.

Chapters Three and Four develop my fictionalist account of mere formalities. This account is primarily a positive metaphysical account of what it is for requirements of etiquette and the like to be merely formal, but it is motivated in part by considerations about our attitudes and language. The basic idea behind this fictionalist proposal is that whatever essential feature makes standards like morality authoritative (reasons, values, or what have you), standards like etiquette are merely formal in virtue of having that feature according to the relevant fiction.

This account is fairly ecumenical insofar as it only relies on one central resource: a distinction between normative reality and normative fiction. But on metaphysically deflationary views about normativity (like expressivism and quietism) there is no good way of understanding that distinction. In Chapters Five and Six I show how this results in serious problems for expressivism and quietism, respectively: they don’t have the resources to explain the continuities and the dissimilarities between standards like morality and etiquette, so they end up with unpalatable commitments about meaning, speakers’ attitudes, and reasons for action.

Acknowledgements

In my penultimate year of high school (2003), Luke Russell taught a short, elective philosophy course called Mind and Morality. I wouldn't say that it woke me from my dogmatic slumber, because that was a slow process. (It took until grad school for me to realize that utilitarianism is neither obvious nor true.) But it did wake me from my dejectedness with my other academic pursuits. While the philosopher is often mocked for having her head in the clouds, from the first my interest in the field was that I was finally able to ask the questions that most concerned me: questions about what we ought to do in a deeply unjust world. My interests have broadened since, and I spend more than my fair share of time fussing over abstractions, but my motivations are still grounded in questions about how we should act and interact with each other.

When I went to the University of Sydney as an undergraduate I was lucky enough take further courses with Luke, who, along with David Braddon-Mitchell and Caroline West, had a significant influence on my philosophical thinking. In my Honours year (2008), Luke and Richard Joyce kindly agreed to supervise my thesis project on Michael Smith's fetishism argument against moral judgment externalism.

I didn't apply for grad school immediately after my Honours year. I completed my law degree first, for two reasons: the sunk cost fallacy, and Kevin Walton. I took all three of his fantastic jurisprudence courses, and he supervised my research project on vagueness and legal interpretation. Kevin not only indulged my near gluttonous appetite for legal philosophy, but was a supportive and challenging mentor.

It is no exaggeration to say that I wouldn't be where I am today without the feedback, support, and advice that Luke, David, Caroline, Richard, and Kevin offered while I was an undergraduate student. I am deeply grateful to all of them.

Once I left Sydney for Princeton I found myself studying with a mixture of philosophers whose work I knew and greatly admired (namely, Philip Pettit, Peter Singer, and Michael Smith) and philosophers whose work I would have admired if I wasn't so inexplicably igno-

rant of it (namely, everyone else). I'm not sure which group was more intimidating. But thankfully the faculty managed to generate an excellent intellectual environment—one that was stimulating and challenging, yet supportive. I learned a great deal from their seminars and supervision, and perhaps even more so from chance encounters in the lounge and casual conversations over tea.

In fact, on reflection my dissertation project was shaped most significantly by a confluence of chance events. Moving to the United States prompted me to read more about the prison-industrial complex. I became increasingly outraged about the behavior of the judiciary who, mostly in fulfilling their official duties, gravely wrong defendants and marginalized groups. Talking about Anthony Appiah's *The Honour Code* made me appreciate that, for better or worse, people frequently take social norms far more seriously than moral norms in determining how to act. Sitting in on a reading group on philosophy of law with Gideon Rosen made me even more convinced that Bentham and Austin have been unfairly dismissed in jurisprudence since H.L.A. Hart first published *The Concept of Law* (despite Hart's own far more sympathetic discussion in his *Essays on Bentham*)—I started trying to work out what Bentham's fictionalist account of legal obligations would look like if it were coupled with hermeneutic ambitions, and decoupled from his radical empiricist commitments. During a summer seminar on Moral Intuitionism at the Central European University someone asked why if we were metaethical naturalists we should find morality more philosophically important and normatively significant than etiquette. I was not persuaded that this was a particular problem for naturalists, but I was struck that (a) it was unclear why we should assume that standards like etiquette are philosophically uninteresting, especially because (b) providing an account of standards like etiquette should presumably form part of any sufficient explanation of the sense in which they are less normatively significant than morality.

I started thinking that the fictionalist view I was developing about legal obligations could explain social obligations more generally: how they are bona fide obligations, even though they do not have the same essential normative import as morality, prudence, and

epistemology. When I was talking about some of my work with David Plunkett, he suggested that I read Tristram McPherson’s “Against Quietist Normative Realism”, which helped me clarify my concerns considerably, as did many subsequent conversations with Tristram. When I pieced this together with my somewhat polemical interests in judicial obedience, the project started to come together.

The sections on expressivism in particular and philosophy of language more generally developed later. The former emerged after a long dialogue where Adam Lerner played the role of Socrates and I played the role of some Athenian with inchoate ideas about expressivism and univocality. The latter emerged after I returned to the literature in jurisprudence with a better understanding of what linguistically informed philosophers, like Angelika Kratzer, say about legal language.

All of this work was significantly shaped by my interactions with my three advisors: Michael Smith, Gideon Rosen, and Philip Pettit. They have been wonderful to work with. All three have provided so much invaluable feedback—challenging, insightful, clarificatory, and encouraging—on more drafts (or drafty, half-formed ideas) than I’d care to count. They’ve also been willing to let me run wild for a while even when it wasn’t clear what parts of this project would amount to. Unexpectedly, given Michael’s own work, one thing that has made these advisors ideal for me is that they proved to be remarkably tolerant when I did not do what they thought I ought to do. I’m very glad to know them, and to have had the opportunity to learn from them.

I have also been fortunate to get helpful feedback on parts of this project from other faculty at Princeton (especially Shamik Dasgupta, Johann Frick, Elizabeth Harman, Frank Jackson, Tom Kelly, Sarah Jane-Leslie, and Sarah McGrath), various visitors to the University Center for Human Values (especially Mark Budolfson, Jonathan Dancy, Pete Graham, Geoff Sayre-McCord, David Plunkett, and Angie Smith), recent and current graduate students at Princeton (Robin Dembroff, Ben Ewing, Mark Harris, Eric Hubble, Kyle Landrum, Yena Lee, Adam Lerner, Errol Lord, Barry Maguire, Michaela McSweeney, Sam Preston, and

Jack Woods), and other generous philosophers (Hrafn Asgeirsson, Matthew Chrisman, Matti Eklund, Max Etchemendy, Daniel Fogal, Jeff Kaplan, Sari Kisilevsky, Sebastian Koehler, Tristram McPherson, David Plunkett, Douglas Portmore, Michael Ridge, Alex Sarch, T.M. Scanlon, Nic Southwood, and Alex Worsnip). I've also had the pleasure of getting feedback from audiences at Australian National University, The University of Edinburgh, Hamilton University, Humboldt Universitat, The University of Michigan Ann Arbor, The University of Missouri-St Louis, New York University, the University of North Carolina Chapel Hill, and of course Princeton University. This is a fairly long list, but I'm not confident that it is complete. My apologies to anyone inadvertently left off. And my deep thanks to everyone who has helped with this project.

I'd like to also thank all of my family and friends (you know who you are) who were so supportive during the last five years, and especially those of you who reminded me that struggling to develop your philosophical position while doing a PhD at Princeton is a first world problem *par excellence*. When you spend too long with your own ideas it's easy to lose perspective. It's nice to have people who help you stay grounded. I'd especially like to thank my mother, Jo, for proofreading this entire dissertation.

Finally, I'd like to thank my closest friend and favorite person: my partner, Sukaina Hirji. It is with all my heart that I dedicate this work to you.

DW

Princeton April 2016

To Sukaina Hirji

My Hobbes and my Susie Derkins

Contents

Abstract	iii
Acknowledgements	iv
1 Introduction	1
2 The Case Against Judicial Obedience	6
2.1 Introduction	6
2.2 The Linguistic Analogy	8
2.3 Uncontroversial Incompatibilities	11
2.3.1 Political Duties	14
2.3.2 The Normal Justification Thesis	16
2.3.3 The Proper Role of the Judiciary	18
2.4 Status Mistakes	22
2.5 Civil Disobedience	30
2.6 Conclusion	36
3 Mere Formalities	37
3.1 Introduction	37
3.2 Foot's Fugitive Thought	39
3.3 Three Puzzles About Mere Formalities	43
3.3.1 Guidance	43
3.3.2 Internalization	47
3.3.3 Conflict	48

3.4	A Way Forward	50
3.4.1	What Makes Mere Formalities Different?	51
3.4.2	The Five Markers of Fictive Attitudes	54
3.4.3	Fictionalism About Mere Formalities	56
3.5	The Virtues of Fictionalism	58
3.5.1	Guidance, Internalization, and Conflict	59
3.5.2	Normativity <i>Sans</i> “Authority”	62
3.6	Miss Manners & <i>Downton Abbey</i>	64
3.7	Conclusion	66
4	How We Create Obligations	67
4.1	Why Is Legal Normativity Puzzling?	68
4.2	The Fictionalist Solution	75
4.3	How Fictionalism Fits Into the Familiar Landscape	80
4.4	What “Legal Obligation” Means	87
4.5	Conclusion	96
5	Expressivism and the Varieties of Normativity	98
5.1	Introduction	98
5.2	Setting the Stage	100
5.2.1	The Linguistic Phenomenon	100
5.2.2	Content and Character	101
5.2.3	Descriptivism	103
5.2.4	Expressivism	104
5.2.5	Varieties of Normativity	106
5.2.6	Taking Stock	107
5.3	Expressivism <i>Sans</i> Character	108
5.4	Semantic Expressivism	112

5.4.1	Objects of Hyperplans	113
5.4.2	Flavors of Approval	117
5.5	Meta-semantic Expressivism	122
5.5.1	A Disunified Meta-semantics	123
5.5.2	A Unified Meta-semantics	125
5.6	The Expressivist's Dilemma	126
5.7	Pragmatic Expressivism	129
5.8	Generalizability	132
5.9	Conclusion	135
6	Why Realists Must Reject Normative Quietism	137
6.1	Introduction	137
6.2	Normative Quietism	137
6.3	Creating Reasons Out of Thin Air	142
6.4	Separate, Conflicting Domains	146
6.5	One Domain, One Standard	160
6.6	Conclusion	166

Chapter 1

Introduction

The distinctions between varieties of normativity are intuitive: actions can be recommended, required, and permitted by morality, prudence, law, and etiquette. This list is far from exhaustive, as humans can and do create an enormous number of standards of correctness, from the requirements of prescriptive grammar to the rules governing gladiatorial combat.

Once we recognize that there is a wide variety of normative standards, interesting questions arise. Some of these questions concern which standards are *normatively authoritative*. When morality and etiquette conflict—when what etiquette requires is immoral and what morality requires is indecorous—which standard should we prioritize in determining what to do? Other questions concern *semantics*. Do we explain the meanings of our normative terms the same way when they are used in relation to all of these various standards? Other questions involve the *explanatory roles* that these standards play. What are the similarities and differences in how requirements from these various standards actually guide our actions, attitudes, and interactions? Finally, and to my mind most importantly, we can ask questions about the *metaphysics* of these varieties of normativity. There are some important continuities between morality and etiquette: they are both systems of recommendations, requirements and permissions, after all. But there are also intuitive metaphysical differences between these standards. Facts about what morality requires seem to be independent of us, whereas facts about what etiquette requires seem to be dependent on us: requirements

of etiquette, unlike requirements of morality, are something we collectively create with our shared attitudes.

My central aim in this project is to explore and answer these questions. This is first and foremost a project in an under-theorizing area of metaethics: the nature of social standards, like etiquette. But like most interesting projects in metaethics it makes forays into a range of other fields, such as social and political philosophy, philosophy of law, metaphysics, and philosophy of language.

My answers to these questions are as follows. Some varieties of normativity are indeed authoritative, and some are not. I take this to be highly intuitive even if we sometimes disagree about how to characterize some particular standards of correctness. Some varieties of normativity, like morality, have automatic justificatory force. With other standards, like law etiquette, the fact that we are required to perform some act does not settle whether we have any reason to act in accordance with that requirement. I argue for this in Chapter Two, focusing on how failures to heed this difference in kind can result in an important type of mistake in normative reasoning: judges, for instance, may well know the content of morality and law, and yet do the wrong thing because they have mistakenly prioritized their legal requirements over their uncontroversially incompatible moral duties. Instead, the judiciary should be engaged in widespread (and, I argue, collective) civil disobedience.

If varieties of normativity can differ in this way, what explains this difference? I take this question up in Chapter Three, offering an account of what it is for a standard like etiquette to be non-authoritative, or “merely formal”. The account I offer is a version of fictionalism. It explains the difference between standards like morality and standards like etiquette in terms of the distinction between reality and fiction. There is some essential feature that morality has in virtue of which it is authoritative; standards like etiquette, by contrast, have that feature according to a fiction, and in that sense they merely purport to be authoritative. In Chapter Three I show how this fictionalist account explains the important continuities

and differences between the explanatory roles that these standards play in how they guide our actions, attitudes, and interactions.

Chapter Four further supports this fictionalist account, showing how it has the resources to explain a further problem: how do we create systems of obligations and permissions whole cloth? My proposed answer is that we create systems of obligations and permissions, like etiquette, by creating a fiction about what we have reason to do. And we create such a fiction in the same way that we create any old fiction: by accepting its content. I show how this account both fills a lacuna left by the three main theories in contemporary philosophy of law, and can be easily coupled with orthodox views on which we have a unified semantics for normative terms like “obligation” when they are used in, for example, moral and legal contexts.

In Chapters Three and Four I strive to be as ecumenical as possible. All we need is a distinction between reality and fiction in contemporary metaethics to get fictionalism off the ground: that distinction is compatible with all forms of realism about standards like morality (so long as realism is understood in a non-deflationary sense). But anti-realists (and deflationary realists) make no such distinction: for them, if there is a normative reality, it’s one that comes cheaply. Chapters Five and Six show how such views—in particular, expressivism and quietism—face serious, neglected challenges with respect to the varieties of normativity. They don’t have the resources to explain the continuities and the dissimilarities between standards like morality and etiquette, so they end up having unpalatable commitments about the meaning of normative language, speakers’ attitudes, and reasons.

The point of these challenges is two-fold. First, to further bolster the positive view: fictionalism. In Chapters Three and Four all I do is show that it would be nice if fictionalism were true, because it would explain a great deal. I don’t systematically engage with all alternative views—there are too many possibilities to consider—in order to argue that we need to be fictionalists. So a good way to put fictionalism to the test is to consider the same issues from points of view on which there is no important metaphysical distinction

between normative reality and normative fiction. It turns out that if this were so, and hence if fictionalism were false, serious problems arise. This is no proof that fictionalism is true. I do not know how to offer such a proof in philosophy. But it does add plausibility to the positive picture: if a theory offers significant explanatory resources only at the cost of positing a distinction (between reality and fiction) that we cannot do without, that's a fairly strong case for accepting that theory, at least until a better competitor arises.

The second point of these challenges is that, coupled with the positive proposal, they offer a roundabout argument for the claim that we are committed to some form of realism about standards like morality. The notion of realism I have in mind here is an intuitive and comparative one, rather than a technical and stipulative one. This is, to my mind, as it should be. J.L. Austin (1962) once wrote that “real” is a “weasel word”: “a definite sense attaches to the assertion that something is real, a real such-and-such, only in the light of a specific way in which it might be, or might have been, not real”. To illustrate his point, Austin wrote “I don’t know just how to take the assertion that it’s a real duck unless I know just what, on that particular occasion, the speaker has it in mind to exclude” (p. 70). That’s the kind of intuitive and comparative use of “real” and “realism” I have in mind. So my view is *not* that standards like morality are real and standards like etiquette are not in the sense that the former exist while the latter do not, just as Austin’s view is not that real ducks exist while and fake ducks do not. My view is that standards like morality are real in the sense of being the genuine article (like real ducks), while standards like etiquette are artifacts that purport to be the genuine article (like fake ducks). It is in this sense that I claim we are committed to the thought that some standards like morality are real while other standards like etiquette are not. And this is the pre-theoretical commitment that, I argue, can be explained and vindicated by appealing to a distinction between normative reality and normative fictions, and that cannot be explained and vindicated without a distinction between normative reality and normative fictions.

If this is right, it reveals something interesting (and by no means unprecedented) about philosophical methodology. If we want to know the nature of real ducks, we should not *just* think about real ducks. Otherwise we might end up thinking that superficial features like swimming and quacking suffice, when those features are in fact shared with fake ducks. We need to engage in a comparative project to avoid mistaking superficial features for deep explanations. Philosophers might not be interested in the nature of ducks, but the same methodological point applies to other phenomena. If we want to know the deep nature of standards like morality, we need to engage in a comparative project to avoid making that same type of philosophical mistake. I doubt that this dissertation will offer the final word on this comparative project. But I hope that it at least sets us off to a good start.

Chapter 2

The Case Against Judicial Obedience

2.1 Introduction

In this Chapter I will defend three main claims. First, judges face uncontroversially incompatible moral and legal duties (§2.3). Second, judges perform immoral acts because they mistakenly give priority to their legal duties (§2.4). And third, the judiciary should instead engage in acts of civil disobedience to change the law (§2.5).

Such criticisms of the judiciary are surprisingly rare, in public discourse as well as in philosophy. Legal ethics continues to be dominated by questions about lawyers, as if judges who sit above the adversarial fray and apply the law are thereby exempt from moral scrutiny. But such criticisms are not unprecedented. In 1968, Robert Cover reviewed “a book long out of print, over 110 years old and almost certainly not soon to be reissued in paperback”. The book was described by its editor, Richard Hilbreth, as being “designed to hold up a mirror to magistrates now sitting on the American bench” for their role in maintaining slavery, especially by enforcing the Fugitive Slave Act of 1850.¹ Cover’s purpose was similar. His review was “a polemic against the American judiciary” for “remain[ing] faithful to its long tradition as executors of immoral law” and thereby being complicit in the “genocidal horror” of the Vietnam war.² After Cover wrote his 1968 polemic against the judiciary for remaining faithful to immoral laws, his colleagues “were more upset” about his “judgments

¹Hilbreth, as cited by Cover (1968), p. 1004.

²Cover (1968), p. 1005–1006 (footnotes omitted).

passed upon the revered and honored dead ... than at the attack upon the sitting bench for its part in Vietnam.”³ This prompted him to write *Justice Accused*, a monograph on antebellum judges; but it also sidelined his charge against the contemporary courts.

I do not want our discussion to be similarly derailed into an exclusive focus on historical or hypothetical cases. So let’s start by highlighting just some of the most serious injustices that are being inflicted by the U.S. legal system, in part by contemporary judges who remain faithful to their long tradition as executors of immoral laws.

In many respects, the situation is worse than it was when Cover wrote his polemical review. The U.S. is no longer engaged in Vietnam. But it is still fighting two equally nefarious wars: the War on Crime that President Johnson declared in 1965 and the War on Drugs that President Nixon declared in 1971. Partly as a result of these wars, the American legal system has become far more punitive. “In the span of little more than three decades,” Stuntz notes, “Americans first embraced punishment levels lower than Sweden’s, then built a justice system more punitive than Russia’s.”⁴ With 5% of the world’s population, the U.S. now has 25% of the world’s prisoners.

The burdens of mass incarceration disproportionately fall on marginalized groups—the poor, the mentally ill, and racial minorities. Black Americans make up 40% of America’s prison population and almost 10% of the world’s prison population, despite being 13% of America’s total population. The police and courts entrench a racial caste system, which Michelle Alexander (2010) calls “The New Jim Crow”.

The conditions of American prisons are also a cause for serious concern. “Too often,” Justice Kennedy noted in the recent U.S. Supreme Court case of *Davis v. Ayala*, “discussion in the legal academy and among practitioners and policymakers concentrates simply on the adjudication of guilt or innocence. Too easily ignored is the question of what comes next. Prisoners are shut away—out of sight, out of mind.” Many American prisons are overcrowded

³Cover (1975), “Acknowledgements”.

⁴Stuntz (2011), p. 34.

to a degree that causes, according to the US Supreme Court, “needless suffering and death”;⁵ according to the Justice Department’s most recent report,⁶ 80,000 prisoners were sexually abused in 2011 and 2012, mostly by custodial officers, only a minuscule proportion of whom were punished; and thousands of U.S. prisoners are subjected to years of solitary confinement despite the clear evidence that it causes and exacerbates serious mental illnesses.⁷

In some important respects, the situation today is better than it was in the past. Abolitionist judges were a rare breed in the antebellum era, and at least in many states a majority of voters supported returning fugitive slaves. Today, U.S. judges are united in their opposition to at least one of the central components of the wars on drugs and crime: the Judicial Conference of the United States has opposed mandatory minimum sentencing provisions regularly and consistently since 1953, and as a matter of “established policy” since at least 1962.⁸ And the particular sentences that mandatory minimums require are widely out of step with Americans’ judgments about just sentencing. Despite this, U.S. judges continue to apply mandatory minimums even when they know that the particular sentences they require are unjust. In a much lauded speech in 2013, then U.S. Attorney General Eric Holder said that “[t]oo many people go to too many prisons for far too long for no good law enforcement reason”. It is rarely mentioned that too many people go to prison for far too long in part because they are convicted and sentenced by judges. This raises our central questions. Should judges continue to apply obviously unjust laws like mandatory minimums? And if not, why do they continue to do so, and what should they do instead?

2.2 The Linguistic Analogy

Before we explore these issues in a legal context, it will be helpful to make a linguistic analogy. Similar issues arise with respect to conflicts between the requirements of prescriptive

⁵*Brown v. Plata* 4 131 S. Ct. 1910 (2011) (Kennedy J.).

⁶See Beck and et al (2013).

⁷See Browne and et al (2011).

⁸See Cassell (2007).

grammar and morality, and the comparative simplicity of these contexts makes it easier to see the structure and plausibility of the view I am proposing.

Consider gender pronouns. In contexts where the antecedent is of an unknown or indeterminate gender identity, it is grammatical to use “they” as a singular pronoun: consider “Each to their own” and “Someone is at the door, but I don’t know their name.” But in other contexts, the only grammatical singular pronouns are “he”, “she”, and “it”. This is morally problematic for several reasons. Some have to do with systematic features of these grammatical norms, which encode and transmit essentialist beliefs about gender.⁹ Others have to do with the offensiveness of using gendered pronouns like “he” or “she” (or the impersonal “it”) in relation to individuals who identify as gender-queer rather than masculine or feminine. If we were discussing the rapper Angel Haze, it would be ungrammatical to say “Their new album is great”, even if we know that Haze identifies as gender-queer and prefers the pronoun “they”.

This illustrates the first central point: it can be uncontroversial that the duties we have according to different standards, like morality and prescriptive grammar, conflict. That should not be surprising. We create the standards of prescriptive grammar, so it stands to reason that we might create grammatical requirements that do not comport with our moral duties. Nonetheless, some might resist this by pointing to the moral reasons we have to speak grammatically: we need grammatical norms to coordinate, communicate, and collect information. While this is true and important, we should not forget that even if *in general* we have good moral reasons to comply with *many or most* grammatical norms, it does not follow that *in every case* we have good moral reasons to comply with *this particular* norm. In some contexts, it might be clear from the context that Haze is a single person even though “they” is typically used as a plural pronoun; and in all contexts to use “he” or “she” would communicate false (and harmful) information about Haze’s gender identity.

⁹See, *inter alia*, Beit-Hallahmi et al. (1974), Wayne (2005), Senden et al. (2015), and Wodak et al. (2015).

At very least, we can imagine people who agree to all of the above, and so agree that their moral and grammatical requirements are incompatible. For those people, this incompatibility is obvious and uncontroversial. Nonetheless, these people might still disagree about what to do; in particular, they might disagree about whether to say things like “Their new album is great” rather than “Her new album is great”.

Such disagreements are not just possible, but actual. Consider an exchange in Ariel Levy’s profile of Jill Soloway.¹⁰ Soloway, famous for creating the television show *Transparent* (about a transgender parent), claims that morally we should use “they” in the third person. Levy responds: “I pointed out that strict grammar forbids using a plural pronoun for a single person; it would sound crazy, for instance, to describe Soloway by saying, ‘They are my favorite director’.” In the context, it is clear that Levy is not objecting to the claim that morally we should use “they”; rather, the objection is that we should not use “they” for a single person because “strict grammar forbids” this. In other words, the disagreement between Soloway and Levy does not concern the content of morality or grammar. It concerns which standard we should give priority to in determining how to write and speak. Plausibly, Soloway and Levy disagree about the status, rather than the contents, of morality and grammar.

I’m inclined to think that Soloway is on the right side of this disagreement. I won’t argue for this, but it is a view that many share. For instance, consider this excerpt from an article published around the same time in the *The Washington Post*.

For many years, I’ve been rooting for—but stopping short of employing—what is known as the singular *they* as the only sensible solution to English’s lack of a gender-neutral third-person singular personal pronoun. (Everyone has *their* own opinion about this.) *He* once filled that role, but a male default hasn’t been palatable for decades. Using *she* in a sort of linguistic affirmative action strikes me as patronizing. Alternating *he* and *she* is silly, as are *he/she*, *(s)he* and

¹⁰‘Dolls and Feelings’, *New Yorker*, December 14, 2015.

attempts at made-up pronouns. The only thing standing in the way of *they* has been the appearance of incorrectness—the lack of acceptance among educated readers.¹¹

Here Bill Walsh is announcing that the *Post* will now use the singular “they” as a singular third person pronoun, at least in relation to “gender-neutral people”. Note that Walsh claims that the “only thing standing in the way of *they*” is not a moral objection, but its (grammatical) incorrectness. If this is right, it illustrates the second central point: that we sometimes act immorally because we mistakenly give priority to a non-moral standard of correctness, in this case grammaticality.

What about the third central point? If we should use “they” rather than “he” or “she” in relation to gender-queer people, we should deliberately violate the conventional standards of prescriptive grammar. Not only does doing so show more respect to individuals like Angel Haze, but it can help change what is grammatical. In that sense, it is a form of civil disobedience. Admittedly, there is not always a great deal that lone speakers can do. But those who are taken to be authorities on language, such as prominent newspaper editors like Walsh, have more power to effect change, and hence a greater responsibility to do so. And when we shift our focus from individuals to groups, the case for civil disobedience becomes even stronger: so much the better, then, that the American Dialect Society voted to make the singular “they” the 2015 Word of the Year, explicitly on the basis that it should be used to refer to known persons who reject “the traditional gender binary of he and she”.

2.3 Uncontroversial Incompatibilities

Conflicts between moral and legal requirements can be just as stark as conflicts between moral and grammatical requirements. But the relative complexity that legal contexts introduce may make this initially harder to see.

¹¹Bill Walsh, ‘The Post drops the ‘mike’—and the hyphen in ‘e-mail’’, *The Washington Post*, December 4, 2015.

To focus matters, let's turn to what will be our central example. In 2002, Weldon Angelos—a 22-year-old man with no criminal record—sold \$350 worth of marijuana to a police informant on three occasions. During two of these deals he was in possession of a Glock. When he was arrested, marijuana and further handguns were found in his apartment, and in police searches of other locations.

Angelos was offered a plea bargain that recommended a prison sentence of fifteen years. He was told that if he refused the offer he would be charged with 20 counts, including five counts of possessing a firearm in furtherance of a drug trafficking crime in violation of §924(c). He refused the plea. At trial, he was convicted of 13 drug, firearm, and money laundering charges, and three violations of §924(c). Our concern is not his conviction, but his sentence.

In November 2004, Angelos was sentenced to a 55–63 year prison term by Judge Paul Cassell of the U.S. Court for the District of Utah.¹² Cassell himself described this sentence as “unjust, cruel, and irrational”. So why did he impose it? Following the Sentencing Guidelines (which will be discussed later), Angelos was sentenced to eight years in prison for 13 drug, firearm, and money laundering charges. This leaves §924(c) of the federal code, which provides that

any person who, during and in relation to any crime of violence for which the person may be prosecuted in a court of the United States, uses or carries a firearm, or who, in furtherance of any such crime, possesses a firearm, shall, in addition to the punishment provided for such crime of violence . . . be sentenced to a term of imprisonment of not less than 5 years. In the case of a second or subsequent conviction under this subsection, the person shall . . . be sentenced to a term of imprisonment of not less than 25 years.

Three violations of §924(c) carry an additional minimum sentence of 55 years. If he had been convicted of five §924(c) counts, it would have been 105 years.

¹²*United States v. Angelos*, 345 F. Supp. 2d 1227 (D. Utah 2004). Hereafter, “*Angelos*”.

The plain meaning of this very precise text, and the clear intent of the statute, clearly require Cassell to sentence Angelos to a minimum of 55 years in prison. These sources of law do not give Cassell scope for what Hart (1994) called “interstitial discretion”, or for what Dworkin (1986) called “constructive interpretation”, either of which would otherwise legally allow Cassell to impose a shorter sentence: “a term of imprisonment of not less than 25 years” cannot be constructively interpreted to mean “a term of imprisonment of less than 25 years”. This is especially clear since, as a trial judge, Cassell must follow binding precedents from higher courts that decisively resolve all relevant ambiguities about the interpretation of §924(c),¹³ as well as the constitutionality of severe, mandatory minimum sentences.¹⁴ It may be that justices in higher level courts are not legally required to affirm this sentence; perhaps the U.S.S.C. could depart from its recent decisions and declare that §924(c) imposes “cruel and unusual punishments” in violation of the Eighth Amendment. But unless or until the U.S.S.C. does so, trial judges are legally required to apply §924(c).

Is this legal requirement uncontroversially incompatible with morality? At very least, it is obvious that Cassell had a strong *pro tanto* moral duty not to impose a punishment that was, in his words, “unjust, cruel and irrational”. Half a century of incarceration is obviously disproportionate to the severity of Angelos’ offense—as Cassell noted,¹⁵ Angelos’ minimum sentence is much more severe than the maximum sentences for far more serious offenders—and this is plausibly sufficient to render his sentence unjust.¹⁶ Moreover, during this 55–63 year prison term, Angelos will be deprived of liberty, social standing and suffrage,

¹³See, *inter alia*, *Smith v. United States* 508 U.S. 223 (1993) (on the meaning of “use”), *Muscarello v. United States* 524 U.S. 125 (1998) and *Deal v. United States* 508 U.S. 129 (1993) (on “count stacking”), and *United States v. Gonzales* 520 U.S. 1 (1997) (on non-concurrence).

¹⁴See *Harmelin v. Michigan*, 501 U. S. 957 (1991); *Ewing v. California*, 538 U. S. 11 (2003); *Lockyer v. Andrade*, 538 U. S. 63 (2003); *Rummel v. Estelle*, 445 U. S. 263 (1980); *Hutto v. Davis*, 454 U. S. 370 (1982). Cf. *Solem v. Helm*, 463 U. S. 277 (1983).

¹⁵See Cassell (2007), who compares the minimum sentence for Angelos (738 months), to the maximum sentences for: a kingpin of a major drug trafficking ring in which death resulted (293 months), an aircraft hijacker (293 months), a terrorist who detonates a bomb in public intending to kill bystanders (235 months), & a second-degree murderer (168 months).

¹⁶Such proportionality principles are widely endorsed. See Duff (2013).

and exposed to unacceptably high risks of being sexual abused,¹⁷ and of being subjected to torturous terms of solitary confinement for arbitrary reasons. This is not to mention the harms Angelos' punishment imposes on others (such as the \$1.3 million bill to the taxpayer).¹⁸ It is very unlikely these harms are counterbalanced by, say, the deterrence effects of Angelos' punishment, especially since drug trafficking is notoriously difficult to deter.¹⁹

Of course, this only shows us that the law is unjust. That much may be uncontroversial.²⁰ But it is possible that Cassell has a moral duty to apply the law even when the law is unjust. We will consider three possible sources for such a duty below. But before we do so, it is important to acknowledge the same point we made above with respect to grammar: even if *in general* judges have good moral reasons to apply *many or most* laws, it does not follow that *in every case* they have good moral reasons to apply *this particular* unjust law. It is at least possible that mandatory minimums are one of the cases where, in Dworkin's words, "judges may have a moral obligation to ignore the law when it is very unjust".²¹ In fact, this turns out to be very plausible.

2.3.1 Political Duties

We have seen that Cassell was legally required to apply §924(c) in *Angelos*, and that Cassell had a *pro tanto* moral duty not to do so because the mandatory minimum sentence for Angelos was unjust, cruel and irrational. So far, this is just like the linguistic analogy above—a case of incompatible requirements. But of course, legal contexts introduce greater

¹⁷For the sake of simplicity, let's crudely assume that the probability of Angelos' being sexually abused in any given year is equal to proportion of prisoners who were sexually abused in 2011–2012 (4%), and that if Angelos is abused in any given year he will be abused as many times as the average victim in 2011–2012 (4 times). 4% of 63 multiplied by 4 is 10.08.

¹⁸As Judge Myron Bright of the Eighth Circuit Court of Appeals wrote in *United States v. Hively* 61 F.3d 1358, 1363 (8th Cir. 1995), mandatory minimums require "excessively long" sentences that "not only ruin lives of prisoners and often their family members, but also drain the American taxpayers of funds which can be measured in billions of dollars".

¹⁹Estimates of this effect "are all over the place" (Rakoff (2015)). But evidence does show that deterrence is ineffective *apropos* drug trafficking. See (e.g.) Felbab-Brown (2013).

²⁰Notably, this is uncontroversial for legal positivists and anti-positivists alike. Fuller (1969) and Finnis (1989, 2002) go to great pains to reject Augustine's slogan that an unjust law is no law at all. For more discussion about legal positivism and anti-positivism, see Chapter Three.

²¹Dworkin (2008), p. 18. See also Dworkin (1986), p. 262.

complexity because there are several reasons why Cassell might have a moral duty to apply §924(c) even though it is unjust. We will now consider the three most plausible sources for such a moral duty, starting with Cassell's duty as a legal official to respect democratic processes and outcomes.²²

Would applying §924(c) show respect for democratic processes and outcomes? The people elected Congress. Congress passed §924(c). But consider how this sausage was made. §924(c) was first introduced in 1968 as a benign sentencing enhancement in a floor amendment;²³ the enhancement was made more and more draconian in a long series of similar amendments, following congressional debates that were rudimentary at best.²⁴ The sentences now required by §924(c) are wildly out of step with Americans' judgments about just sentencing in particular cases.²⁵ And the burdens that mandatory minimums impose largely fall on marginalized communities and disenfranchised individuals who have little democratic power, partly due to mandatory minimums themselves: felony disenfranchisement laws "have been more effective in eliminating black voters in the age of mass incarceration than they were during Jim Crow", and since prisoners are counted as (non-voting) residents in the jurisdictions in which they are incarcerated, "[w]hite rural communities that house prisons wind up with more people in state legislatures representing them, while poor communities of color lose representatives because it appears their population has declined".²⁶

Philosophers may be right to wax lyrical about the respect that citizens and officials owe to democratic laws. But laws and legal systems can be more or less democratic,²⁷ and §924(c) is surely not on the high end of that spectrum: it does not result from informed collective

²²For different accounts of the basis for this political duty, cf. Rawls (1971) and Waldron (1999).

²³To the Gun Control Act, after the assassinations of JFK and Martin Luther King, Jr.

²⁴Cassell (2007) cites former Chief Justice William H. Rehnquist's statement that "Mandatory minimums ... are frequently the result of floor amendments to demonstrate emphatically that legislators want to 'get tough on crime', without any careful consideration."

²⁵Cassell (2007) cites several studies to support this claim. In one "support [for mandatory minimums] dropped precipitously from 88 percent to a mere 17 percent when subjects are asked to apply mandatory sentences to specific hypothetical cases". When asked to evaluate mandatory minimums without considering specific cases, subjects' judgments are distorted by various well-known biases, like the availability heuristic. See Stalans (1993).

²⁶Alexander (2010), p. 193.

²⁷See Goldman (2015).

decision-making, reflect the will of the *demos*, or manifest equal respect for all citizens. Whatever moral reasons this gives Cassell to apply §924(c) are quite weak, which makes it very plausible that they will be outweighed by the strong moral reasons that Cassell has not to subject Angelos to disproportionate and excessively harmful punishments. By analogy, some similar considerations may count in favor of soldiers obeying their commanders; but when a platoon leader repeatedly orders his subordinates to fire on unarmed citizens in circumstances that they know amount to “straight murder”, they should refuse.²⁸

2.3.2 The Normal Justification Thesis

So it does not seem that §924(c)’s democratic pedigree gives Cassell a moral duty to impose an unjust, cruel, and irrational sentence on Angelos. But there are other possible sources for such a duty to obey unjust laws. On Joseph Raz’s view, for instance, the normal justification for obeying the law is based on epistemic considerations.²⁹

What epistemic considerations might give Cassell a moral duty to apply §924(c)? Cassell’s judgment of what punishment is reasonable is inconsistent with Congress’s assessment of what is reasonable. This gives Cassell two options: he could accept Congress’s directives “as authoritatively binding” and try to follow them; or he could try “to follow the reasons which apply to him directly”. Which of these options is he justified in believing would make him “likely better to comply with reasons which apply to him” independently of those directives? It seems plausible that Cassell should believe that as an individual, he is more likely to be wrong about what is reasonable than an institution like Congress.³⁰ If so, that Congress directs him to sentence Angelos to 55–63 years in prison is a reason to do so “which is not to be added to all other relevant reasons when assessing what to do, but should exclude and take the place of some of them”.³¹ In other words, this epistemic consideration need

²⁸For a recent example, see “Cause Célèbre, Scorned by Troops”, *The New York Times*, February 24, 2015. See also McMahan (2009), p. 71–73.

²⁹See Raz (1996). For excellent discussion of the relation between epistemic considerations and the normal justification thesis, see Melissa Lane (1999).

³⁰See Estlund (2007) on institutions’ “epistemic superiority”.

³¹Raz (1996), p. 46.

not *outweigh* the moral reasons not to subject Angelos to unjust punishments. Instead, it provides a moral duty that *excludes* any consideration of these other reasons.

Even if this normally justifies judicial obedience, it does not justify obedience in cases like *Angelos*. We can see why by noting an ambiguity in the options Raz describes. Should Cassell consider the set of Congress's directives, or each law individually? On the first interpretation, the question is whether Cassell should believe that his judgments are more likely than Congress's directives to be wrong about what is reasonable *in any arbitrary case*. On the second, the question is whether Cassell should believe that his judgment is more likely than §924(c) to be wrong about what is reasonable in *Angelos*. To illustrate the difference, consider GPS directions. I can know that I am less likely than Google Maps to be right about what route would be best at any arbitrary time while knowing that it is best not to follow Google's specific direction to drive down a closed-off road.

If we consider Congress's directives as a set, Cassell should not obey §924(c). After all, Congress is at best an internally inconsistent epistemic authority. It provides two radically inconsistent standards for determining what punishments are reasonable: the Sentencing Guidelines, and mandatory minimums.

Whereas the guidelines permit a degree of individualization in determining the appropriate sentence, mandatory minimums employ a relatively narrow approach under which the same sentence may be mandated for widely divergent cases. Whereas the guidelines provide for graduated increases in sentence severity for additional wrongdoing or for prior convictions, mandatory minimums often result in sharp variations in sentences based on what are often only minimal differences in criminal conduct or prior record. Finally, whereas the guidelines incorporate a "real offense" approach to sentencing, mandatory minimums are basically a "charge-specific" approach wherein the sentence is triggered only if the prosecutor chooses to charge the defendant with a certain offense.³²

³²Senator Orrin Hatch, cited by Cassell (2004), p. 1045. See also *United States v. Harris*, 536 U.S. 545, 570 (2002) (Breyer, J., concurring) ("Mandatory minimum statutes are fundamentally inconsistent with

Based on the prosecutor's choice to charge Angelos with §924(c), Cassell was legally required to sentence Angelos to an additional 660 months under §924(c), rather than 24 months under the Sentencing Guidelines, for the exact same criminal conduct (his firearm possession). Knowing this, how could Cassell reasonably conclude that an additional 660 months is just and reasonable?

What if we consider §924(c) on its own? Cassell still should not obey §924(c). To sentence Angelos to 55–63 years in prison is *obviously* unjust and irrational. It is obviously not supported by the balance of reasons that apply independently of §924(c). Cassell knew this was so; and his judgment was supported by the jury, the public, the Guidelines, and numerous amicus briefs from legal experts. To defer to one of Congress's directives here because as an institution it is generally more reliable would be as foolhardy as following Google's directions to drive down an obviously closed-off road because you trust GPS systems, even when Apple Maps confirms your view that this is a radically mistaken route.

2.3.3 The Proper Role of the Judiciary

We have seen that even if *in general* judges have good moral reasons to apply *many or most* laws on the basis of political and epistemic considerations, it does not follow that Cassell had a moral duty to apply §924(c) in *Angelos*. The same holds, I will now argue, for the final putative source for judge's moral duties to apply unjust laws: the judiciary's role-based duties. To focus matters, let's first outline David Luban's influential framework for identifying the role-based moral obligations of agents in complex institutions. On this view, the agent

- (1) justifies the institution by demonstrating its moral goodness; (2) justifies the role by appealing to the structure of the institution; (3) justifies the role

Congress' simultaneous effort to create a fair, honest, and rational sentencing system through the use of Sentencing Guidelines").

obligations by showing that they are essential to the role; and (4) justifies the role act by showing that the obligations require it.³³

Let's grant that the U.S. legal system is justified (or morally good). And let's grant that the structure of this system must involve some division of moral labor. The proper role of the legislature is to make informed, democratic laws; and the proper role of the judiciary is to apply those laws in particular cases.

This suggests that judges have a role-based moral duty to apply the law, which is quite plausible, and supported by both the oaths of judicial office and the American Bar Association's Model Code of Judicial Conduct.³⁴ In Luban's terminology, then, for Cassell to apply §924(c) is a "role act" that is required by Cassell's "role obligation" as a judge to apply the law, and that role obligation is justified by the structure and moral goodness of the U.S. legal system.

Of course, as Luban notes, an agent's role-based duties must "be balanced against the moral reasons for breaking the role".³⁵ So does this duty to apply §924(c) outweigh Cassell's moral reasons not to unjustly punish Angelos?

It might seem that this question turns on a general theoretical issue about the comparative weight of role-based and role-independent moral duties. For some, like Luban, in particularly unjust cases legal officials should be "willing to deviate" from their role-based duties on the ground that doing so is required by "moral common sense".³⁶ For others, like Bradley Wendell (2010), the duty to be faithful to the law must be stronger than officials' duties to do justice in particular cases, so that the legal system fulfills its fundamental roles of resolving social disagreements and providing equal treatment of citizens before the law.

Fortunately, we can set aside such theoretical disagreements. In *Angelos*, Cassell's role-based duties on balance militate against applying §924(c). So however we weigh role-based

³³Luban (1988), p. 131.

³⁴28 U.S. Code §453; A.B.A. MODEL CODE OF JUDICIAL CONDUCT (2011), Rule 1.1.

³⁵Luban (1988), p. 125. See also Applbaum (1999), *passim*.

³⁶Luban (2012), p. 688. See also Luban (1990), *passim*.

duties against other moral reasons, judges like Cassell should not apply mandatory minimums in cases like *Angelos*.

This is in part due to concerns that arise out of the unjust treatment of *Angelos*. One role-based moral duty of the judiciary is to administer justice.³⁷ This is also reflected in the content of judicial oaths: under federal law,³⁸ “[e]ach justice or judge of the United States” solemnly swears to “administer justice without respect to persons, and do equal right to the poor and to the rich”. And since the legitimacy of the judiciary partly depends on its acceptance by the people, another role-based duty of judges is to promote confidence in the judiciary. This is reflected in the Model Code of Judicial Conduct:³⁹ judges must be perceived to administer justice. As Cassell was well aware, *Angelos*’ minimum sentence is so obviously unjust that it brings the law into disrepute.

What about Wendell’s more systematic concerns? The first can be set aside quite easily. Judicial obedience to mandatory minimums is not essential to the capacity of the legal system to fulfill its function to resolve social disagreements in a pluralistic society. After all, the sentences required by laws like §924(c) are uncontroversially unjust and have a questionable democratic pedigree.

The second—ensuring equality before the law—militates against applying §924(c). The application of mandatory minimum sentences generates stark inequalities before the law. To see why, recall that due to the prosecutor’s choice to charge *Angelos* with §924(c), Cassell was legally required to sentence *Angelos* to an additional 660 months under §924(c) rather than 24 months under the Sentencing Guidelines for the exact same conduct (his firearm possession). By enacting mandatory minimums like §924(c), Congress has transferred control over sentencing from judges to prosecutors. That transfer is concerning in itself: sentencing ought to be the responsibility of accountable judges rather than largely unaccountable

³⁷See Greenawalt (1987), p. 281.

³⁸28 U.S. Code §453.

³⁹A.B.A., *supra* note 38, Rule 1.2.

prosecutors. But it becomes all the more concerning when we consider the inconsistencies in how prosecutors exercise this discretion.

From 1990 to 2000 between 10 and 45% of offenders who qualified for a §924(c) enhancement of their sentence were charged under the statute. To the extent that there are patterns in who prosecutors subject to mandatory minimums, they are disturbing. For one, mandatory minimum sentences are disproportionately applied to racial minorities. In 2012, of the 2,189 offenders who were convicted under §924(c), 52.2% were Black, while 18.2% were White. For another, severe mandatory minimums are applied *most* frequently to the *least* culpable offenders. In 2012, 31.9% of those convicted under §924(c) belonged to the least serious category of offender (Category I); this is almost as many as the three most serious categories of offenders combined (Categories IV–VI).⁴⁰

A partial explanation for these patterns is that prosecutors use mandatory minimums as a threat in negotiations over plea bargains. After Angelos rejected a plea bargain that recommended a total term of incarceration of 15 years, he was charged with seven counts of §924(c); had he been convicted on all counts, Cassell would have been required to impose a 105-year sentencing enhancement. Plea bargains are typically accepted by “flagrantly guilty repeat offenders”; they are most likely to be rejected by “first offenders in borderline situations”, who “have plausible defenses and are more likely to insist upon trial”.⁴¹ This results in a criminal justice system where first time offenders like Angelos face far more severe punishments for exercising their constitutional right to a trial.⁴² It also results in a criminal justice system where close to 95% of cases are resolved by plea bargains: even innocent

⁴⁰The breakdown was: Category 1, 31.9%; Category II, 12.0%; Category III, 19.4%; Category IV, 10.9%; Category V, 5.9%; and Category VI, 19.8%. So 63.3% of those convicted of a charge carrying a severe sentencing enhancement belonged to one of the three least serious categories of offenders, and only 36.7% to one of the three most serious categories of offenders. These data are all from The United States Sentencing Commission, *Quick Facts On Section 924(c)* (May 2015).

⁴¹Schulhofer (1993), p. 203.

⁴²For a detailed study of this “trial penalty”, see Hum (2003).

individuals will, and do,⁴³ plead guilty to avoid the considerable risk of being sentenced to decades (or worse) behind bars.

So by enacting mandatory minimums like §924(c), Congress has transferred discretion in sentencing from judges to prosecutors, and allowed the latter to use that discretion inconsistently and abuse it wantonly. As a result, the criminal justice system systematically punishes the innocent (who are bullied into accepting pleas), severely punishes first-time, low-level offenders like *Angelos* (who don't plead guilty), and punishes equally culpable offenders differently (due to the inconsistent use of prosecutorial discretion). Given how laws like §924(c) produce particular and systematic injustices in the U.S. legal system, it is highly implausible that judges' role-based duties justify applying mandatory minimums in cases like *Angelos*. For the judiciary to fulfill its proper roles, it cannot simply rubber stamp inconsistent and irresponsible exercises of prosecutorial discretion no matter how immoral they may be.

2.4 Status Mistakes

That concludes my defense of the claim that judges face uncontroversially incompatible moral and legal duties in cases like *Angelos*. I take it that such conflicts are just as stark as the conflicts between moral and grammatical requirements. This is just initially harder to see due to the relative complexity of legal contexts.

Such conflicts are recognized by judges—hence the judiciary's longstanding opposition to mandatory minimums. Cassell “reached out to the jury, the legal community and beyond in an effort to resolve the conflict between his sense of justice and the law”.⁴⁴ He even wrote that he was “ethically obligated to bring this injustice to the attention of those who are in a position to do something about it”.⁴⁵ (As a result, he was criticized by an appellate court.⁴⁶) But Cassell applied §924(c) regardless. Why?

⁴³2–8% of convicted felons who plead guilty are estimated to be innocent: see Rakoff (2014) and references therein.

⁴⁴Nilsen (2006), p. 539.

⁴⁵*Angelos*, at 1261.

⁴⁶See *United States v. Angelos*, 433 F.3d 738, 753 (10th Cir. 2006).

It is hard to be highly confident about diagnosing why a particular individual performed an act in complex circumstances. But the interesting question here is not about a particular individual, but about the general pattern that this particular individual represents. While the judiciary have been united in their opposition to mandatory minimums since at least 1962, they have also been united by their willingness to apply mandatory minimums even in egregious cases like *Angelos*. This holds for judges who span the political spectrum from conservative to liberal. Cassell—a conservative Bush-appointee—applied §924(c) to *Angelos*. Likewise, Justice Reinhardt—a liberal Carter-appointee—affirmed its application to Marion Hungerford, who received what he regarded as an “irrational, inhumane, and absurd” minimum sentence of 155 years.⁴⁷ The question that concerns us is what explains this pattern.

My answer is that judges perform immoral acts because they mistakenly give priority to their legal duties. In attempting to tacitly justify their judicial fidelity, judges like Cassell make a type of mistake that has not been clearly diagnosed. They make a mistake about the status rather than content of incompatible standards.

Of course, this type of mistake will typically only be made implicitly. Few explicitly state that when morality and the law conflict they give priority to the latter. The most explicit instance of such reasoning I have found is in a much-neglected article in the literature on punishment. “Under what circumstances”, J.D. Mabbott asked, “is the punishment of some particular person justified and why?”⁴⁸ He answered: it is justified when and because a person is guilty of a crime, where this is understood to be “a purely legal matter.”⁴⁹ His

⁴⁷See *United States v. Hungerford* 465 F.3d 1113 (2006). Hungerford was an impecunious, unemployed 48-year-old woman who had no criminal record and suffered from a severe form of Borderline Personality Disorder. She was convicted of seven violations of §924(c) for her “extremely limited” role in conspiring in and aiding and abetting armed robberies. She never even held a gun. See also HLR (2007), pp. 1989, 1991 (“To any student of precedent, this was an easy case [but] the Ninth Circuit should have refused to apply [§924(c)]. It should have reversed Marion Hungerford’s sentence.”)

⁴⁸J.D. Mabbott (1939), p. 152.

⁴⁹Mabbott (1939), p. 154.

view requires only a connection “between punishment and crime”, rather than between punishment and *any* moral consideration.⁵⁰

Mabbott argues that his view “explains”, where “all others do not”,

the case of punishment imposed by an authority who believes the law in question is a bad law. I was myself for some time disciplinary officer of a college whose rules included a rule compelling attendance at chapel. Many of those who broke this rule broke it on principle. I punished them. I certainly did not want to reform them; I respected their characters and their views. I certainly did not want to drive others into chapel through fear of penalties. Nor did I think there had been a wrong done which merited retribution. I wished I could have believed that I would have done the same myself. My position was clear. They had broken a rule; they knew it and I knew it. Nothing more was necessary to make punishment proper.⁵¹

Recall that Mabbott’s initial question was about when and why punishment is justified, not when and why it is *morally* justified. For Mabbott, no such moral question “is relevant”, as punishment is a legal matter: a judge “upholds a bad law ... simply because he has no [legal] right to dispense from punishment.”⁵²

Mabbott’s view does not involve any mistake about the content of morality. How could it? He makes no claims about the content of morality. He does not appeal to political or epistemic or role-based moral reasons to apply bad laws. On his view, legal standards should take priority over moral standards in determining when a particular person will be punished. So Mabbott imposed punishments that he knew to be immoral because he made a mistake about the *status* of morality and law.

⁵⁰Mabbott (1939), p. 155 and *passim*.

⁵¹Mabbott (1939), p. 155.

⁵²Mabbott (1939), p. 157.

This diagnosis raises two crucial questions. First, what is it to make a mistake about the status of standards like morality and law? And second, why should we believe that judges make this mistake, as opposed to other diagnoses?

I will only offer a partial answer to the first question here. To provide a full answer would require us to take up the much larger task of explaining what it is for there to be such a difference in the normative status of standards like morality and law; I defer that task until Chapter Three. For now, it will be enough to offer an intuitive gloss on what it is to make a mistake about the status rather than content of normative standards, then offer further examples that distinguish this from nearby phenomena.

Here's the intuitive gloss. We are subject to different standards, or varieties of normativity, many of which issue in requirements. Sometimes those requirements conflict: actions that are forbidden by one standard are demanded by another. We will flout one requirement (and hence one standard) no matter what we do. And since requirements are 'strict', and issue from different standards, this is not a context wherein we can weigh these requirements against each other to determine what to do. So instead, we must prioritize one standard over another in determining what to do. There are right and wrong ways to resolve such conflicts. We make a mistake about the status rather than content of standards when we resolve them the wrong way.

This provides a simple way to explicate Mabbott's mistake. He took morality and the rules of his college to issue in incompatible requirements. And he gave priority to the latter: the requirements issued by the rules of the college settled what to do, such that the requirements of morality were no longer even "relevant". This is strikingly similar to Levy's implicit reasoning about the singular "they" from §2.2: that "strict grammar forbids" such usage settles what to do, regardless of the force of Soloway's arguments about what morality requires. This helps us see how Mabbott and Levy make the same distinctive type of mistake. Neither believed any false claims about what is required by morality, law, grammar, or the

college. Rather, both resolved conflicts between morality and these non-moral standards the wrong way. They ought to have given priority to morality in determining what to do.

Two further examples can help us distinguish this distinctive type of mistake from nearby phenomena. First, consider a familiar circumstance in which there is some tension between what is moral and what is prudent: perhaps morality would require you to resign from your stimulating, well-paid job in an awful institution. Someone might resolve this tension the wrong way by ignoring what's moral and doing what's prudent instead. In some respects, this is similar to the type of mistake described above. But there are crucial differences. Prudence seems to issue in recommendations rather than requirements, and recommendations do not even purport to settle what we should do all things considered; so to take the recommendations of prudence to settle what to do is to make some mistake about the content of prudence. Moreover, it is natural to think that the content of morality already incorporates prudential considerations: whether an action is morally required may depend on the degree of self-sacrifice it involves. So sticking with your job to avoid a degree of self-sacrifice below that threshold may also involve a mistake about the content of morality.

Second, consider a tension between prudence and a conventional standard. For instance, take the “severe restrictions”, “rules” and “standards of behavior” that governed gladiatorial combat.⁵³ Gladiators were required to abide by these restrictions even when bloodthirsty viewers and emperors preferred more anarchic bouts. Many gladiators hesitated to break the rules even when they knew that “such behavior was likely to infuriate the emperor”, resulting in a serious risk to their lives.⁵⁴ While this case does not involve conflicting requirements, it seems much closer to the distinctive type of mistake described above. For gladiators to risk life and limb in order to comply with the severe requirements imposed on them by their captors seems unreasonable; the recommendations of prudence warrant more weight in their deliberations than the requirement that *murmillones* fight like *murmillones*.

⁵³For historical discussion, see Carter (2006).

⁵⁴Carter (2006), p. 108.

If this is right, it suggests that this distinctive type of mistake arises not because of something about the nature of morality alone (for instance, that moral considerations override all considerations supplied by all other normative standards), but because of something about how standards like morality and prudence compare to conventional standards like the requirements of games and grammar. The former seem to be essentially important, or necessarily reason-giving, in a way that the latter do not.⁵⁵

With this gloss on what it is to make a mistake about the status of standards in mind, let's turn to why we should believe that judges make this type of mistake in cases like *Angelos*. In effect, this was already suggested by Mabbott. He may be wrong that his view succeeds as a *justification* of judicial obedience. But it may still succeed as the best *explanation* of why judges apply laws that they know to be immoral. I will offer five factors that favor this explanation against alternative diagnoses on which judges are tacitly appealing to some moral justification for applying laws like §924(c).

First, my diagnosis provides a compelling explanation for the distinctive language that is present in justifications of judicial obedience. The distinctive language that is present in such justifications typically consists of bare modals. According to Mabbott, a judge upholds a bad law simply because he must do so—he has “no right to dispense from punishment”. Mabbott does not qualify the relevant modals, or other normative terms, in ways that signal their source: that judges *legally* must do so, or have no *legal* right to dispense from punishment. In the context, these bare modals are offered as a complete justification for imposing the punishment.

Cover noted the use of similar language by antebellum abolitionists. He called this “The Judicial ‘Can’t’”, and listed it as the first “paradigm of legal rhetoric” used to justify judicial obedience: judges bemoaned the injustice of slavery, but justified applying the law by simply stating that they “can’t” free fugitive slaves.⁵⁶ As with Mabbott, it is left implicit

⁵⁵Similar observations have been made by Hart (1994), p. 175 and Foot (1972), p. 309, among many others. See Chapter Three for further discussion.

⁵⁶Cover (1975), pp. 119–123.

that these judges *legally* can't free slaves, and this language is typically offered as a complete justification for applying the *Fugitive Slave Act*.

U.S. judges use similar language in justifying their fidelity to bad laws like §924(c). Reinhardt claimed that he had “no choice or discretion”: he “was forced to” affirm Hungerford’s sentence. Likewise, Cassell claimed that “[o]ur mandatory minimum sentencing scheme . . . forced me” to sentence Angelos to 55-63 years: “I had no choice”; “the mandatory nature of the punishment forbade me”; “I had to”.⁵⁷ Similar language is used in justifications offered by their contemporaries on the bench.⁵⁸ And in the context, this language is again offered as a complete justification.

Of course, in some contexts it is ambiguous what is being justified. Mabbott was justifying his conduct. But in legal decisions, judges might be justifying the result, rather than the agent who reached it. At least in some contexts, however, the justification is clearly of the judge. We find the same language used in public extra-judicial writings and private communication. As Cover notes in relation to the use of similar language by abolitionist judges, in many contexts the “rhetorical purpose” was “not the justification of the result and of the underlying principles, but the justification of the judge.” In Cover’s gloss, the judge is telling the audience: “I know that the result reached is morally indefensible and I wish primarily that you understand the sense in which I have been compelled to reach it.”⁵⁹

A good explanation of this language is that judges like Cassell take their legal requirements to settle what to do, despite knowing that what the law requires is obviously incompatible with what morality requires. In that sense, they seem to be prioritizing their legal duties—that they are *legally* compelled to punish—and thereby making a mistake about the status rather than content of morality and law.

Second, my diagnosis also explains what is strikingly absent in justifications of judicial obedience. Consider the competing diagnosis that judges are tacitly appealing to some moral

⁵⁷See *Angelos*, and see also remarks throughout Cassell (2007, 2004), and Cassell and Luna (2010, 2011).

⁵⁸See references in Cassell (2007), n.26.

⁵⁹Cover (1975), p. 119.

justification for applying unjust laws—perhaps one of the justifications canvassed in §2.3. One issue here is that the contemporary judiciary are well aware of the systematic injustices caused by mandatory minimums, which serve to undermine these moral justifications. This makes it doubtful that Cassell was tacitly appealing to something like Raz’s normal justification thesis. After all, Cassell (2007) himself cited the systematic inconsistency between mandatory minimums and “congressionally-approved Guidelines”, noting that the latter only “provide an additional 24 month penalty for the firearms at issue” in *Angelos*, before asking rhetorically, “could a judge conclude that an additional 660 months is a “just punishment”?” If judges are tacitly appealing to some moral justification for applying unjust laws, you would expect them to defend rather than reject the content of that justification.

Third, if judges were tacitly appealing to some moral consideration that favors applying unjust laws, you would expect them to weigh this consideration against the weighty moral reasons not to apply the law. “Competing, inconsistent, moral demands”, Cover noted, “do not resolve but define moral choices.”⁶⁰ So if judges are tacitly appealing to some competing moral consideration, you would not expect them to take this consideration alone to resolve their choice about what to do. It seems more likely, then, that judges are not making any claim about—and hence, not making any mistake about—the content of morality in justifying judicial obedience.

Fourth, the diagnosis offered here explains where others do not why the judiciary converge in applying mandatory minimums that they know to be unjust. Conservatives like Cassell and liberals like Reinhardt have deep disagreements about political morality. If some appeal to political morality explained why judges obey the law, you would expect much more divergence in patterns of judicial obedience between those who have deep disagreements about political morality. This point applies with equal force to disagreements among the current judiciary and to differences between the prevalent beliefs about political morality today and during the antebellum era. If judicial obedience to unjust laws like §924(c) and the *Fugitive*

⁶⁰Cover (1975), p. 122.

Slave Act were driven by tacit beliefs about what morality requires, you would expect much less convergence in judicial obedience to such laws as moral beliefs change over time.

Finally, the diagnosis that agents make mistakes about the status of standards can provide a unified explanation of similar behavior in legal and non-legal contexts. It can explain the relevant similarities between the rule-abiding conduct of Mabbott and Levy and Cassell. It can even explain why similar mistakes are possible when prudential rather than moral considerations are most salient. The hypothesis that judges are always tacitly appealing to moral considerations does not fare well in this regard. We either need some explanation for why similar behavior in legal contexts must be explained differently, or some account of how similar beliefs about the content of morality are somehow driving similar conduct in contexts where such an explanation seems strained at best. Until either explanation is offered, and shown to be adequate, we should prefer the diagnosis offered here: that judges, like plenty of other agents, sometimes do the wrong thing because they mistakenly prioritize the wrong standard.

If this diagnosis is right, we should be deeply dissatisfied with how judges justify judicial obedience. Individuals like Angelos who were wronged by the legal system are entitled to some justification from those who wronged them in the course of fulfilling their official duties. The only justification forthcoming from the bench is that the judiciary are “compelled” to act in such a way. But to quote Cover once more, these “formal restraints are under challenge and cannot be justified by mere reiteration.”⁶¹

2.5 Civil Disobedience

If judges like Cassell shouldn’t apply laws like §924(c), what should they do instead? The third main claim that I will defend is that just as English speakers (especially those in positions of authority) should flout grammatical requirements to change those requirements, the judiciary should engage in civil disobedience to change the law.⁶²

⁶¹Cover (1975), p. 129.

⁶²Following Joel Feinberg (1979), some call officials’ acts of disobedience “rule departures”.

Some philosophers have similarly argued that judges should refuse to discharge their legal duty to apply the law.⁶³ This position is typically defended by first setting out the relevant alternatives as the choices that can be made by individual judges like Cassell: Cassell can apply the law, resign, or engage in civil disobedience. In the relevant cases judicial obedience is ruled out as morally impermissible, and resignation is disparaged as a futile symbolic gesture. So civil disobedience is the last option standing.

I do not wish to defend civil disobedience in this way for three reasons. First, I am not convinced by the objection that resigning in protest is a futile gesture. To the best of my knowledge, there are no recent examples of judges resigning in protest of mandatory minimums.⁶⁴ In the absence of hard data, it is not unreasonable to be optimistic that by resigning rather than applying §924(c), Cassell could have avoided lending legitimacy to §924(c),⁶⁵ and galvanized opposition to mandatory minimums.⁶⁶ So we should not assume that symbolic gestures like resignation would be futile.

Nor should we assume that futility justifies or exculpates. Perhaps Cassell knew that he was replaceable: if he resigned, he would be replaced by another judge who would impose the same unjust punishment on Angelos. But is one justified or excused for committing an immoral act merely when one knows that if one doesn't do it another person would? This is pejoratively called the "arms dealer's defense".⁶⁷ Arms dealers are replaceable: each dealer knows that if he or she doesn't trade arms with ISIS, some other scumbag will instead. So too are drug dealers: if Angelos did not sell marijuana, some other low-level trafficker would take his place. Like Angelos, Cassell's futile participation in collective wrong-doing still makes him complicit,⁶⁸ and concerns about complicity have deep roots in criminal law. If we did

⁶³See Farnsworth (2001), Butler (2007), and especially Brand-Ballard (2010).

⁶⁴There are recent examples of prominent judges resigning in protest of other laws that they took to be unjust: Judge Robert Utter of the Washington Supreme Court (protesting against death penalty), and Judges Lawrence Irving of the District Court of the Southern District of California and John S. Martin of the United States District Court for the Southern District of New York (protesting against the imposition of the Sentencing Guidelines).

⁶⁵Hill (1979).

⁶⁶Kahan (2000).

⁶⁷Gardner (2007).

⁶⁸Lepora and Goodin (2013).

not condemn complicity, individuals like Hungerford—who was convicted and sentenced to over 150 years of incarceration for being an accessory to armed robberies that her boyfriend would probably have pulled off without her—would presumably all walk free.

Moreover, if Cassell had resigned rather than applied §924(c) in 2004, who would have been in a position to object? Perhaps some might object on the basis of the consequences of competent and conscientious judges resigning. But then they should have spoken up when Cassell decided to resign in 2007 due to “the uncertainty surrounding judicial pay”. When judges resign for the sake of their careers without incurring criticism, it is curious that resigning in protest should be so widely derided.

Second, I think there are serious objections to the form of civil disobedience that several philosophers have recently advocated—what Brand-Ballard (2010) calls “lawless judging.” If Cassell simply stated that he refused to apply §924(c) he would be censured and his decision reversed by appellate courts.⁶⁹ So candid disobedience would seem to be as futile as resignation. This leads philosophers to call for acts of surreptitious disobedience: Cassell could have imposed a less severe sentence, or no sentence at all, while offering a spurious and insincere legal rationalization.⁷⁰ Due to the unambiguity of §924(c), this would most likely require a strategic early intervention in the trial, such as a specious ruling on crucial facts or evidence.⁷¹

Surreptitious disobedience is a common, long-standing judicial practice.⁷² Roscoe Pound described judges who, wishing to make just verdicts that did not fit the law, “put a meaning

⁶⁹For discussion, see Brand-Ballard (2010), pp. 187–88, 271–85.

⁷⁰Relatedly, judges can engage in judicial nullification. See Saks (1993), p. 1281 (judges leave “jurors free to decide cases using their own intuitions about justice” by “effectively and persistently offering juries instructions that cannot be understood . . . Judges have done so for centuries, and they are unlikely to stop.”)

⁷¹As a result, this option is much more readily available to trial judges like Cassell than appellate judges like Reinhardt. For discussion, see Smith (1999), p. 1660, and Butler (2007), p. 1792.

⁷²Posner (1995), p. 133 (judges “sometimes bend or break the rules for the sake of other values, such violations being in fact rather common because detection and sanctioning are difficult”). See also Goldsworthy (2011), p. 316 (citing Judge Posner: “if disingenuous judicial opinions are illegitimate, we have had a crisis of judicial legitimacy since *Marbury v. Madison*” [1803]. Goldsworthy continues: “It is widely believed that Chief Justice John Marshall lied in the famous American case of *Marbury v. Madison*, in which the doctrine of judicial review of legislation was firmly established for the first time. It is not that he lied about that doctrine; rather, the claim is that he lied about the meaning of a statute, by adopting an absurd interpretation of it, in order to raise the question of judicial review.”)

into the [legal] text as the juggler puts coins ... into a dummy's hair, to be pulled forth presently with an air of discovery."⁷³ Most relevantly, judges have used spurious legal rationalizations to "evade laws mandating punishments that they regard as excessive"; consider, for example, the way British judges interpreted the legal rule known as "benefit of clergy",

whereby clerics accused of crimes could only be tried in church courts and not in common law ones. Common law judges used this to evade a law requiring all felons to be sentenced to death, by holding that anyone who could read a psalm was deemed to be a cleric, and then extending this to illiterates who could recite a psalm.⁷⁴

There are similar cases where U.S. judges have nullified unambiguous laws that were "ill-conceived, badly drafted and obsolete", and "very difficult to reform by legislative means."⁷⁵ But no recent judges have nullified mandatory minimums on these bases.

So what are the serious objections to this proposal? I am somewhat persuaded by some of the familiar responses to lawless judging: that it flouts the rule of law, violates a publicity constraint on civil disobedience,⁷⁶ and would increase the politicization of the judiciary, resulting in (for instance) activist pro-life judges surreptitiously disobeying *Roe v Wade*.⁷⁷ But I believe that there is a deeper objection in the vicinity. Individual acts of civil disobedience will always be insufficient given the systematic injustices that result from mandatory minimums like §924(c). For Cassell to save Angelos from an unjust, cruel and irrational punishment would help one victim of the legal system, but it would not help other offenders charged with multiple violations of §924(c). A lone judge cannot create anything like the legal fiction that British judges used to avoid subjecting psalm-reciting offenders to a mandatory death sentence. Nor would saving Angelos prevent prosecutors from using the threat of mandatory minimums like §924(c) to bully innocent individuals into pleading guilty.

⁷³Pound (1907), p. 382.

⁷⁴Goldsworthy (2011), p. 311.

⁷⁵Atiyah and Summers (1987), p. 185; see also their ch. 11 and pp. 178–85.

⁷⁶There is an extensive literature on this issue, beginning with Rawls (1971), p. 366.

⁷⁷For discussion, see Brand-Ballard (2010), pp. 182 ff. (on "mimetic failure").

This leads directly to the third reason why I will not defend civil disobedience in this way. It is a mistake to focus on the options available to individual judges. We should not just ask Cassell to justify his conduct, but ask the judiciary as a collective—which is already organized in groups like the federal Judicial Conference—why they continue to apply the law *en masse* rather than engage in civil disobedience.

The most obvious way in which the judiciary could engage in collective civil disobedience is by going on strike. This strike could be general: the judiciary could refuse to hear any cases whatsoever. Or it could be targeted: the judiciary could refuse to hear cases that involve mandatory minimums or accept pleas that result from the threatened application of laws like §924(c).

To the best of my knowledge, the U.S. judiciary has never gone on strike. But judges have recently gone on strike in 30 Spanish provinces (2009), in Malawi (2012), and in Colombia (2015); a city-wide strike in Egypt (2012) and a nation-wide strike in Tunisia (2015) specifically protested laws that restricted the power of the judiciary. A general or targeted strike by U.S. judges is feasible since the Judicial Conference has strongly and consistently opposed mandatory minimums as a matter of established policy for over half a century, precisely because laws like §924(c) unduly restrict the power of the judiciary. Collective action, in other words, is viable because mandatory minimums are so obviously immoral that they are opposed by judges who span the entire political spectrum.

Like individual acts of disobedience, striking allows the judiciary to avoid complicity in injustice in cases like *Angelos*. But only collective disobedience also allows the judiciary to avoid complicity in, *inter alia*, unjust plea bargains.

Moreover, unlike individual resignation or candid disobedience, striking is not futile. Strikes may not prompt Congress to repeal mandatory minimums, but they could achieve other important, efficacious aims.

For one, justices could demand new memos from the Attorney General. On August 12 2013, Eric Holder issued a memo directing federal prosecutors that, *inter alia*, they should not

specify drug quantities which trigger certain mandatory minimums, except for “serious, high-level, or violent drug traffickers”; and on September 24 2014, Holder issued a more narrowly focussed memo that prosecutors should not file charges of §851 (which provides mandatory minimum sentences for repeat offenders) “simply to exert leverage to induce a plea”. These were important improvements on the lax, imprecise and punitive directives that had been previously issued by John Ashcroft, but they are far from enough. The judiciary could strike until Loretta Lynch places more stringent limits on the use and abuse of prosecutorial discretion: prosecutors should be directed not to use *any* mandatory minimums, including §924(c), in either charges against low-level offenders like *Angelos* or negotiations to induce a guilty plea.

For another, judges should demand reform of their judicial powers. Cassell’s position in *Angelos* is tragic: his role-based obligations to administer justice and apply the law are obviously incompatible. And Cassell faced this position because Congress enact and enhance sweeping mandatory minimums without regard for their effects on particular offenders like *Angelos*. Jeremy Bentham held that, in certain conditions, judges should remit decisions to the legislature.⁷⁸ Judges should demand institutional reforms that grant them such powers. If Cassell could have remitted his decision in *Angelos* to Congress, he would have both reconciled the conflict between his obligations and forced the legislature to confront and take responsibility for the injustices of mandatory minimums. This is important, as one reason that mandatory minimums provisions persist—and are introduced and enhanced—is that legislators can defend such policies in the abstract as part of their posturing about being “tough on crime”, then distance themselves from the applications of such policies in particular cases. Forcing legislators to decide whether to impose a 150-year prison term on a low-level offender like Marion Hungerford would make support for mandatory minimums evaporate: legislators must endorse brutal punishments, or be hypocrites.

⁷⁸See Hart (1994), p. 272. See also Postema (1989), pp. 403–435.

Importantly, striking can be efficacious without encountering the pitfalls of surreptitious disobedience. Striking involves a public act of disobedience that neither violates the rule of law nor risks the further politicization of the judiciary: since the Judicial Conference remain divided on controversial moral issues like abortion there is obviously no hope for collective action on, say, *Roe v. Wade*.

Of course, a judicial strike would not be costless. But we routinely accept strikes by other central parts of the justice system (like the NYPD in 2014), and by other providers of essential public services (like teachers). We should be more than willing to accept such temporary costs for the sake of justice.

2.6 Conclusion

In this Chapter I have argued that: judges face uncontroversially incompatible moral and legal duties; judges perform immoral acts because they mistakenly give priority to their legal duties; and judges should instead engage in acts of collective civil disobedience. I have argued for these claims because I take them to be both politically and philosophically significant. I take them to be politically significant because they show how a demand for fairly radical changes in the conduct of currently exalted legal officials can be quite ecumenical. And I take them to be philosophically significant because they open up a different way of looking at the normative landscape.

Philosophers often write and speak as if the most pressing way that we can prevent injustice is to settle contentious questions about the content of morality. I'm not convinced. Plenty of intelligent, well-intentioned people perform acts that they know to be immoral because they make mistakes about the status of morality and conventional standards like law, grammar, games and etiquette. So a neglected question should be thrust into the spotlight: what is it for there to be such a difference in the status of normative standards? It is to this question that we now turn.

Chapter 3

Mere Formalities

3.1 Introduction

Consider these two plausibly true normative judgments:

1. I must give ten percent of my income to charity.
2. I must set the table with forks on the left and knives on the right.

Both concern requirements, and categorical ones to boot: they apply to me regardless of the content of my desires. The first is a requirement of morality. The second is a requirement of etiquette. There's a sense in which morality and etiquette are both "normative": both provide standards of correctness. There's also an intuitive sense in which morality is authoritative in a way that etiquette is not. If (1) is true, I really have to tithe. When (2) is true, it is an open question whether I really have to set knives on the right. We can put this by saying that etiquette, unlike morality, is a *merely formal* normative standard.

To date, philosophical interest in mere formalities has resembled a mercantile interest in fool's gold: once we ascertain that we are not dealing with the genuine article, it is cast aside, ignored. My interest in mere formalities is closer to a chemist's interest in fool's gold: I want to understand what they are, and understand why they look and feel so similar to the genuine article.

In §3.2, I outline the mercantile interest in mere formalities. From Philippa Foot (1972) to Derek Parfit (2011a,b), philosophers have offered an intuitive explanation of the sense in

which requirements of morality are authoritative in a way that requirements of etiquette are not. This explanation only characterizes mere formalities *via negativa*: it only tells us what mere formalities are not.

In §3.3, I make a case for adopting the chemist's interest in mere formalities. I introduce three independently interesting puzzles about mere formalities like etiquette. These puzzles concern how requirements of morality and etiquette look and feel similar in how they guide, are internalized, and conflict. To solve these puzzles, we need to put mere formalities under the microscope.

In §3.4, I make a preliminary case for a positive view about mere formalities. This preliminary case starts by arguing that mere formalities are things that we create with our shared attitudes. I then argue that the relevant attitudes bear the five main markers of fictive attitudes. This suggests that we should consider a kind of fictionalism about mere formalities. On this view, what goes on when I set the table with knives on the right because it is required by etiquette is the same sort of thing that goes on when a child runs away from a tree stump because she is playing a game of make-believe in which tree stumps are bears: in both cases, the person acts because they have accepted a fiction.

In §3.5, I argue that fictionalism solves the puzzles about how mere formalities guide, are internalized, and conflict. Better still, it offers an illuminating, generalizable and ecumenical explanation of the sense in which requirements of morality are authoritative and mere formalities like etiquette are not.

In §3.6, I respond to the most obvious and serious objection to fictionalism: that we sometimes have non-fictive attitudes towards mere formalities, and at least in some cases having such attitudes involves no apparent error.

My central aim is to show that fictionalism is a very plausible theory about what it is to be a mere formality. But my aim is not to show that fictionalism is the only way to go. There will be no systematic comparison between fictionalism and every possible competitor. We're sailing into fairly uncharted waters. So a subsidiary aim is to clarify what we want a

theory of mere formalities for. If the fictionalist's foes can at least agree on that, progress will have been made.

3.2 Foot's Fugitive Thought

It may seem contentious to describe both morality and etiquette as “normative”. But all that I have in mind is what Steve Finlay (2014) describes as “the orthodox view” that “the moral sense of ‘ought’ is only one of many, including prudential and instrumental senses, and those of institutions like legal systems, etiquette, and games.”¹ The same goes, of course, for the moral sense of ‘requirement’.

Here I am using “normative” in a thin sense that is close to its etymological roots.² No ink should be spilled on a terminological debate about whether morality and etiquette are both “normative”. What's interesting is the intuitive claim that morality is normative in some special sense in which etiquette is not. Morality, as we are often told, is authoritative. It must be taken seriously.

Depending on how it is expressed, this claim can seem obscure. (Richard Joyce (2006) often refers to morality's “practical oomph”.) It need not be. Philippa Foot (1972) explained the point modally, rather than metaphorically, in her attempt to capture the commonsensical view that “morality is supposed to be inescapable in some special way” that etiquette is not:

Very roughly the idea seems to be that one may reasonably ask why anyone should bother about what should_e (should from the point of view of etiquette) be done, and that such considerations deserve no notice unless reason is shown. So although people give as their reason for doing something the fact that it is required by etiquette, we do not take this consideration as *in itself giving us reason to act*. Considerations of etiquette do not have any automatic reason-giving

¹Finlay (2014), p. 235. More precisely, these “senses” are all flavors of deontic modals.

²These roots trace “back to the Latin *norma*, a builder's square.” When one uses a square, they “treat gaps between the actual cut and the square ... to show there is something to be “corrected” in the cut rather than the tool. So we have here a seemingly concrete example of ‘action-guidingness’ and an associated ‘standard of correctness’” (Railton (2002), p. 321).

force, and a man might be right if he denied that he had reason to do “what’s done.” ... [B]y contrast, it is supposed that moral considerations necessarily give reasons for acting to any man.³

This was Foot’s “fugitive thought”. It involves two modal claims. First, necessarily, if I should_m ϕ (if I should ϕ “from the point of view of morality”), there is reason for me to ϕ . Secondly, it is not the case that necessarily, if I should_e ϕ , there is reason for me to ϕ . Therein lies the authority (or “practical oomph”) that moral requirements have, and that requirements of etiquette lack.

Importantly, we should not confuse this second claim for the stronger modal claim that necessarily, if I should_e ϕ , it is *not* the case that there is some reason for me to ϕ .⁴ That modal claim is obviously false. After all, the requirements of morality and etiquette can coincide. Killing one’s dinner party guests is both immoral and indecorous. In such cases there are good reasons to do what etiquette requires. Good manners may be morally significant in other ways. For instance, there may be a moral virtue of civility that “has intimate associations with following socially established rules,” including “rules of etiquette”,⁵ such that to be uncouth is often to be uncivil. This is consistent with Foot’s claim.

Foot’s fugitive thought is plausible because it does not seem possible for morality—or, for that matter, prudence or epistemology—to require one to ϕ when there is no reason to ϕ ; but it does seem possible for etiquette to require one to ϕ when there is no reason to ϕ . Etiquette requires that my partner and I set the table with forks on the left and knives on the right, even in the privacy of our own home. Neither of us is terribly bothered about this. My partner grew up in a culture where people ate with their hands; the requirements of etiquette that apply to us seem foreign to her. So for me to violate this requirement may be uncouth, but it would not be rude or disrespectful. And setting knives on the right

³Foot (1972), p. 309. It is also worth remembering Foot’s often neglected point that “behavior does not cease to offend against either morality or etiquette because the agent is indifferent to their purposes and to the disapproval he will incur by flouting them.”

⁴Foot’s fugitive thought involves the claim that $\neg\Box(O_e\phi\rightarrow R\phi)$, not that $\Box(O_e\phi\rightarrow\neg R\phi)$. (Here “ $O_e\phi$ ” signifies “I should_e ϕ ” and “ $R\phi$ ” signifies “there is some reason for me to ϕ ”.)

⁵Calhoun (2000), p. 252. For more on the moral significance of manners, see Buss (1999).

often serves no practical purpose: we're vegan, so we typically eat our measly meals of soft vegetables with our forks in our right hands, without needing to cut anything. We lack even a weak reason to do "what's done".

Foot's fugitive thought is widely accepted in metaethics. Consider Derek Parfit (2011a,b)'s distinction between two "conceptions of normativity":

On the *reason-involving* conception, normativity involves reasons or apparent reasons. On the *rule-involving* conception, normativity involves requirements, or rules, that distinguish between what is *correct* and *incorrect*, or what is *allowed* and *disallowed*. Certain acts are required, for example, by the law, or by the code of honour, or by etiquette, or by certain linguistic rules. It is illegal not to pay our taxes, dishonourable not to pay our gambling debts, and incorrect to eat peas with a spoon, to spell 'committee' with only one 't', and to use 'refute' to mean 'deny'.⁶

Parfit claims that requirements of etiquette, and the other mere formalities he mentions, are not necessarily reason-involving: there may not be any reason to do what is required by etiquette or the law or linguistic rules or codes of honour. By contrast, authoritative standards like morality, prudence and epistemology are necessarily reason-involving.⁷ Let's call these "substantive normativity".

For our purposes, what's important is that Foot and Parfit have a mercantile interest in mere formalities. They characterize mere formalities *via negativa*. At most, they tell us what etiquette et al are *not*—they are *not* necessarily reason-involving. This does not answer what I will call our Core Question:

Core Question: What is it for a property or relation to be a mere formality?

⁶Parfit (2011a), pp. 144-5, emphasis his.

⁷Plenty of philosophers have made similar remarks. for instance, John Broome (2013) claims that "law, morality, rationality, prudence, convention, fashion, etiquette, honour, [and] the rules of chess" all provide examples of "source requirements", but some of these requirements are "normative" (i.e., "we necessarily have reason to satisfy them") while others are not (§7.3).

CORE QUESTION is analogous to a more familiar question from metaethics: What is it for a property or relation to be substantively normative? Perhaps the most well known and widely endorsed answer is REASONS FIRST:

Reasons First A property or relation is substantively normative iff and because it is, or is analyzable in terms of, the reasons relation.⁸

To illustrate: morality is substantively normative because the fact that morality requires me to tithe is analyzable in terms of the reasons for me to do so.

Why is Foot's fugitive thought not enough to answer CORE QUESTION? Because it is not sufficiently illuminating. At most, it tells us why requirements of etiquette et al are not authoritative. It does not tell us why they are *bona fide* requirements all the same, or help us distinguish requirements of etiquette from chairs and numbers, which are not "reason-involving" either.

Parfit's claim that etiquette et al are "rule-involving" does not help much.⁹ Nor does it help to add that etiquette et al are standards of correctness.¹⁰ This just pushes the question back. In virtue of what is etiquette a standard of correctness if it is not authoritative, or normative in the reason-involving sense?

Ultimately, I will defend a fictionalist answer to our CORE QUESTION:

Fictionalism A property or relation is merely formally normative iff and because according to the relevant fiction it is substantively normative.

To illustrate: if we accept REASONS FIRST, the fact that etiquette requires me to set knives on the right is a mere formality because it is analyzable in terms of the reasons for me to do so that exist according to the relevant fiction. I will argue that FICTIONALISM offers

⁸See Parfit (1984, 2011a,b), Scanlon (1998), Smith (1994), Schroeder (2007), Raz (2000) and Skorupski (2010). REASONS FIRST is neutral on whether the reasons relation is analyzable.

⁹It is also a strange way to distinguish morality from etiquette et al. Morality might *also* be rule-involving. Indeed, it *must* be if Parfit's rule-consequentialist theory of morality is true.

¹⁰As McPherson (2011) does in distinguishing "robust" and "merely formal" normativity.

an illuminating, generalizable and ecumenical characterization of what mere formalities are (§3.5). But let's not jump ahead.

Before we get to the positive view, it's worth considering why we should care. Why should we adopt a chemist's interest in mere formalities? If we know that they're not the genuine article, why worry about what they're made of?

The most immediate answer is that there are three independently interesting puzzles about mere formalities. To solve them, we must put standards like etiquette under the microscope and understand what it is to be a mere formality.

3.3 Three Puzzles About Mere Formalities

The three puzzles concern the respects in which requirements of morality and etiquette look and feel similar. These similarities aren't obviously puzzling when we believe there's good reason to do "what's done". The puzzles come into focus when we believe (or know) that there's no reason to do what etiquette requires.

3.3.1 Guidance

Requirements of morality and etiquette look and feel similar in how they guide us, even when we believe that we have no reason to do what etiquette requires. This is puzzling because requirements of morality seem to guide us via our "taking there to be reasons" to do what morality requires. However requirements of etiquette guide us, it must be via some mechanism that looks and feels very similar to taking there to be reasons to do what etiquette requires, but is compatible with our taking there to be no reason to do what etiquette requires.

Let's start by noting the respects in which being guided to deliberate, act, and criticize by requirements of morality and etiquette look and feel similar.

Moral requirements often guide us by structuring our deliberation. This may involve consciously aiming to do what is morally required. I confess that I do this because I take it

to be morally mandatory. That may be an atypical case.¹¹ Moral requirements may more frequently guide our deliberation not by featuring as conscious aims, but by excluding certain options and considerations. I doubt most of us even consider trampling over the slow and elderly when running to get a bus; because we implicitly know that doing so would violate a moral requirement, that action never makes it onto the table as a viable option.

There are other ways that moral requirements may guide us to act. Subjects in experiments who are led to perform acts that they (implicitly) take to be immoral manifest palpable psychological discomfort: they sweat, tremble, stutter, bite their lips, groan, and dig their fingernails into their skin.¹² These behaviors reveal significant, but defeasible, psychological pressures to do what's right.

When we take others, or ourselves, to violate moral requirements, we are also guided to criticize. Such criticism typically involves what Strawson (1962) called the reactive attitudes: guilt, blame and resentment. I am (defeasibly) disposed to blame those who have a similar or greater disposable income to myself yet do not give to charity, because I believe that they violate a moral requirement.

Etiquette guides us to act, deliberate and criticize in remarkably similar ways. Consider (2), the requirement of etiquette that I set tables with forks on the left and knives on the right. I am disposed to set the table in such a manner because it is required by etiquette. If I am cooking a stew and notice that we have no clean knives, I'll deliberate about whether to wash up before or after I finish cooking; even though I know that setting knives on the right serves no purpose, the option of not setting knives never even makes it onto the table. (Sorry.) And if my partner sets only forks, or sets knives on the left, I am disposed towards subtle forms of criticism. Not blame or resentment, of course, but critical attitudes all the same. I take her to have made a mistake, even though I don't take her to have acted contrary to any reasons.

¹¹Indeed, it may involve a "fetishistic" concern for moral rightness: see Smith (1994).

¹²These behaviors are reported in Milgram (1963)'s famous "obedience" experiments.

Why is this puzzling? A common, plausible view is that moral requirements guide us via our taking there to be reasons. What is it to take there to be reasons? Sometimes this involves occurrent beliefs that deploy the concept of a reason (as in my deliberation about tithing). Sometimes it involves something simpler. Perhaps a child could take steps to avoid trampling over the slow and elderly because she takes there to be good reasons to do so, without having any occurrent thoughts at the time, or even possessing the concept of a reason.¹³

Whatever it is to take there to be a reason, it is rationally incompatible with taking there to be no reason. We can home in on this by appealing to reflective avowal. If I would reflectively avow that there is no reason to tithe, I take there to be no reason to tithe, and do *not* take there to be reason to tithe (lest I be incoherent). For our purposes, what's crucial is that I cannot coherently tithe because it is morally required, reflectively avow that it is morally required, *and* reflectively avow that there is no reason to tithe.¹⁴ The believed absence of reasons seems to be a barrier to being guided by moral requirements. Some even claim that it is a barrier to engaging in practical deliberation whatsoever.¹⁵

By contrast, it does seem that I can set forks on the left and knives on the right because it is required by etiquette, reflectively avow that it is required by etiquette, and reflectively avow that there is no reason to do so. I take this to not only be possible, but actual. I can be guided by requirements of etiquette, and engage in practical deliberation, despite the believed absence of reasons.

Some might be skeptical of this. Won't I give as my reason that this way of setting the table is required by etiquette, and if so, in what sense is there a believed absence of reasons? Is my behavior idiosyncratic, or inaccurately self-reported, and if so, what stock should we put in such anecdotal evidence?

¹³See Parfit (2011a), p. 118. See also Sylvan (2015).

¹⁴Unless I am using "morality" in an "inverted commas sense". See Hare (1952, 1981).

¹⁵Enoch (2011a) claims that "we cannot, it seems, *deliberate* in the face of a believed absence of reasons", p. 74 (emphasis his). This plays a central role in his case for Robust Realism.

To answer both concerns, recall Foot's observation that although people give as their reason for doing something the fact that it is required by etiquette, they do not take this to be a reason. Michael Ridge (2014) similarly notes that we often use 'really' "to modify words like 'ought', 'must', 'reason'." He offers the following as an example: "I know that etiquette demands that I put the fork here, but do I really have to put the fork here?" I take it that here the modifier 'really' communicates that we do not reflectively avow the content we assert. We assert that our reason for doing something is the fact that it is required by etiquette, but we do not really take this (or anything else) to be a reason to do it; rather, we reflectively avow, and believe, that there is no reason to do what etiquette requires. And yet we are guided by etiquette all the same. I set the table a certain way even when doing so increases the effort and stress of preparing dinner, and even though I reflectively avow that there is no reason to do so, because that's what's done. This seems prosaic, and far from incoherent.

Moreover, even when we take there to be reasons to do as etiquette requires, it seems that we can still be guided by some other psychological mechanism. People in psychological experiments who are led to perform actions that they take to violate requirements of etiquette manifest palpable psychological discomfort: they feel tension, embarrassment, and sometimes debilitating nausea; and they do so even when they know that the stakes are relatively trivial.¹⁶ These behaviors reveal significant, though defeasible, psychological pressures to do what's done, which are independent of the strength of the relevant reasons.

There are many possible explanations for why the requirements of morality and etiquette guide us to act, deliberate and criticize in ways that look and feel strikingly similar, even when we believe there are no reasons to do what's done. But a particularly promising solution to the puzzle would identify a psychological mechanism that looks and feels similar to taking there to be reasons to do what etiquette requires, but is rationally compatible with taking

¹⁶These behaviors are observed in Milgram (1977)'s famous "subway" experiment.

there to be no reasons to do what etiquette requires (and reflectively avowing that this is so).

3.3.2 Internalization

Some might offer a simpler explanation of how the requirements of morality and etiquette guide: they do so via our adopting the attitude of norm-acceptance.¹⁷ We internalize moral norms and norms of etiquette. Sometimes we do so through conscious deliberation (such as when I adopted moral norms that favor tithing). Sometimes we do so via osmosis (such as when I adopted the moral norm not to trample over the slow and elderly). Either way, by adopting such norms we become disposed to act, deliberate and criticize in the ways described above.

This is another striking similarity. But it belies a puzzling difference between how we deliberately internalize moral norms and norms of etiquette: the former is rationally constrained in (at least) four ways that the latter is not.¹⁸

Consider my deliberate internalization of the moral norm in favor of tithing. To determine what the correct norm was, I had to consider evidence of the merits of tithing. This is our first constraint. Having determined what morality requires, there was only one permissible point of view for me to adopt: I must take the norm seriously. I could not say “I must tithe, but so what?” This is our second constraint.¹⁹ And once I internalized this moral norm, I was under rational pressure to reassess my other attitudes to those who violate this norm. Of course, I could still be pals with *Wolf of Wall Street*, Gordon Gecko types, but by doing so I risked attitudinal incoherence. This is our third constraint. Finally, I cannot directly change my attitude towards this moral norm at will, or on the basis of pragmatic considerations. Nor can you. If Peter Singer offered you \$10,000 to tithe, you could rationally do so. If he offered you the same sum to internalize this moral norm, you could not rationally do so.

¹⁷See Gibbard (1990). On its own, I doubt that this is much of an explanation.

¹⁸I focus on deliberate cases as the constraints on internalizing via osmosis are fuzzier.

¹⁹Some might insist that it is a *moral*, rather than *rational*, constraint. I disagree. (Note that we have the same constraint with prudence and epistemology.) But little hangs on this.

By comparison, the internalization of norms of etiquette is unconstrained. Consider poor Pip from *Great Expectations*. After he is suddenly thrust into the midst of the British aristocracy, he is subject to new, unfamiliar norms of etiquette. Pip can rationally decide what the correct, applicable norms of etiquette are without considering a shred of evidence about their merits. The first constraint does not apply. After deciding what the correct norms are, Pip internalized them. Pip took etiquette seriously. But this was not the only permissible point of view for Pip to adopt. He could have pretended to adopt them, or rejected them: he could say “Etiquette requires me to address people in such-and-such a manner, but so what?” The second constraint does not apply. Once Pip adopted the aristocratic norms, he did reassess his attitudes towards those who violate them, but this also seems optional. Pip could have coherently continued to value his uncouth kin in exactly the same way. The third constraint does not apply. Nor does the fourth. Pip chose to internalize etiquette norms as a direct exercise of his will, and on the basis of pragmatic reasons: he did so to be (or at least pass as) a “gentleman”. This is not incoherent or idiosyncratic. What’s more normal than adopting new social norms in order to fit in?

The strikingly similar internalization of moral norms and norms of etiquette is puzzling since the former is rationally constrained in four ways that the latter is not. Again, there are many possible explanations of this. But a very promising solution to the puzzle would identify a psychological mechanism that is just like the internalization of moral norms, yet floats free of these four constraints.

3.3.3 Conflict

The third and final puzzle is that in practical deliberation the requirements of morality and etiquette seem to conflict, and appear to be commensurable, even in cases where there is a (believed or known) absence of a common currency.

Let’s home in on this phenomenon indirectly. Sometimes what is moral is imprudent, and what is prudent is immoral. Perhaps it is imprudent for me to stick my neck out for

a stigmatized minority lest I be tarred with the same brush; to promote my own prospects involves turning a blind eye to their plight, and to do what morality requires involves some sacrifice of self-interest. I can know this, and weigh these moral and prudential norms against each other in practical deliberation. And there are right and wrong ways to resolve such conflicts. In this case, or perhaps in general, morality may trump prudence.

We have a good explanation for why morality and prudence can conflict and be commensurable: they are both reason-giving. Reasons provide a common currency. When morality and prudence conflict, we have reasons for and against the same action: the (believed or known) reasons for and against sticking my neck out explain why I can weigh moral and prudential norms against each other in practical deliberation; and when there is a determinate answer to which reasons are weightiest, there is one right way to resolve that conflict.

Now let's return to morality and etiquette. Sometimes what is moral is uncouth, and what is proper is impermissible. Norms of etiquette can require that a minority group be segregated, spoken down to, hurt, or humiliated. To stick my neck out for the stigmatized might be indecorous. Here the requirements of morality and etiquette seem to weigh against each other in practical deliberation, and there seem to be right and wrong ways to resolve such conflicts.

This appearance persists even in the face of the (believed or known) absence of reasons. I may have no reason at all to follow morally pernicious social norms. Despite believing this, it can still be intelligible for me to think and speak as if morality and etiquette conflict and are commensurable. I can weigh these requirements against each other in practical deliberation, and I certainly take there to be right and wrong ways to resolve that conflict. But when there is a known absence of reasons we have no common currency, and hence no explanation for how morality and etiquette can seem to conflict and be commensurable.

Some might not find this puzzling. Perhaps the absence of reasons can be weighed against and outweighed by reasons, just as zero can be weighed against and outweighed by one. But this does not account for what is distinctive about how morality and etiquette seem to

conflict. The absence of reasons can *always* be weighed against and outweighed by reasons. Humdrum descriptive propositions can be jointly unsatisfiable with moral requirements. Perhaps all the worlds in which some moral requirement is met are not worlds in which Wall St bankers get bonuses. Say that there is no reason at all for Wall St bankers to get bonuses. Knowing this, I doubt we could intelligibly ask “Should I do what’s right, or see to it that Wall St bankers get bonuses?” Yet even while we believe (or know) that there is no reason to do what’s done, it is intelligible for us to ask “Should_m I do what’s right, or should_e I do what’s done?” or even “Should I do what’s right, or what’s done?” Requirements of morality and etiquette seem to conflict in an interesting way that goes beyond mere joint unsatisfiability.

Again, there may be many possible explanations of this, however it would be ideal if we could find some view that solves each of these three puzzles at once. Moreover, it would be ideal if this solution was generalizable. The same puzzles, I believe, arise for other paradigmatic mere formalities, such as grammar. So we should seek an answer to CORE QUESTION that provides a general solution.

This is why characterizing mere formalities *via negativa* is not sufficient. We need to know *what it is* to be a mere formality so that we can explain how they look and feel so similar to the genuine article—how they guide, are internalized, and can seem to conflict with substantively normative standards like morality—even when we know that there’s no reason to do “what’s done”.

3.4 A Way Forward

We have focussed on requirements of morality and etiquette as representatives of the broader categories of “mere formalities” and “substantive normativity”. There are well-known theories—like REASONS FIRST—of what it is for standards like morality, prudence and epistemology to be substantively normative. What we want is a plausible theory of

what it is for standards like etiquette, law, grammar, codes of honour, and the rules of games to be mere formalities.

There are many possible theories of mere formalities, some of which have been developed more than others. I will not undertake the important but time-consuming task of surveying the range of actual (let alone possible) options.²⁰ Nor will I argue for my preferred view by eliminating all alternatives. Instead, I'll argue that mere formalities have distinctive features that suggest a possible way forward, then follow this path and see how far it leads us. One reason to prefer this approach is that it leads, I argue, to a view that is highly ecumenical, so there may be no deep tension between fictionalism and some other views.

3.4.1 What Makes Mere Formalities Different?

The most fundamental respect in which mere formalities differ from standards like morality is that they are psycho-social features of the world. Requirements of morality, prudence and epistemology may be part of the fabric of the universe. But requirements of etiquette, law, grammar, codes of honour, and games are things that we create with our shared attitudes.²¹ As Parfit wrote, on the so-called “rule-involving conception of normativity” we can “create new normative truths merely by introducing, or getting some people to accept, some rule.”²² Changing shared attitudes (getting people to “accept some rule”) alone creates new requirements.²³ The best explanation for this is that mere formalities are identical to, or are to be analyzed in terms of, certain shared attitudes.

²⁰This task is particularly time-consuming since extant theories are developed in (or can be extrapolated from) a broad range of domains, and depend heavily on domain-specific concerns and authors' controversial and complex background commitments. This is true of views in metaethics (Schroeder (2010) suggests a Humean REASONS FIRST view; Copp (2004, 2009) suggests a “pluralistic and teleological” view), social theory (Anderson (2000) offers an excellent critical overview of different approaches taken within rational choice theory, evolutionary theory, and “social rationality”), philosophy of law (to explain “legal normativity”, some legal positivists like Marmor (2011) appeal to conditional reasons, while others like Shapiro (2011) take legal obligations to be moral obligations according to a “legal theory”), and so on.

²¹This distinction is even accepted by anti-realists like Blackburn (1984): if we all changed our attitudes towards kicking dogs, it would still be immoral, but it would *not* be indecorous.

²²Parfit (2011a), p. 145.

²³Some may prefer to say that all possible mere formalities exist, but which mere formality *applies to us* depends on our shared attitudes. (Cf. Lewis (1975).) Nothing hangs on this.

Mere formalities are not the only psycho-social features of the world. An anthropologist can understand a community's moral norms by understanding their shared moral attitudes, since the latter create the former. But intuitively, those moral norms still differ from mere formalities like etiquette and grammar.²⁴ If mere formalities are identical to or analyzed in terms of shared attitudes, but differ from other psycho-social phenomena, then we need to determine what is fundamentally distinctive about our attitudes to standards like etiquette.

So far, this is not terribly original or controversial. In fact, perhaps the best developed contemporary view about requirements of etiquette et al agrees that they “are constellations of shared normative attitudes”,²⁵ then proceeds by attempting to identify what is distinctive about those attitudes. On this view, which is developed by Nicholas Southwood (2010, 2011), the crucial difference between our attitudes towards requirements of morality and etiquette lies in their “grounds”, where “grounds” are understood to be the reasons that non-derivatively “justify, in one's mind, the normative principle that one accepts”.²⁶

Southwood claims that the grounds for our attitudes towards requirements of etiquette, but not for our attitudes towards moral requirements, include social practices. This seems intuitive. I take it that we must set knives on the right because that's “what's done”, but I do not offer any similar justification for why we must tithe. This motivates Southwood's view that what “makes moral norms and social norms fundamentally distinct” is that “moral judgements” like (1) “are essentially *practice-independent*”, while judgements about mere formalities like (2) “are essentially *practice-dependent* in the sense that they are necessarily grounded, at least in part, in presumed social practices.”²⁷ So there is a significant continuity between our attitudes towards substantive and merely formal standards: “in my mind”, I must tithe because of the reasons to tithe, and I must set knives on the right because of the

²⁴This does not entail that a community's false moral views are substantively normative. (This is one reason for doubting that the merely formal–substantive distinction is *exhaustive*.)

²⁵Southwood (2011), pp. 762, 763, and *supra* nota 4.

²⁶Brennan et al. (2013), p. 68. I have reconstructed this view with all references to its most recent exposition in Brennan et al. (2013) for the sake of charity, and in terms of reasons rather than “justifications” or “considerations” for the sake of continuity with the preceding prose.

²⁷Brennan et al. (2013), pp. 69–72.

reasons to do so. But there is also a significant discontinuity between these attitudes: “in my mind”, among the reasons for me to set knives on the right is that it is “what’s done”.

This proposal faces several challenges. Here’s one. In responding to Philippa Foot, Miss Manners held that we have reason to do what etiquette requires because of “communal harmony, dignity of the person, a need for cultural coherence, and an aesthetic sense.”²⁸ Miss Manners believed this that we only have practice-independent reasons to do “what’s done”. Entire communities—like the one depicted in *Downton Abbey*—might be like Miss Manners in this regard. Their beliefs may be false, but that is beside the point: the grounds for “our” shared attitudes towards morality and etiquette may not always differ.²⁹

We will return to Miss Manners’ beliefs in §3.6. For now, let’s focus on another challenge. For Southwood, “in one’s mind” something must be a reason to do what is required by morality or etiquette. What does this mean? He writes that one “obviously” need not *believe* that it is a reason: “[p]erhaps it is enough” that it “*appears* to one” to be a reason, or that one “*sees*” it as “legitimizing a principle”, or that one “*feels* that it renders certain conduct obligatory.”³⁰ So on Southwood’s view, we have the same eclectic collection of types of attitudes towards morality and etiquette: appearances, seeings, and feelings. The difference lies in their objects: practice-independent and -dependent reasons.

This does not explain the puzzling differences between our paradigmatic attitudes towards requirements of morality and etiquette, which we saw in §3.3. When we are guided by mere formalities, we *assert claims that we would not reflectively avow*: we assert that our reason for doing something is that it is “what’s done”, even when we do not really take this or anything else to be a reason to do it. We can rationally decide what merely formal requirements there are *without considering evidence* about their merits. When we decide what the requirements

²⁸Martin and Stent (1990), p. 244.

²⁹Southwood may face counterexamples in the other direction. Kate Manne (2013) argues that moral reasons “have their source in social practices”. I am not sure whether these reasons are “practice-dependent”: Manne says practices “provide reasons directly”, but “only when certain background conditions hold” (p. 50). Still, a community with beliefs like Manne’s may have moral attitudes that have the same grounds as their attitudes towards etiquette.

³⁰Brennan et al. (2013), pp. 68–69 (emphasis theirs).

of etiquette are, there are *a plurality of permissible points of view*: we can internalize them, pretend to do so, or reject them. We can adopt these attitudes *by a conscious act of will, and for pragmatic reasons*. And decisions to adopt these attitudes are *unconstrained by norms of attitudinal coherence*.

What would explain all of this? I will argue that these differences neatly map on to the five markers that distinguish *belief* from fictive *acceptance*, which suggests that we have different types of attitudes towards morality and etiquette.

3.4.2 The Five Markers of Fictive Attitudes

There are three fictive attitudes: acceptance, pretense and rejection.³¹ We can illustrate them by considering an example from Kendall Walton (1990).

Imagine that a family is playing a game of make-believe in which tree stumps are bears. The young child accepts the fiction: when she sees a tree stump, she yells “Bear!” and runs away in fear. The parent engages in the pretense that she accepts the fiction for the sake of entertaining the young child: when she sees a tree stump she yells “Bear!” and runs away, acting *as if* she were afraid. Her surly teenager rejects the fiction: when she sees a tree stump, she mumbles “It’s just a bloody tree stump. Can we go to the mall?”

Since the attitude of acceptance is most important for our purposes, let’s focus on the attitudes of the child. It’s clear that *accepting the fiction* that <That’s a bear> is very similar to paradigm cases of *believing* that <That’s a bear>.³² But they differ in five interesting and important respects.

First, when you believe that you are seeing a bear, you would reflectively avow that this is so. By contrast, when the child accepts the fiction that <Tree stumps are bears>, she typically would not reflectively avow that this is so. If the surly teenager pointed to a tree

³¹See Eklund (2011). Some contend that acceptance and pretense differ in degree rather than in kind. I disagree, but for our purposes nothing hangs on this.

³²Here I follow the tradition of using <*p*> to designate a proposition: the angle brackets will soon prove helpful for easily distinguishing embedded and unembedded propositions.

and asked “Is that a bear?”, the child would answer “Yes!”; if asked “But is it *really* a bear?”, she will respond differently.³³

Secondly, reasonable beliefs about bears are paradigmatically “shaped primarily by evidence for what is believed and concern for the truth of what is believed. Thus the slogan: belief aims at truth.”³⁴ By contrast, it can clearly be reasonable for the child to accept the fiction that <Tree stumps are bears> in the absence of any evidence for thinking that this proposition is true.

Thirdly, when faced with actual bears, the only permissible point of view is belief. We must believe what’s obviously true. But when faced with fictional bears, several different points of view are typically permissible: the child can accept, pretend to accept, or reject the fiction, without making any mistake. Of course, in some atypical fictions—such as a game of make-believe in which kicking babies is good for them—acceptance and pretense may be impermissible.

Fourthly, belief is not subject to direct voluntary control.³⁵ You cannot directly choose to believe that there is a bear in front of you, just for kicks. By contrast, for the sake of having fun, the child can choose to accept, pretend to accept, or reject a fiction according to which there is a bear in front of her.

Finally, in paradigm cases it is incoherent and criticizable to believe that tree stumps are bears while knowing that they’re not.³⁶ By contrast, accepting a fiction is not constrained by norms of coherence. The child is not criticizable for accepting that <Tree stumps are bears> while knowing that they’re not. Accepting the fiction does not require her to reassess her beliefs about biology.

³³Daly (2008) takes this to be the best evidence for an acceptance–belief distinction.

³⁴Bratman (1992), p. 3.

³⁵See Williams (1973). His further claim that this is a conceptual truth, rather than not a contingent limitation of the will, is much disputed, but orthogonal to our purposes.

³⁶As Bratman (1992) writes (on p. 4): “An agent’s beliefs are subject to an ideal of integration. Other things equal one should be able to agglomerate one’s various beliefs into a larger, overall view; and this larger view should satisfy demands for consistency and coherence.”

Paradigm cases of believing $\langle \text{That's a bear} \rangle$ and accepting a fiction with that content are very similar, but for these five markers of fictive attitudes. Paradigm cases of believing $\langle \text{There is reason for } S \text{ to } \phi \rangle$ and accepting a fiction with that content will also be very similar. As we saw in §3.3, when I believe $\langle \text{There is reason for me to tithe} \rangle$ I am disposed to tithe, rule out not tithing in deliberation, and criticize others for not tithing. Accepting a fiction that $\langle \text{There is reason for me to set forks on the left and knives on the right} \rangle$ will result in similar dispositions towards table-setting;³⁷ but it will also bear the five markers of such fictive attitudes. This should sound quite familiar.

Indeed, paradigmatic attitudes towards mere formalities seem *exactly* like the fictive attitudes we just described. My explanation for this is quite simple, though perhaps somewhat surprising:³⁸ they *are* fictive attitudes. That is, what goes on when I set the table with knives on the right because it is required by etiquette is the same sort of thing that goes on when a child runs away from a tree stump because she is playing a game of make-believe in which tree stumps are bears. In both cases, the person acts because they have accepted a fiction.

3.4.3 Fictionalism About Mere Formalities

If mere formalities are analyzed in terms of shared attitudes, and the relevant attitudes are fictive attitudes, this suggests that we should be fictionalists about mere formalities. On this view, the distinction between substantive normativity and mere formalities reduces to the distinction between reality and fiction.

³⁷Velleman (2000) argues, quite compellingly, that believing that p and imagining that p share the same motivational roles (pp. 255–277). See also Shah and Velleman (2005).

³⁸Should we worry that the explanation is surprising? Should we worry that we do not take our attitudes towards mere formalities to be fictive? Some (such as Cuneo and Christy (2010) and Stanley (2001), pp. 46 ff) have suggested that this is a problem for hermetic fictionalisms. I disagree. Even when I am engaged in literary fictions, the fact that my attitudes are fictive is not luminous (in Williamson (2000)'s sense). And as Joyce (2005) notes, uncommon abstract theoretical reflection is required for agents to form conscious second-order beliefs that classify their attitudes as fictive (p. 292). I take our attitudes toward mere formalities to be fictive because they bear the five markers of fictive attitudes. These markers do *not* include the possession of an occurrent second-order belief: “Gee whiz, my attitudes are fictive!”

Fictionalism can be developed in various ways. An error theoretic fictionalist would hold that there are no requirements of etiquette or law or codes of honour: mere formalities are fictions. But this view is too radical and revisionary for my tastes. The form of fictionalism I favor is quite modest and hermeneutic:

Fictionalism A property or relation is merely formally normative iff and because according to the relevant fiction it is substantively normative.

Let's walk through this proposal in three steps. First, what are the *contents* of the relevant fictions? Our paradigm example of a mere formality is the requirement of etiquette that I set knives on the right. According to the relevant fiction, this requirement is substantively normative. What does that mean? It depends. Recall REASONS FIRST: to be substantively normative is to be analyzed in terms of reasons; so the fact that morality requires me to tithe is analyzed in terms of the reasons for me to tithe. If we accept this account, we will fill in the contents of the relevant fictions with propositions like <There is reason for me to set knives on the right>, or with more complex propositions to match the structure of our account of substantive normativity.

Secondly, what are *mere formalities*? According to FICTIONALISM, mere formalities are facts grounded in fictions. The proposal is *not* that according to the relevant fiction, <I should_e set knives on the right>. (That's the error theoretic view.) The proposal is that I should_e set knives on the right *because* according to the relevant fiction <There is reason for me to set knives on the right>. This distinction is important. The objects of our fictive attitudes do not seem to be requirements of etiquette themselves: I *do* reflectively avow that etiquette requires me to set forks on the left and knives on the right. Rather, the objects of our fictive attitudes seem to be reasons to do what etiquette requires. That's why FICTIONALISM claims that requirements of etiquette and other mere formalities are *facts* grounded in *fictions* about what there is reason to do.

It is worth noting that there are plenty of precedents for the proposal that mere formalities are facts grounded in fictions. Some have argued that facts about modality and mathematics

are grounded in modal and mathematical fictions.³⁹ But we needn't look to such esoteric philosophical doctrines to find precedents for the proposal. Legal facts are often grounded in legal fictions. For example, say a child enters someone's property without permission. The legal fact that the child is not a trespasser is grounded in the fiction that children are invitees even under circumstances in which they have clearly not been invited.⁴⁰

Thirdly, what are *fictions*? FICTIONALISM is compatible with different views here. On the view I like best, fictions are abstract artifacts whose contents are sets of propositions.⁴¹ This makes fictions no more metaphysically mysterious than theories. What makes fictions distinct from theories is that the nature and content of the former is explained in terms of fictive attitudes and the nature and content of the latter is explained in terms of non-fictive attitudes.⁴² These brief remarks should not be read to suggest that the ontology of fictions is simple and uncontroversial. Nothing could be further from the truth. Rather, these brief remarks are intended to convey that fictions, like mere formalities, are psycho-social features of the world: they are abstract artifacts that we create through our (fictive) attitudes. The nature and creation of the fictions posited by FICTIONALISM about mere formalities is no more mysterious than the nature and creation of fictions like *Hamlet* or the game of make-believe described above.

3.5 The Virtues of Fictionalism

That was my preliminary case for and development of FICTIONALISM. It is a good way forward given that mere formalities are analyzed in terms of shared attitudes, and the relevant attitudes bear the five markers of fictive attitudes.

³⁹See Nolan and O'Leary-Hawthorne (1996). Cf. Joyce (2001) on moral fictionalism.

⁴⁰See Fuller (1967) for discussion of legal fictions, including this example (p. 66).

⁴¹Those who oppose this view should not make the mistake of identifying a fiction like *Hamlet* with a concrete object, such as the copy of *Hamlet* that sits on my bookshelf; see Rosen (2009) for discussion. Nor should they make the mistake of identifying fictions with possible worlds (as Lewis (1978) once did); that view encounters problems with fictions that intentionally or unintentionally represent impossibilities (see Priest (1997) for an excellent example). For such reasons, Nolan (2013) and others identify fictions with sets of propositions.

⁴²This view is widely accepted, largely thanks to the work of Walton (1990).

Of course, a good way forward can still lead to a dead end. So let's consider the virtues of FICTIONALISM to see whether this path can lead us far enough.

3.5.1 Guidance, Internalization, and Conflict

A central desideratum for a theory of mere formalities should be that it solves the puzzling respects in which requirements of morality and etiquette look and feel similar even when we believe that there is no reason to do "what's done". Recall the puzzle about guidance: requirements of etiquette and morality guide us to act, deliberate and criticize in strikingly similar ways; yet the believed absence of reasons is a barrier to being guided by moral requirements, but not to being guided by requirements of etiquette. When we are guided by morality, this is because we take there to be reasons. When we are guided by etiquette, we are guided by some mechanism that just mimics taking there to be reasons.

The fictionalist offers a plausible candidate for what that mechanism is: accepting fictional reasons. As we saw above, accepting the fiction that <There are reasons for me to set forks on the left and knives on the right> will closely mimic believing that there are reasons for me to do so—paradigmatically, it results in dispositions to act, deliberate and criticize—but is consistent with my reflectively avowing that there are no reasons for me to do so.

What about the puzzle about internalization? Internalizing moral norms is constrained in four ways that internalizing norms of etiquette is not. These are the exact asymmetries we should expect to find if internalizing moral norms involves adopting beliefs about reasons while internalizing norms of etiquette involves accepting fictional reasons. Take my internalizing moral norms that favor tithing. I must: do so on the basis of evidence about reasons to tithe (because reasonable beliefs are shaped by evidence for what is believed); take this norm seriously (because there is only one permissible point of view towards reality); and reassess my attitudes to those who violate this norm (because beliefs are subject to norms of coherence). Moreover, I cannot directly choose to take this norm seriously for profit (because

beliefs are not subject to direct control by the will, and paradigmatically are not responsive to pragmatic considerations).

By contrast, if Pip's internalizing aristocratic norms of etiquette involves his accepting a fiction according to which $\langle$ There is reason to ϕ $\rangle$, then Pip: need not consider any evidence about reasons to ϕ (because accepting a fiction that p is reasonable in the absence of evidence that p); is permitted to internalize, pretend to internalize, or reject the norm (because typically many points of view are permissible towards fictions); can internalize the norm without reassessing his attitudes towards those who violate it (because accepting a fiction that p is consistent with knowing not p); and can change his attitudes towards the norm by a direct exercise of will in order to fit in (because we can choose to accept fictions, and do so on the basis of pragmatic considerations). So far, so good.

Of course, we are not *always* permitted to take any of these attitudes towards etiquette. Just as it is impermissible to engage in a game of make-believe in which kicking babies is good for them, one should not accept deeply inequalitarian social norms.⁴³ This reflects a general phenomenon whereby we “imaginatively resist”—i.e., refuse to accept—fictions that we take to be “morally deviant.”⁴⁴

The fictionalist's solutions to the first two puzzles are straightforward. The solution to the third is more complex. This puzzle concerns the apparent commensurability of morality and etiquette: even when I know there is no reason to do what etiquette requires, it still seems that requirements of etiquette and morality can conflict, as they can be intelligibly taken to weigh against each other in practical deliberation. How does the fictionalist explain this?

Direct comparisons between relevantly similar real and fictional entities are common, and quite intelligible. If you were to imagine a five-hundred-pound turkey, you could intelligibly

⁴³As the sociologist James Loewen (1988) wrote, “In its everyday operation, segregation consists in a pervasive system of etiquette. Every movement, every point of contact, every interaction, has carefully detailed alternative sets of norms which can be applied to it, depending on the race of the other party.” These norms “express[] and codify[] the relationship of dominance and subservience” between racial majorities and minorities (pp. 58–59).

⁴⁴For excellent discussion, see Hume (1757/1985), Walton (1994), and Gendler (2000).

say “This turkey is heavier than any real turkey”, and thereby directly compare the weights of real and imagined turkeys.⁴⁵ We can also directly compare the weights of real and fictional reasons. In doing so we exercise our general capacity to compare reality with representation.

While comparisons between the real and imagined are intelligible, there is an obvious sense in which it is true that a real turkey is heavier than any imagined turkey. Only real turkeys have weights. This perspective might seem persnickety in some contexts, such as during a game of make-believe. But it is privileged whenever one is engaged in practical deliberation about what to do. If you were asked to get the biggest turkey for Thanksgiving and came home with a drawing of your imagined turkey, you would have made a mistake. When the requirements of morality and etiquette seem to conflict, and we are engaged in practical deliberation about what to do, a similar perspective is privileged and far from persnickety: real reasons beat unreal reasons by default, because only real reasons have weights. This explains the sense in which the question “Should I do what’s right, or what’s done?” is intelligible even when we know there is no reason to do what’s done, and even when we know that though there is obviously only one right answer. So it explains the apparent commensurability of morality and etiquette, despite the known absence of a common currency.

FICTIONALISM offers a simple explanation of why mere formalities look and feel like the genuine article—fictive acceptance looks and feels just like beliefs—that solves the puzzles about guidance, internalization, and conflict in one fell swoop. Better yet, the fictionalist’s solutions to these puzzles can easily generalize from our central example, etiquette, to the myriad of mere formalities.

⁴⁵This example and discussion are inspired by Simon Blackburn (1999b)’s discussion of the ontological argument. Note that the comparison is *direct* in that you can simply compare the weights of real and imagined turkeys; you need not consider what would happen if you exported all real turkeys to a fictional world in which there was a five-hundred-pound turkey.

3.5.2 Normativity *Sans* “Authority”

Another desideratum for a theory of mere formalities is that it should explain the sense in which morality and etiquette are both standards of correctness even though the former is authoritative in a way that the latter is not.

FICTIONALISM explains how mere formalities are *bona fide* standards of correctness: we can act contrary to what we have reason to do, or act contrary to what we have reason to do according to a fiction. This explanation is perfectly compatible with Foot’s fugitive thought. The requirements of etiquette are analyzed in terms of fictional reasons, and fictional reasons are not necessarily real. So it is not necessarily the case that if etiquette requires me to set knives on the right, there is reason for me to set knives on the right. By comparison, if the requirements of morality are analyzed in terms of reasons, then it is necessarily the case that if morality requires me to do so, there is reason for me to do so.

This makes FICTIONALISM an illuminating supplement to Foot’s fugitive thought. It provides a positive account of what it is to be a mere formality which explains the respects in which etiquette et al differ from both substantive normativity (like morality) and descriptive phenomena like chairs and numbers.

While my focus has been on the requirements of etiquette, FICTIONALISM is also easily generalizable to other standards, such as legal requirements. Indeed, a common view in philosophy of law is that legal requirements can be analyzed in terms of “legal reasons”,⁴⁶ where “legal reasons” are only reasons in some unspecified “qualified” sense.⁴⁷ As Enoch (2011b) put the idea, “legal reasons” are like “imaginary friends”: the fact that there is a legal reason to ϕ does not entail that there is a reason to ϕ , just as the fact that one has an imaginary friend does not entail that one has a friend. FICTIONALISM helpfully develops this common view: legal requirements are analyzed in terms of legal reasons, where legal reasons are reasons that exist according to the legal fiction.

⁴⁶See especially Horta (2011).

⁴⁷See Shapiro (2006)’s distinction between adjectival and qualified readings of “legal reason”.

That FICTIONALISM is generalizable is also dialectically helpful since some may resist the view, and this project, because they quibble with the classification of particular standards of correctness. Some claim that morality is on par with etiquette: it is not necessarily reason-giving.⁴⁸ Some also claim that law is on par with morality: it is necessarily reason-giving.⁴⁹ But these debates need not detain us. As long as we accept that some standards (such as prudence) are authoritative, and others (such as the requirements of games) are not, we can distinguish between substantive normativity and mere formalities, and explain this general distinction in terms of the distinction between reality and fiction.

A final virtue of FICTIONALISM is that it is a highly ecumenical view about both normative authority and mere formalities. So far we have focussed on views that explain authority and substantive normativity in terms of reasons. While REASONS FIRST is currently in vogue, FICTIONALISM is by no means tethered to this view, and should not cling too tightly to its coattails. For instance, other philosophers explain substantive normativity in terms of value.⁵⁰ Some even reject Foot's fugitive thought, but endorse similar modal claims about values: necessarily, if morality requires me to ϕ , it is good that I ϕ , but it is not necessarily the case that if etiquette requires me to ϕ , it is good that I ϕ .⁵¹ However these philosophers can still accept FICTIONALISM as an answer to CORE QUESTION. They just fill in the details of the view in terms of fictional values rather than fictional reasons, to match their account of normative authority.

Perhaps more surprisingly, FICTIONALISM is compatible with many other views about mere formalities. Recall Southwood's view that the objects of our attitudes towards morality and etiquette differ: the former are about practice-independent reasons, and the latter are about practice-dependent reasons. We saw that Southwood denies that these attitudes are always beliefs, but is otherwise non-committal about what *types* of attitude we take towards

⁴⁸Foot (1972) took this view (until Foot (1997)), as did Singer (1993) (p. 286). Their opponents—moral rationalists—often commit to something stronger than Foot's first claim: that necessarily, if I should_m ϕ , there is *most* reason for me to ϕ (see Portmore (2011), p. 6).

⁴⁹Defenders of this view, like Greenberg (2014), take law to give us moral reasons for action.

⁵⁰See Moore (1903). Other philosophers appeal to "fittingness": e.g., Chappell (2012).

⁵¹I take it that Railton (1986) would endorse these two modal claims about value.

morality and etiquette; as a result, his view does not explain the puzzles in §3.3. It is open to Southwood to fix this by embracing the view that our attitudes towards morality involve beliefs (about practice-independent reasons) while our attitudes towards etiquette involve acceptance (of a fiction about practice-dependent reasons). So Southwood's view can be developed as a version of FICTIONALISM.⁵²

3.6 Miss Manners & *Downton Abbey*

I've argued that FICTIONALISM satisfies two central desiderata for an answer to our CORE QUESTION: it solves the puzzles about guidance, internalization and conflict; and it offers an illuminating, generalizable and ecumenical explanation of how mere formalities are standards of correctness *sans* normative authority.

But of course, FICTIONALISM will have detractors. For reasons of brevity, I will only consider the most obvious and serious objection it faces: sometimes we have non-fictive attitudes towards mere formalities. The problem here is subtle. On my view, mere formalities are facts grounded in fictions. So it is of no concern that we have non-fictive attitudes towards the *requirements* of etiquette or law or grammar. However, it is concerning that we often have non-fictive attitudes towards the *reasons* to comply with such requirements. Worse yet, those attitudes do not always seem to involve any kind of error: when Miss Manners believes that she has reasons to do “what’s done”, she does not seem like a Lord of the Rings literalist who has mistaken fiction for reality. And worse still, folks like Miss Manners are not, or not always, outliers: perhaps in the era depicted by *Downton Abbey*, most of “us” took this attitude towards etiquette.

The problem with this objection is that it assumes that the presence of belief entails the absence of fictive attitudes. But that can't be right. After all, parts of the content of fictions can be true, and not just true-according-to-the-fiction. Consider the proposition: <According to Saramago's *The History of the Siege of Lisbon* <The Portuguese laid siege

⁵²Southwood agreed with this assessment in personal correspondence.

to the Moorish city of Lisbon>>: this is true, because the embedded proposition is true according to the fiction. But the unembedded proposition <The Portuguese laid siege to the Moorish city of Lisbon> is also true; in that sense, part of the fiction is true. Indeed, in historical fictions most of the unembedded propositions may well be true.

Given this, it is neither surprising nor problematic that when reading such fictions we both believe and accept the same propositions. This may make it harder to *detect* the fictive attitudes, as the five markers we considered earlier will not be present in our actual conduct: we will, for instance, assert and reflectively avow that the Portuguese laid siege to the Moorish city of Lisbon. But we clearly still *have* fictive attitudes all the same. Perhaps this is because of our counterfactual conduct: if Saramago had written “the Peruvians” instead, we would have asserted but not reflectively avowed the relevant propositions.

These same lessons carry over to FICTIONALISM about mere formalities. <There is reason for me to set knives on the right> can be true, and not just true according to the relevant fiction. If setting the table incorrectly would be hurtful or uncivil, the content of the fiction may be true. For some standards, most of the content of the relevant fiction may be true: in almost all cases, there may be reason to do what the law requires. And this can be believed by a particular individual (Miss Manners) or most individuals in a society (like the one depicted by *Downton Abbey*). However, that does not show that such individuals lack the relevant fictive attitudes. It just makes it harder to detect their fictive attitudes by considering their actual conduct. We may need to consider their counterfactual conduct. Plausibly, if those individuals were in a case where the relevant social norm was transparently trivial or terrible they would continue to assert, but not reflectively avow, that there is reason to do “what’s done”.

This is why, methodologically, it is most helpful to focus on transparently trivial or terrible social norms. That our shared attitudes towards mere formalities are fictive is easily obscured in cases where we believe or know that we have reason to do “what’s done”, just

as our fictive attitudes towards literature are harder to detect when we believe or know that the relevant content is true.

3.7 Conclusion

We need a systematic theory of mere formalities. We need that theory to explain how mere formalities guide, are internalized, and conflict with other standards despite the believed absence of reasons. And we need that theory to explain how mere formalities are fundamentally distinct from “substantive” standards like morality and epistemology, yet still provide *bona fide* standards of correctness.

My ambition was not just to clarify this problem, but to offer a solution. The distinctive features of our attitudes towards mere formalities suggest a form of fictionalism, which turns out to offer a very plausible way forward. So long as we accept a distinction between reality and fiction, we can offer an illuminating, generalizable and ecumenical analysis of mere formalities. This captures what all mere formalities have in common. Often this category is introduced ostensively: mere formalities include requirements of etiquette, law, linguistic rules, games, clubs, and codes of honor. You know, things like that. FICTIONALISM captures in what respects things must be *like that* if they are to be mere formalities.

However, this is not a complete case for FICTIONALISM. I have made no systematic comparison between my preferred view and all possible competitors. Nor have I considered all of the motivations and complications for this view about mere formalities, such as the question of how we *create* varieties of normativity like etiquette, and whether FICTIONALISM coheres with orthodox semantic views about deontic modals. I take up both of these tasks in the next chapter.

Chapter 4

How We Create Obligations

Whatever business the law may be conversant about, may be reduced to one sort of operation viz that of creating duties (Bentham (1782/1970), p. 249).

[A]n obligation seems a fact of a kind which it is impossible to create (Prichard (1940/2002), p. 257).

Many contemporary legal theorists purport to ‘explain legal normativity’, but often fail to articulate what it means to say that the law is normative or in what way that property requires explanation (Bix (2015), p. 125).

On a commonsensical view about the law, when the legislature passes bills they create new laws, and thereby create new legal obligations and permissions. Sometimes normative terms like “duty” are used in the content of the legislation.¹ Even when such terms are not used, we can still describe the content of the law as “the set of rights, obligations, privileges, and powers in force in a legal system.”² Or we might say that the content of the law is the set of rules in the legal system, which *generates* the relevant set of legal obligations. On either description we create legal obligations.

We can call this set of legal rights, obligations, privileges, and powers “legal normativity”. Like many, I think legal normativity warrants an explanation. Unlike most, I think that the most striking feature of legal normativity that warrants an explanation is that it is something

¹See, e.g., Australia’s *Commonwealth Electoral Act 1918*, 245(1): “It shall be the duty of every elector to vote at each election.” Deontic modals like (“must”, “can”) are more common.

²Hershovitz (2015) describes this as an orthodox view, before rejecting it (p. 1163).

that we can and do *create*. This phenomenon is not unique to the legal domain. We can also create obligations and permissions of other artificial, or “merely formal”, standards, such as prescriptive grammar, games, and etiquette. In this Chapter I aim to provide a general explanation of how we do so.

I will first offer an account of why the commonsensical view that we can create legal obligations is puzzling (§4.1), then offer a fictionalist solution to the puzzle (§4.2). I will then argue that this solution turns out to be surprisingly ecumenical (§4.3), and where it is inconsistent with common commitments in philosophy of law it has the theoretical advantage that it makes better sense of legal language (§4.4).

4.1 Why Is Legal Normativity Puzzling?

One of the most striking features of legal normativity that warrants an explanation is that it is something that we can and do create. In some cases, this is easy to explain. The creation of some legal obligations is explained in terms of the exercise of legal powers—the legislature has the power to create new laws.³ But at some point there were no prior legal rights, obligations, privileges, powers—at some point we created the system of legal obligations whole cloth. That’s where our puzzle starts.

As I said, this phenomenon is not distinctive to the law. But it is distinctive to “merely formal” standards like law. At some point we created the system of requirements of prescriptive grammar whole cloth, and likewise for the requirements of games and etiquette. But at no point did we create the system of requirements of morality whole cloth, and ditto for the requirements of prudence and epistemology.

This point is sometimes obscured in analogies drawn between legal obligations and promissory obligations:⁴ that is, the moral obligations we have because of the promises

³To put this point in the terms used by Hart (1994), the legislature has the legal power to create and change the “primary rules” in the legal system because there are “secondary rules” in the legal system that grant them this power.

⁴Talk of “promissory obligations” is ambiguous. It can refer to conventional norms about promise-keeping; this is a merely formal standard. It can also refer to objective moral norms about promise-keeping. This

we have made. You think that the actions and attitudes involved in making a promise *create* a promissory obligation. (Many doubt this,⁵ but let's leave that aside.) If we can create some particular moral obligation to keep a promise, it is because the act of promising—so long as it meets certain conditions—triggers a prior general moral obligation.⁶ What happens with promises, in other words, is just a more complicated instance of what happens with pain. Given the prior general moral obligation not to cause pain, by making it the case that ϕ ing would cause pain we can make it the case that one has a particular moral obligation not to ϕ . Intuitively, even if we can “create” the particular moral obligation, we cannot create the prior general moral obligation. Even leading contemporary anti-realists about morality defend such claims.⁷ And those that deny this intuitive claim about morality typically accept a similar claim about some normative standards, like practical rationality.⁸

So how do we create systems of legal obligations and the like whole cloth? At a high level of abstraction, there seems to be a near-consensus in the literature. Social practices—i.e., regularities in behavior—may be necessary,⁹ but they are not sufficient.¹⁰ The attitude of *acceptance* plays the most crucial role in the explanation.¹¹

The trouble emerges once we start asking what *type* of attitude acceptance is, and what the *object* of this attitude is. The trouble with the latter is that many accounts of what we accept presuppose the very thing that we want to explain: it is no good saying that we accept a system of obligations, permissions, privileges and powers, because we wanted to appeal to the attitude of acceptance in order to explain how we create that system of normative properties and relations in the first place.

is the use I have in mind here. To see how these can come apart, note that it is perfectly possible that according to the norms in some community we have an obligation to keep promises to perform horrendously immoral acts, even though we have no moral obligation to do so.

⁵See Prichard (1940/2002) and Heuer (2012) (on “Prichard’s axiom”).

⁶I take this talk of triggering from David Enoch (2011b).

⁷See, e.g., Blackburn (1993, 1999a)’s “quasi-realism” on the mind-independence of morality.

⁸For instance, Gauthier (1986)’s contractarianism. Street (2006)’s constructivism is an exception.

⁹Some doubt this, and with good reason: see Brennan et al. (2013) (pp. 20–22).

¹⁰See Hart (1994)’s famous “gunmen-writ-large” thought experiment (p. 56).

¹¹“What makes a normative principle a norm of a particular group or community is that it is somehow *accepted* in that group or community.” Brennan et al. (2013), p. 3.

The trouble with the former is more subtle. There are three intuitive desiderata for any account of what type of attitude acceptance is, and it is hard to see how we can satisfy all three at once because they seem to pull in different directions.

The first desideratum is *intensional adequacy*. This attitude of acceptance must allow us to construct a new standard, transforming *what we do* into *what we ought to do*. In that sense, this type of attitude must have a certain constructive power.¹²

The second desideratum is *extensional adequacy*. This type of attitude must be present whenever we create new systems of obligations, and must be distinct from the attitudes that feature in other nearby social practices. Otherwise our account of how we create obligations will have either false negatives or false positives.

The final desideratum is *psychological plausibility*. We often create standards of correctness with little fanfare. Children create new games—with complex systems of rules, requirements, permissions, and powers—and we intuitively regard this as a fun, innocuous activity. So if an account of how we create obligations predicts that it involves a type of attitude that is too often unreasonable for individuals to have when engaged in such activities, that account is psychologically implausible.

To see why these three desiderata seem to be difficult to satisfy at once, consider the three most common candidates for the attitude of acceptance. The first appeals to a cluster of dispositions. On many interpretations,¹³ this was the path Hart took in *The Concept of Law*. He argued that seeing a pattern of behavior “as a general standard to be followed by the group as a whole” played a pivotal role in the creation of social rules in general. While he described this attitude in various ways,¹⁴ his view seemed to be that we accept a practice when we are disposed to: (a) be guided by that practice; (b) criticize deviations from it; (c)

¹²In case this is obscure, consider the following analogy. To turn metal into money involves not only the physical transformation that takes place in the mint, but the social transformation that turns coins into currency. The second is the type of construction I have in mind here.

¹³See Perry (2014). Cf. Shapiro (2006), who takes acceptance to “explain” our dispositions.

¹⁴For instance, as the “critical reflective attitude” of viewing a pattern of behavior as: a “reason and justification” for action; the “basis for claims, demands, admissions, criticisms or punishment”; and establishing the “legitimacy” of such criticism: see Hart (1994), pp. 90, 56.

regard such criticisms as legitimate; and (d) express such criticism using normative language, such as deontic modals.¹⁵ If enough of us take this attitude the practice is transformed into a social rule.¹⁶

How does this dispositional account fare with respect to our desiderata? It is hard to see how the account is intensionally adequate. How does adopting these dispositions transform what we do into what we ought to do in any interesting sense? Some may answer: “because that’s what social rules are.” But this does not offer a satisfactory explanation; it simply insists that the account is explanatory.

The problems with extensional adequacy are even more stark. Consider social practices in sports. In cricket there is a practice (i.e., a general pattern of behavior) of drawing near the batsmen when slow bowlers are bowling. Most players are guided by this practice, criticize deviations from this practice, regard such criticism as legitimate, and express such criticism using normative language: “You should have drawn nearer to the batsman!” But players are not *required* to draw near.¹⁷ Even if everyone has these dispositions the practice remains a strategy; it is not transformed into a rule. Hart responded to this objection by restricting the account further. Along with (a)–(d), for a practice to be transformed into a social rule it must: (e) be enforced by serious pressure to conform; (f) be believed to be important to social life; and (g) possibly conflict with the interests of subjects.¹⁸ However, this makes the account too restrictive. Children create games, which involve social rules on Hart’s view, without (e)–(f) being present. So this dispositional account does not seem promising.¹⁹

The second most common candidate for acceptance is a non-cognitive attitude. In his *Essays on Bentham* Hart argued that to accept a practice is to approve of punishments

¹⁵Hart (1994), p. 57.

¹⁶See, e.g., Hart (1994), p. 58: “The acceptance, and so the existence, of such a rule...” It is worth noting that, as Green (2002) discusses, there were some evolutions in Hart’s theoretical ambitions: he settled on the aim of providing a general account of “conventional obligations”, p. 518.

¹⁷This point was first made by Warnock (1971). See also Shapiro (2011) and Perry (2014).

¹⁸Hart (1994), pp. 85–88.

¹⁹For further criticism of this view, see Raz (1999), pp. 53–58 and Dworkin (1977), pp. 48–58.

for deviations from that practice.²⁰ In *Legality* Scott Shapiro argued that acceptance is an intention to act in accordance with a social practice.²¹

How would such non-cognitive attitudes—approving or intending—transform patterns of behavior into social rules with accompanying obligations and permissions? Hart’s answer appeals to unorthodox expressivist commitments about the semantics of legal language which we will discuss later (in §4.4 and Chapter Five). Shapiro’s answer is less clear. He appeals to the claim that the adoption of plans triggers rational requirements: “the adoption of plans by legal officials will induce a rational requirement for those individuals to comply with the law”.²² Shapiro recognizes that this account does not explain how citizens who do not accept the law also have legal requirements;²³ but he does not recognize that this account struggles to even explain officials’ legal obligations (as legal obligations have a narrow-scope, but rational requirements do not),²⁴ or their legal permissions, privileges and powers (as forming new intentions does not trigger new rational permissions, privileges or powers).

It is also unclear whether these non-cognitivist accounts of acceptance are extensionally adequate. Hart’s appeals to approval of punishments result in far too many false negatives. For many social rules, including many legal rules,²⁵ we do not provide (and would not approve of) a punishment for deviations from the relevant pattern of behavior. This is especially obvious in games. There is no punishment for violating the rule that one must

²⁰Hart (1982), pp. 159–160. See Toh (2005) for discussion of this interpretation of Hart.

²¹Shapiro (2011)’s account of intentions seems to be quite idiosyncratic. On his view, the objects of intentions are plans, and “plans are abstract entities that require, permit, or authorize agents to act, or not act, in various ways under certain conditions” (p. 127). So on Shapiro’s view the object of my intention to eat breakfast is not some (propositional representation of an) action, but some (propositional representation of a) requirement or permission to perform this action.

²²Shapiro (2011), p. 181.

²³Shapiro (2011) tries to fix this problem by adopting a Razian semantics for legal language: such citizens’ claims about having a legal obligation to perform some action “may be understood to mean that from the legal point of view one is (morally) obligated to perform that action”, where the “legal point of view” is understood to be “the perspective of a certain normative theory” that holds that the laws of the relevant legal system “are morally legitimate and obligating” but is “silent on how to answer” all other moral questions (pp. 185–186). For objections to this Razian view, see §4.4.

²⁴Plausibly, the relevant rational requirement has the form [if you intend to ϕ , then ϕ], which is equivalent to the disjunction [ϕ or do not intend to ϕ]; the legal requirement is simply to ϕ .

²⁵See Kramer (2001) on “nominal duties”, pp. 65–78. As Green (2002) notes (p. 517, n.8), after Hart’s objections to Austin’s “sanction theory” of legal obligations in *The Concept of Law*, his *Essays on Bentham* offer a “puzzling endorsement of something like a sanction theory”.

not castle in check, because this is a constitutive rule of the game.²⁶ Shapiro's appeal to intentions, on the other hand, will result in false positives. After all, cricketers presumably do intend to draw near slow bowlers.

Non-cognitivist views also struggle to satisfy the final desideratum: they are not psychologically plausible. This is because attitudes like intentions are subject to strict norms of coherence. If you intend to eat and intend to fast, your attitudes are incoherent, irrational. So what happens when we accept a standard that has internally inconsistent dictates?²⁷ Or accept multiple standards that contain jointly non-compossible requirements? These are familiar phenomena. But the non-cognitivist suggests that we are unreasonable in all such cases, as we have adopted an incoherent set of plans (or states of approval, or what have you). So these views fail to explain why accepting merely formal standards is typically an innocuous affair.

The final candidate for the attitude of acceptance is some form of cognitivism. We might identify acceptance with moral beliefs: the belief that social rules are morally justified;²⁸ the belief that we are morally required to act in accordance with the rules; or, the belief that it is morally justified to accept the rules.²⁹ Or we might identify acceptance with beliefs about generic normative notions like reasons (for rules, for complying with rules, or for accepting rules).³⁰ Both forms of the view face similar problems with the three desiderata we identified above.

Again, these accounts are not intensionally adequate. A community of Ayn Rand's acolytes might accept her "objectivist" moral egoism—they might all form all of the requisite moral beliefs, as well as the requisite generic normative beliefs—but this does not make it true that they ought to be egoists in any interesting sense. Believing that we have decisive (moral) reasons to act in accordance with whatever rules are laid out in *The Fountainhead*

²⁶See Rawls (1955).

²⁷For an interesting real-world example, see Cohen (2004) on American baseball.

²⁸See MacCormick (1978).

²⁹Holton (1998), p. 6.

³⁰For general discussion, see Bix (2006).

and *Atlas Shrugged* does not generate a new normative standard. It simply involves false beliefs about an existing normative standard.³¹

Nor are these accounts extensionally adequate. If we appeal to moral beliefs, we face false negatives. This is partly because this account is overly-intellectualized: children can create games without any moral beliefs whatsoever. But it is also because this account is overly optimistic about our motivations. Mischievous children might accept the rules of the game because they believe that (complying with) the rules is immoral: anyone who thinks that schoolyard games are never designed to be vicious had an unusually blessed childhood. We might not face these same problems if we appeal to beliefs about generic normative notions. However, this move results in false positives. Cricketers presumably believe that they have good reason to draw near slow bowlers, and that they have those reasons because of how the game is played.

Finally, cognitivists face even worse problems with the final desideratum. Beliefs are subject to a strict norm of coherence, and this norm will be violated when we accept an internally inconsistent standard, or jointly inconsistent standards. Moreover, beliefs are subject to strict norms about responsiveness to evidence: reasonable beliefs are paradigmatically “shaped primarily by evidence for what is believed and concern for the truth of what is believed. Thus the slogan: belief aims at truth.”³² However, it seems perfectly common, and *prima facie* reasonable, to accept requirements of games, grammar, etiquette, and the like without a shred of evidence that there are (moral) reasons for the requirements, for compliance, or for acceptance. So cognitivism doubly fails to explain why accepting and creating standards of correctness are typically an innocuous affair, rather than an epistemically irresponsible activity.

This cluster of problems for dispositional, non-cognitive, and cognitive accounts of acceptance is troubling. In part, the problems are troubling because it is not clear what other candidates for acceptance we could consider. And in part they are troubling because they

³¹Similar points have been made by, for instance, Coleman and Leiter (1996), p. 241.

³²Bratman (1992), p. 3.

show how our initially plausible desiderata seem to pull in different directions. For our account to have intensional adequacy we need to identify an attitude of acceptance that has a certain constructive power. For our account to be psychologically plausible we must identify an attitude of acceptance that is not subject to strict norms of coherence (and responsiveness to evidence). For our account to be extensionally adequate we need to appeal to attitudes that are not present in nearby activities like strategic behavior in sport. These all count against appealing to familiar attitudes like intentions and beliefs. But if we appeal to unfamiliar attitudes our account will also face problems with extensional adequacy, as it will be doubtful that children have such exotic attitudes when they create and participate in games.

4.2 The Fictionalist Solution

This puzzle should seem soluble, because we do create normative standards. So I propose that we should search for a solution by identifying the most familiar candidate for acceptance that does seem to have the requisite constructive power. That is, we should start by satisfying the neglected desideratum of intensional adequacy, then consider whether the account is extensionally adequate and psychologically plausible.

It is commonsensical that we can create fictions.³³ Some also argue that we can create fictional characters, which seems just as plausible: Shakespeare, we might say, created both *Hamlet* and Hamlet. This so-called “creationist” view takes many forms,³⁴ but the explanation of *how* we create fictional characters typically appeals to facts about fictive acceptance.³⁵ More generally, we might say, we create fictions by accepting their contents. Creationists typically appeal to the attitudes of the author that are present in certain intentional activi-

³³“The term “fiction” derives from the Latin *fingere* meaning “to form,” and this linguistic root is still evident in treating fictional characters as entities formed by the work of an author or authors in composing a work of fiction. We do not describe authors of fictional works as discovering their characters or selecting them from an ever-present set of abstract, nonexistent, or possible objects. Instead, we describe authors as inventing their characters, making them up, or creating them, so that before being written about by an author, there is no fictional object. Taking authors to be genuinely creating as they make up fictional characters is central to our understanding of fiction” (Thomasson (1999), p.6).

³⁴See Thomasson (1999, 2003), Salmon (1998), Searle (1975), Schiffer (1996), & Kripke (2013).

³⁵For instance, Searle (1975) appeals to acceptance (or “pretense”) and illocutionary acts: (“by pretending to refer” to a person, the author “creates a fictional character”, p. 330).

ties, but stated at this level of abstraction the view being proposed here is also compatible with literary theories that put less emphasis on authorial intentions and more emphasis on the creative acts of the audience. (In Barthes (1977)’s famous slogan, “the author is dead”.) On such versions of the view, audiences create the fiction by accepting its content.³⁶

Amie Thomasson has argued that “the ontology of fictions can ... serve as a model for the ontology of other social and cultural objects in the everyday world”, including legal obligations.³⁷ Thomasson argues for this on the ground that these objects all fall into the category of *abstract artifacts*: legal obligations, like fictions, are abstract rather than concrete entities, yet also artificial rather than eternal Platonic forms. But I want to explore the idea that the ontology of fictions provides a model for the ontology of legal obligations for a more specific reason. Fictive acceptance is a familiar attitude that has the constructive power we are looking for: it explains how we create abstract artifacts in general. So if it is plausible that the creation of merely formal standards is a particular instance of this more general phenomenon, we will have an intensionally adequate account of how we create legal obligations.³⁸

I think it is plausible that the creation of merely formal standards is a particular instance of this more general phenomenon. To see why, imagine that some children see a cluster of

³⁶This is important for two reasons. First, because this move easily avoids Brock (2010)’s objections about the role of authorial intention. And second, because in at least some cases, such as customary law, legal obligations may be created by a community (rather than a separate legislator), and without any overt intention to create anything like a law or obligation. Hart (1994) made both of these points. Thomasson (1999) does not seem to appreciate them: “a law does not exist unless it is brought into existence by acts of writing, raising hands, or marking ballots (each of which, to be a genuine piece of voting behavior, requires that it be done with the intention of casting a vote)” (p. 130). She does write that the creative acts can sometimes be “of a community”, but always insists that they are “intentional” (pp. 13, 31).

³⁷Thomasson (1999), p. 14, and pp. 13, 41 and 133 (on the creation of laws).

³⁸It is worth acknowledging that some philosophers are suspicious of the notion of abstract artifacts (“Can there really be abstract things that we make? Some might find it implausible that even God could literally create an abstract object” van Inwagen (2003), pp. 153-154). and of proffered explanations for how we create fictional characters (such explanations suffer “from the same problem as theological creationism: the purported explanation is more mysterious than the data it seeks to explain” Brock (2010), p. 336). But strangely, such suspicions are only taken to impugn an ontological commitment to fictional characters, and are not directed towards other paradigmatic candidates for abstract artifacts like fictions and theories. (Thomasson (1999) makes this same point.) Two particularly striking instances of this are Walton (1990) and Brock (2002, 2010). The former explicitly accepts that we create fictions but denies that we create fictional characters. The latter is a “fictionalist about fictional characters”, which raises hard questions for his view. If fictional characters exist according to fictions, do those fictions exist, and did we create them?

tree stumps and start to act and speak *as if* tree stumps are bears. Upon seeing a stump they might point, scream “Bear!”, and run away. As good creationists, let’s say that the group of children has created a fiction by accepting its contents, which include propositions like “Tree stumps are bears”. Now imagine that the children’s practices evolve in various ways. Perhaps they will start sneaking up on stumps, tying ropes around them, and declaring thereafter that this “captures” the bears. They might start acting and speaking as if they “must” capture the bears, and that whoever captures the most bears wins. Sometimes the practices will evolve in response to unforeseen events: after one kid tries to untie another’s stump, the others might (i) act and speak as if one “can” release bears that someone else captured, rather than (ii) acting and speaking as if one “cannot” do so. There may be no conscious choice for (i) over (ii)—this might just reflect the relative popularity of the children involved. In any case, as the practice develops, it seems obvious enough that a game has emerged with an object and a system of requirements and permissions.

The basic idea of the fictionalist proposal is that the children created a fiction about bears and a system of obligations and permissions in the same way: via the attitude of acceptance. Initially the children accepted descriptive propositions (“Tree stumps are bears”). They thereby created a fiction with descriptive content. Then they accepted further normative propositions (“One must capture bears”). They thereby added normative propositions to the content of the fiction. And in doing so, they generated a distinct standard of correctness: what one ought, must, and can do according to the relevant fiction. So the creation of a merely formal standard is just a particular instance of a more general phenomenon: the creation of fictions.

How might this fictionalist proposal be developed? In Chapter Three I defended a particular form of fictionalism as a promising account of mere formalities in general:

Fictionalism A property or relation is merely formally normative iff and because according to the relevant fiction it is substantively normative.

To illustrate this proposal, it might be that relations and properties like moral obligations are “substantively normative” iff and because they are analyzable in terms of reasons: roughly, one is morally required to ϕ iff and because one has decisive reason to ϕ . If something like this is true, then relations and properties like legal obligations are merely formally normative iff and because they are analyzable in terms of reasons that exist according to the relevant fiction. So in the game of make-believe we imagined above, players are required to capture bears iff and because according to the relevant fiction players have decisive reason to capture bears, and so on. One nice feature of this way of illustrating the proposal is that it allows us to see how the objects of the attitude of acceptance (a fiction about reasons) are distinct from the properties and relations we want FICTIONALISM to explain (obligations and permissions).

So far I have proposed FICTIONALISM as an explanation of how we create mere formalities like legal obligations. But what about our two other desiderata?

FICTIONALISM is extensionally adequate for two reasons. First, it is designed to explain the simplest case of children creating a game, so it does not rely on any idiosyncratic feature of a particular domain (like legal punishment). This makes it easy for this explanation to generalize to more complex domains like law and grammar. Second, it identifies an attitude that is familiar enough to be plausibly present in activities like the creation of games, but distinctive enough to be absent from nearby phenomena like strategic behavior in sport. Cricketers may well believe that they have reason to draw near slow bowlers, but it is doubtful that they accept a fiction according to which they have reason to draw near; their conduct is unlikely to manifest any of the distinctive markers of fictive attitudes.³⁹

What about our final desideratum? FICTIONALISM does not encounter problems with psychological plausibility because fictive acceptance floats free of the norms of coherence that govern other cognitive and non-cognitive states, as well as the norms of responsiveness to evidence that govern beliefs. Today children might accept a fiction according to which tree stumps are bears even though this is inconsistent with their knowledge and evidence

³⁹For more discussion, see Chapter Three.

about biology. Tomorrow children might accept a different fiction according to which tree stumps are pirates. This is all innocuous. We are permitted to accept fictions for all sorts of reasons, even if their contents are internally or mutually inconsistent—no one, after all, is inclined to label an engrossed reader irrational when she realizes that Watson’s shoulder wound from *A Study in Scarlet* is now a leg wound in *The Sign of the Four*. Similarly, it is innocuous for children to accept a normative fiction according to which they have decisive reason to capture bears even if this is inconsistent with their beliefs, with their evidence, and with the contents of the fiction that they accepted yesterday.

So far all of this discussion has been at the level of the genus, mere formalities. But is FICTIONALISM a plausible view about legal obligations in particular?

One initial worry is that FICTIONALISM trivializes legal practices. Judges’ attitudes towards the law, after all, seem far more solemn than the playfulness we paradigmatically find in children’s games of make-believe. But this rests on a misconception of fictive acceptance. Accepting a fiction can be a solemn affair.⁴⁰ The markers of fictive acceptance do not include frivolity. Rather, they include distinctive features like context-dependence. We can rationally accept a proposition in one context that we do not accept in another, whereas we cannot rationally believe a proposition in one context and not believe it in another (at least, not unless the proposition’s content includes context-sensitive terms like indexicals). Interestingly, Adam Perry (2014) has pointed out that those same features are present in judicial behavior: in that sense, a judge may be “like an actor in a play” insofar as she only accepts certain propositions while on the bench, for a wide range of reasons.

A more general worry, however, concerns how FICTIONALISM can fit into the familiar landscape of contemporary philosophy of law. In the next section, I will argue that FICTIONALISM has the virtue of being surprisingly ecumenical. It is compatible with common

⁴⁰There is an interesting parallel between this fictionalist proposal and views held by Nietzsche (1882/1974a), who appealed to “the seriousness one had as a child, at play” (94) as being essential to the explanation of how we create new values. See also Nietzsche (1882/1974b, 1883/1982).

commitments about legal normativity, and can even clarify core features of views that are often considered obscure or elusive (by non-partisans, of course).

4.3 How Fictionalism Fits Into the Familiar Landscape

The field of jurisprudence has been preoccupied with questions about legal normativity for over half a century. The battleground is well populated, and neatly divided into three warring factions, each of which carries the standard of a towering figure: H.L.A. Hart, Joseph Raz, and Ronald Dworkin. The three views that they championed—inclusive legal positivism (Hart), exclusive legal positivism (Raz), and anti-positivism (Dworkin)—were offered as answers to the same question:

The question is what facts determine the content of the law. In different ways, all three views purport to provide a metaphysical account of our legal rights, obligations, privileges, and powers. In other words, they purport to tell us what makes it the case that we have the legal rights, obligations, privileges, and powers that we do.⁴¹

FICTIONALISM is also a metaphysical account of what makes it the case that we have the legal rights, obligations, privileges, and powers that we do. So it is a surprising claim that FICTIONALISM can be compatible with all of these inconsistent views.

Nonetheless, I think that this claim is true. I do not want to overstate my case here. FICTIONALISM is not compatible with everything that these philosophers have said about legal obligations. Here it is helpful to distinguish *metaphysical* and *semantic* commitments in philosophy of law. The two are often conflated; there is a general “uncertainty in contemporary jurisprudence about whether its subject-matter is (a) the *concept* of law, or rather (b) law as a social reality”.⁴² In this section my aim is to argue that FICTIONALISM is compatible with positivist’s and anti-positivist’s core metaphysical commitments about legal obligations

⁴¹Hershovitz (2015), p. 1167.

⁴²Finnis (2002), p. 15.

considered as a social reality; I will defer discussion of the meaning of concepts and terms until §4.4.

There are three general reasons why FICTIONALISM can be such an ecumenical metaphysical view about legal normativity. First, much of the discussion of “legal normativity” in contemporary philosophy of law is simply about a different (albeit related) subject matter from the one at issue here. Some philosophers offer no account of legal obligations, take it for granted that we create legal obligations, and ask whether and how legal obligations can give us (decisive and/or moral) reasons for action. This is true of some positivists and anti-positivists alike.⁴³ There is no direct tension between these views and FICTIONALISM; the latter is not a view about whether, when, or why we have reason to act in accordance with our legal obligations.

Second, positivists’ and anti-positivists’ core metaphysical commitments about legal obligations concern what *ultimately* explains legal obligations. Gideon Rosen takes the legal positivist’s “distinctive claim” to be “that the grounds of law are wholly social, consisting ultimately in the acts of officials and the social practices in which they are embedded”, and Scott Shapiro defines legal positivism as the view that “all legal facts are ultimately determined by social facts alone.”⁴⁴ Anti-positivism is typically characterized as the view that legal facts are ultimately determined by social and *moral* (or, more broadly, *normative*) facts, not social facts alone. There is an important, neglected methodological problem here. If we are interested in legal obligations as social reality, theorizing about legal obligations should be continuous with theorizing in social ontology more generally. As Elizabeth Barnes (2014) argued, there’s far more to social ontology than identifying what facts are explanatorily fundamental. We might all agree that quarks and bosons ultimately explain everything,

⁴³For instance, see Marmor (2011). Anti-positivists are often assumed to accept the view that “all legal requirements are also moral requirements” (Coleman and Leiter (1996), p. 244); and moral requirements, according to the natural law tradition, are most definitely not artificial. However, it is clear that traditionally anti-positivists have accepted that we do create legal obligations, which are not identical to moral obligations: Finnis (2002) explicitly rejects Coleman and Leiter’s characterization of his view (p. 20), and accepts that we can create “legal requirements” that are not “moral requirements”, but provides no account of what they are.

⁴⁴Rosen (2010), p. 110; Shapiro (2011), p. 27.

but that would not provide an adequate account of *what it is* to be a woman or man, or black or white. The same observation applies to legal obligations: knowing that they can (must) be ultimately explained by descriptive facts alone, as Hart (Raz) held, or that they must be ultimately explained by descriptive and normative facts, as Dworkin held, would not provide an adequate metaphysical account of *what it is* to be a legal obligation. So each of these theories leaves a lacuna that FICTIONALISM, thankfully, can fill. It can provide an account of what it is to be a legal obligation that is compatible with other metaphysical commitments about what ultimately does the explanatory work: that is, facts about legal obligations could be explained in terms of facts about the contents of fictions, and facts about the contents of fictions could in turn be explained in terms of descriptive and/or normative facts.

Third, and perhaps most importantly, FICTIONALISM has been proposed as a view about the genus of mere formalities. It offers no account of the differentia that distinguish legal obligations from the requirements of games and prescriptive grammar. And different views about the differentia can easily motivate positivism or anti-positivism. For instance, a natural thought is that we can distinguish these species by their distinctive roles.⁴⁵ While many think that (attitudes towards) moral obligations play separate explanatory and justificatory roles,⁴⁶ there is a great deal of disagreement about whether the same holds for (attitudes towards) legal obligations. If legal obligations play a justificatory role that distinguishes them from prescriptive grammar et al., that may motivate anti-positivism; if they only play an explanatory role in guiding our conduct, that may motivate legal positivism, or varieties thereof.⁴⁷

To illustrate all three points, let's consider in detail the version of anti-positivism defended by Dworkin (1986). Dworkin provides no clear account of what it is for an act to be legally

⁴⁵Similar claims are often made in essentialist and teleological terms: law's function is partly "constitutive" of what law is (Shapiro (2011)): "most modern theorists agree that [law's] function" partly explains its "essence" (Soper (1984)); "all legal systems necessarily perform ... social functions", including guiding action, as part of their "essence" (Raz (2009)).

⁴⁶See Owens (2012), ch. 5 (on "bare wrongings"). Korsgaard (1996) argued that any view of moral normativity must have "justificatory adequacy" and "explanatory adequacy" (pp. 12–13).

⁴⁷For instance, Shapiro (2001) famously argued from the guiding function of law to exclusive legal positivism.

obligatory. He is more concerned with the justificatory role that legal obligations play: they provide “a general justification for the exercise of coercive power by the state.”⁴⁸ We could square FICTIONALISM with Dworkin’s view by taking facts about fictions to be grounded in descriptive facts alone and appealing to this justificatory role as what distinguishes legal obligations from other mere formalities. But there is a deeper connection between the two views. Dworkin relied extensively on analogies between law and fictions: the contents of both are ultimately determined by facts about what fits and justifies the total set of relevant practices (such as the written text of the statutes or novels). The fictionalist holds that facts about legal obligations are determined by the content of the relevant fiction (about what we have reason to do). Dworkin is committed to facts about the contents of fictions being ultimately determined by facts about fit and justification. Dworkin’s central metaphysical commitment about legal obligations follows from this straightforwardly.

So far I have argued that FICTIONALISM is compatible with common metaphysical commitments about legal normativity. But I also promised to show that it can clarify features of views that otherwise seem obscure. If this is right, FICTIONALISM is an even more welcome addition to the familiar landscape in jurisprudence.

Here are three examples of how FICTIONALISM offers us theoretical resources to demystify obscure claims. First, natural law theorists often contend that we can create legal obligations which are not moral obligations, but these are not “genuine obligations” either.⁴⁹ This sounds odd. If obligations exist, in what sense are they not “genuine”? FICTIONALISM makes this thought perfectly intelligible. We can create Sherlock Holmes, but we do not thereby create a “genuine detective” because Holmes is not the same as a flesh-and-blood investigator. Likewise, we can create legal obligations which are not “genuine” in the sense that they are not metaphysically or normatively equivalent to moral obligations. Moral

⁴⁸Dworkin (1986), p. 90. Hart (1994) famously denied this: it is “quite vain to seek any more specific purpose which law as such serves beyond providing guides to human conduct and standards of criticism of such conduct” (p. 249).

⁴⁹Finnis (2002) notes that in deeply unjust legal systems we might create “legal requirements”, but “cannot create what any self-respecting person would count as a genuine obligation” (pp. 22–3).

obligations, after all, are explained by facts about reasons, not facts about reasons that exist according to a fiction.

Second, many legal positivists suggest that legal obligations are grounded in facts about “legal reasons”.⁵⁰ Some such philosophers at least seriously entertain the thought that legal reasons are “a distinct class of reasons for acting” that count in favor of actions alongside moral and prudential reasons.⁵¹ I find this implausible and have no wish to defend it. (Are there also distinctive classes of grammar reasons and etiquette reasons and chess reasons that compete with moral and prudential reasons?)⁵² But notably, many positivists who agree that “legal reasons” are not “real reasons” in this sense nonetheless appeal to legal reasons, without providing an account of what they are if they are not “real”.⁵³ This view seems quite obscure. As Herskovitz (2015) points out, “if you say that legal reasons are not real reasons, you owe us an account of what they are, since the claim is that they are not what they purport to be.”⁵⁴ But note how easily FICTIONALISM can offer such an account: legal reasons are reasons that exist according to a legal fiction, and since that fiction can be false legal reasons may not be what they purport to be. This gives us a clear account of otherwise obscure philosophical commitments, and one that has the virtue of being quite consonant with legal positivists’ remarks about legal reasons.⁵⁵

Third, many legal positivists’ accounts of legal obligations centrally appeal to “what the law claims”. Such appeals are typically part of an account of the meaning of legal language that was first advanced by Joseph Raz and has received a great deal of support in jurisprudence since: “*A* is legally obligated to ϕ ” means “From the point of view of the law *A* is morally obligated to ϕ ”.⁵⁶ Sometimes this view is presented in a way that is even more

⁵⁰See, e.g., Horty (2012).

⁵¹See, for instance, Coleman (2011), p. 5. See also Raz (2009): legal reasons are a “special and distinct kind of reasons”, like “moral reasons or prudential reasons” (p. 68).

⁵²See also Kramer (1999), p. 379 (“morality and prudence exhaust the realm of reasons”).

⁵³See Shapiro (2006) (on the “qualified” reading of “legal reasons”). And see also Enoch (2011b), pp. 16–19.

⁵⁴Herskovitz (2015), p. 1194, n.55.

⁵⁵Enoch (2011b) suggests a fictionalist view in comparing “legal reasons” to “imaginary friends”.

⁵⁶The two most prominent endorsements of this view are: Raz (1999, 2009); and Shapiro (2011), that one is legally obligated to ϕ “may be understood to mean that from the legal point of view one is (morally

committal.⁵⁷ I will address this as a semantic view shortly, but here I want to focus on the metaphysical commitments of this view.

On this view, what are legal obligations? The answer seems to be that “legal obligations are simply constituted by the linguistic contents of the pronouncements of legal institutions.”⁵⁸ That’s odd, as you might think that there are legal obligations that have not yet featured in the contents of any institutional pronouncements.

On this view, what is it for the law to “make claims”, or “have a point of view”? The best answer on offer attributes to the law claims made by its officials: “Law makes claims only insofar as law-applying officials make those very same claims at the very same time and place. The claims of law are identical to certain claims of its officials.”⁵⁹ This is problematic, as the presence of officials is an idiosyncratic feature of the legal domain; it remains unclear whether or how this metaphysical and semantic view could extend to other merely formal standards like grammar or games.

FICTIONALISM offers a less elusive, and easily generalizable, way to interpret this position. For the law to “make claims” or “have a point of view” is for the relevant legal fiction to contain propositions. These propositions need not be pronounced by officials to be attributed to the law: we did not need to appeal to officials or speech acts to explain how “Tree stumps are bears” became part of the content of a fiction.

obligated to perform that action” (p. 185). See also Coleman (2011), pp. 22, 78. Even Raz’s critics often accept that his view offers a plausible account of the literal meaning of “legal obligation”: see, e.g., Kramer (1999), p. 376 n. 1, and *passim*.

⁵⁷For instance, Leslie Green (2002) writes that “legally speaking there is an obligation to ϕ ” means that from the point of view of the law there are bringing reasons to ϕ and exclusionary reasons not to act on some of the reasons to the contrary” (p. 519), where the latter is a proffered metaphysical analysis of what moral obligations are.

⁵⁸Greenberg (2014), p. 1304. He also notes that this view denies that “a legal obligation is a kind of obligation at all”. This is continuous with how the view is described by its advocates: “a legal obligation or right is none other than an obligation or right that exists according to law. And an obligation of right that exists according to law is none other than an obligation or right, the existence of which law claims” (Gardner (2012), p. 133): “The normativity of law thus involves questions of how the law presents itself to us” (Green (2002), p. 519).

⁵⁹Gardner (2012), p. 131. This is consonant, Gardner notes, with Raz’s view that “The claims that law makes for itself are evident from the language it adopts and from the opinions expressed by its spokesmen” (Raz (1995), p. 295).

Before we move on to considering semantic commitments about legal language, it is worth addressing one final point about the metaphysics of legal normativity. Much of the current literature on legal normativity is framed around what has become known as “Hume’s law”, which is often glossed as follows: no normative fact can be fully grounded in (or “derived from”) descriptive facts alone.⁶⁰ I think that this is problematic for three reasons. First, it’s not clear that “Hume’s law” is actually a *law*, since it is vulnerable to counterexamples.⁶¹ Second, it’s not clear that “Hume’s law” is actually *Hume’s*: Hume only claimed that facts about what we morally ought to do cannot be derived from descriptive propositions alone; he made no analogous claims about what we prudentially or legally ought to do. Thirdly, and most importantly, “Hume’s law” says nothing about the salient feature that distinguishes legal and grammatical standards from moral and prudential standards: that we create the former. So if we only focus on a claim about grounding or derivation we may miss something that is interesting, puzzling, and distinctive about domains like law.

But in case Hume’s law keeps you up at night, it is worth seeing why that should not lead you to reject FICTIONALISM. For one thing, we have already seen that there are plausible views on which facts about the relevant fictions will be partly grounded in normative facts (*a la* Dworkin’s anti-positivism). For another, the facts that these fictions ground could be construed as non-normative depending upon where we stand with respect to some complex commitments about propositions and prefixes.⁶²

⁶⁰Hume (1739)’s point was framed in terms of deriving propositions, not explaining facts. See Maguire (forthcoming) for a gloss of Hume’s law in terms of the ideology of grounding.

Both legal positivists and anti-positivists frame the discussion around Hume’s law. Here’s Coleman and Leiter (1996): “If law is a matter of fact, not value, then what can explain its normative force? How can we derive a normative conclusion (about law’s authority) from a factual premise (it’s legality)?” (p. 244). And here’s Finnis (2002): “No *ought* from a mere *is*. So, since positivism prides itself on dealing only in facts, it deprives itself of the only conceivable *source* of reasons for action (oughts), namely true and intrinsic values” (p. 23).

⁶¹See, for instance, Anscombe (1958), pp. 2-3, and Searle (1964), especially p. 56.

⁶²On the version of fictionalism I presented, there are straightforward facts of the form “Legally, *A* is obligated to ϕ ”. But one might prefer a version of fictionalism according to which any sentence of that form is shorthand for “According to the legal fiction, *A* is obligated to ϕ ”. The latter is a descriptive proposition that contains normative concepts. For discussion, see Lewis (1978) and Raz (2009) on whether “according to...” prefixes form a part of the relevant proposition.

4.4 What “Legal Obligation” Means

I have emphasized how ecumenical FICTIONALISM is with respect to metaphysical commitments about legal normativity. Now let’s turn to consider common semantic commitments about the concept or term “obligation” as it appears in legal contexts. Here I want to contend that FICTIONALISM should be inconsistent with central claims made by Hart, Raz, and Dworkin about what “legal obligation” means, because these claims are poorly motivated and at odds with an orthodox view in linguistics.

Let’s start by outlining that orthodox view about the meaning of normative terms—that is, deontic modals like “ought”, “must”, and “can”, and cognates like “obligation”.⁶³ This view takes “obligation” to be a *standard-relative* term: it has a single “standing meaning”, or “character”, that can be relativized to different standards (“moral”; “legal”). To illustrate this view, consider these sentences:

1. Antigone is legally obligated not to bury Polynices.
2. Antigone is morally obligated to bury Polynices.
3. Antigone is not morally obligated not to bury Polynices.

Does “obligated” mean the same thing in all three of these sentences? If so, what does it mean? And what do the modifiers (“legally”, “morally”) add to its meaning?

On this orthodox view, “obligated” means the same thing in each sentence. Its meaning is its character, which is understood as a function from contexts to contents. Identifying the right function takes a good deal of work that would be quite distracting. We can illustrate the approach using a simple candidate:

Obligated_{fg} $\phi = 1$ iff $\forall w((w \in [[F]] \cap [[G]]) \supset w \in [[\phi]])$.

⁶³This view is most associated with Angelika Kratzer (1977, 1981, 1991). To be clear, the view I describe as orthodox is only committed to (a) taking each modal to mean a function from contexts to contents, and (b) taking a modal base *and* an ordering source to be relevant to each modal. There is plenty of disagreement about how to fill in the specific details.

The “modal base” f provides a set of propositions, F , which represents the relevant circumstances. The “ordering source” g provides a set of propositions, G , which represents the demands of the relevant standard. “ A is obligated to ϕ ” is true iff “ $A \phi$ ” is true at every world in the intersection of F and G .⁶⁴ Let’s apply this to (2): this is true iff “Antigone buries Polynices” is true at every world in the intersection of F and G . F is a set of propositions—like “Polynices is Antigone’s brother”, “Polynices has not been buried”, and so on—that represents the relevant circumstances. G is a set of propositions—like “Everyone buries their family members”—that represents the demands of the relevant (moral) standard. The intersection of F and G is the set of worlds where all of the propositions in F and G are true. It is obviously true that “Antigone buries Polynices” at all of these worlds. This explains why (2) is true. And it is fairly easy, from here, to see how it explains why (3) will be true too.

Of course, there are obstacles for this orthodox semantics for “obligated”, but it remains an orthodoxy in linguistics because it is supported by plenty of linguistic data, is highly parsimonious, and has significant explanatory power. Showing due deference to linguists, our starting point in jurisprudence should be that this is what “ought” and “obligation” mean in English, as well as in other natural languages; this view is robustly cross-linguistic in its application.⁶⁵

Linguists apply this view to “moral obligation” and “legal obligation” in the exact same way. So (1) is true iff “Antigone does not bury Polynices” is true at every world in the intersection of F and G . Since this sentence was uttered in the same context and circumstances as before, F is the same set of propositions above. But G is different: the set of propositions representing the demands of the relevant (legal) standard includes, for instance, “If x dies attacking Thebes, no one buries x ”. “Antigone does not bury Polynices” is true at every one of these worlds, so (1) is true too.

⁶⁴The function above is far too simple, but it neatly illustrates the essential innovation from Kratzer: that *two* conversational backgrounds are relevant to each modal. The modal base determines which worlds are accessible, while the ordering source ranks those worlds.

⁶⁵See especially von Stechow and Iatridou (2008).

We’ve now illustrated how, according to this orthodox semantics, “obligated” means the same thing in moral and legal contexts, and in particular it means the same function from contexts to contents, where contents are understood to be propositions. We’ve also illustrated what the modifiers “legally” and “morally” add to the meaning of the sentence: they explicitly shift the ordering source. According to linguists, there is an “open-ended set” of standards to which “ought” and “obligated” can be relativized.⁶⁶ For some, but not all, of these standards we often use explicit modifiers like “legally” and “morally” (e.g., “grammatically”). In the absence of such modifiers we must pick out the salient ordering source: e.g., “Kobe must not double dribble” is implicitly relativized to the rules of basketball. So this semantics also applies to the other merely formal standards we have considered, like games and etiquette.

With this orthodox view in mind, let’s consider how it is at odds with central claims made by Hart, Raz, and Dworkin about what “legal obligation” means. As was mentioned before, Hart defended an expressivist semantics for “legal obligations”: “Antigone is legally obligated not to bury Polynices” expresses approval of Antigone being punished if she buries Polynices, and we should understand the meaning of the sentence in terms of the expression of this non-cognitive attitude. This view is straightforwardly inconsistent with the orthodox semantics outlined above. (Though one could offer a version of Hart’s expressivism as a meta-semantic interpretation of the formal semantics above, as we’ll discuss in Chapter Five.) Moreover, Hart explicitly denied “the identity of meaning of ‘obligation’ in moral and legal contexts.”⁶⁷ This makes Hart’s view inconsistent with linguistic data. For instance, on this view it is not clear how we should interpret sentences like “Antigone had two obligations: one legal, one moral.” The view predicts that this sentence should read like “Thebes had two banks: one was a financial institution, and the other was a riverside.”

⁶⁶See von Fintel and Heim (2011), p. 36.

⁶⁷Hart (1982) went on to declare that a cognitivist theory was well suited to moral contexts but a “different, non-cognitive theory of duty” is “[f]ar better adapted to the legal case” (pp. 159–160).

By contrast, FICTIONALISM can be compatible with this orthodox semantics, and its commitment to the univocality of normative terms in moral and legal contexts. All that this semantics tells us is that salient features of the context (e.g., that we’re on a basketball court) or explicit lexical material (e.g., “legally”) contribute an ordering on the set of worlds in the modal base. The semantics says nothing about what makes a given ordering the appropriate ordering. It could be the ordering supplied by objective normative facts in moral contexts, and supplied by some salient fiction in legal contexts. To put this point in other terms, the ordering sources supply a *weighted ranking* of (propositions about) actions, like “Antigone buries Polynices”. One way to develop that idea is to take (distinct) ordering sources to consist of (distinct classes of) reasons; we need not be wedded to this, but it is helpful for expository purposes. The important point is that this semantics need not commit us to the *existence* of distinct classes of reasons. In moral contexts, it may be true that “*A* is obligated to ϕ ” because of how “*A* ϕ s” is ranked by the reasons that exist. In legal contexts it may be that “*A* is obligated to ϕ ” because of how “*A* ϕ s” is ranked by the reasons that exist according to the relevant fiction. This difference is perfectly consistent with FICTIONALISM, univocality, and this orthodox semantics.

What about Raz’s commitments with respect to the semantics of legal language? As we saw in the previous section, Raz takes “obligation” to have a distinctively moral meaning in legal contexts: “*A* is legally obligated to ϕ ” literally means “From the point of view of the law *A* is morally obligated to ϕ ”. The tension between Raz’s view and the orthodox semantics is subtle. The views are not strictly incompatible. On Raz’s view we can still take “legally” to contribute a distinct ordering. After all, we can take modifiers like “Kantian” in “Kantian obligation” to contribute a distinct ordering, even though one’s Kantian obligations just are one’s moral obligations according to the point of view of Kant (or Kantians). So where’s the tension?

In two places. First, this orthodox semantics shows that Raz’s view is poorly motivated. Raz and those who endorse or sympathize with his view try to motivate it by appealing to

the univocality of normative terms like “obligation”, which Hart denied. But as was already mentioned above, Raz’s view fails to deliver a unified approach to the meaning of “obligation” in legal and non-legal contexts, because there are no relevant officials (and certainly none making claims about moral obligations) with games and prescriptive grammar.⁶⁸ Moreover, Raz’s view is often motivated by an inference from the premise that law employs the same concept of “obligation” as morality to the conclusion that claims about legal obligations must be claims about moral obligations.⁶⁹ Indeed, the semantic views available to positivists are often framed to suggest that we either deny univocality or accept Raz’s view:

Some people think that there is a distinctively legal domain of normativity, separate from other normative domains like morality ... Other positivists, however, think that law employs the same concept of obligation as morality so ... claims about a person’s legal obligations are really claims about her moral obligations.⁷⁰

Once we have the distinction between the character and content of “obligation” in mind, we can see why this is doubly mistaken. The commitment to a distinctively legal domain of normativity does not commit one to Hart’s denial of the univocality of normative terms. It commits to there being a distinctively legal ordering source to which normative terms can be relativized. Likewise, the claim that law employs the same concept of obligation as morality does not commit one to the view that claims about legal obligations are really claims about moral obligations. This is like saying that all uses of “I” share the same meaning so all uses of “I” are really claims about *me* (the content of *my* uses of “I”). It ignores the plausible view that moral and legal language use the same term, “obligation”, which can be relativized to *standards* in the same way that “I” can be relativized to *speakers*. To put the same point in conceptual terms, moral and legal language can employ the same *generic* (rather than necessarily *moral*) normative concept of “obligation”.

⁶⁸Gardner (2012) seems to agree that this semantics for “legal obligation” is discontinuous with the semantics for uses of the term “obligations” in relation to etiquette and games.

⁶⁹Raz (2009), p. 158; see also his pp. 38–39.

⁷⁰Hershovitz (2015), pp. 1168–1169.

Second, Raz is committed to treating “legal” in the same way that we treat “Kantian”: as modifying “obligation” by taking us to a point of view about moral obligations. Linguists, by contrast, treat “legal” in the same way that they treat “moral”. If we are to depart from linguistic orthodoxy, it would be nice to see some linguistic data that favors Raz’s approach. He has provided none.⁷¹ And when we search for linguistic data that would discriminate between the two approaches, they militate against Raz’s view. Consider modifiers like “Kantian” once more. You might say, “The agent has a Kantian obligation not to lie”. This analogy isn’t perfect: these terms only work as adjectives, whereas “legally” also works as an adverb. But the analogy is close enough since it captures the basic idea behind Raz’s view: “Kantian obligations” are not a distinct domain of normativity from “moral obligations”; rather, one’s “Kantian obligations” are simply the moral obligations one has from the Kantian point of view. This allows us to test Raz’s semantics. “Legal” should pattern like “Kantian”. We can, for instance, say “She has a Kantian obligation not to lie” or “She has a Kantian *moral* obligation not to lie”. The meaning of both sentences is equally clear, and clearly equivalent. If “legal obligation” operates like “Kantian obligation”, we should similarly be able to say “Antigone has a legal obligation not to bury Polynices” or “Antigone has a legal *moral* obligation not to bury Polynices”. But here the latter sentence is very difficult to understand, and not obviously equivalent in meaning to the former. This is strong linguistic data against the view that “legal”, like “Kantian”, modifies “obligation” by taking us to a point of view about moral obligations. Rather, “legal” seems to function the same way as “moral”: by taking us to a distinct domain of normativity. Which is precisely how we should treat “legal” and “grammatical” and so on according to FICTIONALISM.

Now let’s turn to Dworkin’s commitments. This will take a little more work to set up. For most of his career Dworkin accepted the view that there were distinctively legal obligations which we create, and which may even be “too immoral to enforce.”⁷² But then

⁷¹The closest thing to linguistic data that Raz (2009) has offered are an analogy between legal language and “advice” that is “given from a point of view” (p. 156). But the latter phenomenon seems to be better explained in terms of another distinctive ordering source: practical rationality.

⁷²Dworkin (1986), p. 262.

he changed his mind, and declared that he should have accepted a view on which “‘law’ and ‘morals’ describe different collections of norms.”⁷³ He endorsed a metaphysical view according to which our legal obligations are a species of our moral obligations. Similar views have now been defended in more detail by Greenberg (2014) and Hershovitz (2015). These metaphysical views are inconsistent with FICTIONALISM. But they are also at odds with orthodox accounts of how we explain uses of normative terms in legal contexts, and I want to focus on this as a way of showing why we should favor FICTIONALISM over these revisionary views.⁷⁴

Dworkin himself says little about how to square his view with how we use legal language, but Hershovitz takes up that task, and the problems he faces are quite general. The initial problem for this view is that our talk of “legal obligations” looks just like our talk of “moral obligations”. Both suggest that we are talking about a distinct domain of normativity. Hershovitz suggests that instead we should treat talk of “legal obligations” like talk of “familial obligations” or “promissory obligations”: the adjective just signals “the source of the [moral] obligations in question.”⁷⁵

This is not inconsistent with the letter of the orthodox semantics outlined above, but it is in tension with how linguists treat “legally” (as well as other modifiers like “grammatically”—the account, Hershovitz makes clear, is intended to generalize to the other “law-like” standards that I am calling mere formalities). This is because if legal obligations are treated as a species of moral obligations we encounter terrible predictions about collections of claims like (1)–(3). On this view, such claims commit us to moral dilemmas: (1) and (2) entail that Antigone is morally obligated to bury and not bury Polynices. Worse, on this view such

⁷³Dworkin (2011), p. 402. Dworkin goes on to note that he had previously suggested this view in Dworkin (1977), but did not appreciate the significance of the view at the time.

⁷⁴Proponents of this view recognize that *prima facie* it has revisionary implications for how we use legal language: “we talk about our legal obligations; distinguish them from our moral obligations; and recognize that the two can conflict, sometimes in terrible ways. All that would seem to imply that our legal practices generate a distinctively legal domain of normativity,” unless “there are other explanations for our inclination to talk that way”, Hershovitz (2015), pp. 1186–1187.

⁷⁵Hershovitz (2015), p. 1187.

claims commit us to contradictions: (1) and (3) entail that Antigone is and is not morally obligated not to bury Polynices.

Hershovitz recognizes that there are such “hard cases” for his view where we regard people as having “a legal obligation to do something morally repugnant”, which “people could not possibly have a moral obligation to do”.⁷⁶ In such cases Hershovitz claims that we do not actually have a legal obligation, but “we often have reasons to be morally obtuse about our moral obligations”, so “we should sometimes regard ourselves as subject to legal obligations”.⁷⁷ The reasons he identifies have to do with judicial fidelity to the law. We speak as if legal officials always have legal (and hence moral) obligations to enforce repugnant laws because “we want them to approach their jobs with humility and deference.”⁷⁸

This response to these “hard cases” is seriously inadequate for three reasons. First, it has a very limited scope. Hershovitz only identifies reasons to speak falsely in public discourse about officials in the U.S. legal system. What about in private discourse? Shouldn’t my barber and I be able to see through the charade and truly claim that legal officials have no such legal obligation? What about discourse about citizens? My barber and I might well be quite convinced that denizens of the “dark ghetto” owe no such humility and deference to the state, but we would surely still think and say that they are subject to the same legal obligations as we are.⁷⁹ Or what about distant (foreign, historical, fictional) legal systems? If we were sufficiently arrogant, my barber and I might worry that our private speech could have some negative influence on the Ninth Circuit, but we’d have no such concerns about eroding the humility and deference among Nero’s magistrates or Kim Jong-Un’s Central Court, let alone among Creon’s officials as they are characterized in Sophocles’ *Antigone*. All of these concerns are heightened when we shift from talk to thought. Even if we have reason to speak falsely, why should I still falsely believe that we are subject to legal obligations to act in accordance with morally repugnant laws?

⁷⁶Hershovitz (2015), p. 1190.

⁷⁷Hershovitz (2015), p. 1192.

⁷⁸Hershovitz (2015), p. 1192. I am even skeptical of this, for reasons canvassed in Chapter Two.

⁷⁹See Shelby (2007) for a defense of this view.

Second, Hershovitz only identifies some reasons to think and speak falsely. From this he concludes that we “should” think and speak falsely. But of course there will be a long list of weighty reasons *not* to think or say, falsely, that people have legal and moral obligations to perform morally repugnant acts. The latter reasons would frequently outweigh the former, so it would often be unreasonable to think or say that we have legal obligations in these “hard cases”. Moreover, the former reasons are all *pragmatic* (concerning the effects of what is believed or asserted) and many of the latter are *alethic* (concerning the truth of what is believed or asserted); this makes it quite contentious to claim that the former ever or often outweigh the latter. Worse, some of our reasons not to speak falsely will concern incoherence. On Hershovitz’s view, when we say both (1) and (3) we contradict ourselves about whether or not Antigone has a moral obligation not to bury Polynices. We still need a further story about why we have sufficient reasons to assert not just falsehoods, but contradictions.

Perhaps both of these points are compatible with Hershovitz’s qualified claim that “we should *sometimes* regard ourselves as subject to legal obligations” that we do not have. But once we emphasize that qualification the jig is up. His view needs to explain why (almost) all of us agree that (almost) everyone is and has been subject to legal obligations to act in accordance with morally repugnant laws like, to use his example, the Fugitive Slave Act of 1850. If it falls shy of this mark, it remains a highly revisionary view about how we use normative language in legal contexts.

And the final problem with Hershovitz’s view is that even though it is intended as a general account about law and law-like practices, he says nothing about other “law-like” standards. What reason would we have to say that someone is grammatically required to (say) use “she” rather than “they” in relation to a gender-queer individual even when by our own lights no one is morally required to do so?

There’s more to be said about this account of legal obligations and legal language. Hershovitz also tries to motivate the view as a “kind of eliminativism” that is more “ontologically

spare” than the commonsensical view that there are distinctively legal obligations.⁸⁰ But the combination of the orthodox semantics we outlined above and FICTIONALISM show that these arguments fall flat. Considerations of ontological profligacy are best understood as concerning the number of *types* of entities we are committed to, rather than the number of *token* entities we are committed to. Once we are committed to ordering sources, it is not ontologically profligate to claim that there is an open-ended set of ordering sources, as this only introduces more tokens of the same type of entity. And FICTIONALISM shows that we can adopt this orthodox semantics without committing ourselves to any obscure or suspicious entities: we only need to appeal to something like the reasons that exist according to fictions.

4.5 Conclusion

Requirements of law, grammar, games and etiquette exist as a social reality. So we need some account of what they are, and how we create them. FICTIONALISM provides precisely that. In a game of make-believe we “must” capture bears iff and because according to the relevant fiction we have decisive reason to capture bears; and we can make it the case that there is such a fiction, and hence such a requirement, by accepting the contents of that fiction. This proposal provides a fairly ecumenical account of the metaphysics of legal obligations, and has the advantage over alternative views that it provides an orthodox and non-revisionary approach to legal language.

Before concluding, I want to note that this proposal was inspired by, and is indebted to, the form of fictionalism that was first endorsed by Jeremy Bentham (1782/1970) and has been neglected ever since.⁸¹ In large part this neglect is due to problems on Bentham’s

⁸⁰Hershovitz (2015), p. 1193.

⁸¹Bentham is frequently presented as having a sanction-based account of the meaning of “legal obligation”. For instance: “The Benthamite answer was to deny that [claims about legal obligations] are normative utterances. ‘X has a duty to ϕ ’ means ‘X is liable to a sanction if he does not ϕ ’” (Raz (2009), p. 154). Hart (1982)’s treatment is more sympathetic, and accurately notes that Bentham’s sanction-based theory was an intentionally revisionary replacement for our fictive (and arguably meaningless) discourse about legal obligations.

behalf. His view is difficult to disentangle from his radical empiricist commitments about semantics and metaphysics,⁸² and used for quite unusual theoretical ambitions. The central idea in Bentham's view was that our thought and talk about legal obligations are in some sense a fictional discourse. But Bentham had little interest in understanding the content of this fiction, beyond identifying some "emblematic or archetypical" imagery that it seems to involve. Instead, Bentham dismissed our fictional discourse as "syphilitic" and "pestilential", in order to argue that it could be replaced by something more scientific. Yet I hope to have shown how Bentham's central idea can be developed to offer a rich account of legal normativity.

⁸²For discussion, see Rosen (2005), Harrison (1983) and Quinn (2015).

Chapter 5

Expressivism and the Varieties of Normativity

5.1 Introduction

The expressivist advances a view about how we explain the meaning of a fragment of language, such as claims about what we morally ought to do. Critics principally evaluate expressivism on those terms. This is a serious mistake. We don't just use that fragment of language in isolation. We make claims about what we morally, legally, rationally and prudentially ought to do. To account for this linguistic phenomenon, the expressivist owes us an account not just of each fragment of language, but of how they weave together into a broader tapestry.

The linguistic phenomenon in question is the relativity of normative terms to standards. In the same circumstances, one act can count as what we “ought to do” relative to one standard (e.g., morality) and what we “ought not do” relative to another (e.g., legality). We see the same phenomenon with descriptive terms: in the same circumstances, one entity can count as “tall” relative to one standard and “not tall” relative to another. Such claims are consistent, and have different contents. However this does not show that “ought” (or “tall”) is lexically ambiguous. Rather, “ought”, like “tall”, has a single *standing meaning*, or *character*, which systematically determines those different contents.

Historically, the expressivist neglected this phenomenon. She offered an explanation of the contents of terms like “ought” in one fragment of language. She provided no explanation of the character of normative terms like “ought”.

In §5.3, I explain why the expressivist must explain the character of “ought” across all varieties of normativity—that is, across all standards to which “ought” is relativized. Such an explanation is required for two reasons. First, if “ought” had different contents but no single character, it would have a different meaning in moral, legal, rational and prudential claims. There is no evidence for this prediction of widespread lexical ambiguity, and much evidence against it. So the expressivist needs an account of the character of “ought” if it is to be consistent with the linguistic data. Second, without such an explanation, expressivism will lack the parsimony and explanatory power of its “descriptivist” competitors.

As we shall see, there are two ways that the expressivist could explain the character of “ought”: at the semantic level, or at the meta-semantic level. In §5.4 I consider whether the semantic expressivist can explain the character of “ought”. The best option here builds on work by Allan Gibbard (1990, 2003). However, this view gets the wrong results about inconsistency: it makes claims about what we morally ought to do inconsistent with claims about what we legally, prudentially or rationally ought to do. This problem with inconsistency, I argue, also afflicts any plausible alternative expressivist semantics.

In §5.5, I consider whether the meta-semantic expressivist can explain the character of “ought”. The best option here was developed by Michael Ridge (2014). But it leaves the expressivist with an implausibly disunified meta-semantics: there are two different explanations for why “ought” has one meaning. There is no good way for meta-semantic expressivists to avoid this problematic disunity.

If this is right, the expressivist faces a serious challenge. She must provide an explanation of the character of “ought”, but has no good way to do so. My hypothesis is that to explain the univocity of “ought”, the expressivist must posit that “ought” expresses the same type of conative attitude when it is relativized to varieties of normativity; but once she posits that, no matter how we fill in the details of her view we encounter a general problem in explaining inconsistency because attitudinal inconsistency is too coarse-grained. I use this hypothesis to frame the present challenge to expressivism as a dilemma in §5.6. In §5.7 I

note that by construing her view as a pragmatic thesis the expressivist can easily evade both horns of the dilemma; but this is at best a Pyrrhic victory, since it does not vindicate any anti-realist ambitions about normativity.

Before concluding, I sketch two additional difficulties that the expressivist encounters in offering a generalizable view about the character of modals like “ought” (§5.8): one involves embedding; the other involves non-deontic modals.

As Seth Yalcin (2012) noted, expressivism has been “so unmoored from linguistic phenomena [that] it is unsurprising that there is little agreement concerning what the expressivist’s linguistic programme is, or should be.” This can make expressivism an elusive target. Given this, my aim here is to draw our attention to a linguistic phenomenon that the expressivist has neglected in order to throw down a gauntlet, not to offer an impossibility proof. Perhaps some form of expressivism avoids both horns of the dilemma while vindicating anti-realism. If so, progress will still have been made: at very least, we will have reached some agreement about what the expressivist’s linguistic programme should be.

5.2 Setting the Stage

We should not proceed unmoored. So what is this linguistic phenomenon? And what are “content” and “character”, “descriptivism” and “expressivism”?

5.2.1 The Linguistic Phenomenon

Consider Sophocles’ *Antigone*. After her brother Polynices was killed in battle, Creon—the king of Thebes—forbids Antigone to bury him, on pain of death by stoning. Say that Antigone and Creon discuss what Antigone ought to do:

1. Morally, I ought to bury Polynices, but legally I ought not do so.
2. No, morally and legally, you ought not bury Polynices.

In saying (1), Antigone takes morality and legality to be *divergent standards*: the same act in the same circumstances (“burying Polynices”) is what she “ought to do” relative to one standard (morality) and what she “ought not do” relative to another (legality). In saying (2), Creon takes morality and legality to here be *convergent standards*: according to both, Antigone ought not bury Polynices.

This exchange will be our main focus. But it is synecdochic: we are using this example to represent the wider linguistic phenomenon of standard-relativity. We could also substitute legality for another standard, like prudence:

3. Morally, I ought to bury Polynices, but prudentially I ought not do so.

4. No, morally and prudentially, you ought not bury Polynices.

And we could *embed* such relativized claims, as Creon does in the following sentences (after the prophet Tiresias convinces him of the error of his ways):

5. Morally you ought not bury Polynices; but if you bury Polynices, rationally I ought not blame you.

6. Legally, I ought to punish you. But morally, I should not have made it the case that legally I ought to punish you.

It is worth noting a few intuitive and important features of such exchanges. Antigone and Creon disagree about what she morally ought to do, but agree about what she legally and prudentially ought to do. Moreover, they do not disagree with themselves: it is consistent for Antigone to say that “burying Polynices” is what she “ought to do” relative to morality, *and* what she “ought not do” relative to legality and prudence. Ditto for Creon’s claims (5) and (6).

5.2.2 Content and Character

Now let’s distinguish between *content* and what Kaplan (1989) called *character*. These technical terms track an intuitive distinction between context-variant and -invariant meaning.

Theories of meaning provide different explanations of that distinction,¹ as well as broader and narrower interpretations of “context”. All that I mean by the terms is the following intuitive and ecumenical distinction.

Say that James, who is 5', is taking part in a school photo. His classmates are all around 4'5". The teacher says “James, you’re tall: go stand at the back.” Later, James joins his new basketball team. His teammates are all around 5'5". The coach says “Sorry James, you’re not tall: you’ll play guard.” Do these uses of “tall” mean the same thing? One intuitive answer is “No”. The teacher’s use means “of a height is greater than 4'5".” The coach’s use means “of a height greater than 5'5".” This is the sense in which these uses of “tall” have different *contents*. Another intuitive answer is “Yes”: in both uses, “tall” means roughly “of a degree of height greater than the average member of x ”.² This is the *character* of “tall”; in a context, it determines the variable contents of “tall”.

Here “tall” is used in two different circumstances. But the same intuitive distinction can be made for different uses of “tall” in the same circumstance:

7. Creon is tall for a Theban, but he’s not tall for a Greek.

8. No, he’s tall for a Theban *and* for a Greek.

Here the same entity (“Creon”) is taken to be “tall” or “not tall” relative to different standards (“for a Theban”; “for a Greek”). There remains an intuitive sense in which “tall” means the same thing in (7) and (8). Since what changes are the standards, not the circumstances (which is what some mean by “context”), we can take the *character* of “tall” to be its “standard-invariant meaning”.

The standard-relativity of the descriptive term “tall”, in (7)–(8), is analogous to the standard-relativity of the normative term “ought” in (1)–(2). So by analogy, our concern is with the distinction between the content (standard-variant meaning) and character (standard-invariant meaning) of “ought”.

¹For instance, compare Kaplan (1989) and Heck (2001); and see generally Speaks (2014).

²See Barker (2002), p. 2, for a less rough account of the character of “tall”.

The challenge for expressivists will be to explain the standard-invariant meaning of “ought” (or “ A ought to ϕ ”), and how it systematically interact with standards like “morally” and “legally” to determine contents.

5.2.3 Descriptivism

But let’s not jump ahead. Let’s first consider the most orthodox account of the character of “tall” from representationalist, truth-conditional semantic theories. This account appeals to functions: “tall” is taken to mean some function from contexts to contents, where contents are understood to be propositions.

“Descriptivists” offer the same representational explanation for the meaning of descriptive and normative terms alike. So they can also take the character of “ought” to be a function from contexts to propositions. This approach is well-developed, thanks to David Lewis (1973) and Angelika Kratzer (1977, 1981, 1991), among others. To illustrate it, here is one candidate function (which is familiar from Chapter Four):

$$\mathbf{Ought}_{fg}\phi = 1 \text{ iff } \forall w((w \in [[F]] \cap [[G]]) \supset w \in [[\phi]]).$$

The modal base f provides a set of propositions, F , which represents the relevant circumstances. The ordering source g provides a set of propositions, G , which represents the demands of the relevant standard. “ A ought to ϕ ” is true iff “ $A \phi$ ” is true at every world in the intersection of F and G .³

This account of the character of “ought” has many virtues, two of which are dialectically important. First, it explains how the character of “ought” systematically interacts with standards to determine contents. Second, it provides fine-grained inconsistency conditions. (1) and (2) are inconsistent: “Antigone buries Polynices” cannot be true and false at the intersection of the F and G where G represents the same legal demands and F represents the

³The function above is far too simple, but it neatly illustrates the essential innovation from Kratzer: that *two* conversational backgrounds are relevant to each modal. The modal base determines which worlds are accessible, while the ordering source ranks those worlds.

same circumstances. But the conjuncts of (1) are consistent: it can be true that “Antigone buries Polynices” at the intersection of the F and G where G represents the moral demands, and true that “Antigone does not bury Polynices” at the intersection of the F and G where G represents the legal demands. This explains why Antigone and Creon disagree, but Antigone does not disagree with herself.

Of course, there are problems with and puzzles for the function above, and truth-conditional accounts of the character of “ought” generally. Still, such accounts are promising. And perhaps more importantly, even those who reject them should still accept that descriptivists face no special challenge in explaining the character of “ought”. After all, there will be some account of the character of standard-relative descriptive terms, like “tall”. If the descriptivist is right, whatever machinery works for “tall” can also be wheeled out for “ought”.

5.2.4 Expressivism

The expressivist rejects descriptivism. She claims that there is a deep difference between how we should explain the meaning of descriptive terms (like “tall”) and normative terms (like “ought”). Three motivations are offered for this.

The first concerns representation. If (7) and (8) are representational, they commit us to objects instantiating properties like *being tall*. If (1) to (6) are representational, they commit us to objects instantiating properties like *being what one ought to do*. Unlike *being tall*, such properties do not seem to be part of the familiar physical world, so many worry that they would be “of a very strange sort, utterly different from anything else in the universe.”⁴

The second concerns disagreement. If speakers systematically differ in their criteria for applying descriptive terms like “tall”, they don’t really disagree when they assert (7) and (8): they merely talk past each other. Not so if Antigone and Creon systematically differ in their criteria for applying the words “morally ought”. If Antigone is a divine command

⁴Mackie (1977), p. 38.

theorist and Creon is a virtue ethicist, they still really disagree about what she morally ought to do.⁵

The third concerns “internalism”. Roughly, this is the view that there is an internal connection between sincere uses of “ought” and speakers’ conative attitudes.⁶ If a completely dispassionate and indifferent person said “Creon is tall”, we would have no grounds to doubt her sincerity, but not so if she said “Antigone ought to bury Polynices”. Absent some indication of her conative attitudes—that she is *for* Antigone burying Polynices—she would seem to be parroting the normative claims of others. This suggests that speakers’ conative attitudes play importantly different roles in normative and descriptive claims.

Enter the expressivist. She proposes that we “explain the meaning of a term” by “what states of mind the term can be used to express” in sincere speech acts.⁷ Further, she takes descriptive terms like “tall” to express cognitive states, and normative terms like “ought” to express conative states (like plans or approval).

The expressivist aims to vindicate anti-realism. She is “primarily motivated” by the thought that by explaining the meaning of normative but not descriptive terms via conative states we can avoid “unwanted ontological commitments.”⁸ Desires and beliefs, after all, don’t carry the same ontological commitments.⁹

The expressivist also aims to explain disagreement in terms of inconsistent conative attitudes. Antigone and Creon disagree because they express rationally discordant attitudes. The details of the explanation differ, but the form is like the following. Antigone expresses a plan to bury Polynices and Creon expresses a plan for Antigone not to bury Polynices. Since these plans are rationally discordant—a speaker who possessed both would be practically

⁵See especially Horgan and Timmons (1991).

⁶What “internal” connection? This is a vexed question. See Smith (1994) and Ridge (2014).

⁷Gibbard (2003), p. 7; see also Ridge (2014), p. 5. As Bar-On and Sias (2013) note, many take this pattern of explanation “to be more-or-less definitional of expressivism” (p. 702).

⁸Bar-On and Sias (2013), p. 701. See Gibbard (2012) pp. 236–240; Ridge (2014), p. 101. Expressivist views have been offered for some non-normative terms and claims—about, *inter alia*, the *apriori* (Field (2000)), probability (Logue (1995)) and epistemic modals (Yalcin (2012))—typically with the view that such terms and claims are also non-representational.

⁹Beyond this intuitive gloss, I will not attempt to solve the difficult problem of how we should distinguish realism from anti- (and quasi-)realism. For discussion, see Dreier (2004).

irrational—they are inconsistent, and so too are the sentences that express such attitudes.¹⁰ Such attitudinal inconsistency can persist in the face of systematic differences in speakers’ criteria for applying words like “morally ought”. So expressivists claim to explain why virtue ethicists and divine command theorists disagree.

Finally, the expressivist aims to explain internalism. The internal connection between sincere uses of normative terms and speakers’ conative attitudes is explained by taking those terms to *express* speakers’ conative attitudes.

For each of these motivations, the descriptivist has (at least) one response.¹¹ With one exception (see §5.7), I put these to one side. The aim of this paper is not to defend descriptivism, but to explore a novel challenge to expressivism.

5.2.5 Varieties of Normativity

Expressivism is typically advanced as a theory about the meaning of normative language: the expressivist offers an explanation of the meaning of claims about what we “ought to do”, or what we “morally ought to do”. But this ignores the varieties of normativity. Antigone’s and Creon’s claims about what they “ought to do” are relativized to moral, legal, prudential, and rational standards.

In (1) to (6) “ought” is explicitly relativized to these standards. Speakers also implicitly relativize “ought”, and other deontic modals, to standards that are made salient by the conversational context: “You must reply in the third person to an invitation issued in the third person” and “You can’t double dribble!” implicitly relativize deontic modals to standards set by etiquette and basketball. These cases also illustrate that the scale of varieties of normativity goes well beyond morality, legality, prudence and rationality: indeed, there may well be an “open-ended set” of standards to which “ought” can be relativized.¹²

¹⁰See Schroeder (2008) and Woods and Baker (2015) for excellent discussion of this issue.

¹¹Considering just the first motivation, descriptivists can be error theorists, naturalists, non-naturalists or quietists about the ontological commitments of normative language.

¹²See von Fintel and Heim (2011), p. 36.

This phenomenon has been noted before. It is familiar to linguists, who call varieties of normativity “flavors” to which deontic modals are relativized. And it is familiar to philosophers. Michael Zimmerman (1996) wrote that “ought” is often used to “express moral obligation”, as well as “to express what is required not by morality but by prudence, or law, or aesthetics, or the rules of chess, or whatever.” Likewise, Stephen Finlay (2014) says “the orthodox view” is that “the moral sense of ‘ought’ is only one of many, including prudential and instrumental senses, and those of institutions like legal systems, etiquette, and games.”

Historically, this phenomenon has been largely ignored by the expressivist. She has focussed on one (or at most, a narrow subset) of the “senses” of “ought”,¹³ and offered a theory that explains the *content* of that sense of “ought”. She offered no account of the character of “ought”. Nor can such an account easily be found by patching together her accounts of the contents of some particular senses of “ought.” This assessment is not unorthodox. While Jamie Dreier (2009) writes that expressivists “decline to say what” normative terms or claims mean by speaking of their “content, or character”, we “can think of” expressivism as taking (sets of) conative states to be “normative *contents*”; that is, we can think of (sets of) attitudes as the semantic “contents of sentences.”¹⁴ Dreier even contrasts expressivism with relativism, which assigns normative claims “a Kaplanian character”, or “a function from contexts to contents.” Similarly, Mark Schroeder (2008)’s book-length evaluation of expressivism noted that the view is yet to introduce “context-sensitive elements” to yield “a picture of the role of context ... in which individual terms have Kaplanian characters”.¹⁵

5.2.6 Taking Stock

So far, we’ve seen that there is a linguistic phenomenon of standard-relativity, which occurs with descriptive and normative terms alike. One way to explain the meaning of a standard-

¹³Sticking to our central case, the two most prominent expressivists—Blackburn (1984, 1993, 1999a) and Gibbard (1990, 2003)—defend expressivism primarily about the moral “ought” (or “moral rightness”, etc.), and ignore the legal “ought”. Kevin Toh (2005) does the opposite.

¹⁴Dreier (2009), pp. 83, 85 (emphasis his). Why the “sets of” qualification? See §5.4 below.

¹⁵Schroeder (2008), pp. 179–180.

relative term is to provide an account of its character, which determines its contents. The descriptivist can use the same machinery in explaining the characters of descriptive and normative terms. The expressivist, by contrast, has not provided an account of the characters of “ought” et al.

5.3 Expressivism *Sans* Character

These observations do not amount to a challenge to the expressivist. We have not yet seen why she needs to provide an account of the character of “ought.” To do so, let’s continue pressing the analogy between “ought” and “tall”. Does a first-order semantic theory need an account of the character of “tall”?

Say the meaning of any use of “tall” is its semantic content. The contents of uses of “tall” vary depending on the standards to which “tall” is relativized: “tall” (for a Theban) has a different content to “tall” (for a Greek). So different uses of “tall” will differ in meaning: “tall” will be lexically ambiguous.

A semantic theory with this structure would encounter two serious, interrelated problems. First, it would lack the theoretical virtues of parsimony and explanatory power. It would posit a vast number of distinct meanings of “tall”. And it would owe us an explanation of what these different meanings have in common, if they are not systematically determined by the character of “tall”.¹⁶

Second, it would fail to explain the linguistic data. Predictions of lexical ambiguity are testable, and this one fails the relevant tests. Consider the Conjunction Reduction Test, wherein two sentences that contain a purportedly ambiguous term are conjoined, using the term once in contexts where both meanings are encouraged. To illustrate: “Pale colors are light”; “Feathers are light”; “Pale colors and feathers are light.” The third sentence fails the test because “light” seems *zeugmatic*: it has been used with two other parts of the sentence (“colors”, “feathers”) and must be understood differently in relation to each. This

¹⁶See Peter Ludlow (2008) for an ambitious attempt to discharge this explanatory debt.

is evidence for thinking that “light” is lexically ambiguous.¹⁷ By contrast, consider (8): “Creon is tall for a Theban *and* for a Greek”. If “tall” were lexically ambiguous, this should seem zeugmatic. But it doesn’t. These linguistic data are not explained by a theory that takes “tall” to be ambiguous.

Would a first-order semantic theory that provides no account of the character of “ought” face these same problems? Absolutely. Say the meaning of any use of “ought” is its semantic content. The contents of different uses of “ought” vary depending on the standards to which “ought” is relativized. So different uses of “ought” will also differ in meaning: “ought” will be lexically ambiguous.¹⁸

If the expressivist does not provide an account of the character of “ought”, her theory faces two serious, inter-related problems. First, it would lack the virtues of parsimony and explanatory power. It would posit widespread lexical ambiguity, perhaps even attributing to “ought” an open-ended set of meanings (to correspond to the open-ended set of varieties of normativity). This violates the maxim that we should not multiply meanings beyond necessity. Moreover, with no account of the stable character that determines these varying contents, it would not explain what these different contents have in common.¹⁹

Second, the expressivist will not explain the linguistic data. When we have *convergent standards*, we can take sentences in which “ought” is relativized to different standards (“Morally, Antigone ought not bury Polynices” and “Legally, Antigone ought not bury Polynices”) and conjoin them using “ought” once (“Morally and legally, Antigone ought not bury Polynices”). If “ought” were lexically ambiguous, such sentences would seem zeugmatic, as “light” does in “Pale blue and feathers are light”. That “ought” does not seem zeugmatic

¹⁷Sennet (2011). Linguistic evidence from this test—as well as the “and so” and “binding” tests used below in the example of Zoe—is rarely decisive. It can be hard to distinguish between ambiguity and nearby phenomena, like semantic indeterminacy. Further issues are posed when speakers produce zeugma for poetic effect: in “She took her hat, and her leave”, “took” is used literally in relation to “her hat” and figuratively in relation to “her leave”.

¹⁸See Chrisman (2007) pp. 308–9 for helpful discussion of “ought” and lexical ambiguity.

¹⁹For more on what explanatory power is needed, see Swanson (2007) p. 1195 and Finlay (2014) pp. 55 ff. The expressivist has offered no attempt to discharge this explanatory debt. It is also worth noting that known attempts to do so—such as Ludlow (2008)’s—would undercut the disagreement-based motivation for expressivism. See Plunkett and Sundell (2013).

in (2) and (4) and other like cases strongly suggests that it is not lexically ambiguous. This is a straightforward application of the Conjunction Reduction Test.

We get the same results if we apply other standard tests for lexical ambiguity to “ought”, and other deontic modals, that are relativized to different standards. Consider the following case. Zoe has three kids: Joe, Kim and Leo. Joe is in the Evil Villains Club, where it is a rule that members are silent when someone is presenting a scheme, as is happening now. Kim is in a legally designated noise-reduction zone. Leo is near a dangerous predator; if he makes a peep, it’ll kill him and some kids. Now we can consider sentences like “Joe must be quiet, and so too for Kim and Leo”, or “All of Zoe’s kids must be quiet”. In these sentences, “must” is implicitly relativized to different standards. But it does not strike us as zeugmatic, which strongly suggests that it is not ambiguous.²⁰

Perhaps these arguments are not decisive. But at very least they show that the onus is on the expressivist to show that her explanation of the meaning(s) (contents) of “ought” is plausible in the absence of an account of its character. It is difficult to see how the expressivist could discharge this explanatory debt, as a disunified semantics is certainly not *favoured* by the linguistic evidence.

Now we have a challenge for expressivists. However, we still have not seen anything that suggests that it is insurmountable. Before we turn to the complex and interesting strategies that the expressivist can employ in attempting to provide an account for the character of “ought”, it will prove helpful to consider an analogy that exposes the general difficulties she will encounter.

Say I defended an expressivist view about “tall”, taking “ x is tall” to express the speaker’s approval of x ’s height, and explaining the meaning of “tall” in terms of the expression of this conative attitude. This yields a view about the content of claims like “Creon is tall”. But “tall” is a standard-relative term. You should ask me: What is the *character* of “tall”, and what does “... for a Theban” or “... for a Greek” add to the meaning of “Creon is

²⁰Zwicky and Sadock (1975) and Stanley (2005) for discussion of these tests for ambiguity.

tall”? My answer would be subject to several constraints. First, it must be generalizable to the open-ended set of standards to which “tall” can be relativized (“... for a human”, “... for an animal”, and so on). Second, it better not deny that “tall” is univocal when it is relativized to different standards. After all, when standards converge, we can engage in conjunction reduction, as in (8): “Creon is tall for a Theban and for a Greek”. And third, it better explain why “Creon is tall” and “Creon is not tall” are inconsistent when relative to the same standard (“for a Greek”), but consistent when relative to different standards, as in cases of divergent standards like (7): “Creon is tall for a Theban, but not for a Greek”.

These last two constraints would prove difficult to satisfy at the same time. “Creon is tall for a Greek” and “Creon is not tall for a Greek” would have to express discordant attitudes (such as approval and disapproval of Creon’s height) if my view were to explain their inconsistency. “Creon is tall for a Greek” and “Creon is tall for a Theban” would have to express the same conative attitude (approval of Creon’s height) if my view were to explain the univocality of “tall”. But “Creon is tall for a Theban” and “Creon is not tall for a Greek” cannot express inconsistent attitudes lest my view have terrible predictions about (7).

Expressivists about standard-relative normative terms like “ought” face the same question: What is the character of “ought”, and what does “Morally...” or “Legally...” add to the meaning of “Antigone ought to bury Polynices”? Their answers are subject to the same constraints: they must be generalizable; they must explain the univocality of “ought”; and they must explain inconsistency. It will prove difficult to satisfy these last two constraints at the same time.

All forms of expressivism face this challenge. There are two main ways for the expressivist to meet it. She can explain the character of “ought” at the semantic level, or at the meta-semantic level.²¹ Let’s consider these in turn.

²¹For the sake of expository convenience, I’ll leave aside other issues on which forms of expressivism differ. For instance, some take expressivism to explain the meaning of *terms*. Others take expressivism to explain the meaning of *sentences*. This issue is orthogonal to our purposes: all references to the meaning of terms (“ought”) can be translated into the meaning of sentences (“A ought to ϕ ”). Indeed, there had better be such

5.4 Semantic Expressivism

Semantic expressivists answer the first-order question “What does “ought” mean?” They take “ought” to mean the attitude it expresses when uttered sincerely.²² This is a *rival* to the descriptivist view of the meaning of “ought” from §5.2.

So how can the semantic expressivist explain the character of “ought”?²³ I’ll first argue for a minimal condition that any such explanation must satisfy:

Unified Expressivism “Ought” always expresses conative attitude X .²⁴

We can see why this is a plausible minimum condition in two ways. First, consider what an expressivist view would look like without it. Say “ought” expresses a plan in “Morally, Antigone ought not bury Polynices” and a belief in “Legally, Antigone ought not bury Polynices”. If “ought” expresses two different types of mental states and its meaning is identical to the state expressed, it is ambiguous. This is the result that the expressivist must avoid. So when it is relativized to different standards, “ought” must express the same attitude type.

Second, consider how this minimum condition can do explanatory work. The character of “ought” can be played by this attitude *type*, and the contents of “ought” can be played by *tokens* of this type. Take a simple example of conjunction reduction: “Antigone ought not bury Polynices, or deceive Creon”. We can take “ought” to express one type of attitude—such as *being for*—which systematically determines the contents of the conjuncts: being for Antigone

a translation schema, as it is well-known that expressivists must explain the “compositionality of meaning”, which “simply holds that the meaning of a given sentence is a function from the meaning of its parts and their arrangement” (Ridge (2014), p. 137). It is also worth noting Schroeder (2008)’s point that it is much easier to assign Kaplanian characters to terms than to sentences (pp. 179–80).

²²This is roughly Rosen (1998)’s characterization of expressivism: it offers “a mapping from statements ... to the mental states they ‘express’ when uttered sincerely”, p. 387.

²³A semi-technical aside. Expressivists with more technical bearings may wish to claim that “ A ought to ϕ ” means a function from agent–act pairs ($\langle A, \phi \rangle$) to attitudes. I am fairly confident that we can translate all of the following in these terms—for instance, “ought” always expresses X iff this function is always to the same type of attitude, X —but I have avoided doing so in the prose for the sake of accessibility and expository convenience.

²⁴In fact, all we need for present purposes is the weaker negative thesis that “ought” does not express different types of attitudes when it is relativized to different varieties of normativity.

not burying Polynices and being for Antigone not deceiving Creon. And we can offer this same in cases involving standard-relativity, like (2): “Morally and legally, Antigone ought not bury Polynices”. If “ought” expresses and means one type of mental state (character), this type can systematically determine the different attitude tokens (contents) expressed by the conjoined claims “Morally, Antigone ought not bury Polynices” and “Legally, Antigone ought not bury Polynices”. This is a plausible explanation for why “ought” is not zeugmatic in (2).

But it is not a full explanation of the character of “ought”, because it does not explain varieties of normativity. How does this type of mental state interact with the standards to which “ought” is relativized to produce token mental states? What do “morally” and “legally” add to the meaning of (2)?

There are two promising places to look for answer. The semantic expressivist could claim that varieties of normativity modify the objects of this type of attitude, or modify other features of the attitude itself. Let’s take these in turn.

5.4.1 Objects of Hyperplans

The most promising version of the first strategy comes from Allan Gibbard. Gibbard (2003) appeals to the notion of a *hyperplan*, a plan which

covers any occasion for choice one might conceivably be in, and for each alternative open on such an occasion, to adopt the plan involves either rejecting the alternative or rejecting it. In other words, the plan either forbids an alternative or permits it.²⁵

The meaning of normative terms is explained via the expression of hyperplans.²⁶

²⁵Gibbard (2003), p. 56. A different Gibbardian strategy that I don’t consider (since I don’t know how to develop it) would appeal to *simulating* conative states. See Gibbard (2015).

²⁶More accurately, since “real people are undecided about what to do in most possible situations (just as we are undecided about most details of what is the case, which world is ours), our planning states can be represented by sets of hyperplans”, Dreier (2009), p. 84.

There are three reasons why Gibbardian expressivism is a particularly promising place to look to develop an expressivist explanation of the character of “ought”. First, hyperplans have structural features that are similar to the possible worlds employed in the truth-conditional semantics we considered in §5.2. As Yalcin writes, hyperplans can be “formally analogous” to functions.²⁷

Second, Yalcin has argued that this framework can explain a closely related linguistic phenomenon—the relativization of “ought” to bodies of information—without requiring an additional “distinct lexical entry for ‘ought’.”²⁸

Yalcin does not consider how the framework can address the phenomenon with which we are concerned: standard-relativity. But Gibbard himself discusses how “ought” claims that are relative to *moral* and *rational* standards could express and mean the same type of mental state. This is the final reason why Gibbardian expressivism is a particularly promising place to start.

Gibbard (1990) proposed that rational and moral “ought” claims express the same type of mental state with different *objects*.²⁹ “Rationally, *A* ought to ϕ ” expresses the speaker’s hyperplan for *A* to ϕ .³⁰ (More technically: the object of the attitude is the prejacent.) “Morally, *A* ought to ϕ ” expresses the speaker’s hyperplan to blame *A* if *A* does not ϕ . This only addresses two members of the open-ended set of varieties of normativity, but the strategy could be generalized by identifying other act-types which form the objects of conditional hyperplans. For example, “Legally, *A* ought to ϕ ” might express the speaker’s

²⁷Yalcin (2012), p. 148.

²⁸Yalcin (2012), p. 154. The relativization of words like “ought” to standards and to bodies of information are often classed together under the umbrella “context-sensitivity”.

²⁹Like Wedgwood (2010) (p. 121), I take it that these objects can be actions or attitudes. I take it that similar objections arise either way, but I focus on actions in the discussion below.

³⁰For the sake of explanatory convenience, I am eliding over several exegetical issues. First, as mentioned above, speakers express *sets* of hyperplans. Second, *hyperplans* are not employed until Gibbard (2003). Third, unlike Gibbard (1990), Gibbard (2003)’s view is about the technical term “the thing to do”, rather the rational “ought”: i.e., “ ϕ -ing is the thing to do” expresses the speaker’s hyperplan to ϕ . None of this affects the objections that follow.

hyperplan for A to be punished if A does not ϕ .³¹ Similar moves could be made when “ought” is relative to prudence, etiquette, the rules of games, and so on and so forth.

While it is initially promising, this strategy generates the wrong results about inconsistency. It makes claims about what we morally ought to do inconsistent with claims about what we legally, prudentially or rationally ought to do.

Consider how this view explains the inconsistency of (1) and (2): “Morally, Antigone ought not bury Polynices” and “Morally, Antigone ought to bury Polynices” express hyperplans for Antigone to be blamed if she does bury Polynices and for Antigone to be blamed if she does not bury Polynices. These plans are discordant, so the sentences are inconsistent: Creon and Antigone disagree.

But what about (1) itself? “Morally, Antigone ought to bury Polynices, but legally she ought not do so” expresses hyperplans for Antigone to be blamed if she does bury Polynices and for Antigone to be punished if she does not bury Polynices. But on many respectable views, being punished constitutively involves being blamed.³² If this is so, these hyperplans are discordant,³³ the conjuncts are inconsistent, and Antigone disagrees with herself. She is irrational.

This is a highly implausible result. How can the expressivist respond?

She could bite the bullet. Perhaps “Legally Antigone ought to bury Polynices” entails that “Morally, Antigone ought to bury Polynices”, so the conjuncts of (1) are inconsistent. This response tethers expressivism to a controversial form of legal anti-positivism.³⁴ If it is analytic that punishing Antigone involves blaming Antigone, then this form of legal anti-positivism will be an analytic truth, which would be surprising to positivists and anti-

³¹A broadly similar view about “ A is legally obligated to ϕ ” was suggested by Hart (1982).

³²This view was popularized by Feinberg (1970). See also Duff (2003, 2013).

³³That may be too quick. Perhaps “legally ought” expresses (passive) plans for someone *to be* punished (and blamed), which are not discordant with (active) plans *to* blame someone that are expressed by “morally ought”. (See Harman (1973) on passive “ought to be” and active “ought to” claims.) But this should not detain us. I take it that “legally ought” would express plans to be blamed by the state, and so by the relevant legal officials. So any appeal to a passive-active distinction would not help in analogous cases where the speaker is Creon.

³⁴Of the kind defended by Greenberg (2014). Dworkin (1986) accepts no such entailment.

positivists alike. Moreover, (1) and (2) are just one example of a general phenomenon. We can substitute legality for other standards, like etiquette; the same issue will arise, and this anti-positivist response will become increasingly implausible.³⁵

She could claim that it is consistent and rationally compatible for Antigone to express hyperplans to be blamed for burying and for not burying Polynices. There's nothing irrational about planning to be blamed, whatever one does. Perhaps this is plausible. But if so, the expressivist loses her explanation for why (1) and (2) are inconsistent, and why Antigone and Creon disagree. Moreover, we can construct variants of the case where this move will be unavailable.³⁶

The expressivist could appeal to more specific objects of our attitudes that are constitutively independent. However, this will amplify other concerns. It will make it harder for the account to generalize. The expressivist will have to identify a highly specific distinct object of our hyperplans for each of the many varieties of normativity. That is a daunting task. Especially since we must *have* attitudes with such highly specific objects, and conventionally express them when we make sincere claims invoking the relevant standard. It is already doubtful that the pacifists at my yoga studio regularly plan to blame and punish, but I am sure that they nonetheless make sincere claims about what we morally and legally ought to do. If the expressivist appeals to attitudes with even more specific objects, it will become increasingly doubtful that anyone has them, let alone that everyone shares them, and conventionally expresses them.

³⁵To illustrate, consider "According to etiquette, A ought to ϕ ." The expressivist could take it to express the speaker's conditional plan to socially distance herself from A if A does not ϕ . But on Scanlon (2008)'s view, blame constitutively involves social distancing. This will suffice to generate inconsistencies between "ought" claims that are relative to morality and etiquette. Biting the bullet by accepting anti-positivism about etiquette is deeply counter-intuitive.

³⁶Depending on one's stance on moral dilemmas, "I morally ought to ϕ and to not ϕ " may be consistent. But this should not detain us. We can make the same objection with "I am *not* morally obligated to ϕ , but I am morally obligated to ϕ ", and "I am *not* morally obligated to ϕ , but I am legally obligated to ϕ ." The former sentence is definitely inconsistent, so for Gibbardian expressivists it must express a practically irrational pair of attitudes towards my being blamed if I don't ϕ . But if punishment constitutively involves blame, then the latter sentence also expresses a practically irrational pair of attitudes towards that same object.

Moreover, no matter which distinct objects of our attitudes the expressivist appeals to, she is guaranteed to get the wrong results about inconsistency. To see why, recall (5): Creon's claim that "Morally, Antigone ought not bury Polynices; but if she buries Polynices, I rationally ought not blame her". This is obviously consistent. Does the Gibbardian view get this result? The first conjunct expresses Creon's hyperplan to blame Antigone if Antigone buries Polynices. The second conjunct expresses Creon's hyperplan not to blame Antigone if she buries Polynices. (Recall: the object of the hyperplans expressed by rational "ought" claims is the prejacent.) So Creon expresses rationally incompatible attitudes.

Can the expressivist tinker with the details of her view to avoid this result? I have no proof that she can't, but I'm skeptical. The Gibbardian strategy is to explain the character of "ought" in terms of the expression of the same type of attitude, and explain the varieties of normativity as altering the objects of that attitude. But the above paragraph illustrates how the object of the expressed attitude will depend on the variety of normativity *and* the prejacent: in other words, what token attitude " A ought to ϕ " expresses depends on the modifier (e.g., "rationally") and the values of A and ϕ . These can vary independently. So in the infinite possible combinations of varieties of normativity and prejacents, what rules out the possibility of cases where " A ought to ϕ " and " B ought to ψ " are relativized to different standards yet express rationally discordant attitudes?

5.4.2 Flavors of Approval

What if the expressivist gives up on this appeal to hyperplans and their objects, claiming that "ought" expresses a different type of conative attitude? She may appeal to *approval* instead.³⁷ Initially, this might not seem to be a promising move. After all, the reason Yalcin and others prefer hyperplans is that they can play a similar role to possible worlds in truth-conditional semantics. However, there are at least two reasons why the expressivist might wish to make this move.

³⁷Blackburn (1993, 1999a, 1984) typically appeals to approval instead of plans.

The first is not very compelling. It might seem that the Gibbardian's woes can be traced back to her claim that "ought" expresses an "all-in" attitude: hyperplans. Approval, by contrast, is a "pro tanto" attitude. If I plan to drink and not drink my coffee, I'm irrational. If I express those plans, what I express is inconsistent. But I can approve of both drinking and not drinking my coffee without being irrational, and express that approval without inconsistency. This might help the expressivist avoid over-generating inconsistency. Say that by asserting (1)—"Morally, I ought to bury Polynices, but legally, I ought not do so"—Antigone expresses her approval of both burying and not burying Polynices. Those attitudes, and the sentences that express them, are consistent.

Before we issue a hurrah for approval, we should notice that this form of expressivism *under-generates* inconsistency. Consider the disagreement between Antigone and Creon. When he says (2)—"No, morally and legally, you ought to bury Polynices"—he expresses his approval of not burying Polynices. That attitude is consistent with what Antigone expresses. Now the expressivist has lost one of the central motivations for her view: that it explains moral disagreement.

There is, however, a second reason why we might wish to appeal to approval. Perhaps approval is a type of mental state that comes in different "flavors" to match and explain the varieties of normativity. Perhaps there is moral approval and legal approval and rational approval and so on and so forth. This is a different strategy for the semantic expressivist. Since she still takes "ought" to express the same type of attitude which systematically determines the content of "ought" claims, the expressivist still meets the minimum condition we set out above. Moreover, since she claims that varieties of normativity modify the "flavor" of the attitude itself, the expressivist provides a different explanation of the character of "ought" and how it is relativized to standards.

The trick is to explain what this talk of "flavors" of approval amounts to. Björnsson and McPherson (2014) provide an illuminating way to cash this out. To distinguish the "non-cognitive attitudes corresponding to judgements of moral wrongness, for example, from

attitudes involved in aesthetic disapproval”, they appeal to the distinct *grounds* of states of approval.³⁸ This is fairly intuitive. If I witnessed Peter Singer gracefully scoop a drowning child out of a pond, I would morally and aesthetically approve of his act: I would express the same type of attitude (approval), with the same object (Singer’s action). The token attitudes are still distinct, however, as they are based on different features of the object: I would *morally* approve because Singer’s act averted harm, and I would *aesthetically* approve because of the elegant arc of the motion itself. So the same type of attitude, approval, may come in different “flavors”, and tokens of this type with the same objects may be distinguished by their grounds.

This forms the basis for a second promising expressivist strategy. “*A ought to ϕ* ” always expresses approval towards *A ϕ ing*; when “ought” is relativized to standards (“morally”, “legally”, ...), that expresses the distinct “flavors” (read: grounds) of the token state of approval. To pursue this strategy, the expressivist needs to identify these distinct grounds. On Björnsson and McPherson’s view, “Morally, *A ought to ϕ* ” expresses the speaker’s approval of *A ϕ ing*, where her approval is (paradigmatically) based on the perception that for *A* not to ϕ would involve *A* “intentionally harming or risking harm to others” and/or *A* failing “to respect certain boundaries that play a central role in sustaining social cooperation”. All we need, it seems, is the distinct (paradigmatic) grounds for other varieties of normativity: legality, prudence, rationality, and so forth.

One quick and promising way to develop and generalize this strategy would be to borrow a formal feature from the descriptivist’s account of the character of “ought” (from §5.2). The descriptivist appeals to an ordering source, *g*, that represents the demands of the relevant standard. The expressivist can take that same set of propositions to be the grounds for

³⁸Actually, “grounds” are only one of the five features of moral claims that Björnsson and McPherson appeal to, but the other four offer no additional resources here. As they acknowledge, three of these features—aversion, personal-level acceptance, and engagement—“are shared with instances of many other types” of claims. That is also true of the remaining feature: social hostility. When Creon claims that “Antigone ought not ϕ ”, he may well be disposed to be socially hostile to Antigone if she ϕ s regardless of whether “ought” is relativized to morality or legality. (Indeed, he must be in cases where these standards converge.)

the relevant states of approval. So *g* allows for standard-relativity on a descriptivist or expressivist view.³⁹

However, no matter how this strategy is developed, it encounters the same problems with inconsistency that afflicted its predecessor. To see why, let's ask: Does having two tokens of the same type of attitude with inconsistent contents cease to be rationally discordant when those attitudes have different grounds?

Let's say that it does. This is the best case for the expressivist. Initially, she now seems able to deliver the right results about inconsistency. Consider (1) and (2). Antigone says "Morally, I ought to bury Polynices", expressing approval of burying Polynices, and Creon claims "Morally, Antigone ought not bury Polynices", expressing approval of her not burying Polynices. Let's say that these attitudes have the same ground (e.g., harm), and so are inconsistent. We have an explanation of interpersonal disagreement. Now consider the other conjunct of (1): "Legally, I ought not bury Polynices". This expresses Antigone's approval of not burying Polynices. But since it has a different ground (e.g., respect for the law of the land), this does not generate intrapersonal disagreement.

So far, so good. But this view encounters several problems. For one, it is consistent with Björnsson and McPherson's view that Antigone could express moral approval of burying Polynices based on one ground (harm) and Creon could express moral approval of her not burying Polynices based a different ground (social cooperation). But this would undergenerate inconsistency.

For another, it is consistent with Björnsson and McPherson's view that Antigone could express moral approval of burying Polynices based on one ground (social cooperation) and legal disapproval of not burying Polynices based on *that same ground*. If it is plausible that claims about what we morally and legally ought to do both express approval, it is also plausible that they both sometimes express approval grounded in the perception that acts involve "a failure to respect certain boundaries that play a central role in sustaining social

³⁹I'm grateful to Alex Kaiserman and Hrafn Asgeirsson for suggesting this idea.

cooperation.” But then the two conjuncts of (1) express moral approval of burying Polynices and legal approval of not burying Polynices based on the same ground. If the grounds of moral and legal approval can overlap, this will over-generate inconsistency. This problem is amplified if the expressivist develops this strategy by taking g to individuate standards, as it is well-known and quite obvious that different modal flavors (morality, legality) can share many of the same demands.

The expressivist could attempt to solve both of these problems by positing that there is a unique, non-disjunctive ground for all tokens of moral approval (e.g. harm), and making a similar move for every other variety of normativity. But that would be a very counter-intuitive commitment. It is already doubtful that Antigone (our divine command theorist) and Creon (our virtue ethicist) express moral approval grounded in the perception that acts are harmful or fail to respect boundaries that play a role in social cooperation. Making the distinct ground of moral approval even more specific will be far too discriminatory.⁴⁰

Moreover, even this fix would not help when we consider cases involving embedding. Take (6), Creon’s claim that “Legally, I ought to punish Antigone. But morally, I should not have made it the case that legally I ought to punish her”. *Ex hypothesi*, this speech act expresses Creon’s disapproval of his own approval of punishing Antigone.⁴¹ Given their commitments elsewhere, expressivists are hard-pressed to deny that these first-order and higher-order attitudes are rationally discordant,⁴² regardless of any difference in their grounds.

What if having two tokens of the same type of attitude with inconsistent contents does *not* cease to be rationally discordant when those attitudes are based on different grounds? This seems more plausible. If I plan to drink my coffee because it is delicious and plan not

⁴⁰The expressivist, like Björnsson and McPherson (2014), might take distinct grounds to only be features of *paradigmatic* moral claims. But this is no help. Now we need to know how to discriminate between non-paradigmatic moral claims and paradigmatic non-moral claims.

⁴¹Or rather, disapproval of *having made it the case that* he approves of punishing. I’m not sure if this injects enough space between the higher- and lower-order attitudes to avoid discordance. We could remove that space if we switched from “ A ought to ϕ ” to “It ought to be that p ” (as in: “Morally, it ought not be the case that I legally ought to ϕ .” But this would introduce other complications regarding how we should treat “ought to be” constructions.

⁴²At least, Gibbard and Blackburn are hard-pressed to deny this. On Gibbard, see Wedgwood (2010), p. 121; on Blackburn, see the discussion of quasi-realism below in §5.8.

to drink my coffee because I am over-caffeinated, these plans are inconsistent and render me irrational.⁴³ But if the expressivist accepts this, the second strategy is a non-starter. Identifying distinct grounds for our attitudes will not generate the more fine-grained inconsistency conditions that the expressivist needs: she will still lack any explanation for why moral “ought” claims can be inconsistent with each other, but not with legal “ought” claims; so she will still lack an explanation for why Antigone and Creon disagree, but Antigone does not disagree with herself.

5.5 Meta-semantic Expressivism

So far we’ve seen that the expressivist owes us an explanation of the character of “ought” across all varieties of normativity, and that an expressivist semantics struggles to deliver that while also explaining why inconsistency is standard-relative. Once the expressivist must posit that “ought” always expresses the same type of mental state (to explain its character), any way of explaining the relativization of “ought” to standards allows for inconsistency between moral, rational, legal and prudential “ought” claims. Equally implausibly, it requires us to posit that speakers like Antigone and Creon are practically irrational.

This gives us due cause to be skeptical that the expressivist should pursue a semantic programme that rivals well-developed truth-conditional accounts of the character of “ought”. Instead, she could pursue a meta-semantic programme.

Meta-semantic theories answer the second-order question: “In virtue of what does the phonetic and orthographic sequence ‘ought’ mean what it does?” Michael Ridge (2014) offers the most well-developed form of meta-semantic expressivism, and his view seems particularly promising in this context. Ridge explicitly accepts that “the key terms with which normative claims are made (‘ought’, ‘must’, ‘reason’, ‘good’, ‘bad’, and cognates) are context-sensitive”,

⁴³Some might think that it is consistent to *approve* of drinking coffee and not drinking coffee for different grounds. But the “pro tanto” nature of approval is what makes this seem plausible, not the difference in grounds. It also seems consistent to approve of drinking coffee and not drinking coffee for the same reason: because of how it would affect my mood.

just like “tall”; but they are “not brutally ambiguous”. Rather, they “have a unified meaning”, insofar as they have a single Kaplanian character.⁴⁴

Ridge *adopts* a truth-conditional account of the character of “ought” (like Kratzer’s from §5.2).⁴⁵ The character of “ought” is a function from contexts to contents: call it “ Z ”. Ridge argues that meta-semantic expressivism is consistent with this first-order semantic theory. Let’s grant this, and focus on how we should answer to the meta-semantic question: “in virtue of what does the orthographic and phonetic sequence “ought” mean Z ”? The expressivist answer is that “ought” means Z in virtue of expressing some conative state.

5.5.1 A Disunified Meta-semantics

Perhaps surprisingly, Ridge thinks that this answer is only true *in some cases*. He explicitly adopts a disunified meta-semantics: “Expressivism provides the right meta-semantic theory for *only certain uses*” of terms like “ought”.⁴⁶ Some uses of “ought” mean Z in virtue of expressing a non-representative conative attitude. Other uses of “ought” mean Z in virtue of expressing a “robustly representational belief”, a belief which carries ontological commitments.⁴⁷ We have one *explanandum* (“ought” means Z) and two different *explanantia*.⁴⁸

Does anything mark the uses of “ought” which express conative attitudes rather than representational beliefs? Ridge suggests that the flavor or standard is a relevant marker. He takes there to be, for example, a “deep discontinuity between the ‘ought’ of etiquette, on the

⁴⁴Ridge (2014), p. 9.

⁴⁵This is not to say that there are no important differences in the precise details of their semantics for modals. For example, Ridge argues that his view better addresses concerns about inconsistent sets of requirements. These details are simply orthogonal to our concerns.

⁴⁶Ridge (2014), p. 9 (emphasis added).

⁴⁷Ridge (2014), p. 93.

⁴⁸In our quite extensive personal communication, Ridge notes that this passage “must be handled with great care”, as he had “in mind the meta-semantic story one tells about the semantic *contents* (not *characters*) of sentence *tokens* (not *types*)”. Ridge contends that his meta-semantic explanation of why “ought” means Z is unified: “Linguistic conventions . . . dictate that a speaker who asserts ‘ A ought to ϕ in context C ’ in a context of utterance U has made the judgment” that is fixed in a certain way, where “the judgment” is simply a variable. “In *some* contexts, the relevant judgment will simply be an ordinary garden-variety representational belief”, while “in *other* contexts, the relevant judgment will be” non-cognitive. I take it that this explanation is only superficially unified: when we replace “the judgment” with a disjunction of radically dissimilar disjuncts (“cognitive or non-cognitive attitudes”), it is clearly a disunified meta-semantic explanation of why “ought” means function Z .

one hand, and the ‘ought’ of morality,” on the other; and Ridge takes his disunified meta-semantic expressivism to “privilege standards”, such as morality, “which we think deserve to be taken seriously”.⁴⁹

For the sake of brevity, we can ignore the rest of Ridge’s rich discussion of how the expression of conative attitudes plausibly explains why “ought” means *Z* when it is relativized to standards like morality. Instead, let’s note that Ridge provides very little discussion of two pressing issues: how the expression of “robustly representational” cognitive attitudes plausibly explains why “ought” means *Z* when it is relativized to other standards like etiquette; and how these two distinct expressivist and non-expressivist explanations interrelate.

Let’s consider these issues in turn. On the first, it is worth emphasizing that Ridge must concede that there is a viable non-expressivist explanation of why “ought” means *Z* in a wide variety of uses, such as when relativized to etiquette and (let’s assume) law. This explanation is non-expressivist insofar as it appeals to robustly representational beliefs. And it is viable in that it must explain: why “ought” means *Z*; how context selects the relevant standard; how competent speakers use “ought” to communicate, coordinate, and collect information; and, how speakers disagree even in the face of systematic differences in their criteria for applying words like “legally ought.”⁵⁰ Once that viable non-expressivist explanation is on the table for some uses, why not offer it across the board? A unified meta-semantics is preferable, if only for the sake of parsimony.

On the second, it seems that Ridge is committed to an unexplained linguistic coincidence. There is one *explanandum*: that the word “ought” means *Z*. There are two radically different *explanantia*; the expressivist, after all, is emphatic about the differences between representational beliefs and non-representational conative states. If the *explanantia* are radically different, why is the *explanandum* exactly the same? Why don’t the radical differences be-

⁴⁹Ridge (2014), p. 40.

⁵⁰Dworkin (1986) famously called this “theoretical disagreement” about the law.

tween the states that we are expressing translate to differences in meaning?⁵¹ And relatedly, why would we employ one word to express such radically different mental states?

This combination of a unified semantics and disunified meta-semantics seems quite implausible. Importantly, there are no companions in guilt to whom Ridge can appeal. We take “tall” to have a single meaning, even though it can be relativized to different standards. Once we agree upon that, surely we should hope for a unified meta-semantics for “tall”, rather than different meta-semantic explanations for why different uses of “tall” have the same meaning? It may be reasonable to offer very different meta-semantic explanations for very different fragments of language, such as proper names as logical connectives, but it is unprecedented to offer very different meta-semantic explanations for a fragment of language for which we have a unified semantics: deontic modals.

5.5.2 A Unified Meta-semantics

At this point, the meta-semantic expressivist might wish to part company from Ridge and offer a unified explanation for why “ought” means Z : in every use “ought” means Z because it expresses the speaker’s non-cognitive states. For this to truly be a unified explanation, we need something more precise. (After all, would we have a unified explanation if in some uses “ought” meant Z because it expresses approval and in other uses “ought” meant Z because it expressed hyperplans?) At the very least, we need “ought” to mean what it does in claims like (1)–(4) because it expresses the same *type* of non-cognitive state.

Notice that this is the same constraint that we reached when we considered whether the semantic expressivist could provide an explanation of the character of “ought”. And once that constraint was in place, we saw that the expressivist encountered serious problems with getting the right results about inconsistency. She either failed to generate inconsistency between claims about what we morally ought to do, or over-generated inconsistencies between claims about what we morally, legally, prudentially and rationally ought to do. Now the

⁵¹For a related objection, see Chrisman (2015), pp. 179–80, regarding how expressivists treat logical connectives that embed normative components (e.g. “I am tall or pain is bad”).

meta-semantic expressivist will face the same problems. She will either predict that speakers like Antigone and Creon either fail to disagree with each other, or predict that they disagree with themselves. Neither option is very palatable.

5.6 The Expressivist's Dilemma

So far the challenge to expressivism has been framed in terms of relatively highfalutin distinctions (content and character, semantics and meta-semantics). We can now abstract away from such details and reframe the challenge.

First, let's make two simple and plausible observations about "ought" et al:

Inconsistency "A ought to ϕ " and "A ought not ϕ " are always and only inconsistent when "ought" is relative to the same standard.

Univocality Deontic modals are not lexically ambiguous when relativized.

INCONSISTENCY explains why the conjuncts of (1)—"Morally, I ought to bury Polynices, but legally I ought not do so"—are consistent. It also explains why Antigone disagrees with Creon when she says (1) and he says (2): "No, morally and legally, you ought not bury Polynices". It is a general truth about the inconsistency conditions for modals that are relativized to standards.⁵² UNIVOCALITY is supported by a lot of linguistic evidence (applying standard tests for ambiguity), and considerations of parsimony and explanatory power. These are data that our theories about the meaning of deontic modals must explain. My conjecture is that no form of expressivism plausibly explains both.

Why is that? The answer, I venture, does not depend on how the expressivist (a) identifies the conative state that deontic modals express, or (b) specifies what it is to understand the

⁵²I am not sure how this general truth is best stated for two reasons. First, because it needs to be broader: "A ought to ϕ " and "A ought not ψ " are consistent when "ought" is relative to different standards. (We saw that Gibbard's view does not seem to explain this.) And second, because it may need to be narrower to accommodate possible exceptions where "A ought to ϕ " and "A ought not ϕ " are inconsistent even though they are relative to different standards, due to an interesting relation between those standards. By analogy, "Creon is tall (for an α)" and "Creon is not tall (for a β)" may be inconsistent due to some relation between α s and β s.

meaning of those modals in terms of those states. It concerns whether “ought” always expresses the same type of mental state. UNIFIED EXPRESSIVISM, recall, is simply the position that it does.

Now we can state the dilemma. UNIFIED EXPRESSIVISM is either true or false. If it is true, the expressivist fails to plausibly explain INCONSISTENCY. If it is false, the expressivist fails to plausibly explain UNIVOCALITY. So the expressivist fails to plausibly explain either INCONSISTENCY or UNIVOCALITY.

Consider the first horn of the dilemma. In (1)–(4), Antigone and Creon disagree about what she morally ought to do, but agree about what she legally and prudentially ought to do. The unified expressivist has no good way to explain this. If “Antigone morally ought not bury Polynices”, “Antigone morally ought to bury Polynices” and “Antigone legally ought not bury Polynices” all express the same type of attitude, and the first two express inconsistent attitudes, why do the latter two express consistent attitudes? The expressivist needs a general, systematic explanation of this. None of the candidates we considered succeeded.

All that matters for this argument is that “ought” expresses the same type of mental state, X , regardless of the relevant standard. We can substitute X for hyperplans or approval or “being for” or whatnot. The issue with inconsistency remains the same. This is why the identification of the relevant conative state is downstream from the choice between unified and disunified expressivism.

Now consider the second horn of the dilemma. The expressivist proposes that we are to explain the meaning of “ought” in terms of the mental state(s) that “ought” expresses. If this is so, then we can offer the same explanation of the meaning of “ought” *only if* “ought” expresses the same type of mental state. If different uses of “ought” express plans and preferences, or approval and beliefs, or what have you, then the expressivist is committed to offering different explanations of the meaning of “ought”. I take it that this is true on *any* specification of what it is to explain the meaning of “ought”, which is why the semantic–meta-semantic distinction is also downstream. However, it is worth noting that a disunified

meta-semantic view is on slightly better footing than a disunified semantic view. Both lack similar theoretical virtues (especially parsimony), but the latter is also in tension with linguistic data about zeugma.

This way of framing the challenge reveals why some of the responses that are available to expressivists at most offer partial solutions to the problem. First, note that there are *two* distinct, related issues with inconsistency: if UNIFIED EXPRESSIVISM is true and (1) and (2) express rationally discordant attitudes then (1) does too, which is sufficient to make the conjuncts of (1) inconsistent. This reveals why it won't help to move to a (semantic or meta-semantic) *hybrid* expressivism on which normative terms express conative and cognitive states.⁵³ If the expression of rationally discordant conative states is sufficient to make the conjuncts of (1) inconsistent, it is no help to add that (1) also expresses beliefs. And even if it isn't—which is a big concession—it is still a problem to predict that by asserting (1) Antigone expresses rationally discordant attitudes.

It's not clear whether the unified expressivist can avoid such predictions. She may lean on her account of sincerity, conceding that by asserting (1) Antigone conventionally expresses rationally discordant attitudes, but arguing that Antigone need not *have* those attitudes; she is only *committed* to having those attitudes, such that she either has those attitudes or is *insincere*. But these predictions are no better. Antigone does not seem practically irrational, insincere, or committed to practical irrationality (whatever that may mean).

Second, note that to explain univocality it is not enough to get the result that “ought” means the same thing in (2). It must also mean the same thing in both uses in (1). This reveals why it won't help to move to a view on which some legal claims (like Creon's) express conative states while others (like Antigone's) express cognitive states. That move could seem well-motivated. Hart (1994) famously distinguished “internal” and “external” legal claims: the former are used to guide, recommend, and criticize, while the latter are not. Perhaps, then, we should offer an expressivist treatment of Creon's internal legal claim

⁵³Ridge (2006) is a hybrid expressivist; I ignored this complication in §5.5 to simplify things.

and a descriptivist treatment of Antigone's external legal claim. This could seem to solve the problems with conjunction reduction: if Creon's conjoined uses of "morally ought" and "legally ought" both express (say) a plan, there is no ambiguity in (2). Meanwhile, Antigone's use of "morally ought" expresses a plan and her use of "legally ought" expresses a belief, so (1) does not express practically irrational attitudes. However, if this were so, Antigone's and Creon's uses of "legally ought" would differ in meaning, which would make it hard to explain why Antigone and Creon agree about what she "legally ought" to do.

Third, note that expressivists face the same challenge with all varieties of normativity. The same issues arise when we swap "legally" for "prudentially" in (3) and (4). So even if the expressivist can legitimately offer a disunified approach to "morally ought" and (external uses of) "legally ought", she still faces the burden of explaining why there is no inconsistency in claims like (3): "Morally I ought to bury Polynices, but prudentially I ought not do so".

5.7 Pragmatic Expressivism

This dilemma is not meant to be an impossibility proof. It is a challenge: the onus is on the expressivist to deliver a form of her view that clearly avoids being impaled on both horns of the dilemma. Interestingly, there is a view that can achieve that quite easily, but it offers the expressivist a Pyrrhic victory: it wins the battle but loses the war, as it fails to vindicate any anti-realist ambitions.

The view in question appeals to pragmatics, or communication beyond literal meaning. "Smith is punctual" literally means that Smith is punctual, but if it is the main content of a letter of recommendation it communicates—in particular, it *con conversationally implicates*—that Smith is a mediocre candidate.⁵⁴ This oft-repeated example illustrates how speakers use more than semantic meaning to communicate information about their mental states and external circumstances.

⁵⁴Grice (1957, 1975, 1989).

Speakers may use normative terms to communicate information about their conative states: “I ought to bury Polynices” might *conversationally implicate*, but not *mean*, Antigone’s approval of burying Polynices. This pragmatic thesis could be considered a form of expressivism; it could offer a different specification of what it is to understand meaning in terms of the expression of attitudes.⁵⁵

This pragmatic thesis is consistent with UNIVOCALITY and INCONSISTENCY. It does not require positing lexical ambiguity or offering disunified meta-semantic explanations, for the simple reason that pragmatics is distinct from semantics and meta-semantics. Nor does it lead to implausible predictions of inconsistency between sentences or discordance in speakers’ attitudes. Pragmatic implicatures are flexible and cancelable,⁵⁶ and preserve as far as possible the presumption that speakers are rational.⁵⁷ When Antigone says “Morally I ought to bury Polynices”, this may implicate that she is for burying Polynices; but when she says “Morally I ought to bury Polynices but legally I ought not do so”, this need not implicate that she is irrationally for and against burying Polynices.

This pragmatic thesis can play a role in explaining the meanings (contents) of uses of “ought”. Standard-relative terms can be used to form sentences that are syntactically complete but semantically incomplete. Consider “Creon is tall”. *Being tall* is a relational property, and this sentence fails to specify one of the *relata*. It may be that conversational implicatures about speakers’ intentions help determine the relevant comparison class—and hence the standard to which “tall” is relativized—to generate a logically complete proposition. This is what King and Stanley (2005) call a “weak” pragmatic effect: pragmatics plays a role in assigning contents to uses of standard-relative terms. The same may hold for “Antigone ought to bury Polynices”: it might be syntactically complete but semantically incomplete, as the standard is left unspecified.⁵⁸ Conversational implicatures about speak-

⁵⁵See Yalcin (2007, 2012) for an expressivist theory that is predominantly pragmatic.

⁵⁶For excellent evidence that the so-called “internal connection” between sincere uses of normative terms and speakers’ conative attitudes is in fact cancelable, see Woods (2014).

⁵⁷This is central to how pragmatics has been understood since Grice (1957, 1975, 1989).

⁵⁸See Finlay (2014) for related discussion (in terms of unspecified ends).

ers' attitudes can help determine the relevant standards, and so help assign contents and generate logically complete propositions.

I take it that there's much to be said for this pragmatic thesis. But many worry that "pragmatic expressivism" is a form of expressivism in name only. Ridge, for instance, claims that "expressivism is most clearly *not* well understood as a theory in pragmatics", because that would not vindicate anti-realism:

If that were all that expressivism had to add to the theory of meaning then it would be compatible with a fully representational theory of literal meaning. The whole point of going expressivist was to avoid the problems inherent in a representationalist approach.⁵⁹

I will not try to adjudicate whether anti-representationalism is "the whole point" of going expressivist, or whether expressivism should be a broader church that admits merely pragmatic parishioners. Instead, I will make three points. First, given its role in weak pragmatic effects, this plausible pragmatic thesis can easily be misconstrued as a distinct version of meta-semantic expressivism. Second, such pragmatic theses are consistent with a descriptivist semantics and meta-semantics. And third, descriptivists can and do exploit such pragmatic theses to explain phenomena that were thought to favor expressivism. This includes cases of disagreement where speakers (like virtue ethicists and divine command theorists) systematically differ in their criteria for applying words (like "morally ought").⁶⁰ And it includes the so-called "internal" connection between speakers' conative states and their sincere uses of normative terms like "ought".⁶¹ So such pragmatic theses may neutralize two of the three motivations for expressivism. This makes it more difficult for the semantic or meta-semantic expressivist to lean on disagreement or internalism in fending off the present challenge.

⁵⁹Ridge (2014), p. 103, emphasis added.

⁶⁰See Plunkett and Sundell (2013). See also Dowell (Forthcoming) for skepticism about semantic explanations of our intuitions about disagreement in "moral twin Earth" cases.

⁶¹This "quasi-expressivist" strategy is extensively developed by Finlay (2014).

If this is right, I must amend my earlier hypothesis. If the expressivist is to explain the univocality of “ought” in terms of the expression of conative attitudes, she must posit that “ought” expresses the same type of conative attitude when it is relativized to varieties of normativity; but once she posits that, no matter how we fill in the details we encounter a general problem in explaining inconsistency because attitudinal inconsistency is too coarse-grained. The expressivist can avoid both horns of the dilemma by adopting a pragmatic view that accepts, but does not *explain*, the univocality of “ought”: but by doing so, she fails to deliver on a central motivation for expressivism (anti-realism).

5.8 Generalizability

So far we have focussed on the difficulties the expressivist faces in explaining both univocality and inconsistency. Before concluding, it is worth briefly noting two of the additional difficulties that she may face with a further constraint on an expressivist account of the character of “ought”: generalizability.

First, we need an expressivist view about standard-relativity that generalizes to *embedded* relativized deontic modals. Consider these counterfactuals:

9. If we were in a state of nature, it would still be the case that morally we must not kill.
10. If we were in a state of nature, it would not be the case that morally we must not kill.
11. If we were in a state of nature, it would not be the case that legally we must not kill.

Expressivists have insisted that counterfactuals like (9) express either higher-order disapproval of changing one’s attitude towards killing if the antecedent obtains, or hyperplans to not kill that are “not contingent” on whether the antecedent obtains.⁶² Either of these “quasi-realist” strategies can explain why (10) is false: (9) is true, and inconsistent with (10), so (10) is false. But what about counterfactuals like (11), which is one way of stating

⁶²See Blackburn (1984), p. 198, Gibbard (2003), pp. 186, 183.

the important truth that we create legal obligations? If (10) and (11) express the same type of attitude, why is (10) false and (11) true? Why are (9) and (10) inconsistent, while (9) and (11) are consistent? Alternatively, if (10) and (11) express different types of attitudes, why should we accept this disunified (and seemingly *ad hoc*) treatment of embedded relativized deontic modals?

It is also worth noting that quasi-realist commitments about counterfactuals like (9) can generate problems for certain expressivist treatments of (1)–(6). We’ve already seen this once (in §5.4): (6) may express discordant lower- and higher-order attitudes. Similarly, an expressivist who took “morally ought” to express plans and “legally ought” to express contingent plans will encounter problems with inconsistency. Say “I legally ought to ϕ ” expresses a plan to ϕ if p . That sentence would be inconsistent with counterfactuals like “If it were true that p , it would still be the case that morally I ought not ϕ ”, which express (*ex hypothesi*) a plan to ϕ that is “not contingent” on whether p is true.

Second, we may need an expressivist account of standard-relativity that generalizes to *non-deontic* standards.⁶³ Imagine a philosophy seminar where two speakers debate the relationship between morality and metaphysical necessity:

12. Metaphysically, Antigone must bury Polynices. So it is not the case that morally, Antigone must bury Polynices.

13. No! Metaphysically *and* morally, Antigone must bury Polynices.

A speaker who said (13) might sound quaint (as do many other contributors to philosophy seminars). But her use of “must” does not seem zeugmatic, as “light” does in “Pale colors and feathers are light.” If this is right, then this instance of conjunction reduction puts pressure on the expressivist to offer a unified semantics for modals when they are relativized

⁶³Wertheimer (1972) made a similar point *apropos* non-cognitivism, albeit with a different strategy: he did not appeal to the linguistic data from tests like conjunction reduction, or draw connections between concerns about univocality, inconsistency, and internalism.

to moral or metaphysical standards (which is something that descriptivists have already provided).⁶⁴

Expressivists may well balk at the thought that their view must generalize from moral to metaphysical modals. For one thing, many of the paradigmatic motivations for expressivism are not present in metaphysical claims: one may sincerely claim “Metaphysically, Creon must kill Antigone” even if one views this act dispassionately, or with unmitigated sorrow. For another, expressivism is often predicated upon there being a deep difference between how we explain the meaning of *normative* language and *descriptive* language. If they offer a unified treatment of morality and metaphysics, that very distinction looks unstable.⁶⁵

But when we consider standards like legality, it is hard to see how such a distinction cuts nature, or rather natural language, at its joints. Legality is plainly not “non-normative” in the way that metaphysics is non-normative. The legal “must” and the moral “must” are *deontic* modals; the metaphysical “must” is not. The legal “must” and the moral “must” can both guide us in practical deliberation;⁶⁶ the metaphysical “must” cannot. But the legal “must” may not be “normative” in the same way that the moral “must” is normative. Like other artificial standards—etiquette, games, grammar—it may not be *authoritative*.

What the expressivist needs is some clear cut between normative and non-normative natural language. If we merely compare “ought” and “tall”, that cut may seem obvious. But once we recognize the continuities between moral and metaphysical modals, and the deeper continuities between moral and legal deontic modals, it is hard to find a non-*ad hoc* place to draw the line. This is perhaps the most troubling challenge to expressivists. If there

⁶⁴Again, see Kratzer (1977, 1981, 1991).

⁶⁵Notably, when Schroeder (2008) concluded (for different reasons) that the best form of expressivism offers a unified treatment of normative and descriptive language in terms of the same attitude type, he predicted that many people will “conclude this amounts to a *reductio* of expressivism”, pp. 92–93. His prediction was correct: see Woods and Baker (2015).

To be clear, my contention here is different from Schroeder’s conclusion. I am not arguing that the expressivist must offer a unified treatment of *all* language; I am arguing that she must offer a unified treatment of all normative and some non-normative language (e.g., modals). The expressivist can deny that the latter is representational and, in that sense, “descriptive”. But it is still not normative. That’s why the normative–descriptive distinction looks unstable.

⁶⁶See Hart (1994)’s discussion of the “internal point of view”.

is no clear demarcation between normative and descriptive fragments of language, how can there be a deep difference between how we explain their meaning?

5.9 Conclusion

Expressivism, as Yalcin noted, has been unmoored from linguistic phenomena. Given this, the best way to determine whether there is a plausible expressivist programme, and if so what that programme should be, is to focus on neglected linguistic phenomena, like the relativization of “ought” to different standards.

This phenomenon poses a novel challenge to expressivism. Does “ought” express the same type of conative state when relativized to different standards? If so, expressivists struggle to explain inconsistency. If not, expressivists struggle to deliver a plausible explanation of the univocal meaning of “ought”.

It’s too early to tell whether some form of expressivism survives this challenge without abandoning all anti-realist ambitions. My main aim, however, has been to throw down a gauntlet. The expressivist still owes us an account of the semantic character (rather than content) of “ought”. We are yet to see how she can either rival or adopt the truth-conditional first-order semantic theories that take the character of “ought” to mean some function from contexts to contents. We’re used to hearing the mantra that everything descriptivists do, expressivists can do cheaper. For the expressivist to deliver on that front, she must show how she can explain the univocality of “ought” in terms of the expression of a type of conative attitude that has sufficiently fine-grained inconsistency conditions.

Aside from this challenge to expressivism, I hope to have shed light on two things. The first is the methodological point with which we began. It is fine to ask questions about the meaning of fragments of natural language while considering those fragments in isolation. But since we do not use those fragments in isolation, we must not lose sight of the bigger picture. We must ask how those fragments fit together. Like everyone else, the expressivist owes us an explanation not just of moral language, but of how the open-ended set of varieties

of normativity fit together. In other words, everyone—including expressivists of all stripes—must explain the deep continuities between the uses of deontic modals in relation to the varieties of normativity.

The second concerns explaining the deep discontinuities between uses of deontic modals in relation to varieties of normativity. Like Ridge, I think some standards like morality always deserve to be taken seriously, while other standards like etiquette do not. Ridge’s explanation of the difference is that claims about what we ought to do according to morality express judgments that have the functional role of contributing to settling what we will do, while claims about what we ought to do according to etiquette do not. This is false: as we have seen in Chapter Two and Chapter Three, etiquette judgments routinely play complex roles in contributing to settling how we will set tables, and plenty more besides. This specific problem aside, I think expressivism is in general quite poorly placed to explain the distinction between authoritative standards like morality and “merely formal” standards like law and etiquette. Expressivism explains differences in the meanings of terms and sentences via the differences in the mental states they express in sincere speech acts. With regard to *this* distinction, we do not want to posit a difference in the meanings of our terms. What we want, if what I’ve argued in Chapters Two and Three is true, is to posit a distinction between normative reality and normative fiction. As anti-realists about normativity, expressivists cannot give us that. Nor can they give us an analog for the distinction between reality and fiction in terms of the mental states we express, because there is no obvious non-cognitive analog for the distinction between belief and fictive acceptance. In fact, even if there were such a distinction, it is not clear that expressivists could appeal to it; pretending to approve or plan would presumably count as a form of insincerity. This makes the problem expressivists face all the harder to solve: they simultaneously need to explain both the deep continuities and discontinuities between varieties of normativity, but are not well placed to do either.

Chapter 6

Why Realists Must Reject Normative Quietism

6.1 Introduction

The last two decades have seen a surge of support for normative quietism: most notably, from Ronald Dworkin (1996, 2011), Tom Nagel (1996, 1997), Derek Parfit (2011a,b) and T.M. Scanlon (1998, 2014). Detractors like David Enoch (2011a) and Tristram McPherson (2011) object that quietism is incompatible with realism about normativity. The resulting debate has stagnated somewhat. In this chapter I explore and defend a more promising way of developing that objection: I'll argue that if normative quietism is true, we can create reasons out of thin air, so normative realists must reject normative quietism. I'll close by diagnosing where quietists go wrong, arguing that they need but can't provide a distinction between normative reality and normative fiction.

6.2 Normative Quietism

Normative quietism is notoriously difficult to characterize. For the sake of clarity I'll focus on the form of quietism that Scanlon defends in *Being Realistic About Reasons*.¹ Admittedly, I am not sure how readily the objections I will raise generalize to other quietists' views since their details are fairly elusive—I don't know what to make of Parfit (2011b)'s claim that irreducible reasons exist “in a non-ontological sense”, and that this view “has no metaphysical

¹All page references will be to Scanlon (2014), unless otherwise indicated.

implications” and so “cannot be open to metaphysical objections” (747). The three core ideas at the heart of Scanlon’s view, by contrast, are relatively clear and attractive.

Following Scanlon, let’s warm up to normative quietism by first considering mathematical quietism. “Two plus two is four” is true. What makes it true? Is it true because numbers exist as Platonic forms? If so, how do we have epistemic access to these aphysical entities? Or, do facts about numbers reduce to physical facts about quarks and bosons and the like? If so, how does the reduction go? Scanlon rejects both of these views. Claims like “Two plus two is four” are true because they are licensed by standards internal to the mathematical domain. Numbers exist because such claims quantify over numbers. So “questions about the existence of numbers ... are settled by mathematical reasoning” (19), and mathematical claims can’t be metaphysically or epistemologically problematic.

Similarly, claims like “There is reason to give to charity” are true because they are licensed by standards internal to normative discourse, and reasons exist because such claims quantify over reasons. Questions about the existence of reasons are thus settled by first-order normative reasoning. So we need not invoke Platonic forms or reduce facts about reasons to facts about quarks and bosons to explain how reasons exist or how we have epistemic access to them.

This illustrates what I take to be the first core idea of Scanlon’s quietism:

Diversity: Different statements have different subject matters, and what it takes for a statement to be true depends upon its subject matter.

To be clear, there is a single account of what it takes for any statement to be true: it must be licensed by the relevant domain-dependent standard. However, the standards for specific statements differ depending on their subject matters.

I take it that DIVERSITY is the plausible idea that Scanlon intends to model through his somewhat more abstruse appeals to “domains” and “standards” (which will be clarified and discussed further below). Scanlon seems to agree. He writes that his use of the term

“domains” is “just a way of referring to the fact that statements can make claims about different subjects”, which is “a common-sense idea”. Scanlon immediately adds that “the idea that questions about a given domain (i.e. about a given subject) should be settled by the best ways of thinking about that subject is a piece of common sense, even a triviality” (23).

An attractive feature of Scanlon’s quietism is that DIVERSITY is fully general: it applies to all statements about all subjects, including natural science. In this sense, it is a tad misleading for me to describe Scanlon as a *normative quietist*, rather than a *quietist*. Scanlon does not accept a metaphysically inflationary account of what it is for protons to exist and a deflationary account of what it is for reasons to exist; if he did, he could justly be accused of being a realist about science and an anti-realist (or quasi- or ir-realist) about normativity.

As Scanlon recognizes, one could accept DIVERSITY and yet deny quietism. One could insist that for “Two plus two is four” to be true it must be licensed by the standards of mathematics *and* some other domain. That other domain could be some privileged first-order domain(s), like physical science, or some general second-order domain, such as metaphysics. Scanlon emphatically rejects both of these possibilities.

This takes us to the second core idea behind his quietism:

Autonomy: There are no privileged first- or second-order domains.

As Scanlon notes, DIVERSITY “becomes philosophically controversial only when it is combined with the claim that statements about domains other than the natural world should be seen as autonomous” (23). So let’s outline the controversial commitments Scanlon accepts with respect to AUTONOMY.

Start with the first-order domains. For Scanlon, all domains are created equal, and none are more equal than others. On its face, that doesn’t seem problematic. But Scanlon acknowledges that it has striking implications:

[W]e could adopt some way of talking which specified criteria of identity for objects of a certain sort, and truth conditions for sentences containing terms referring to them which allowed for existential generalizations from such sentences. According to my view, as long as this way of talking was well defined [and] internally coherent, ... we would be committed to the existence of things quantified over in the existential statements counted as true in this way of talking (27).

For example, at various points in history we have adopted well defined and internally coherent ways of talking about supernatural phenomena like witches. So there is a supernatural domain with its own domain-specific standards, which will license claims like “Witches caused the illness”. Since such claims quantify over witches, Scanlon is stuck with the conclusion that witches exist, and indeed are no less real than protons or reasons; at least, that is the view so far.

What about second-order domains? Scanlon accepts that there could be a domain like metaphysics whose subject matter is “existence”, or “the universe”. But he holds that the notions of “existence” and “the universe” in this second-order domain would be “entirely parasitic on the [first-order] domains that contribute to it”: there is a “merely disjunctive “universe” in which something exists if it is a physical object, or a number, or is entailed by true statements of some other domain” (22–24). Scanlon very explicitly rejects any “general idea of existence” and holds that “genuine ontological questions are all domain-specific”, which he recognizes is a “radically anti-metaphysical” conclusion (n.10). So Scanlon is right that coupling DIVERSITY with AUTONOMY is philosophically controversial: it has significant extensional implications (about what exists) and intensional implications (about there being no domain-general account of existence or the universe, other than the quietist’s quite uninformative one above).

Scanlon is, of course, quite content with these intensional implications. But he does wish to ameliorate concerns about the extensional implications of his view. And rightly so. After

all, different domains could license inconsistent statements, and his quietism would be a non-starter if it led to contradictions.

One way to reduce the incidence of such contradictions is to insist that only the truth values of “pure”, rather than “mixed”, claims in a domain are properly settled by standards internal to that domain. Claims like “Jones has a reason to leave the burning building now” are mixed, insofar as they “make normative claims but also make or presuppose claims about natural facts” (38). Mixed claims should be factored into pure claims: e.g., a claim about “natural facts” (e.g., “Jones is in a burning building”) and a conditional normative claim.² This helps, but Scanlon recognizes that on its own it is insufficient: “Even pure statements in one domain can entail or presuppose claims in some other domain” (21).

This takes us to the third core idea in Scanlon’s view:

Consistency: If domains α and β contain, entail, or presuppose inconsistent statements—if p is true in α and false in β —either α or β must be rejected.

Scanlon illustrates CONSISTENCY with scientific and supernatural domains. The supernatural domain licenses “claims about events in the physical world and their causes” which “conflict with claims of physics and other empirical sciences, and this conflict provides decisive reason to reject the idea that there are witches” (21). We’ll return to Scanlon’s treatment of this example later. For now, it’s sufficient to note that CONSISTENCY would not, by Scanlon’s own admission, allow us to reject claims in a supernatural domain that only quantified over (say) causally inefficacious ghosts. For Scanlon, so long as a first-order domain does not conflict with any other domain then that domain is entirely autonomous.

At this point, you may wonder why Scanlon is willing to court such controversies in the first place. Why not privilege some first-order domains, or allow that we can ask “genuine ontological questions” that are domain-independent?

²These pure claims all have the following form $R(p, x, c, a)$: if p is true and an agent x is in condition c , p is a reason for x to perform act a (or have attitude a) in c (31, 37).

Scanlon argues that the mathematical and normative domains do not contain, entail or presuppose any statements that are false according to any other domain, such as the domain of physical science (22). On his view, that's enough. For Scanlon, "the truth values of statements about one domain, insofar as they do not conflict with statements of some other domain, are properly settled by the standards of the domain that they are about" (19). If Scanlon gave up on part of this picture—if he denied AUTONOMY—then he would open the door for skeptical arguments about whether there really are numbers and reasons.

6.3 Creating Reasons Out of Thin Air

So far we've outlined the attractive, albeit controversial, core ideas at the heart of Scanlon's quietism. In time we'll delve further into the details of the view—especially about the nature and role of domains and standards—but for now let's introduce the objection that quietism is incompatible with realism.

As I said, I am by no means the first to make an objection in this ballpark. Detractors have objected to quietism's intensional or extensional implications, but dialectically, they have failed to find a point of purchase against the quietist.

Objections to quietism's intensional implications often rely upon metaphysical ideology that Scanlon, *qua* quietist, rejects. Enoch (2011a), for instance, objects that Scanlon's quietism is "a notational variant of a fictionalist view", and so not a form of realism, because quietism does not distinguish between what "exists" in the "soft, licensed-by-internal-standards" sense, and what "EXISTS" in "the metaphysically heavy, Platonistically friendly" sense (124). But as we've already alluded to, Scanlon simply rejects that distinction (28). Scanlon counts himself as a realist, and denies that realists must (or should) be committed to any "metaphysically heavy", domain-independent notion of existence.

Extensional objections are simpler, and more promising. One was illustrated above: if we adopted a certain way of talking about causally inefficacious ghosts, it would follow from Scanlon's quietism that these entities exist. Similarly, McPherson (2011) objects, if a

community adopted a well defined, internally consistent way of speaking about “schmeasons”, it would follow from quietism that schmeasons exist. Indeed, schmeasons would be no less “real” than reasons. Enoch makes a similar point *apropos* “counter-reasons”. These objections seem damning. Yet in response Scanlon has shown a remarkable willingness to bite bullets and soften their bitter aftertaste with a hefty dose of pragmatism. He concedes that we can “adopt some way of talking” and thereby add new entities to our “ontological commitments”, then asks rhetorically, “Can we take seriously an idea of existence that comes so cheaply?” Scanlon answers, oracularly, that

... the question about such entities is not whether they really exist. This question is settled by the standards of the domain ... The question is only whether we have any reason to be concerned with these entities and their properties (27).

If statements about causally inefficacious ghosts or schmeasons or what have you do not conflict with statements about any other domain, Scanlon seems to think, we have no reason to be concerned with such entities. And if we have no reason to be concerned with ghosts or schmeasons, Scanlon thinks, we can be quite cavalier about our ontological commitments with respect to such entities.

Much more can and will be said about Scanlon’s responses to these objections, but for now I want to point out how the argumentative strategy I will pursue here departs from these objections in ways that have a better hope of finding a point of purchase against the quietist. Instead of considering how easily Scanlon can acquire ontological commitments with respect to ghosts or schmeasons, let’s just consider the extensional implications of quietism with respect to reasons.

I will argue that it is possible for a community to adopt a new well defined, internally consistent, way of talking about reasons, statements about which do not conflict with statements in any other domain. If quietism is true, that community thereby creates new reasons to do arbitrary things out of thin air.

If this is right, no dose of pragmatism will soften the blow. In part, this is because of the practical import of these entities: for Scanlon reasons *are* what matters. So we do have reason to be concerned with such entities. But the real kicker is the theoretical import of these entities in Scanlon's "Reasons Fundamentalist" metaethical theory. This is where the intensional upshot of the objection becomes important. Like other quietists, Scanlon clearly accepts certain "realist" commitments about the nature of reasons: that reasons cannot be arbitrary; that reasons cannot be created merely by adopting some way of talking; that reasons do not depend on our will or our flights of fancy; that reasons, to put it simply, are independent of us.³ If Scanlon's quietism entails that we can create new reasons to do arbitrary things out of thin air, Scanlon's quietism cannot vindicate his own realist commitments about reasons.

This is just a general outline of the objection; the devil is obviously in the details. This is unfortunate, since some crucial details of Scanlon's view are fuzzy. It will help to orientate ourselves by focusing on a familiar example.

Consider our discourse about morality and etiquette. It is often observed that morality and etiquette are both self-justifying. Both purport to tell us what we ought, must, can, and have reason to do. According to morality, we have reason to do what morality requires us to do. According to etiquette, we have reason to do what etiquette requires us to do. But requirements of etiquette are things we can clearly create with our shared attitudes. If quietism is true, why don't we create new reasons when we create new requirements of etiquette?

To make the comparison vivid, consider two particular requirements:

Tithing Norm According to morality, the wealthy must give to charity.

Prudish Norm According to etiquette, the unwed must be chaste.

³Such claims are made throughout Scanlon (2014) (e.g.: 4–8, on desire-based theories, and 93–104, on constructivism). Dworkin (1996, 2011) and Parfit (2011a,b) make similar claims.

We can make it the case that Prudish Norm is true merely by changing our attitudes; but we cannot make it the case that Tithing Norm is true merely by changing our attitudes. Quietists accept this (see, e.g., Parfit (2011a) 144–145).

What can quietists say about the corresponding claims about reasons?

Tithing Reason There is a reason for the wealthy to give to charity.

Prudish Reason There is a reason for the unwed to be chaste.

Quietists clearly accept that there is some connection between the truth of Tithing Norm and the truth of Tithing Reason. I claim that by their lights they must accept a similar connection between Prudish Norm and Prudish Reason.

This way of framing the objection avoids taking on too many commitments about how to interpret certain details of Scanlon's view. But we won't be able to develop the objection without at least discussing those details. In particular, how we develop the objection depends on whether discourse about etiquette forms a separate domain with its own domain-specific standard (*a la* the supernatural domain), or is subsumed under the one normative domain.

In §6.4, I'll argue that if Scanlon's quietism is true, discourse about etiquette forms a separate domain; and if quietism is true and discourse about etiquette forms a separate domain, we can create new reasons out of thin air. If realism is true, we cannot create reasons out of thin air, so we must reject quietism.

In §6.5, I'll argue in the alternative that if Scanlon's quietism is true and discourse about etiquette is subsumed under the one normative domain, we can still create new reasons out of thin air. So again, realists must reject quietism.

My central aim is to defend the conclusion that realists must reject quietism. With both arguments there will be several spots where one can get off the boat. In particular, one can resist the arguments by reinterpreting or revising Scanlon's views about domain-individuation, the nature and role of domain-specific standards, inter-domain conflict, indeterminacy, and revisability. We will see, however, that these escape routes do not offer a way to salvage a

view that is both plausible and recognizable as a form of quietism. So a subsidiary aim in this chapter is to clarify and challenge these particular quietist commitments.

6.4 Separate, Conflicting Domains

In this section I'll argue that if Scanlon's quietism is true, discourse about etiquette forms a separate domain; and if quietism is true and discourse about etiquette forms a separate domain, we can create new reasons out of thin air.

Let's start with the premise that if Scanlon's quietism is true, discourse about etiquette forms a separate domain. That is to say, there is a set of statements like Prudish Norm and Prudish Reason that have a distinct subject matter: social propriety. The upshot of this is that such statements are true if they are licensed by a domain-specific standard, subject to the CONSISTENCY constraint.

Quietists might be tempted to reject this first premise on one of two bases. First, quietists could deny that discourse about etiquette includes statements about reasons. Parfit (2011a) suggests this when he distinguishes between two "conceptions of normativity". For Parfit, morality is normative on the "*reason-involving* conception", according to which "normativity involves reasons or apparent reasons", but etiquette is only normative on the "*rule-involving* conception", according to which "normativity involves requirements, or rules, that distinguish between what is *correct* and *incorrect*, or what is *allowed* and *disallowed*" (144–5, emphasis his). Along with etiquette, Parfit offers requirements of law, linguistic rules and codes of honour as paradigm examples of "rule-involving normativity". Discourse about all such requirements is not "reason-involving".

If that were true, it would block the argument. For quietists, there would be a connection between Tithing Norm and Tithing Reason but not between Prudish Norm and Prudish Reason, because we use "reason" in discourse about requirements of morality, but not in discourse about requirements of etiquette.

Unfortunately, this response is unavailable to quietists. *Contra* Parfit, our discourse about all of these requirements is “reason-involving”. This point was implicitly made by Philippa Foot (1972), who noted that “people give as their reason for doing something the fact that it is required by etiquette” (319).⁴ The same point is better developed in philosophy of law, where it is an orthodoxy that according to legal discourse new legal rules give us new “legal reasons”; this feature of legal discourse is well defined, internally coherent, and has been formally modeled by John Horty (2011). Admittedly, this only shows that discourse about requirements of morality and etiquette and law all use the orthographic and phonetic sequence, “reason”. Quietists could try to deny that such discourse all involves the same *concept* of “reason”. But this is hard for Scanlon to defend. While his position on how we attribute this concept is not clear, it seems to turn on the distinctive role that “reason” plays in deliberation and motivation (54). Claims about “legal reasons” et al. play those same distinctive roles in deliberation, motivation and criticism, at least among those who take what H.L.A. Hart (1994) called the “internal point of view” (as we saw in Chapters Two and Three). So this move won’t help quietists.

Moreover, quietists (and realists) should not offer this response. If one claims that discourse about etiquette does not involve the concept “reason”, one encounters serious difficulties in explaining disagreement. The world is full of fogeys and prudes who take the internal point of view towards etiquette norms in general, and Prudish Norm in particular. They take Prudish Reason to be true. I take Prudish Reason to be false. I take it that we disagree about whether there is a reason for the unwed to be chaste. How can we make sense of this disagreement if the prude and I do not use the same concept, “reason”?

Second, quietists could agree that discourse about etiquette includes statements about reasons, and insist that because this is the case etiquette discourse must be subsumed under the one normative domain: that is, the domain of reasons. If this is right, the first argument fails; but the second might succeed (see §6.5).

⁴Notably, on Nic Southwood (2011)’s view, judgments about conventional norms (like etiquette) constitutively involve some cognitive attitude towards “practice-dependent reasons”.

However, this response is awfully *ad hoc*. To see why, let's consider in detail what Scanlon tells us about the nature and role of domains. A domain is not "a distinct realm of objects"; rather, it is a distinct realm of discourse. A "domain is better understood in terms of the kinds of claims it involves, and hence in terms of the concepts that it deals with, such as number, set, physical object, reason, or morally right action" (19). Scanlon provides no criteria for individuating domains, and says that this is intentional: he meant to "leave the question of what domains there are entirely open", including whether some groups of "particular domains" are "in fact single unified domains" (23, and n.6).

This does not worry Scanlon. His appeal to "domains" is "just a way of referring to the fact that statements can make claims about different subjects"—it is meant to model the intuitive idea we called DIVERSITY. So "the question about domains is not whether they exist"; it is whether they are "helpful" (23).

But this does worry me. If we have no principled basis to prefer one way of carving up domains rather than another, but can reach different conclusions depending on which one we choose, then domains are not helpful. They make our ultimate ontological commitments depend on choices that are utterly arbitrary.

Indeed, our choices about domain-individuation can be arbitrary in at least two respects. First, we can require more or less *conceptual overlap* between sets of statements in order for them to fall under the same domain. We know that if sets of statements had no conceptual overlap they would be about entirely different subjects, and hence form separate domains. But we also know that there will be some conceptual overlap between separate domains: all domains must share some basic logical concepts; and the supernatural and scientific domains could not conflict if they did not share the concept of "causation". So what could resolve a dispute about whether x amount of overlap is enough?

Second, we could individuate concepts themselves coarsely (reasons) or finely (moral reasons; epistemic reasons) in determining how to individuate domains. Scanlon seems to

want to do both: he speaks as if there is one normative domain—the domain of reasons—as well as a separate moral domain (e.g., 22).

In the absence of criteria, we should determine whether etiquette forms a separate domain by considering Scanlon’s treatment of examples. Scanlon holds that supernatural and scientific discourses form separate domains even though both involve claims about “causation”, which is as central to the scientific domain as “reason” is to the normative domain. And throughout the book Scanlon allows that there is a normative domain (involving claims about “reason”) *and* a moral domain (involving claims about “reason” and “morally right action”). Surely we should say the same about etiquette, “reason”, and “socially proper action”. We could reach that same conclusion by appealing to fine-grained subject matters: “moral reasons”, “legal reasons”, and (more awkwardly) “reasons of etiquette”.

So the quietist should not insist that since etiquette discourse involves claims about reasons it must be subsumed under the one normative domain. That position is *ad hoc* (given Scanlon’s verdicts about other cases) and arbitrary (given the absence of clear criteria for domain-individuation). And we do not want our ultimate ontological commitments to depend upon arbitrary choices.

I have been defending the first premise—that etiquette discourse, including claims like Prudish Reason, forms a separate domain—by playing defense. There is also a positive argument in favor of the claim that if quietism is true, etiquette discourse forms a separate domain. Scanlon tells us that “the domain of practical reasons is not a unified subject matter” (104). This relates to his position on indeterminacy, which we will address shortly. For now, we should simply note that if talking about “domains” is just a way of talking about statements’ subject matters, and “practical reasons” is not a unified subject matter, then it seems that the natural conclusion to draw is that *there is no normative domain*. Or if there is, it is like the domain of metaphysics: it is “entirely parasitic” on the more unified domains (“moral reasons”, “legal reasons”) that contribute to it.

Now let's turn to the second premise: if quietism is true and discourse about etiquette forms a separate domain, we can create new reasons out of thin air. If etiquette forms a separate domain, statements like Prudish Reason are true if and because they are licensed by the domain-specific standard of etiquette. (I'm setting aside the CONSISTENCY constraint, which we'll return to shortly.) That would allow us to make reasons exist by merely changing our attitudes.

Let's walk through this step by step. First, what are standards? Scanlon tells us that a standard consists in part "of substantive principles about the domain", which are in turn "justified by less explicitly codified reasoning about the subject matter in question," and "can be revised in the light of further reasoning of this kind." That reasoning "proceeds by appeal to our best general understanding of the nature of the concepts basic to the domain in question and to the most obvious particular truths within it" (20). To illustrate this, Scanlon claims that the standards of the empirical sciences are justified in part by "our conception of the subject matter of science" as "comprised of objects in space, causally interacting with us" (20), while the standards for set theory are justified in part by our best understanding of its central concept: "set".

Second, would the standard of the etiquette domain license Prudish Norm and Prudish Reason? Yes they would. Prudish Norm and Prudish Reason can be among the "obvious particular truths" in some local etiquette system, and can be licensed by the standard of that domain in the same way that Tithing Norm and Tithing Reason are licensed by the standard of the moral domain.

Third, if the standard of the etiquette domain can license Prudish Reason, can we create reasons out of thin air? Yes we can. Reasons exist if and because statements that quantify over reasons are licensed by the relevant standard. And a crucial difference between morality and etiquette is that we can *create* the relevant standard of etiquette via our practices and attitudes. So when we create the relevant standard of etiquette that licenses new statements about reasons, like Prudish Reason, we create new reasons out of thin air.

Some might dispute this third step by offering an alternative interpretation of the role of domain-specific standards. Scanlon describes their role as follows: a claim about what exists is “properly settled” by the standards of the relevant domain. This is ambiguous. On one reading, standards only play an epistemic role. The best way for us to reach justified beliefs about what exists is to reason in accordance with the standards for the relevant domain. That is, standards settle what we are justified in believing to be true. I agree that, for Scanlon, standards play this role.⁵ And I agree that if this is all that standards do within Scanlon’s quietism, the above argument would not go through: by creating a standard of etiquette we would change what statements about reasons we’re justified in believing, but not what statements about reasons are true.

In §6.5 we’ll see why even this has problematic implications. But for now, I will dispute that standards *only* play this epistemic role.⁶ On my reading, standards also play an ontological role: domain-specific standards settle what is true, and hence what exists. (To exist just is to be quantified over by a true statement.) Why accept this second reading? Because this is how Scanlon describes his view: standards settle “truth values” (not just justification), and standards provide “the conditions required in order for objects in different domains to exist”; this fits with his minimalist conceptions of facts and properties (20–22, 45, and n. 12). Moreover, if standards only played an epistemic role, Scanlon’s view would not be recognizably quietist; it would be compatible with a range of inflationary metaphysical views. Scanlon needs standards to play an ontological role to explain and support his other metaphysical commitments (or lack thereof).⁷

⁵See, for example, Scanlon’s references to standards *justifying* existential claims (25, 123).

⁶Maguire (2014) defends this interpretation: “The text is richly ambiguous between two interpretations of this “settling” relation, as between the constructivist view that standards for settling truth-values *ground* those truth values, and the realist view that even the best employment of the standards for settling truth-values may get things wrong. As the last sentences of the book suggest, the balance of Scanlon’s conviction lies with the realist position.” Wedgwood (2015) suggests a similar interpretation: “Given that it is plausible that different sorts of beliefs are justified in different ways, there is no reason for accepting the Quinean view that there is a single method for answering all ontological questions”. It is hard to see how to square this interpretation with the explicit textual evidence, or with Scanlon’s ambitions.

⁷These include, *inter alia*, (a) Scanlon’s refusal to commit to naturalism or non-naturalist Platonism, (b) his views about indeterminacy, and (c) his responses to objections (below).

I don't mean to suggest that it is unproblematic for Scanlon to hold that standards play this ontological role. For instance, it immediately raises the spectre of an explanatory circle. Scanlon writes that the standards for a domain are partly set by appealing to "obvious truths" within that domain, *and* that standards settle truth-values. How can standards determine truths while the truths also determine the standards? I have no idea. But that's what Scanlon has said in print. So this is a place where one can get off the boat, though doing so won't salvage a view that's still Scanlon's, or recognizably quietist.

We've seen that if Scanlon's quietism is true, discourse about etiquette forms a separate domain, and so we can create new reasons out of thin air. But so far we've set aside the CONSISTENCY constraint, which offers the quietist a way of resisting this conclusion. You might think that just as the supernatural and scientific domains conflict, the etiquette domain will contain, entail, or presuppose claims that are false according to some other normative domain. So just as the inter-domain conflict between the supernatural and scientific domains "provides decisive reason to reject the idea that there are witches", the inter-domain conflict between etiquette and (say) morality will provide decisive reason to reject the idea that there is reason for the unwed to be chaste.⁸

There are two problems with this response. First, if etiquette conflicts with some other normative domain like morality, why does that "provide decisive reason" to reject Prudish Reason rather than Tithing Reason? Why privilege morality over etiquette? The problem here is fully general: that domains α and β conflict is a symmetric condition, and so cannot itself provide decisive reason to reject α over β rather than β over α . In his discussion of witches, Scanlon simply asserts that the conflict between the supernatural and scientific domains provides decisive reason to reject the former. But why, if quietism is true, should we reject the idea that there are witches rather than reject the idea that there are quarks?

⁸This was Scanlon's response to the argument in personal correspondence: "In order for us to be free to "set up" etiquette or some other domain however we like, it must be understood in a way that does not involve interdomain conflicts ... [but] in order for such a domain to entail conclusions about reasons it must be understood in a way that does involve such conflicts".

If all domains are equal—as DIVERSITY and AUTONOMY claim—why privilege science over the supernatural, or morality over etiquette?

Scanlon never gives us an asymmetric condition to serve as a tie-breaker. Indeed, it's not clear how he could. He may say that scientific discourse has “demonstrated success” in “predicting and explaining what seem to be obvious facts about this world” (20), and witch discourse has not.⁹ But how are *quietists* entitled to make such comparisons? What domain are such comparisons internal to? If they are internal to the scientific domain, they add nothing: we already knew that scientific claims meet scientific standards but supernatural claims don't, and vice versa. So this is still a symmetric condition, and cannot serve as an explanation for why quietists are justified in privileging science.

One might suggest that comparisons between scientific and supernatural standards are internal to some third domain. But what third domain could that be? It cannot be another *first-order* domain. No first-order domain can license second-order pronouncements upon the scientific and supernatural domains. Moreover, having three conflicting first-order domains is no better than having two. We should not resolve existential questions by counting the relevant domains and going with the majority,¹⁰ especially in the absence of non-*ad hoc* criteria for domain-individuation (as we saw above). More plausibly, such comparisons could be internal to a *second-order* domain. But if so, this re-introduces the very metaphysical vantage point, with its domain-independent standards for existence, that quietists were in the business of maligning.

Plus, even if quietists can legitimately make inter-domain comparisons without re-introducing a metaphysical vantage point, how would help this help with the case at hand?

⁹Scanlon confirmed that this was his view in personal correspondence: “in cases of inter-domain conflict, we need some reason for giving priority to one domain. The success of science in delivering systematic, accurate predictions gives us ample reason to give it priority.”

¹⁰I take it to be self-evident that we should not do metaphysics by plebiscite. But if one does wish to resolve existential questions by counting the relevant domains and going with the majority, they must resolve what Pettit (2001) has called “the discursive dilemma”. Say α contains p and q and $(p \wedge q)$, β contains p and $\neg q$ and $\neg(p \wedge q)$, and γ contains $\neg p$ and q and $\neg(p \wedge q)$. The majority of first-order domains yield an inconsistent set of commitments: to p and q and $\neg(p \wedge q)$. It is unclear what resources quietists would have to resolve this.

Say predictive accuracy justifies privileging science over the supernatural. What similar comparison would justify privileging morality over etiquette? Neither normative domain is in the business of predicting anything.

What if Scanlon can justify resolving inter-domain conflicts asymmetrically? The second problem still remains: there may be no conflict between etiquette and any other domain in the first place. This may sound surprising. Say Prudish Reason is true according to etiquette, and is currently not true according to any other normative domain. You might think it follows that Prudish Reason is false according to some other normative domain, like morality. But that's a mistake. Prudish Reason can be neither-true-nor-false according to other normative domains. This is plausible given Scanlon's stance on indeterminacy, which we touched on above. Scanlon takes it to be a fact, or at least a plausible conjecture, that the standards of normative domains provide determinate truth-values for some claims about reasons and indeterminate truth-values for others.

Admittedly, Scanlon's stance on indeterminacy is not easy to interpret.¹¹ He barely mentions familiar concerns about vague predicates.¹² His concern here is different: "the domain of practical reasons is not a unified subject matter", in that "our thinking about reasons depends on too many disparate starting points", so "the range of non-derived elements in the domain of reasons for action is too varied to be plausibly explained by a systematic overall account" (104, 101–102).

If quietism is true, there are at least two ways that this could generate indeterminacy. First, it may be that our standards contain one fixed set of discrete starting points that

¹¹First, it is unclear whether Scanlon is committed to indeterminacy. He starts his discussion claiming that determinacy is a "much less plausible hope" for normative domains than for the mathematical domain (85), and ends by taking it to be a *fact* "that the domain of practical reasons is not a unified subject matter like the domain of sets, the content of which we should expect to be determined by overall principles characterizing this domain" (104). Is determinacy an implausible hope, or is indeterminacy a fact? Secondly, Scanlon discusses the determinacy of mathematics at length, but does not discuss the determinacy of any other domain, including that of physical science. Thirdly, it is "utterly unclear" (as Wedgwood (2015) says, 5) why Scanlon takes indeterminacy about normativity be plausible, or a fact. More on this below.

¹²While "it seems clearly true that the fact that some action is necessary to avoid serious physical pain is, in most circumstances, a reason to do it", it can be "uncertain at the edges", in that there can be borderline cases of pain, or of pain that is not "worth bearing" (85–86).

yield determinate truth values for some claims and indeterminate truth values for others. As an illustrative analogy, consider a bureaucratic process that admits people who have 60 points of identification or more and rejects people who have 40 points of identification or less; those with 41–59 points of identification have an indeterminate status, but not because any particular source of identification (e.g., “driver’s license”) involves vagueness.

Secondly, our standards may not contain *one* fixed set of discrete starting points. This seems to match Scanlon’s concerns in the text: he accepts that it is possible that while impeccably reasoning about reasons, different people “may reach different reflective equilibria”, where “the best explanation of the plausibility of [their] different answers to ... particular questions is that they flow from different overall accounts of this subject matter, and that the subject in question can equally well be understood in either of these two ways” (79–80).

As an illustrative analogy, consider athletes. While impeccably reasoning about who or what is and is not an athlete different people may reach different reflective equilibria, where the best explanation of why they do so is that their incompatible answers to particular questions flow from different but equally good overall accounts or understandings of the subject matter. For instance, some take Secretariat, a racehorse, to be obviously one of the best athletes of all time; others take it to be equally obvious that racehorses are not athletes.¹³ Since all overall accounts of athletes still converge on some claims, those claims have determinate truth values: “Usain Bolt is an athlete” is true; “Rocks are athletes” is false. But since overall accounts sometimes diverge, claims like “Secretariat is an athlete” have indeterminate truth values. This indeterminacy does not result from a vague predicate, but from a subject matter that is not sufficiently unified to be best explained by one unique systematic overall account.

If this is how we should understand Scanlon’s view about indeterminacy, Prudish Reason is a particularly suitable candidate for a claim that has an indeterminate truth value according to normative domains. Many people accept as “a discrete starting point” a kind

¹³See Ludlow (2008) and Plunkett and Sundell (2013) for discussion of this example.

of moral puritanism about sex; they take there to be underived *moral* reasons for the unwed to be chaste. “Liberals” tend to reject this claim. But for many “conservatives”, it is foundational to their moral thought.¹⁴ Given this, it is plausible that while reasoning impeccably about morality, different people may reach different (“liberal” or “conservative”) reflective equilibria. The best explanation of the plausibility of their different answers (e.g., to whether according to morality the unwed have reason to be chaste) may be that they flow from different overall accounts of morality.

However, little hangs on this example, or on this understanding of Scanlon’s stance on indeterminacy. Whatever claims about reasons have indeterminate truth-values according to the standards of other normative domains, there could be a local standard of etiquette according to which only those claims are true. After all, we can create such standards, so we can gerrymander their contents. So if there is indeterminacy in normative domains, we can create an etiquette domain (via mere changes in our attitudes) that licenses new claims about reasons without entailing claims that are false according to other normative domain. So we can create new reasons without creating inter-domain conflicts.

There are several ways that one could get off the boat here. One could deny the indeterminacy of normativity. It is unclear if the view that would emerge is recognizably quietist. But if it is, the view would be vulnerable on other fronts. Scanlon appeals to the indeterminacy of normativity in fending off various non-quietist views which are inspired by Hume and Kant: in claiming that practical reasons “depend on too many disparate starting points” to be a unified subject matter, Scanlon is denying that we can explain practical reasons in terms of a single starting point like actual or ideal desires, or rational agency. So parting ways with Scanlon on the indeterminacy of normativity is far from cost-free.

Alternatively, one could revise CONSISTENCY. Recall that CONSISTENCY holds that domains α and β conflict when p is true in α and false in β . We could go for something weaker: domains α and β conflict when p is true in α and not true in β . Now etiquette

¹⁴See Jonathan Haidt (2012) for evidence for and discussion of these empirical claims.

conflicts with other normative domains. But this makes inter-domain conflicts ubiquitous. Just take any claim about reasons that is true according to normative domains but neither-true-nor-false according to science. (The standards of science are silent about reasons. So statements about reasons are not true according to science, and they had better not be false either.) This is a very unwelcome result for Scanlon, who claims that there is in fact no tension between scientific and normative discourse.

Could we revise Scanlon's view in order to generate a conflict between etiquette and morality without generating a conflict between morality and science? Perhaps one could restrict the scope of this revised CONSISTENCY principle such that it only applied to normative domains. This may not be *ad hoc*: perhaps the conditions for conflict between indeterminate domains should be weaker. On this view, a statement about reasons is true only if it is true according to the standard of every normative domain. Licensing claims about reasons is a bit like launching missiles from a submarine: everyone needs to turn their key. This is sufficient to generate a conflict between etiquette and other normative domains. And it reduces unwanted inter-domain conflicts. But it does not eliminate them entirely. Plausibly, epistemology is silent about moral reasons, and morality is silent about epistemic reasons. So even a revised and restricted CONSISTENCY constraint cannot generate the needed conflict between etiquette and morality without generating unwelcome conflicts between epistemology and morality.

A final possibility is more promising. The quietist could concede that the claims that etiquette contains and entails are not false according to other domains, but argue that etiquette still *presupposes* claims that are false according to other normative domains. This would generate an inter-domain conflict.

Scanlon offered this response in personal correspondence. He suggested that morality and etiquette have similar normative presuppositions, and that those presuppositions are true in the former case and false in the latter. So to understand the response, let's consider the normative presuppositions of morality:

The claims that we make about moral right and wrong generally presuppose *that there are moral standards that everyone has good reason to take seriously as guides to conduct and as standards for objecting to what others do*. But the ordinary ways of understanding morality, and ordinary ways of arguing for moral conclusions do not make clear what these reasons are, or establish that we have such reasons. There is therefore a question, external to morality, *whether the usual ways of establishing that a form of conduct is wrong also guarantee that there are good reasons not to engage in it*. This question is not scientific or metaphysical but normative—a question about what we have reason to do (22, emphasis mine).

I take the basic idea behind this response to be that all normative domains—morality, etiquette, epistemology, etc.—have such normative presuppositions. Does this provide a basis for generating the needed conflict between etiquette and some other domain without generating unwelcome inter-domain conflicts?

No. These normative presuppositions are either overly demanding, or trivial.

Scanlon offers two quite different formulations of the relevant normative presuppositions (which I have emphasized above). Roughly, they correspond to the following: *that everyone has good reason to take Tithing Norm seriously as a guide to conduct*; and *that the usual ways of establishing that Tithing Norm is true guarantee that there are reasons to tithe*. Let's consider them in turn.

The first formulation is overly demanding. It may be false that everyone has good reason to take norms of etiquette seriously as guides to conduct. However, if morality is esoteric, it is also false that everyone *has* good reason to take moral norms seriously as guides to conduct. And if morality is self-effacing, there may not *be* good reason for everyone to be guided by what morality requires. So it is too demanding to require that all normative domains—morality, etiquette, etc.—are such that everyone has good reason to be guided by their dictates. If Scanlon takes this route, his quietism threatens to herald moral nihilism.

What about the second formulation? On a natural interpretation, it is trivial. Do the usual ways of establishing that Tithing Norm is true guarantee that there is reason to tithe? Sure. Our ways of establishing that Tithing Norm is true (the standards of the moral domain) guarantee that Tithing Reason is true too. But by parity of reasoning, our ways of establishing that Prudish Norm is true (the standards of the etiquette domain) also guarantee that Prudish Reason is true, and hence that there is reason to do what etiquette requires.

Perhaps this natural interpretation is mistaken. After all, Scanlon says that morality's presuppositions are "external", so perhaps they must be made true in a domain other than morality (and *mutatis mutandis* for etiquette et al.). This makes the presupposition non-trivial. But it also has the surprising consequence that moral nihilism is true, and all moral claims like Tithing Norm are false, unless it is guaranteed that there are non-moral reasons to be moral. This is a very controversial, even counter-orthodox, position. Since Prichard (1912), many have thought that it is a mistake to demand that there must be some non-moral source of reasons to act in accordance with the requirements of morality.

Moreover, the relevant presupposition is now far too demanding, even for Scanlon. Scanlon (1998) argued that if his contractualism is true, it is guaranteed that there are non-moral reasons to do what "we owe to each other". But he also indicated that "what we owe to each other" may not exhaust morality; it may not include our moral requirements regarding animals and the natural environment (179). So it is an open question whether, even for Scanlon, our usual ways of establishing that conduct that mistreats animals or the environment is wrong guarantee that there are "good (non-moral reasons not to engage in it."

Another way to see why this presupposition is too demanding is to consider other normative domains, like prudence and epistemology. To undermine my prospects is imprudent even if the usual ways of establishing this do not guarantee that I have non-prudential reasons to act prudently. To believe contrary to my evidence is unjustified even if the usual ways of establishing this do not guarantee that I have non-epistemic reasons to form justified beliefs.

The upshot of all of this is that if quietists are to insist that it is an external presupposition of any normative domain that we are guaranteed to have good reason to act in accordance with the dictates of that domain, then quietism opens the door for nihilism about morality, prudence and epistemology. And if quietists do not insist that morality et al. have such external presuppositions, then by parity of reasoning they must say the same about etiquette et al. too.

Where does this leave us? We first saw that if Scanlon's quietism is true, discourse about etiquette forms a separate domain, so we can create new reasons out of thin air. To resist this conclusion the quietist must respond that etiquette conflicts with other normative domains—it contains, entails, or presupposes claims that are false according to other normative domains—and this conflict provides decisive reason to reject the etiquette domain. However, neither step in this response succeeds. The only ways to generate a conflict between etiquette and other normative domains massively over-generate inter-domain conflicts. And the quietist has no resources to explain why symmetric inter-domain conflicts should be resolved asymmetrically by rejecting etiquette. So we're left with the conclusion that if quietism is true we can create new reasons out of thin air.

6.5 One Domain, One Standard

The first argument rests on a way of individuating domains that is especially plausible given Scanlon's stance on indeterminacy and his commitment to the separateness of the particular domains he discusses (such as the scientific and supernatural domains, and the normative and moral domains). But let's assume, *ex hypothesi*, that one supplemented Scanlon's view with clear criteria for domain-individuation which justified the stance that etiquette discourse does not form a separate domain, but instead falls under one and only normative domain. The upshot of this is that claims like Prudish Reason must be licensed by the standards of the one normative domain in order to be true.¹⁵

¹⁵This mirrors Scanlon's response (29) to Enoch (2011a)'s "counter-reasons" example (124).

Does it follow from this that Prudish Reason is not true? Or, even more pertinently, does it follow that we cannot make Prudish Reason true via mere changes in our attitudes? I don't think so. In this section I'll argue that if Scanlon's quietism is true and discourse about etiquette is subsumed under the one normative domain, we can still create new reasons out of thin air.

This argument rests on what is perhaps the most elusive detail of Scanlon's view: the revisability of standards. Let's warm up to the basic idea by considering scientific and supernatural discourse once more. Recall that they can contain inconsistent claims about causation, such as "Bacteria caused the illness" and "Witches caused the illness", where either cause would exclude the other. If these discourses form separate domains, they conflict. But what if they are subsumed under a single domain, such as the domain of causation? We cannot appeal to CONSISTENCY. So we must determine whether it's true that "Witches caused the illness" by appealing to the standards of the causation domain. And this is where things get interesting. It would be a mistake to assume that the standards of the broader causation domain would be identical to the standards of the narrower scientific domain. After all, a domain is just a set of statements with a single subject matter, and the standards for a domain are determined in part by the "obvious particular truths" within that set of statements. So the broader set will contain obvious particular truths that are absent from the narrower set, and hence the broader domain may well have a different standard than the narrower domain. I am not sure if that would occur here, as there is a significant tension between the candidate "obvious particular truths" within the broader set of statements, and that tension would need to be resolved.

Still, this example illustrates how the content of the standard for a domain may change depending on what statements are subsumed under that domain. The example above illustrates this point synchronically: we considered how we could carve out a broader or narrower domain at a single point in time. But we could make the same point diachronically by considering what happens when we add to the relevant set of statements over time. By adding

to the “obvious particular truths” we could revise the relevant standard, and thereby change what statements are true in the domain. If we could make “Witches caused the illness” true in such a manner, we could create witches out of thin air.

Now let’s return to the normative domain. Say that discourse about morality, prudence, and epistemology all falls under the normative domain. Then discourse about etiquette—statements like Prudish Norm and Prudish Reason—emerges and is subsumed under that domain too. This adds new “obvious particular truths” to the domain. And unlike in the example of supernatural discourse, it does not introduce any tension between the obvious particular truths in the domain: these “obvious particular truths” were previously neither-true-nor-false, not false. So it seems probable that by adding a set of statements like Prudish Reason to the normative domain we will revise the standards of the latter so that they license the former. In other words, by subsuming etiquette discourse under the one normative domain we can still create new reasons out of thin air. If this is right, realists must still reject normative quietism.

Quietist could try to resist this argument by rejecting either of its two crucial steps: that we can change what statements are true by revising the standards of a domain; and that adding statements like Prudish Reason to the normative domain would prompt such a revision. Let’s consider each in turn.

According to Scanlon’s quietism, can we change what statements are true by revising the relevant domain-specific standards? I believe so. We saw that, for Scanlon, domain-specific standards determine the truth-values of statements. From when he first discusses standards Scanlon notes that they can be “revised” (20). He even states that “nothing beyond the (*evolving*) standards of a domain can be relevant to the *truth* of statements within it” (21, emphasis mine).

It is surprising that Scanlon does not discuss the controversial implications that seem to follow from this view. The following, for instance, is at least possible on Scanlon’s view: scientific standards “evolved” between the ancient, medieval, and modern eras, such that

quantum field theory was false (or indeterminate) but is now true, while Humorism was true but is now false.¹⁶ This commits Scanlon to a form of relativism about truth. Worse yet, it commits Scanlon to a form of relativism about existence. Something exists iff it is quantified over by true statements. So entities like phlogiston and quantum particles can pop in and out of existence with “evolutions” in our scientific standards.

These implications seem to be absurd and paradigmatically anti-realist; and there is no evidence that they were explicitly accepted by Scanlon. So perhaps we should charitably reinterpret or revise Scanlon’s stance on the revisability of standards. Recall that Scanlon accepts that standards play epistemic and ontological roles. We could distinguish between two types of standards—truth-value-determining ontological standards and justification-determining epistemic standards—one of which is immutable, and the other revisable.¹⁷ So revised, Scanlon’s view could yield the more intuitive verdict that phlogiston never existed: claims like “Phlogiston caused the fire” were always false (according to the relevant fixed ontological standards), but at times we were justified in believing these falsehoods (according to the relevant fluctuating epistemic standards).

This proposal might seem promising, but it fits neither the letter nor the spirit of Scanlon’s view. If ontological standards are immutable and epistemic standards are revisable, the two can come apart. This raises many of the troublesome epistemological problems that quietists like Scanlon want to dissolve rather than solve. How we could know that what we’re justified in believing to be true according to our current, changeable epistemic standards—for math, normativity, science, or any domain—conforms to what’s actually true according to the immutable ontological standards? The quietist fails at her own game if she opens the door for skepticism about epistemic access to the truth.

Moreover, the proposal that ontological standards are immutable but epistemic standards about normativity are revisable is still problematic for realists about normativity. It follows

¹⁶Humorism is the historically influential view that physical and mental qualities are caused by the proportions of the four cardinal humors: blood, phlegm, yellow bile, and black bile.

¹⁷I thank Tristram McPherson for this suggestion. It is, to be clear, not Scanlon’s view.

from this proposal that by adding statements like Prudish Norm and Prudish Reason to the normative domain we cannot make Prudish Reason true, but we can make everyone justified in believing that Prudish Reason is true. This is a smaller bullet to bite, but it still leaves a bitter aftertaste. Even the “liberals” in a society with “conservative” social mores would become justified in falsely believing that there are reasons for the unwed to be chaste. So this is a view on which we cannot create *practical reasons* by changing our attitudes, but we can we create epistemic justification (i.e., *epistemic reasons*) to believe that there are practical reasons when there aren’t, merely by changing our discourse and our attitudes. This still leaves our central conclusion in tact: if quietism is true, we can create reasons out of thin air.

Since it does not seem to be very promising, let’s set aside the first response and consider the second. For Scanlon, we can change what statements are true by revising standards. But would subsuming etiquette discourse under the normative domain prompt such a revision and make Prudish Reason true?

I don’t know if we can reach a decisive conclusion about this issue either way, given how little Scanlon says about the nature of standards, or their revisability. But based on the little we know from the text, I think it’s more plausible that subsuming etiquette discourse would prompt such a revision. We know that standards are not revised due to changes in our linguistic practices alone (19 n.3). More helpfully, we know that domain-specific standards are set by reasoning that appeals to “our best general understanding of the subject matter in question” and “the most obvious particular truths” within the domain (20). We know that according to Scanlon our best general understanding of reasons (the subject matter in question) is one on which there are many “disparate starting points” for reasoning impeccably about reasons. And we know that in a sufficiently conservative community claims like “According to etiquette there is reason for the unwed to be chaste” can be at least as obvious as truths like “According to morality there is reason for the wealthy to give to charity”.

If we subsume etiquette discourse under the one normative domain, we should treat etiquette as another disparate starting point for reasoning impeccably about reasons (along with morality, prudence, and so on). And from that starting point we can reach claims like Prudish Reason, which will be at least as obvious as truths reached from other starting points, like Tithing Reason. This is a strong *prima facie* case for thinking that adding such starting points could prompt a revision, and it need not appeal to our linguistic practices at all.

This case is strengthened by considering an analogy to other points in history where we added new statements to our broader normative discourse. Discourse about distributive justice—including, perhaps, statements like Tithing Reason—is relatively recent. Presumably, distributive justice became another disparate starting point for reasoning impeccably about reasons, which could revise the standards for the normative domain to license different claims about reasons. However it is that this worked for distributive justice and Tithing Reason, it should work the same way for social propriety and Prudish Reason.

We've now seen two arguments for the conclusion that if quietism is true, we can create new reasons out of thin air, which is incompatible with realism. The arguments depend on different ways of individuating domains. So to block the objection, the quietist needs to show that *either* both arguments fail, *or* one argument fails and the other rests on an unacceptable way of individuating domains. If one succeeds and both ways of individuating domains are acceptable, quietism makes our ontological commitments depend on arbitrary choices.

Several challenges to Scanlon's quietist commitments about the nature and role of domain-specific standards, inter-domain conflict, indeterminacy, and revisability have also been considered. Even if both arguments fail, I believe these challenges still provide ample grounds for realists to reject normative quietism.

6.6 Conclusion

The debate over whether realism is compatible with quietism has stagnated because of the realist's difficulties in finding a point of purchase against quietists: Scanlon gladly rejects the metaphysical distinction between what "exists" and what "EXISTS", and cavalierly accepts the implication we can make entities like ghosts and schmeasons exist. But Scanlon cannot maintain such nonchalance about quietism's intensional and extensional implications about reasons. So if Scanlon's quietism entails that we can create new reasons out of thin air, Scanlon's own realist commitments about reasons militate against quietism.

Two final questions are worth considering. What is the source of the problem for Scanlon's quietism? And is it a problem for quietism, or just a problem for Scanlon's version thereof? The two problems are obviously related. Diagnosing the source of the problem will reveal whether its roots are deep or superficial.

We can diagnose the source of the problem for quietists in two ways. First, we can consider how the problem stems from what is present in the account: the initially plausible ideas that lie at the core of Scanlon's view, DIVERSITY and AUTONOMY, allow us to create new unfamiliar entities like ghosts by adopting a new way of talking about them. I've shown that the same holds for familiar entities like reasons.

The same may also hold for other familiar entities like numbers and quarks. This depends on whether the indeterminacy of the normative domain is central to our diagnosis of the source of the problem. I suspect that the problem remains even if we are committed to determinacy, due to the independent problems with the CONSISTENCY constraint: quietists need a principled basis for resolving inter-domain conflicts asymmetrically, but none is available. Even if I'm wrong about this, it is cold comfort to quietists. That their view yields such unwanted results about reasons should make realists reject the general enterprise.

If this is right, the roots of the problem run deep, and affect quietism writ large. To avoid the objection, quietists will have to get off the boat early. Perhaps they will abandon DIVERSITY and AUTONOMY entirely, in which case I confess that I'm not sure what lies at

the core of quietism, and why it is even initially attractive. Perhaps they will restrict the view to a handful of privileged domains—like normativity, mathematics, and physical science. But that will be awfully *ad hoc*. As we’ve seen, it is hard to justify privileging some first-order domains over others without appealing to comparisons in some general second-order domain (like metaphysics) that quietists are in the business of maligning. Or, perhaps quietists will abandon any aspiration to provide a general view. Perhaps they will accept a general, metaphysically inflationary account of why quarks exist, and a metaphysically deflationary account of why reasons exist. However, if reasons are held to be less real than quarks, it is hard to see how normative quietism amounts to more than normative anti-realism by another name.

Second, we can consider how the problem stems from what is absent in the account: an appeal to the metaphysical status of morality. This is a long-standing commitment of Scanlon’s. In the opening pages of *What We Owe To Each Other*, he wrote that

it does not seem that the reason-giving force of facts about right and wrong derives from their metaphysical status. This is shown by the fact that it is not clear how an account of this status ... would, simply in virtue of the “reality” it would give to the subject matter of judgments of right and wrong, bolster their reason-giving force.

If what I’ve argued in Chapters Two and Three is true, this is wrong-headed. Our discourse about morality and etiquette may both involve truths about reasons, but the former involves truths about reasons that exist while the latter involves truths about reasons that exist according to a fiction. This provides a clear account of how the metaphysical status of reality and etiquette—the “reality” of the subject matter of judgments about what is right and wrong, as opposed to judgments about what is decorous and indecorous—explain why morality has (and etiquette lacks) authority, or “reason-giving force”. It also provides a clear account of how when we create new norms of etiquette we create new requirements without

necessarily creating new reasons. We only need to change the reasons that exist according to a fiction.

In that sense, it seems that we can explain the normative asymmetry between morality and etiquette in terms of the metaphysical distinction between reality and fiction. But this explanation is unavailable to Scanlon because of his deflationary metaphysical commitments. For Scanlon, there is no domain-independent characterization of reality, but only a “merely disjunctive universe”; what’s true “according to etiquette” cannot fail to represent a domain-independent reality in the way that most of us would say that what’s true according to *Sherlock* fails to represent some domain-independent reality. This diagnosis may sound surprising. After all, Scanlon distances his view from fictionalism—“a practice of merely speaking as if there were such things” (24). But he never offers or even appeals to a distinction between reality and fiction, including in the place where one would most expect him to: even when considering “fictional characters” like Sherlock, Scanlon suggests that they do not exist because they “lack some important property that might be claimed for them because they do not correspond to parts of the (causal) world” (24, n. 9), rather than because discourse about Sherlock is embedded under prefixes that cancel assertoric force (“according to *Sherlock*”). Even here, Scanlon implicitly settles the ontological implications of quietism only by using its distinctive resources: domains and standards. Implicitly, Scanlon takes Sherlock to lack some important property according to science, rather than according to *Sherlock*.¹⁸

These two diagnoses are complementary, not competing. It is because he eschews any appeal to the metaphysical status of morality that Scanlon must appeal instead to the initially attractive resources of domains and standards. And it is because these resources are not fit for the task at hand that Scanlon ends up not being able to vindicate his own realist commitments about reasons.

¹⁸Scanlon was kind enough to confirm this in personal correspondence: in the absence of inter-domain conflicts, “there is never any ontological commitment to distance oneself from”.

Bibliography

Recent case: Civil Disobedience - The Role of Judges - Ninth Circuit Affirms Mandatory Sentence. - United States v. Hungerford, 465 f.3d 1113 (9th cir. 2006). Harvard Law Review, 120(7), 2007.

Michelle Alexander. The New Jim Crow. The New Press, 2010.

Elizabeth Anderson. Beyond homo economicus: New developments in theories of social norms. Philosophy and Public Affairs, 29(2):170–200, 2000.

G. E. M. Anscombe. Modern Moral Philosophy. Philosophy, 33(124), 1958.

Arthur Isak Applbaum. Ethics for Adversaries: The Morality of Roles in Public and Professional Life. Princeton University Press, 1999.

P.S. Atiyah and R.S. Summers. Form and Substance in Anglo-American Law: A Comparative Study in Legal Reasoning, Legal Theory, and Legal Institutions. Clarendon Press, 1987.

J.L. Austin. Sense and Sensibilia. Oxford University Press, 2nd edition, 1962.

D. Bar-On and James Sias. Varieties of expressivism. Philosophy Compass, 8(8):699–713, 2013.

Chris Barker. The dynamics of vagueness. Linguistics and Philosophy, 25:1–36, 2002.

Elizabeth Barnes. Going beyond the fundamental: Feminism in contemporary metaphysics. Proceedings of the Aristotelian Society, 64, 2014.

- Roland Barthes. Image / Music / Text, chapter Death of the Author, pages 142–147. Hill and Wang, 1977.
- Allen Beck and et al. Sexual Victimization in Prisons and Jails Reported by Inmates, 2011-2012, 2013.
- Benjamin Beit-Hallahmi, John C. Catford, Ralph E. Cooley, Cecelia Yoder Dull, Alexander Z. Guiora, and Maria Raluszny. Grammatical gender and gender identity development: Cross-cultural and cross-lingual implications. American Journal of Orthopsychiatry, 44 (3):424–431, 1974.
- Jeremy Bentham. Of Laws in General. Athlone Press, 1782/1970.
- Brian Bix. Legal positivism and “explaining” normativity and authority. APA Newsletter on Philosophy and Law, 5(2), 2006.
- H. Brian Bix. Problems of Normativity, Rules and Rule-Following, chapter Rules and Normativity in Law, pages 125–146. Springer International Publishing, 2015.
- Gunnar Björnsson and Tristram McPherson. Moral attitudes for non-cognitivists: Solving the specification problem. Mind, 123(489), 2014.
- Simon Blackburn. Spreading The Word. Clarendon, Oxford, 1984.
- Simon Blackburn. Essays in Quasi-Realism. Oxford University Press, Oxford, 1993.
- Simon Blackburn. Ruling Passions. Oxford University Press, Oxford, 1999a.
- Simon Blackburn. Think: A compelling introduction to philosophy. Oxford University Press, 1999b.
- Jeffrey Brand-Ballard. Limits of Legality: The Ethics of Lawless Judging. Oxford University Press, 2010.
- Michael Bratman. Practical reasoning and acceptance in a context. Mind, 101:1–15, 1992.

- Geoffrey Brennan, Lina Eriksson, Robert E. Goodin, and Nicholas Southwood. Explaining Norms. Oxford University Press, 2013.
- Stuart Brock. Fictionalism about fictional characters. Noûs, 36(1):1–21, 2002.
- Stuart Brock. The creationist fiction: the case against creationism about fictional characters. Philosophical Review, 119:337–364, 2010.
- John Broome. Rationality Through Reasoning. Wiley, 2013.
- Angela Browne and et al. Prisons Within Prisons: The Use of Segregation in the United States. Federal Sentencing Reporter, 24(1), 2011.
- Sarah Buss. Appearing respectful: The moral significance of manners. Ethics, 109:795–826, 1999.
- Paul Butler. When judges lie (and when they should). Minnesota Law Review, 91, 2007.
- Cheshire Calhoun. The virtue of civility. Philosophy and Public Affairs, 29(3):251–275, 2000.
- M. J. Carter. Gladiatorial combat: The rules of engagement. The Classical Journal, 102(2): 97–114, 2006.
- Paul Cassell. Statement of Judge Paul Cassell on Behalf of the Judicial Conference of the United States (June 2007). Federal Sentencing Reporter, 19(5), 2007.
- Paul Cassell and Erik Luna. Mandatory minimalism. Cardozo Law Review, 32(1), 2010.
- Paul G. Cassell. Too severe? A defense of the Federal Sentencing Guidelines (and a critique of the federal mandatory minimums). Stanford Law Review, 56, 2004.
- Paul G. Cassell and Erik Luna. Sense and sensibility in mandatory minimum sentencing. Federal Sentencing Reporter, 2011.

- Richard Yetter Chappell. Fittingness: The sole normative primitive. The Philosophical Quarterly, 62(249):684–704, 2012.
- Matthew Chrisman. On the meaning of “ought”. Oxford Studies in Metaethics, 7, 2007.
- Matthew Chrisman. The Meaning of “Ought”: Beyond Descriptivism and Expressivism in Metaethics. Oxford University Press, 2015.
- Ted Cohen. Baseball and Philosophy, chapter There are No Ties at First Base. Carus-Open Court, 2004.
- Jules L. Coleman. The architecture of jurisprudence. Yale Law Journal, 2(121), 2011.
- Jules L. Coleman and Brian Leiter. A Companion to Philosophy of Law and Legal Theory, chapter Legal Positivism. Blackwell, 1996.
- David Copp. Normativity and naturalism, chapter Moral naturalism and three grades of normativity. Ontos-Verlag, 2004.
- David Copp. Toward a pluralistic and teleological theory of normativity. Philosophical Issues, 19(1):21–37, 2009.
- Robert Cover. Justice Accused: Antislavery and the Judicial Process. Yale University Press, 1975.
- Robert M. Cover. Book review: Atrocious judges: Lives of judges infamous as tools of tyrants and instruments of oppression. Columbia Law Review, 1968.
- Terence Cuneo and Sean Christy. New Waves in Metaethics, chapter The Myth of Moral Fictionalism. Palgrave Macmillan, 2010.
- Chris Daly. Fictionalism and the attitudes. Philosophical Studies, pages 423–40, 2008.
- Janice Dowell. The metaethical insignificance of moral twin earth. Oxford Studies in Metaethics, Forthcoming.

- James Dreier. Meta-ethics and the problem of creeping minimalism. Philosophical Perspectives, 18:23–44, 2004.
- James Dreier. Relativism (and expressivism) and the problem of disagreement. Philosophical Perspectives, 23(1):79–110, 2009.
- Antony Duff. Legal Punishment. The Stanford Encyclopedia of Philosophy, 2013.
- R. A. Duff. Punishment, Communication, and Community. Oxford University Press, 2003.
- Ronald Dworkin. Taking Rights Seriously. Harvard University Press, 1977.
- Ronald Dworkin. Law’s Empire. Belknap Press, 1986.
- Ronald Dworkin. Objectivity and truth: You’d better believe it. Philosophy and Public Affairs, 25(2):87–139, 1996.
- Ronald Dworkin. Justice in Robes. Belknap Press, 2008.
- Ronald Dworkin. Justice for Hedgehogs. Belknap Press, 2011.
- Matti Eklund. Fictionalism. The Stanford Encyclopedia of Philosophy, 2011.
- David Enoch. Taking Morality Seriously: A Defense of Robust Realism. Oxford University Press, 2011a.
- David Enoch. Oxford Studies in Philosophy of Law, chapter Reason-Giving and the Law. Oxford University Press, 2011b.
- D. Estlund. Democratic authority. Princeton University Press, 2007.
- Ward Farnsworth. ‘To Do a Great Right, Do a Little Wrong’: A User’s Guide to Judicial Lawlessness. Minnesota Law Review, 86, 2001.
- Joel Feinberg. Doing and Deserving, chapter The Expressive Function of Punishment. Princeton University Press, 1970.

- Joel Feinberg. Civil disobedience in the modern world. Humanities in Society, 2(1):37–60, 1979.
- Vanda Felbab-Brown. Focused Deterrence, Selective Targeting, Drug Trafficking and Organized Crime: Concepts and Practicalities. International Drug Policy Consortium, February 2013.
- Hartry Field. Apriority as an evaluative notion. In Paul Boghossian and Christopher Peacocke, editors, New Essays on the A Priori, pages 117–149. Oxford University Press, Oxford, 2000.
- Stephen Finlay. Confusion of Tongues: A Theory of Normative Language. Oxford University Press, 2014.
- John Finnis. Natural law and natural rights. Oxford University Press, 1989.
- John Finnis. The Oxford Handbook of Jurisprudence and Philosophy of Law, chapter Natural Law: The Classical Tradition. Oxford University Press, 2002.
- Phillipa Foot. Morality as a system of hypothetical imperatives. Philosophical Review, 81(3):305–316, 1972.
- Phillipa Foot. Recantation, page 322. Moral Discourse and Practice, 1997.
- Lon L. Fuller. Legal Fictions. Stanford University Press, 1967.
- Lon L. Fuller. The Morality of Law. Yale University Press, 1969.
- John Gardner. Complicity and causality. Criminal Law and Philosophy, 2007.
- John Gardner. Law as a Leap of Faith: Essays on Law in General, chapter How Law Claims, What Law Claims. Oxford University Press, 2012.
- D. Gauthier. Morals by Agreement. Oxford University Press, 1986.

- Tamar Gendler. The puzzle of imaginative resistance. The Journal of Philosophy, 97(2): 55–81, 2000.
- Allan Gibbard. Wise Choices, Apt Feelings. Harvard University Press, Cambridge, Mass., 1990.
- Allan Gibbard. Thinking How to Live. Harvard University Press, 2003.
- Allan Gibbard. Meaning and Normativity. Oxford University Press, 2012.
- Allan Gibbard. Passions and Projections: Themes from the Philosophy of Simon Blackburn, chapter Improving Sensibilities. Oxford University Press, 2015.
- Alvin Goldman. What is democracy (and what is its raison d’etre)? Journal of the American Philosophical Association, 1(2):233–256, 2015.
- Jeffrey Denys Goldsworthy. The Limits of Judicial Fidelity to Law: The Coxford Lecture. Canadian Journal of Law and Jurisprudence, 24(1), 2011.
- Leslie Green. The Oxford Handbook of Jurisprudence and Philosophy of Law, chapter Law and Obligation. Oxford University Press, 2002.
- Kent Greenawalt. Conflicts of Law and Morality. Clarendon Press, 1987.
- Mark Greenberg. The moral impact theory of law. Yale Law Journal, 2014.
- H.P. Grice. Logic and conversation. In Peter Cole and Jerry Morgan, editors, Syntax and Semantics, Volume 3: Speech Acts, pages 41–58. Academic Press, New York, 1975.
- H.P. Grice. Studies in the Way of Words. Harvard University Press, Cambridge, Mass., 1989.
- Paul Grice. Meaning. Philosophical Review, 66, 1957.

- Jonathan Haidt. The Righteous Mind: Why Good People are Divided by Politics and Religion. Pantheon Books, 2012.
- Richard Mervyn Hare. The Language of Morals. Clarendon Press, 1952.
- R.M. Hare. Moral Thinking: Its Levels, Method and Point. Oxford University Press, 1981.
- Gil Harman. Review of The Significance of Sense: Meaning, Modality, and Morality. The Philosophical Review, 82(2):235–239, 1973.
- Ross Harrison. Jeremy Bentham. Routledge & Kegan Paul, 1983.
- H.L.A. Hart. Essays on Bentham: Jurisprudence and Political Theory. Oxford University Press, 1982.
- H.L.A. Hart. The Concept of Law. Clarendon Press, 2nd edition, 1994.
- Richard Heck. Do demonstratives have senses? Philosophers' Imprint, 2(2):1–33, 2001.
- Scott Hershovitz. The end of jurisprudence. Yale Law Journal, 124, 2015.
- Ulrike Heuer. Promising - part 2. Philosophy Compass, 7(12):842–851, 2012.
- Thomas E. Hill. Symbolic protest and calculated silence. Philosophy and Public Affairs, 9(1):83–102, 1979.
- Richard Holton. Positivism and the internal point of view. Law and Philosophy, 17(5/6), 1998.
- T. Horgan and M. Timmons. New wave moral realism meets moral twin earth. Journal of Philosophical Research, 16:447–65, 1991.
- John Horty. Rules and reasons in the theory of precedent. Legal Theory, 17:1–33, 2011.
- John Horty. Reasons as Defaults. Oxford University Press, 2012.

An Offer You Can't Refuse: How US Federal Prosecutors Force Drug Defendants to Plead Guilty. Human Rights Watch, 2003.

David Hume. A Treatise of Human Nature. Clarendon, Oxford, 2nd edition, 1739. First edition by Selby-Bigge published in 1888.

David Hume. Essays: Moral, Political and Literary, chapter Of the Standard Taste, pages 226–49. Liberty Fund, 1757/1985.

Richard Joyce. The Myth of Morality. Cambridge University Press, 2001.

Richard Joyce. Fictionalism in Metaphysics, chapter Moral Fictionalism. Oxford University Press, 2005.

Richard Joyce. The Evolution of Morality. MIT Press, 2006.

Dan Kahan. Gentle nudges vs. hard shoves: Solving the sticky norms problem. Yale Law Journal, 2000.

David Kaplan. Demonstratives. In Joseph Almog, John Perry, and Howard Wettstein, editors, Themes from Kaplan. Oxford University Press, New York, 1989.

Jeffrey C. King and Jason Stanley. Semantics vs. Pragmatics, chapter Semantics, Pragmatics, and The Role of Semantic Content. Oxford University Press, 2005.

Christine Korsgaard. The Sources of Normativity. Cambridge University Press, 1996.

Matthew Kramer. Requirements, Reasons, and Raz: Legal positivism and legal duties. Ethics, 109, 1999.

Matthew Kramer. Rights, Wrongs, and Responsibilities, chapter Getting rights right, pages 28–95. Palgrave, 2001.

Angelika Kratzer. What 'must' and 'can' must and can mean. Linguistics and Philosophy, 1:337–355, 1977.

- Angelika Kratzer. Words, Worlds, and Contexts: New Approaches in Word Semantics. Eds. . B, chapter The Notional Category of Modality, pages 38–74. W. de Gruyter, 1981.
- Angelika Kratzer. Semantics: An International Handbook of Contemporary Research, chapter Modality, pages 639–50. W. de Gruyter, 1991.
- Saul Kripke. Reference and Existence. Oxford University Press, 2013. Unpublished manuscript; delivered as the Locke Lectures in 1973.
- Melissa Lane. States of nature, epistemic and political. Proceedings of the Aristotelian Society, 99(2):211–224, 1999.
- Chiara Lepora and Robert E. Goodin. On Complicity and Compromise. Oxford University Press, 2013.
- David Lewis. Counterfactuals. Blackwell, Oxford, 1973.
- David Lewis. Languages and language. In Keith Gunderson, editor, Language, Mind and Knowledge, volume 7 of Minnesota Studies in the Philosophy of Science, pages 3–35. University of Minnesota Press, Minneapolis, 1975.
- David Lewis. Truth in fiction. American Philosophical Quarterly, 15:37–46, 1978.
- James W. Loewen. The Mississippi Chinese: Between Black and White. Waveland Press, 1988.
- James Logue. Projective Probability. Oxford University Press, 1995.
- David Luban. Lawyers and Justice: an Ethical Study. Princeton University Press, 1988.
- David Luban. Freedom and constraint in legal ethics: Some mid-course corrections to lawyers and justice. Maryland Law Review, 1990.
- David Luban. Misplaced fidelity. Texas Law Review, 2012.

- Peter Ludlow. Cheap contextualism. Philosophical Issues, 18:104–129, 2008.
- J.D. Mabbott. Punishment. Mind, 48(152-167), 1939.
- Neil MacCormick. Legal Reasoning and Legal Theory. Clarendon Press, 1978.
- J. L. Mackie. Ethics: Inventing Right and Wrong. Penguin, London, 1977.
- B. Maguire. Grounding the autonomy of ethics. Oxford Studies in Metaethics, forthcoming.
- Barry Maguire. T.M. Scanlon’s Being Realistic About Reasons. Times Literary Supplement, 2014.
- Kate Manne. On being social in metaethics. Oxford Studies in Metaethics, 8:50–73, 2013.
- Andrei Marmor. Philosophy of Law. Princeton University Press, 2011.
- Judith Martin and Gunther S. Stent. I think; therefore I thank: A philosophy of etiquette. The American Scholar, 59(2):237–254, 1990.
- Jeff McMahan. Killing in War. Oxford University Press, 2009.
- Tristram McPherson. Against quietist normative realism. Philosophical Studies, 154:223–240, 2011.
- Stanley Milgram. Behavioral study of obedience. Journal of Abnormal and Social Psychology, 67(4):371–8, 1963.
- Stanley Milgram. The Individual in a Social World: Essays and Experiments. Addison-Wesley, 1977.
- G.E. Moore. Principia Ethica. Cambridge University Press, 1903.
- T Nagel. The Sources of Normativity, chapter Universality and the Reflective Self, pages 200–209. Cambridge University Press, 1996.

- T. Nagel. The Last Word. Oxford University Press, 1997.
- Frederick Nietzsche. Beyond Good and Evil. Random House, 1882/1974a.
- Frederick Nietzsche. The Portable Nietzsche, chapter Thus Spoke Zarathustra, pages 103–439. Viking Penguin, 1883/1982.
- Friedrich Nietzsche. The Gay Science. Random House, 1882/1974b.
- Eva S. Nilsen. Indecent standards: The case of u.s. versus weldon angelos. Roger Williams University Law Review, 11(2):537–563, 2006.
- Daniel Nolan. Impossible worlds. Philosophy Compass, 8(4):360–372, 2013.
- Daniel Nolan and John O’Leary-Hawthorne. Reflexive fictionalisms. Analysis, 56:26–32, 1996.
- David Owens. Shaping the Normative Landscape. Oxford University Press, 2012.
- Derek Parfit. Reasons and Persons. Clarendon, Oxford, 1984.
- Derek Parfit. On What Matters, volume 1. Oxford University Press, 2011a.
- Derek Parfit. On What Matters, volume 2. Oxford University Press, 2011b.
- Adam Perry. The internal aspect of social rules. Oxford Journal of Legal Studies, 2014.
- Philip Pettit. Deliberative democracy and the discursive dilemma. Philosophical Issues, 11: 268–99, 2001.
- David Plunkett and Timothy Sundell. Disagreement and the semantics of normative and evaluative terms. Philosophers’ Imprint, 13:1–37, 2013.
- Douglas W. Portmore. Commonsense Consequentialism: Wherein Morality Meets Rationality. Oxford University Press, 2011.

- Richard Posner. Overcoming Law. Harvard University Press, 1995.
- Gerald J. Postema. Bentham and the Common Law Tradition. Clarendon Press, 1989.
- Roscoe Pound. Spurious interpretation. Columbia Law Review, 7(6):379–386, 1907.
- H.A. Prichard. Does moral philosophy rest on a mistake? Mind, 21, 1912.
- H.A. Prichard. Moral Writings, chapter The Obligation to Keep a Promise, pages 257–65. Clarendon, 1940/2002.
- Graham Priest. Sylvan’s box: A short story and ten morals. Notre Dame Journal of Formal Logic, 38(4):573–82, 1997.
- Michael Quinn. Legal Fictions in Theory and Practice, chapter Fuller on Legal Fictions: A Benthamite Perspective. Springer, 2015.
- Peter Railton. Facts and value. Philosophical Topics, 14:5–31, 1986.
- Peter Railton. Normative force and normative freedom: Hume and Kant, but not Hume versus Kant. Ratio, 12(4), 2002.
- Jed S. Rakoff. Why innocent people plead guilty. New York Review of Books, November 20, 2014.
- Jed S. Rakoff. Mass incarceration: The silence of the judges. New York Review of Books, May 21, 2015.
- John Rawls. Two concepts of rules. Philosophical Review, 1955.
- John Rawls. A Theory of Justice. Belknap, 1971.
- Joseph Raz. Ethics in the Public Domain: Essays in the Morality of Law and Politics, chapter Authority, Law and Morality. Oxford University Press, 1995.
- Joseph Raz. The Morality of Freedom. Oxford University Press, 1996.

- Joseph Raz. Practical Reason and Norms. Oxford University Press, 1999.
- Joseph Raz. Normativity, chapter Explaining Normativity: On Rationality and the Justification of Reason. Blackwell, 2000.
- Joseph Raz. The Authority of Law: Essays on Law and Morality. Oxford University Press, 2009.
- Michael Ridge. Ecumenical Expressivism: Finessing Frege. Ethics, 116:pp. 302–36, 2006.
- Michael Ridge. Impassioned Belief. Oxford University Press, 2014.
- Gideon Rosen. Blackburn’s Essays in Quasi-Realism. Noûs, 32(3):386–405, 1998.
- Gideon Rosen. Fictionalism in Metaphysics, chapter Problems in the History of Fictionalism. Oxford University Press, 2005.
- Gideon Rosen. Abstract objects. The Stanford Encyclopedia of Philosophy, 2009.
- Gideon Rosen. Modality: Metaphysics, Logic, and Epistemology, chapter Metaphysical Dependence: Grounding and Reduction. Oxford University Press, 2010.
- Michael J. Saks. Judicial nullification. Indiana Law Journal, 68, 1993.
- Nathan Salmon. Nonexistence. Noûs, 32:277–319, 1998.
- Thomas Scanlon. What We Owe to Each Other. Harvard University Press, 1998.
- Timothy Scanlon. Moral Dimensions: Permissibility, Meaning, Blame. Harvard University Press, 2008.
- T.M. Scanlon. Being Realistic about Reasons. Oxford University Press, 2014.
- Stephen Schiffer. Language-created, language-independent entities. Philosophical Topics, 24:149–167, 1996.

Mark Schroeder. Slaves of the Passions. Oxford University Press, 2007.

Mark Schroeder. Being For: Evaluating the Semantic Program of Expressivism. Oxford University Press, 2008.

Mark Schroeder. Value and the right kind of reason. Oxford Studies in Metaethics, 5:25–55, 2010.

Stephen J. Schulhofer. Rethinking mandatory minimums. Wake Forest Law Review, 1993.

John Searle. How to derive “ought” from “is”. The Philosophical Review, 73(1):43–58, 1964.

John Searle. The logical status of fictional discourse. New Literary History, 6(319–332), 1975.

Gustafsson Senden, Marie, Back, Emma A, and Annd Lindqvist. Introducing a gender-neutral pronoun in a natural gender language: The influence of time on attitudes and behavior. Frontiers in Psychology, 6(893), 2015.

Adam Sennet. Ambiguity. The Stanford Encyclopedia of Philosophy, 2011.

Nishi Shah and David Velleman. Doxastic deliberation. The Philosophical Review, 144(4): 497–534, 2005.

Scott Shapiro. Hart’s Postscript: Essays on the Postscript to the Concept of Law, chapter On Hart’s Way Out. Oxford University Press, 2001.

Scott Shapiro. What is the internal point of view? Fordham Law Review, 75, 2006.

Scott Shapiro. Legality. Harvard University Press, 2011.

Tommie Shelby. Justice, deviance and the dark ghetto. Philosophy and Public Affairs, 35 (2):126–160, 2007.

Peter Singer. Practical Ethics. Cambridge University Press, 2nd edition, 1993.

- John Skorupski. The Domain of Reasons. Oxford University Press, 2010.
- M. B. Smith. May Judges Ever Nullify the Law. Notre Dame Law Review, 74, 1999.
- Michael Smith. The Moral Problem. Blackwell, 1994.
- Philip Soper. A Theory of Law. Harvard University Press, 1984.
- Nicholas Southwood. New waves in meta-ethics, chapter The authority of social norms, pages 234–48. Palgrave MacMillan, 2010.
- Nicholas Southwood. The moral/conventional distinction. Mind, 120, 2011.
- Jeff Speaks. Theories of meaning. The Stanford Encyclopedia of Philosophy, 2014.
- Loretta J. Stalans. Citizens’ crime stereotypes, biased recall, and punishment preferences in abstract cases: The educative role of interpersonal sources. Law and Human Behavior, 17 (4):451–470, 1993.
- Jason Stanley. Hermeneutic fictionalism. In Peter French and Howard K. Wettstein, editors, Midwest Studies in Philosophy XXV: Figurative Language, pages 36–71. Basil Blackwell, Oxford, 2001.
- Jason Stanley. Contextualism in Philosophy: Knowledge, Meaning, and Truth, chapter Semantics in Context, pages 221–54. Oxford University Press, 2005.
- P. F. Strawson. Freedom and resentment. Proceedings of the British Academy, 48, 1962.
- Sharon Street. A darwinian dilemma for realist theories of value. Philosophical Studies, 127 (1):109–166, 2006.
- William J. Stuntz. The Collapse of American Criminal Justice. Harvard University Press, 2011.
- Eric Swanson. Modality in language. Philosophy Compass, 3(6):1193–1207, 2007.

- Kurt Sylvan. What apparent reasons appear to be. Philosophical Studies, 172, 2015.
- Amie L. Thomasson. Fiction and metaphysics. Cambridge University Press, 1999.
- Amie L. Thomasson. Speaking of fictional characters. Dialectical, 57:207–226, 2003.
- Kevin Toh. Hart’s expressivism and his Benthamite project. Legal Theory, 11:75–123, 2005.
- Peter van Inwagen. Oxford Handbook of Metaphysics, chapter Existence, Ontological Commitment, and Fictional Entities. Oxford University Press, 2003.
- David Velleman. The Possibility of Practical Reason, chapter The Aim of Belief. Oxford University Press, 2000.
- Kai von Fintel and Irene Heim. Intensional Semantics. 2011.
- Kai von Fintel and Sabine Iatridou. Time and Modality. Studies in Natural Language and Linguistic Theory, chapter How to Say Ought in Foreign: The Composition of Weak Necessity Modals, pages 115–41. Springer, 2008.
- Jeremy Waldron. The Dignity of Legislation. Cambridge University Press, 1999.
- Kendall Walton. Mimesis as Make-Believe. Harvard University Press, Cambridge, Mass., 1990.
- Kendall Walton. Morals in fiction and fictional morality/i. Proceedings of the Aristotelian Society, LXVIII:27–50, 1994.
- G. Warnock. The Object of Morality. Methuen and Co, 1971.
- Linda D Wayne. Neutral pronouns: A modest proposal whose time has come. Canadian Woman Studies, 24, 2005.
- Ralph Wedgwood. Schroeder on expressivism: For – or against? Analysis, 70(1):117–129, 2010.

- Ralph Wedgwood. Book review: Being realistic about reasons. The Philosophical Quarterly, 2015.
- Bradley Wendell. Lawyers and Fidelity to Law. Princeton University Press, 2010.
- Roger Wertheimer. The Significance of Sense: Meaning, Modality, and Morality. Cornell University Press, 1972.
- Bernard Williams. Problems of the Self. Cambridge University Press, Cambridge, 1973.
- Timothy Williamson. Knowledge and its Limits. Oxford University Press, Oxford, 2000.
- Daniel Wodak, Sarah-Jane Leslie, and Marjorie Rhodes. What a loaded generalization: Generics and social cognition. Philosophy Compass, 10(9):625–635, 2015.
- Jack Woods. Expressivism and Moore’s paradox. Philosopher’s Imprint, 14(5), 2014.
- Jack Woods and Derek Baker. How expressivists can and should explain inconsistency. Ethics, 125(2), 2015.
- Seth Yalcin. Epistemic modals. Mind, 116:983–1026, 2007.
- Seth Yalcin. Bayesian expressivism. Proceedings of the Aristotelian Society, 42, 2012.
- Michael Zimmerman. The Concept of Moral Obligation. Cambridge University Press, 1996.
- A. Zwicky and J. Sadock. Syntax and Semantics, chapter Ambiguity tests and how to fail them. Academic Press, 1975.