

UC Riverside

UC Riverside Electronic Theses and Dissertations

Title

Reasons and Resentment

Permalink

<https://escholarship.org/uc/item/6t06c7qx>

Author

Coates, Daniel Justin

Publication Date

2012

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA
RIVERSIDE

Reasons and Resentment

A Dissertation submitted in partial satisfaction
of the requirements for the degree of

Doctor of Philosophy

in

Philosophy

by

Daniel Justin Coates

June 2012

Dissertation Committee:

Dr. John Martin Fischer, Chairperson

Dr. Michael Nelson

Dr. John Perry

Dr. Andrews Reath

Dr. Gary Watson

Copyright by
Daniel Justin Coates
2012

The Dissertation of Daniel Justin Coates is approved:

Committee Chairperson

University of California, Riverside

Acknowledgments

More people than I can remember have helped me get here. So, to the best of my ability and in something approaching chronological order (with certain, more significant figures appearing closer to the end), I want to thank those who have helped me get here today.

First, my undergraduate advisors—Sandra Chaney, David Grier, and John Wingard—who instilled in me an appreciation for good questions. I also want to thank Trevor Kvaran, Sean Martin, and Jim Sias for being my first philosophical friends. I learned a lot from drinking beer with you guys. On that note, I want to thank the Brick Store Pub and Taco Mac, for hosting those great times. I also want to thank Stephen Jacobson and Tim O’Keefe for their help in my growth as a philosopher. And I especially want to thank Eddy Nahmias for the hard work he put into my philosophical development.

Upon arriving at the University of California, Riverside, I found myself at home, philosophically speaking. During my first two years at UCR, there was a large group of us working on issues related to free will and moral responsibility: Christopher Franklin, Ben Mitchell-Yellin, Garrett Pendergraft, Philip Swenson, Patrick Todd, and Neal Tognazzini. What I owe to these guys philosophically is immeasurable. But what I owe to them personally is also significant: they were good friends in good times and in bad. Despite thinking philosophically about the sentiments, I still find it hard to express them. So thanks guys.

Of course, at a place like UCR, these are hardly the only people who helped me think through issues related to human agency. So I also want to thank: Zac Bachman, Joshua

Hollowell, Agnieszka Jaworska, Coleen Macnamara, Meredith McFadden, Jonah Nagashima, and Monique Wonderly.

I also became good friends with a number of people who didn't work on free will and responsibility: Matt Braich, Joe Cressotti, Dan Ehrlich, Megs Gendreau, Michael Goerger, Mark Johnson, Samantha Matherne, Luis Montes, Courtney Morris, John Ramsey, and Patrick Ryan. There are a lot of good philosophers in that list (and good people too!) So again, I owe a great deal to this (diverse) group, both intellectually and personally. I especially want to thank Michael and Sam, who, whether they know it or not, helped me decide to stay at UCR.

Finally, I want to thank the faculty of UCR. It probably goes without saying that my philosophical accomplishments (such as they are) are due to their training. Among the faculty who are not on my committee, I want to especially thank Maudemarie Clark and Pierre Keller. I ended up taking a lot of history classes from Maude and Pierre—*German history!*—and I am undeniably a better philosopher for it. I'd also like to thank John Fischer, Peter Graham, Coleen Macnamara, Michael Nelson, Andy Reath, Erich Reck, Eric Schwitzgebel, Charles Siewert, and Gary Watson for their seminars, and for things I learned from them outside of class.

Of course, I am extremely grateful to my committee. Many thanks to John Perry for agreeing to participate in the defense, even though I am guilty of *unstructured* procrastination and didn't get him a draft until the last minute. Thanks also to Gary Watson for providing a model of how to think about these and related issues. I always thought Strawson was probably philosophically important, but then I read Gary's "Responsibility and the Limits of Evil," which just sealed it for me. I am especially grateful then, that the author

of the best, by my lights, discussion of “Freedom and Resentment” is on my committee. And nervous. Also, I should probably thank Gary doubly (or tripelley—as the Belgians might say of their beers), since he’s the one that introduced me to Hollingshead’s Deli. I also appreciate Andy Reath’s help in writing this dissertation and more importantly, for demonstrating the virtues of careful scholarship. He was also a good sport about my relatively new-found interest in Kant, which resulted in what I suspect were a lot of bad questions. Thanks as well to Michael Nelson. Michael spent a lot of time helping me to develop these ideas, and although I know he’s not convinced by them, they are much better than they would have been without his help.

And then there’s John Fischer, who I couldn’t thank enough for all he’s done. I came to UCR in large part to work with John, but at the time, I don’t think I could’ve imagined how good he’d be to work with. I’ve recently re-read some of my early work, and why he believed in me, I can’t be sure. But I’m glad that he did. Mostly however, I am especially grateful that at some point, while he was teaching me how to be a professional philosopher, we became friends.

Lastly, I’d just like to note that it’s almost certain that had things been different, Paul Hoffman would have been on this committee. And I suspect that he would’ve come to the defense with his copy of Hume’s *Treatise* and some serious questions that I probably wouldn’t be able to answer satisfactorily. I am sad that he will not be there.

* * *

For allowing me to present versions of various chapters, I want to thank the participants of the UCR Agency Writing Workshop, the “Ready or Not” (RON) dissertation writing workshop, the UCR faculty and students that attended my practice job talk, which was a presentation of the ideas that became Chapter 6. For helpful discussions on these and related ideas, I’d especially like to thank John Fischer, Pierre Keller, Samantha Matherne, Ben Mitchell-Yellin, Michael Nelson, Andy Reath, and Neal Tognazzini. And for comments on earlier drafts of various chapters, I’d like to thank—in addition to my committee, Zac Bachman, Martha Nussbaum, and Neal Tognazzini.

Dedication

For my family, and most of all, for Steph

ABSTRACT OF THE DISSERTATION

Reasons and Resentment

by

Daniel Justin Coates

Doctor of Philosophy, Graduate Program in Philosophy
University of California, Riverside, June 2012
Dr. John Martin Fischer, Chairperson

In this dissertation I develop a theory of practical reasons as such, and then I extend that theory to specifically moral reasons. According to the theory of practical reasons that I develop in Part I, the existence and weight of an agent's reason to act in a particular way depends on an agent's motivational states—specifically those motivational states issuing from practical orientations that play some role in structuring the agent's practical identity. I then argue that this account of practical reasons is a well-motivated reductive account of practical reasons that is extensionally adequate.

In Part II, I turn to questions of specifically moral reasons, which putatively have the properties of being categorical and of having practical priority. I argue that the theory of practical reasons that I develop in Part I doesn't rule out the possibility that moral reasons would have these two properties, and I suggest (following Peter Strawson) that such properties could be grounded in our orientation towards others as participants in certain forms of meaningful interpersonal relationships—an orientation that is perhaps internal to our practical identity as relational agents. I argue that if this is correct, then moral reasons will be categorical and have practical priority.

Table of Contents

Preface	xi
1. Introducing Structuralism	1
Part I: Practical Reasons	7
2. Reasons and Structure	8
3. Towards a Well-Motivated Humean Theory of Reasons	40
4. The Extensional Adequacy of Structuralism Part I: <i>Too Many Reasons</i>	60
Part II: Moral Reasons	83
5. The Extensional Adequacy of Structuralism Part II: <i>Too Few Reasons</i>	84
6. Structuralism and the Practical Priority of Moral Reasons	109
7. The Participant Attitude and Moral Reasons	130
Selected Bibliography	168

Preface

About the time that I was convincing myself of a fairly idiosyncratic interpretation of Strawson's "Freedom and Resentment" (which I discuss in Chapter 7), I read Mark Schroeder's *Slaves of the Passions* with Matt Braich, Maude Clark, Meredith McFadden, Michael Nelson, and Eric Walker. This dissertation therefore, can be seen as a marriage of ideas found in each of these sources. And although Strawson is a kind of Kantian (at least as I read him) and Schroeder is a kind of Humean, I think that it's a marriage that works well—a bit like one of those cute YouTube videos of a cat sleeping peacefully with a dog.

I should say though, that I'm not sure where I come down on the ideas developed in the dissertation. I do find Structuralism to be promising. But I'm not yet sure whether moral reasons really are categorical and practically authoritative. The world would undoubtedly be neater if they were. But I worry that this sort of view fails to take seriously—I mean *really* seriously—the unfortunately messy complexity of human life. I guess I just don't know what to say about such matters. Of course, since admissions of ignorance are frowned upon in philosophy, I defend the orthodox view of moral reasons, that they are categorical and that they have practical priority. And I hope that my own uncertainty does not render my arguments less clear or compelling.

Justin Coates
University of California, Riverside
May 2012

Chapter 1: Introducing Structuralism

§1.

Reasons are said to be facts that bear on practical questions that arise in our lives.¹ That it's raining is a reason to wear galoshes and carry an umbrella *because* this fact bears on the practical question I have of what to wear and carry as I'm about to leave my apartment. Likewise, the fact that standard factory farming practices cause a significant amount of unnecessary pain and suffering is a reason for me to refrain from eating bacon, even though bacon is especially tasty. After all, the fact that factory farming causes a significant amount of unnecessary pain and suffering bears on the practical question I ask myself several times a day: What should I eat?

But these examples, illustrative as they no doubt are, do not tell us what it is for some fact to *bear on* a practical question? That is, these examples don't go very far in providing us an *analysis* of this bearing on relation that seems constitutive of some fact's being a reason at all. So: what is it for some fact to bear on a practical question?

In this dissertation, I hope to answer this question.

Actually, that's too strong. So let me try again.

In this dissertation, I'll begin exploring a possible answer to this question. The view I explore—a view I call “Structuralism”—is a view that I have a good deal of affinity for. But despite my affinity for Structuralism, it's still not clear to me that it is wholly adequate as an analysis of the bearing on relation. Nevertheless, for the purposes of the dissertation, I will develop and defend Structuralism not only as an adequate analysis of practical reasons,

¹ For a helpful discussion of this account of reasons, see Pamela Hieronymi [2005].

but as a particularly promising way of understanding what it is for some fact to bear on a practical question.

In the remainder of this introduction then, I will introduce you to the dissertation. Consider this your roadmap to Structuralism.

§2.

I divide my discussion of Structuralism into two parts, the first dealing with practical reasons as such and the second focusing on specifically moral reasons. In Part I, I introduce Structuralism as a theory of practical reasons as such. I then argue that Structuralism offers a plausible reductive base for a theory of reasons; unlike some other theories of practical reasons, it offers us the *right kind* of analysis. I conclude Part I by discussing the extensional adequacy of Structuralism. In Part II, I return to the question of Structuralism's extensional adequacy. I then argue in the concluding chapters that not only is Structuralism compatible with a conception of moral reasons according to which moral reasons are categorical and practical authoritative, but that it can plausibly explain why moral reasons have these properties.

Having seen the forest, let's briefly look at the trees.²

In Chapter 2 "Reasons and Structure" I motivate and articulate a Humean theory of practical reasons: Structuralism. Structuralism explains the existence of S 's reason R to φ in terms the motivational states that (i) have an object that would be promoted by S 's φ -ing if R is true and that (ii) issue from practical orientations that play some role in structuring S 's

² Actually, we'll be getting to the trees in the chapters themselves. So if there's something in between the forest and the trees (in the relevant sense of "in between" of course), then let's briefly look at that.

practical identity (e.g., *S*'s loves). When the practical orientations in question play a more significant role in structuring *S*'s practical identity, on Structuralism, reasons explained by this orientation are comparative weightier. Thus, Structuralism also provides the Humean with an account of the weights of reasons. I go on to argue that Structuralism offers us the *right kind* of reductive analysis by comparing it to Proportionalism—a Humean theory that offers the *wrong kind* of reductive analysis—and by further explicating the notion of a “practical orientation.”

In Chapter 3 “Towards a Well-Motivated Humean Theory of Reasons” I consider an objection to Humean accounts (like Structuralism, but also like Mark Schroeder’s Hypotheticalism) that reject Proportionalism as an analysis of the weights of reasons.³ According to this objection, which is due to David Enoch, any Humean account that rejects Proportionalism is poorly motivated and ad hoc, since by Enoch’s lights all Humeans are fundamentally committed to the truth of Proportionalism in at least some cases.⁴ I argue that Enoch’s objections to anti-Proportionalist accounts of Humeanism rest on a mistake, and that Humean accounts like Structuralism (but unlike Schroeder’s Hypotheticalism) are able to motivate their rejection of Proportionalism on uncontroversially Humean grounds. Thus, *pace* Enoch, it is possible that some anti-Proportionalist Humean theories of reasons are in fact well-motivated.

In Chapter 4 “The Extensional Adequacy of Structuralism Part I: *Too Many Reasons*,” I consider how Structuralism might respond to the charge that it would entail the existence of too many reasons. I begin by arguing that because Structuralism circumscribes the set of

³ For more on Hypotheticalism, see Schroeder [2007a].

⁴ Enoch [2011].

motivational states with reasons-giving authority to those that issue from practical orientations that play a role in structuring the agent's practical identity, it does not entail that just any motivational state (e.g., Davidson's yen to touch a woman's elbow) will be the source of reasons. I then go on to distinguish between two versions of the Too Many Reasons problem. First I consider the *Way Too Many Reasons* problem, and I argue that this problem poses no threat to Structuralism. Second I consider the *Still Too Many Reasons* problem. Unlike the Way Too Many Reasons problem, *this* problem might show Structuralism to be extensionally inadequate if there is good reason to think that agents will never have reasons to act immorally that are sufficiently weighty so as to rationalize that action. I conclude this chapter (and Part I) with a discussion of modest Humeanism, which is a form of Structuralism that rejects Rationalism—the thesis that all immoral action is also irrational—and so denies that Structuralism's entailment of “still too many” weighty reasons is really a problem.

In Chapter 5 “The Extensional Adequacy of Structuralism Part II: *Too Few Reasons*,” I consider a related objection—viz., that Structuralism and other Humean theories of reasons would entail that agents have too few reasons, since there are some reasons (e.g., moral reasons) that are reasons for all agents. Following Mark Schroeder I argue that on Humean theories of Reasons, an agent's reasons can be overdetermined by a range of motivational states whose objects would be promoted by acting in ways consistent with moral principles. Thus, I argue that moral reasons will be incredibly cheap on Structuralism. So cheap that plausibly all agents will have moral reasons. In an effort to defend this claim, I consider objections to two recent (and promising) probabilistic accounts of promotion. Although I agree that extant accounts are subject to counterexamples, I suggest a new probabilistic

account of the promotion relation that is not subject to such counterexamples and that will rationalize moral reasons for all agents, thus securing the categoricity of moral reasons.

In Chapter 6 “Structuralism and the Practical Priority of Moral Reasons,” I turn to the issue of whether a conception of moral reasons as having practical priority is compatible with the truth of Structuralism. I argue that it is if we accept that reasons explained by motivational states issuing from practical orientations that are *internal* to an agent’s practical identity. To say that some practical orientation is internal (in the relevant sense) to an agent’s practical identity is to say that the role it plays in structuring her practical identity is a necessary one, such that she cannot dissociate herself from that practical orientation without doing irreparable violence to her own practical identity. Since such practical orientations would play a necessary role in structuring *who we are* as agents, it would plausibly follow that reasons explained by motivational states issuing from these practical orientations would also be necessary for the agent—*practically* necessary. Thus if moral reasons can be explained by motivational states issuing from an internal practical orientation, they will have practical priority. I admit though that simply showing that Structuralism is compatible with the practical priority of moral reasons is weak, since it doesn’t follow from this that there are any internal practical orientations upon which we can ground moral reasons. I turn to this task in the final chapter.

In Chapter 7 “The Participant Attitude and Moral Reasons,” I argue that the Participant Attitude—a practical orientation described by Peter Strawson in “Freedom and Resentment”—is an *internal* practical orientation because it plays a necessary role in structuring our identities as rational agents. Since it plays a necessary role in structuring our rational agency as such, it would be a practical orientation that plays a necessary role in

structuring the identity of all rational agents. Thus it could explain not only the practical priority of moral reasons but also the categoricity of moral reasons. I conclude by offering an explanation of how the Participant Attitude can explain the categoricity and practical priority of moral reasons. This explanation is a form of “Constitutivism,” but it is a distinctively Strawsonian form, since it is grounded in our engagement with others in meaningful reciprocal interpersonal relationships.

* * *

Of course, the devil is in the details. And given the scope of this project—even the fairly unlimited scope of a dissertation—I unfortunately had to omit many crucial details. But I hope that you come away from *Reasons and Resentment* thinking that, even if everything isn’t quite right, or quite worked out sufficiently, there’s something promising about understanding our practical reasons as being grounded in those practical orientations that play a role in structuring our identities as agents.

Part I

Practical Reasons

“Reason is, and ought only to be the slave of the passions, and can never pretend to have any other office than to serve and obey them.”

David Hume [T B2.3.3]

Chapter 2: Reasons and Structure⁵

§1.

Dan and Joe are browsing at the local record store. Dan sees a relatively obscure jazz record from the early 1960s, and he gets really excited. You see, Dan loves obscure, hard-to-listen-to jazz. Joe on the other hand, prefers more accessible music, and as he always does when Dan gets excited about some unknown album, Joe rolls his eyes.⁶

In this case, the *fact* that the store has the obscure jazz album on sale is a *reason* for Dan to buy the album. And conversely, this fact isn't a reason for Joe to buy the album. This much seems obvious. But what's less obvious is the explanation of why in this case, Dan has a reason to buy the album, but Joe doesn't.

A natural answer to this question appeals to the differences in Dan and Joe's psychological makeup. On the one hand, Dan loves obscure jazz records, and in general, buying them when he finds them on sale at reasonable prices appeals to him. And because buying and listening to such music will promote these preferences and concerns, the fact that the record store has this album on sale is plausibly thought to be a *normative* reason for Dan to buy it—i.e., a reason that not only explains Dan's behavior but that can also (in the right circumstances) justify Dan's decision to buy the record.

On the other hand, Joe doesn't really like obscure jazz records (in general, the hard work involved in listening to them doesn't appeal to Joe). So, Joe's preferences and

⁵ I appreciate helpful comments on these ideas and on earlier drafts of this paper from an audience at the University of California, Riverside, and from John Martin Fischer, Michael Nelson, Andrews Reath, Philip Swenson, and Neal Tognazzini.

⁶ This case borrows its structure from Mark Schroeder's [2007] case of Ronnie and Bradley. I changed the case primarily because I can relate to Dan much better than I can relate to Ronnie, who loves dancing—an interest that seems utterly foreign to me.

concerns won't be promoted by buying and listening to such music, and this (putatively) explains why the fact that the record store has the record on sale isn't a reason for Joe to buy the record (or at least, this suggests that the fact that the record store has the record on sale isn't as weighty of a reason for Joe as it is for Dan).⁷ Thus, it's natural to conclude that Dan has a reason that Joe lacks in this case because Dan cares about different things.

§2.

As I said, this is a natural answer to the question, "Why is the fact the record store has an obscure jazz album on sale a reason for Dan (but not for Joe) to buy the album?" At first blush, it would seem that appealing to non-normative features of Dan's and Joe's respective conative states is an elegant way of explaining why some facts are reasons for Dan but not for Joe. And given the explanatory power of this answer in the case of Dan and Joe, it's natural to ask, "But can *all* of an agent's practical reasons be explained in this way?" That is, we might wonder whether it is generally true that for any normative reason, the existence of that reason ontologically depends on an agent's conative states (e.g., desires, loves, cares, concerns, hopes, etc.).

Of course, to suggest that all reasons do depend on agents' conative states is certainly a bold claim. For although it's plausible that some reasons depend on our conative states, it is by no means obvious that all normative reasons depend on the actual conative states of agents. For example, even if I don't *enjoy*, *want to*, or *care about* helping the poor or being nice to others, most of us would still think that I have some reason(s) to do these

⁷ I will say more concerning the content of this parenthetical aside in §3 of this Chapter and later in Chapters 4-6, but for now, I simply leave the point as is.

things. And if I had these reasons independently of my actual conative states, then this would undermine a general explanation of normative reasons in terms of the actual conative states of agents.⁸

In the face of such worries, the “Humean”—inspired by Hume’s claim that “reason is, and ought only to be the slave of the passions” [B2.3.3]—embraces this bold claim.⁹ And following Mark Schroeder [2007], we can say more precisely that the Humean Theory of Reasons (HTR)—an account of reasons that is committed to the “Humean” thought described above—is committed to nothing more than a parity thesis. More perspicuously, Schroeder claims that HTR holds that “every reason is explained by the kind of psychological state that explains [Dan’s] reason in the same way as [Dan’s] is” [Schroeder 2007, 2]. And because Dan’s reason is explained by his conative states, HTR holds that all reasons will be explained in this way. Thus all reasons will indeed be slaves to the passions.¹⁰

In what follows, I hope to develop and defend a version of HTR as a viable theory of practical reasons. This specific theory of reasons, which I call Structuralism, is a Humean account of (i) the existence and (ii) the weight of reasons. In its favor, I shall argue that Structuralism is a well-motivated account of practical reasons that takes seriously the Humean thought that if we want to explain an agent’s normative reasons we need to look no further than her conative states. Moreover, I will also argue that Structuralism can

⁸ I discuss this worry—the Too Few Reasons problem—in Chapter 5.

⁹ Perhaps “embraces” is too strong, since some who take this claim seriously (like me), are merely sympathetic towards it and interested in exploring its prospects as an account of practical reasons.

¹⁰ And for my purposes, this account of HTR (even if it is ahistorical or anachronistic) will suffice. That is, for the purposes of defending an account of reasons based on the parity thesis, whether Schroeder’s characterization of HTR is ultimately attributable to Hume is inconsequential. Of course, I don’t mean to suggest that historical work on Hume is inconsequential. Rather, I only mean to suggest that even if the theory of reasons that Schroeder loosely attributes to Hume is something Hume would not accept, it’s nevertheless a theory is worth taking seriously in its own right. For more on this point, see Schroeder [2007a].

accommodate many of the intuitions that motivate traditional *objections* to HTR—intuitions concerning (i) the proper extension of practical reasons and (ii) the putative practical authority of specifically moral reasons. As such, Structuralism isn't simply a viable position for dyed-in-the-wool Humeans; it should also be attractive to those who, like many in the Kantian moral tradition, take moral reasons to be categorical and practically authoritative.¹¹

But before we can get to questions concerning moral reasons specifically, I want to first offer a general statement of Structuralism. This chapter will serve as such. Also in this chapter, I will motivate Structuralism by contrasting it with an alternative form of HTR called Proportionalism. I claim that Structuralism, but not Proportionalism, can provide the Humean with the *right kind* of analysis of the *bearing* relationship that is constitutive of some fact being a practical reason.¹² What I mean is that what it is for some fact to be a reason is for it to bear on a practical question, and that Structuralism provides the right kind of analysis of this bearing on relationship. Thus, Structuralism provides the right kind of analysis of practical reasons. Once I argue that Structuralism can plausibly serve as the right kind of analysis, I will further clarify its content and commitments as a theory of practical reasons.

In subsequent chapters, I will argue that Structuralism answers a recent challenge to anti-Proportionalist Humean theories (Chapter 3), that Structuralism is extensionally adequate (Chapters 4-5), that Structuralism is compatible with thinking that specifically moral reasons have practical priority (Chapter 6), and that Structuralism, in conjunction with

¹¹ As we'll also see, Structuralism should also be attractive to those in the contractualist tradition. Indeed, because Structuralism is simply a metaphysical theory about *what it is* for some fact to reason, it is compatible with a wide range of first-order normative ethical theories.

¹² I say a bit more about what I mean by the "bearing" relationship in §3 of this chapter. For more on this point, see Hieronymi [2005].

some plausible assumptions, provides us with a compelling account of specifically moral reasons (Chapters 7). But first, I tackle a much more preliminary question—viz., what is Structuralism?

§3.

As I've said, HTR holds that all normative reasons can be explained in the same way that Dan's reasons are explained. And plausibly, Dan's reasons are explained by his motivational states.¹³ So, on HTR, all of an agent's normative reasons will be explained by her motivational states. Of course, this bare parity thesis doesn't tell us *which* motivational states are relevant to an agent's reasons. After all, lots of states are putatively motivating—e.g., desires, loves, cares, hopes, perhaps certain evaluative judgments, etc. Nor does the bare parity thesis tell us anything about the comparative weights of an agent's reasons. For example, if Dan has only enough money to buy the record on sale or to buy his food for his daughter, we think that the reasons for buying the record are outweighed by the comparatively weightier reasons for taking care of his daughter. But again, HTR as such tells us nothing concerning how we should explain this property of reasons, viz., that they have comparative deliberative weights.

Thus, as it stands, HTR is too thin. It is compatible with several different accounts of the existence and weight of an agent's reasons, so long as such accounts each explain all of an agent's reasons in the same way they explain Dan's reasons. So if we want HTR to be more than simply a hypothesis about the general basis of reasons, we need a specific account

¹³ Obviously, explanation is a notoriously vexed topic, but as I'm understanding it here, HTR claims that normative reasons are explained by an agent's motivational states because an agent's particular motivational states (e.g., Dan's desire to own the record) explain why some particular fact (e.g., that the record is on sale) *bears on* Dan's practical question—viz., what to do.

of which conative states explain the existence and weight of an agent's reasons and how they do so. And this is precisely where Structuralism comes in.

3.1.

First and foremost, Structuralism holds that an agent's reasons will depend only on those motivational states and practical orientations that *play some role in structuring an agent's practical identity* (hence, "Structuralism"). In other words, on Structuralism, the motivational states and practical orientations that play a role in constituting our identities as agents (where "our identities as agents" is understood as distinct from our personal identity) are the sources of our practical reasons.¹⁴

But more precisely, Structuralism holds that an agent S 's reason R to ϕ depends on the role O plays in structuring S 's practical identity, where O is the practical orientation that issues in S 's motivational state m whose object is promoted by ϕ -ing if R is true.¹⁵ Thus, because reasons are considerations that bear on practical questions,¹⁶ we can understand

¹⁴ As a helpful heuristic, we can say that some state s plays a role in structuring an agent's practical identity just in case the agent would alter or otherwise damage her practical identity by *dissociating* herself from s . Indeed, when I was first developing Structuralism, I actually thought the amount of damage an agent would do to her practical identity by dissociating herself from some s explained the existence and (particularly) the weight of her reasons. But then I realized that the possibility of dissociation couldn't be normatively fundamental. After all, insofar as an agent's ability to dissociate from some particular s has any normative significance, it is only because the notion of 'dissociation' gives us better purchase on what it would mean for s to play some role in structuring an agent's practical identity.

¹⁵ This initial statement of Structuralism is meant as a conceptual refinement of Mark Schroeder's account of reasons according to which "for R to be a reason for X to do A is for there to be some p such that X has a desire whose object is p , and the truth of R is part of what explains why X 's doing A promotes p " [Schroeder 2007a, 59]. As we'll see, Structuralism is tied to a principled reason for thinking not just any desires will be part of the explanation of an agent's reasons. But more on this to come.

¹⁶ As I've already alluded to (cf. fn. 5), Hieronymi [2005] introduces this way of characterizing reasons as an alternative to a more traditional characterization of reasons due to T. M. Scanlon according to which reasons are considerations (or facts) that favor particular actions or attitudes. For the reasons Hieronymi

Structuralism as a reductive analysis of what it is for some fact to *bear on* a particular practical question. Some fact bears on a practical question just in case it is connected to an agent's motivational states and practical orientations in just the way Structuralism describes.

To unpack this, recall Dan's reason to buy the obscure jazz record (i.e., the fact that it was on sale). According to Structuralism, Dan has this reason to buy the record because Dan has a special concern for that form of music; that is, he has a reason to buy the record because he is practically oriented towards jazz in a way that shapes and directs his engagement with music.¹⁷ In other words, Dan's actual practical orientation towards jazz gives rise to a complex set of conative dispositions—dispositions that constitute his concern for jazz and explain why, in particular contexts, music has the significance for Dan that it does. But according to Structuralism, to fully explain Dan's reason, we need to appeal to specific motivational states Dan has—e.g., his desire for this particular album or his desire to find a good deal on a record he didn't own.¹⁸ And it is because the objects of these desires—"that he own this particular album" or "that he get a good jazz album cheaply"—are promoted by his buying the album, that the fact that it is on sale is a reason for Dan to buy the album.

Notice that on Structuralism, the mere fact that Dan wants the album doesn't explain his reason to do so. It's only because this desire issues from a practical orientation that Dan has—a practical orientation that plays some role in structuring Dan's practical identity—that

adduces in her paper, as well as ones I offer in Chapter 4, I think we should prefer Hieronymi's characterization.

¹⁷ I say more about the notion of "practical orientation" in §6 of this chapter. See also Fischer [2010] on "normative orientations" for a related notion (albeit one that appears in a slightly different context).

¹⁸ As we'll see in §5 of this chapter, we could also appeal to other motivational states (e.g., evaluative judgments or beliefs) to help explain Dan's reason in this case.

the desire plays some role in grounding the existence of Dan's reason to buy the album. So if Joe has a fleeting, transient desire to buy the album too—say, the cover art simply strikes him in the right way—on Structuralism, this impulse doesn't necessarily ground a similar reason for Joe to buy the record since it isn't (by hypothesis) connected in the right way to Joe's practical orientations. In other words, because such a desire doesn't issue from one of the practical orientations that structures his identity as an agent, that desire cannot *express* Joe's practical self, and so, this desire would lack reasons-giving authority.¹⁹

3.2.

But Structuralism doesn't simply provide us with an account of the conditions under which reasons exist; it also provides an account of the comparative weights of reasons. And this is important because what an agent *ought* to do all things considered is a function of the weights of her reasons. Because of the significance of the weighting of reasons to normative theory, it is a desideratum on any theory of practical reasons that it not only explains the existence of an agent's reasons, but also the comparative weights of her reasons as well.²⁰ And on Structuralism, the weightiness of S 's reason R to φ will depend on:

- (1) The degree O plays a role in structuring S 's practical identity, where O is one of the practical orientations that play some role in structuring S 's practical identity,
- and

¹⁹ Obviously, the assumption here is that reasons-giving authority issues from our practical identities. I can't argue for this claim here, but it has been widely accepted throughout the history of moral philosophy (e.g., in Plato, Hume, Kant, Nietzsche, etc.).

²⁰ As we will see in Chapter 3, Mark Schroeder's otherwise plausible Humean theory "Hypotheticalism" fails to provide an analysis of the weights of an agent's reasons. This is just one reason to think that Structuralism is an improvement on Schroeder's theory of practical reasons.

(2) *How well* φ -ing will promote m , where m is the motivational state that issues from O .²¹

To help clarify this, consider the following analogy. When constructing a building—say, a Gothic cathedral—the most important piece of the foundation is the cornerstone. When the mason sets the cornerstone, the foundation of the cathedral is built around it. Thus, its location and attitude determine the location and attitude of the rest of the foundation.²²

Likewise, by redirecting the lateral forces of the cathedral’s walls to the ground, the cathedral’s flying buttresses play an important role in structuring it. Of course, the flying buttresses play a *different* role than the cornerstone, but like the cornerstone, their role is very important to the structural integrity of the cathedral. However, compare the cornerstone or the flying buttresses to the beautiful stained glass windows. Though they are great to look at, the windows play no architectural role in structuring the cathedral. Thus, in some important sense, they have less architectural significance.

Structuralism makes similar claims, though it’s our practical identities rather than Gothic cathedrals that are relevant. Thus, if some practical orientation plays a larger role in structuring an agent’s practical identity, motivational states issuing from it will have normative significance—i.e., weight. To illustrate, consider Laura. Laura loves her daughter Evelyn greatly. Laura also finds great enjoyment from her garden. And let’s say that Laura’s love for Evelyn and for gardening both issue from practical orientations Laura has towards Evelyn and gardening. Moreover, let’s say that, though each of these practical orientations

²¹ For the purposes of this chapter, I focus on (1), since it is distinctive of Structuralism. But (2) will be important in Chapter 4 where I argue that Structuralism is consistent with our best theories of practical reasons.

²² Keith Lehrer [1997] uses a similar analogy—that of the keystone in an arch—to explain to role self-trust plays in structuring our capacity for reason, knowledge, and wisdom.

plays important but distinctive roles in structuring Laura's practical identity, her love for her daughter plays a more significant role.

Now suppose that these orientations unfortunately sometimes issue in motivational states that will be promoted by inconsistent courses of action. In such cases, Laura has to deliberate or *weigh* her reasons about what to do. On Structuralism, because Laura's practical orientation toward her daughter plays a more significant role in structuring her practical identity than does her practical orientation towards gardening, the reasons that are explained by the former practical orientation will be comparatively weightier. This accords, I think, with our judgment that she should take care of her daughter before she takes care of her garden in cases in which each of these courses of actions precludes the other. Of course, things are a bit more complicated, but such intricacies can be similarly explained by Structuralism's commitment to (2) as a component of what explains the weight of an agent's reasons.

Perhaps I can further clarify Structuralism's account of the weighting of reasons by relying on the heuristic I introduced earlier—viz., dissociation.²³ Accordingly, we can say that on Structuralism, if the damage an agent would do to her practical identity by dissociating herself from some practical orientation is relatively minor, then those dispositions play little role in structuring the agent's practical identity. Accordingly, none of the agent's reasons that depend on motivational states issuing from those practical orientations will be particularly weighty. Conversely, if the damage the agent would do to her practical identity by dissociating herself from some practical orientation is relatively significant, then we should conclude that that orientation plays an important role in

²³ See fn. 14.

structuring the agent's practical identity. As a result, reasons that depend on motivational states issuing from these significant practical orientations will be comparatively weighty. So understood, Structuralism is a variant of HTR that explains both the existence of and the (deliberative) weight of an agent's reasons in terms of the properties of her motivational states (and the practical orientations from which those states issue).

3.3.

Now that you've got a sense for Structuralism, we are almost ready to start exploring Structuralism's prospects as a general analysis of practical reasons (a task that will only be complete at the conclusion of the dissertation). But before I begin the business of defending Structuralism as a promising analysis of practical reasons, I want to further clarify the content and commitments of Structuralism. To begin this task, (in §4) I'll contrast Structuralism with another version of HTR that also grounds the existence and weight of an agent's normative reasons in properties of her non-normative motivational states: Proportionalism.²⁴ After briefly discussing Proportionalism, I'll characterize those motivational states (in §5). Then, I'll say a bit more about Structuralism's claim that motivations states have reasons-giving authority only if they issue from practical orientations that structure an agent's practical identity (in §6). And finally, I'll conclude with a brief look forward (in §7).

²⁴ I borrow this term from Schroeder [2007].

§4.

Like Structuralism, Proportionalism is a Humean theory of practical reasons. But unlike Structuralism, Proportionalism holds that the existence of an agent's normative reasons is wholly explained by her desires.²⁵ Moreover, on Proportionalism, the role that those desires play in structuring an agent's practical identity is irrelevant to the weight of reasons. Instead, Proportionalism claims that the weight of an agent's reason R to ϕ is proportional to (i) the *strength* of the desire that explains R , and (ii) how well ϕ -ing on the basis of R promotes the object of the explaining desire.²⁶

Notice that both Proportionalism and Structuralism explain the weight of an agent's reasons, which is itself a normative property of the reasons, by appealing to non-normative properties of an agent's motivational states. At this level of magnification Proportionalism and Structuralism seem to be on par. However, in my view, Proportionalism appeals to the *wrong kind* of non-normative property, viz., the strength of those motivational states. To see this, consider the thought (frequently implicit in the reasoning of playground bullies and neo-Conservative strategists alike) that "might *makes* right." Taken purely descriptively,

²⁵ Of course, it would be possible to develop a version of proportionalism according to which the existence of an agent's normative reasons depends on those motivational states (of which desires are but one example) that issue from the practical orientations that play a role in structuring the agent's practical identity. And at that level of descriptions, there would be no difference between this Proportionalist account of the existence of reasons and Structuralism. The difference would lie in their competing explanations of the weight of an agent's reasons.

²⁶ Perhaps you've noticed that Structuralism itself might plausibly be thought to be a kind of proportionalism since it holds that the weight of a reason is proportional to the role that broad sets of conative dispositions play in structuring our practical identities. So it turns out that the main difference between Proportionalism and Structuralism is what non-normative properties of motivational states (and in the case of Structuralism, the practical orientations from which those states issue) are relevant for grounding an agent's reasons. However, for reasons of clarity, I prefer not to label Structuralism as a form of Proportionalism, since following Schroeder [2007] (from whom we have the label 'Proportionalism'), Proportionalism tends to be specifically identified with the view that the weights of reasons are proportional to the strengths of the explaining desires.

perhaps there is something to this thought—the winners tend to write the history books, and so the mightiest tend to have more opportunity to rationalize their side of the story. However, taken normatively as an account of what *makes* some action right, this thought is utterly implausible. But this implausible connection between “might” and “right” is precisely what Proportionalism seems committed to. After all, on Proportionalism causally *strong* (or “mighty”) desires make *weighty* reasons, which just looks like a specific application of “might making right.” Thus, the fundamental problem with Proportionalism is that it accounts for the weight of reasons in a way that too quickly confuses causal strength with normative import in the form of reasons-giving authority.

Alternatively, Structuralism appeals to the kind of non-normative property of an agent’s motivational states that might plausibly serve as a basis for the weighting of her reasons—i.e., it offers the *right kind* of analysis. On Structuralism, the weights of an agent’s reasons depend on the role that particular practical orientations play in structuring an agent’s practical identity. And reasons explained by motivational states that issue from a practical orientation that plays a major role in structuring an agent’s practical identity are accordingly weighty. Now, you might worry that this analysis leaves Structuralism in no better position than Proportionalism. However, unlike the causal strength of an agent’s desires, which intuitively seems to be irrelevant to considerations of rightness or of reasons-giving authority, the role a practical orientation plays in structuring an agent’s practical identity seems like a perfectly legitimate basis upon which to ground the weight of reasons. After all, it’s natural to think that what reasons we have and their relative weights ultimately depend

on *who we are* as agents.²⁷ Thus, when some practical orientation plays a significant role in structuring my identity as an agent, it should come as no surprise that any reasons explained by motivational states issuing from this practical orientation will be weighty.

So while Proportionalism grounds the weight of an agent's reasons in a property of her desires that is normatively insignificant, Structuralism grounds the weight of agent's reasons in a property of her practical orientations that is intuitively tied to the existence of her reasons. Hence Structuralism does plausibly ground the weight of an agent's reasons in a property of her practical orientations that matters normatively. Not only is this a significant advantage of Structuralism over Proportionalism, it also gives us a reason to think that Structuralism is promising in its own right. Because Structuralism grounds the existence and weight of reasons in properties of motivational states that issue from those practical orientations that structure an agent's practical identity, it offers us an account of reasons that could possibly explain the normative significance of reasons as facts that bear on practical questions.

And although it seems to me that Proportionalism's failure to offer such an account is its most fundamental problem, it's hardly the only problem for Proportionalism. Most notably: Proportionalism seems to have a two-way extensional inadequacy because it predicts both that agents have *too many* and *too few* reasons.²⁸ Moreover, on no set of plausibly

²⁷ Again, here Structuralism's commitment to Reasons Internalism (i.e., the view according to which R is a reason for S to φ only if S could be motivated to R) becomes evident. Perhaps surprisingly however, I will say little to defend Reasons Internalism in the dissertation. Indeed, I suspect that there are no decisive arguments for or against Reasons Internalism. Accordingly, if the reader is sympathetic to Reasons Externalism, I would invite him or her to, upon completing *Reasons and Resentment*, simply weigh the theoretical advantages of Structuralism against the theoretical costs of abandoning Reasons Externalism. When I conduct such a weighing, I find that I am not particularly troubled by the pull of Reasons Externalism.

²⁸ This point is most clearly made in Schroeder [2007].

assumptions is Proportionalism compatible with the categoricity or practical priority of moral reasons.²⁹ But on plausible assumptions, we'll see that Structuralism is extensionally adequate and compatible with the categoricity and practical priority of moral reasons (more on this to come!). Later in the dissertation (cf. Chapters 4-7) I'll explore these features of Structuralism in more detail. But first, I continue the task of clarifying the content and commitments of Structuralism as a Humean theory of practical reasons, turning now to the issues of which motivational states might have reasons-giving authority and the nature and normative significance of our practical orientations.

§5.

A theory of practical reasons is *Humean* if it purports to explain every practical reason in the same way it explains Dan's and Joe's practical reasons—i.e., if every practical reason is explained by actual motivational states of the agent. And Structuralism is an improvement on the bare statement of HTR because it gives us a better sense of *which* motivational states can have reasons-giving authority. Specifically, Structuralism claims that the existence and weight of an agent *S*'s reason *R* to ϕ is explained by motivational states that issue from those practical orientations that play some role in structuring the agent's practical identity. And this is illuminating. After all, if Structuralism is true, then presumably, the man who in a functional state that motivates him to turn on every radio he sees—a functional state that is wholly disconnected from the rest of his practical identity that comes and goes

²⁹ This fact is, I think, ultimately explained by Proportionalism's commitment to a "might-makes-right"-like principle.

unexpectedly—does not have a reason to turn on every radio he sees (at least not on the basis of this urge).³⁰

Thus Structuralism offers the Humean a principled way of accounting for the fact that not every motivational state is the sort that has reasons-giving authority. So understood, Structuralism offers us a plausible constraint on motivational states with reasons-giving authority—viz., only those that issue from those practical orientations that play some role in structuring an agent’s practical identity. But this is merely a negative limit on which motivational states have reasons-giving authority. And it’s worth considering whether Structuralism can offer a positive account—perhaps highlighting the specific motivational states that do have reasons-giving authority.

Traditionally—and this can be seen in Proportionalism as well as Mark Schroeder’s Hypotheticalism (a more sophisticated Humean theory)³¹—desires have been the motivational state most frequently thought to have reasons-giving authority. And this would certainly fit with Structuralism. On such a view, those desires that issue from a practical orientation that plays some role in structuring our practical identities would have reasons-giving authority (i.e., they could explain the existence of an agent’s reasons). And while I want to agree that this class of desires does have reasons-giving authority, as I see it, Structuralism is in a position to embrace a much more capacious characterization of those motivational states with reasons-giving authority. Thus, officially, on Structuralism, any mental state that issues from a practical orientation that structures an agent’s practical

³⁰ This case is due to Warren Quinn [1995]. I take Structuralism’s ability to correctly assess that in such a case, the functional state in question lacks reasons-giving authority to be a feather in its cap.

³¹ I discuss Schroeder’s Hypotheticalism in greater detail in Chapter 3.

identity which could also serve to rationalize an agent's action will be a "motivational state" of the relevant sort.

And indeed this is capacious, since on this view not only desires, but other pro-attitudes like likes and cares can have reasons-giving authority. But the class of motivational states with reasons-giving authority doesn't stop there. After all, we often judge that things are good or bad, and these evaluative judgments are frequently cited as the explanation of why we are motivated to act in particular ways. So perhaps, even evaluative judgments and beliefs could be motivational states with reasons-giving authority on Structuralism.

To see this, consider an alternate explanation of why Dan has a reason to buy the obscure record that Joe lacks. On this alternative, the explanation for why Dan has a reason that Joe lacks isn't fundamentally that Dan wants it and that Joe doesn't want it. Rather, it's that Dan judges that it is good and Joe judges that it is bad. No doubt, judgments of this sort often give rise to desires, but these desires' role in explaining why Dan has a reason that Joe lacks ultimately depends on Dan's and Joe's evaluative judgments. In other words, such desires are "motivated," and any putative reasons-giving authority they might have rests in the judgments that motivate them.³²

It seems to me that this is a perfectly adequate explanation of Dan and Joe's case. But if it is, then doesn't it fly in the face of the Humean theory of motivation according to which judgments or beliefs—really any states that have mind-to-world direction of fit—are insufficient to explain motivation?³³ After all, if, as Russ Shafer-Landau suggests, we

³² Here I have in mind Thomas Nagel's [1970] discussion of motivated desires here.

³³ The Humean theory of motivation is not entailed by Humean theories of reasons per se, but following Hume himself (and for pretty obvious reasons) they often go hand-in-hand. A note on nomenclature: since this dissertation seeks to provide a theory of practical reasons, if I call a view "Humean" or

understand the Humean theory of motivation as claiming that “when taken by themselves, beliefs neither motivate nor generate any motivationally efficacious states” [Shafer-Landau 2003, 121], then the putative adequacy of the above explanation of Dan’s and Joe’s motivations seems problematic for theories of motivation that explain all motivation in terms of purely conative (rather than cognitive) states. And indeed, this is precisely the sort of case that Shafer-Landau uses to motivate his *anti*-Humean theory of motivation.³⁴

So far, this discussion of the Humean theory of motivation might seem interesting, but neither here nor there when it comes to a theory of practical reasons as such. But if an anti-Humean theory of motivation is true, it gives us reason to doubt that the practical orientations of an agent could be the basis of her reasons. After all, if we suppose that reasons are necessarily tied to motivational states—as all versions of HTR claim—and that evaluative judgments can motivate, then if evaluative judgments as such can be independent of an agent’s practical orientations (and many philosophers have suggested that they can be), we have reason to doubt that all of my reasons can be explained by appeal to my practical orientations. Thus, the falsity of a Humean theory of motivation would put pressure on Structuralism, which takes an agent’s practical orientations (understood in a particular way to be described in §6) to be the basis for her reasons.

But we must be careful here. While it’s certainly true that Dan’s judgment that the album is good can explain why he was motivated to act in particular ways (either by directly

label it a kind of “Humeanism,” I mean to refer to the Humean theory of reasons. Any other Humean views that I introduce will always be labeled in more descriptive ways (e.g., the “Humean theory of motivation”).

³⁴ Shafer-Landau [2003] pgs. 124-25. There Shafer-Landau considers a case in which it’s a false belief about what you want that best explains your motivation to endure difficult times because in this case, you discover that you didn’t actually want the thing you were pursuing. So plausibly, since the desire couldn’t explain your motivation (it didn’t exist), your belief must.

motivating him or by generating a motivated desire that in turn motivates him), I don't think this vitiates the Humean claim that all motivation is ultimately tied to our fundamentally conative practical orientations.

To illustrate, consider the following analogy. My striking the match is sufficient in ordinary contexts to explain why a fire starts at the end of the match. In such contexts, nothing more needs to be said. In other words, it's *true* that my striking the match explains the fire. But it's also true that if I had struck the match in a vacuum, no fire would have ignited. And as we all know from eighth-grade physical science, fire needs oxygen. So while it's true that my striking the match explains why it catches on fire, it's also true that it only does so because the room I'm standing in is full of O₂. This suggests that whether the act of striking the match can be said to explain the fire is ultimately context-dependent.³⁵ In ordinary contexts of conversation, it's true to say that my act of striking the match generates the fire. And this is because ordinary contexts presuppose that there the natural conditions are roughly standard. But when we want a picture of the underlying physics, we see that it's really my act of striking the match in conjunction with the presence of O₂ that explains the genesis of the fire.³⁶ That is, when we shift contexts, we'll find that we very often shift explanans as well.

I think something similar is going on when it comes to explaining an agent's motivations in terms of (i) her beliefs or judgments or (ii) in terms of desires that are

³⁵ Some would deny that explanation is context-dependent in this way. However, I find the idea of an invariantist account of explanation to be puzzling. Explanatory relations are relations that allow us to *make sense* of one relatum in terms of the other relatum, and clearly, I think, the *making sense* relation is context-dependent. For more, see van Fraassen [1980] and Larry Wright [MS].

³⁶ It's actually much more complicated than this! But hey, I'm doing the *metaphysics* of reasons here, not the *physics* of fire.

generated by those beliefs or judgments. I suspect that we can perfectly well explain Dan's motivation to buy the obscure jazz record simply by appealing to his judgment that the record is good.³⁷ Thus, in ordinary contexts of action explanation and evaluation, we don't need to appeal to any sets of conative dispositions or more specifically, to any of the agent's practical orientations; in these contexts, purely cognitive states like beliefs are sufficient to explain an agent's action because we presuppose that such beliefs are grounded in the agent's practical orientations. But just as the match lights only if there is O₂ in the room, plausibly, Dan's judgment that the record is good will motivate him (in a way that will rationalize his action) only if this judgment ultimately issues from Dan's prior engagement with the world. After all, Dan's judgment that the record is good is not itself *sui generis*,³⁸ but instead reflects Dan's complex concern for and enjoyment of music—a concern that is reflected in a practical orientation that plays some role in structuring Dan's practical identity.

Indeed, just as every lit match will ignite in the presence of O₂,³⁹ so too will every judgment that motivates (in a way that will rationalize the agent's action) do so only against the background of an agent who is practically oriented towards her environment. Thus, on this view it will still be possible to explain a particular instance of motivation or action by appeal to a judgment. And in ordinary contexts, this might be the best way to explain why an agent was motivated to act. But again, in a philosophical investigation of such judgments,

³⁷ Likewise, we could plausibly explain Dan's motivation to buy the obscure jazz record by appealing to his desire to buy the record which is itself generated by his judgment that the record is good.

³⁸ Of course, if such judgments truly were *sui generis* and therefore wholly disconnected from an agent's practical orientations, then it's unclear why they should have any more reasons-giving authority than functional states of the sort Quinn [1995] discusses.

³⁹ Strictly speaking, this might not be quite true (cf. fn 18). But if it's not, it only points to the limits of the analogy, not to the limits of the view I'm developing.

it's worth pointing out (as I do here) that they do not motivate independently of the practical orientations from which they issue.

So on the capacious view sketched above, the motivational states that Structuralism takes to be relevant to an agent's reasons are any mental states that could rationalize the agent's action in virtue of issuing from the agent's practical orientations. And this includes both desires (and other desire-like states) and evaluative judgments. But though this includes a considerable number of candidate states, it does eliminate some motivational states as potentially explaining an agent's reasons. For instance, while Quinn's [1995] desire to turn on radios or Davidson's [2001] yen to drink a can of paint, are certainly motivational states, they are not connected to an agent's practical orientations in the right sort of way. After all, unlike motivational states issuing from those dispositions that structure her practical identity, Quinn's desire or Davidson's yen can only help to causally explain an agent's actions. They are impotent as *rationalizations* since they do not allow the agent to make sense of that action as her own. For example, the man who acts on his yen to drink a can of paint cannot make sense of that action as an expression of who he is (in any sense), since it does issue from who he is *qua* agent. Thus, these states lack the reasons-giving authority of motivational states issuing from an agent's practical orientations.

Hopefully, having settled which atomic motivational states might have reasons-giving authority, I've clarified the content of Structuralism to some degree. But it also makes clear that more needs to be said. Specifically, what are the practical orientations that structure an agent's practical identity? Are they her values? Loves? Concerns? Or something else? And why do our orientations get to serve as the basis of our practical reasons? These are important questions, which I turn to below.

§6.

As I've said, those practical orientations that structure agents' practical identities are the oxygen to motivation's fire—not always invoked in ordinary explanations of how our motivational states are generated, but always in the background.⁴⁰ Indeed, the thought was that just as a world that lacked oxygen would lack fire (and all its important concomitants like heat, light, safety, etc.), so too an entity that lacked these practical orientations would lack motivations with reasons-giving authority (and all their important concomitants like an ability to make sense of yourself, an ability to meaningfully engage with others as equal persons, etc.). Needless to say, on Structuralism, the practical orientations that structure our practical identities are important.

But what are practical orientations?

Well, for starters, I label the attitudes I have in mind “practical *orientations*” because I want it to be clear that they aren't simply atomic mental states.⁴¹ Atomic mental states, like say, my desire that the Oklahoma City Thunder win the NBA Championship this season (Kevin Durant is a very exciting player!) or my belief that Finnegan (my cat) is lounging in the living room, aren't themselves closely enough linked to my practical identity to generate reasons. Moreover, because atomic mental states as such come and go (plausibly, they don't

⁴⁰ However note another similarity: just as the oxygen will always be invoked in physical explanations of a fire, so too will the practical orientations that play some role in structuring an agent's practical identity need to be invoked in metaphysical explanations of an agent's reasons.

⁴¹ This will be true even if you have a conception of atomic beliefs and desires being dispositions to act in particular ways. Unless those dispositions are sufficiently broad so as to structure behavior across a wide range of counterfactual scenarios, they don't play enough of a role in structuring an agent's practical identity to generate reasons themselves.

endure in the same way *we* endure, nor can they plausibly serve as the Lockean ties of personal identity), they can't play a significant role in constituting or expressing our natures. After all, whether I have this atomic desire or that atomic belief doesn't affect my practical identity at all; I could dissociate myself from these states without loss. In other words, my identity as an agent is resilient with respect to these atomic mental states.

Thus I'm not talking about atomic mental states. But *this* still doesn't tell us much about the practical orientations that play some role in structuring our practical identities. So positively, as I understand these practical orientations, they are stances, or attitudes (in the broadest sense) that we take towards some particular person(s), object(s), event(s), or state(s) of affairs.⁴² But what is it to be oriented in this way? And why are the stances we take on the world fundamental constituents of our practical identities in a way that putatively gives them reasons-giving authority?

6.1.

First (and perhaps foremost), an orientation is a stance or an attitude that shapes and directs our engagement with the world. In virtue of having a particular orientation, some things will be particularly salient or significant, while other things will be relegated to the background. This is obvious in cases of mundane perception. How I am oriented with respect to you will dramatically affect how I see you; if I am facing you, and I see you in good light, I might notice (among other things) your height, your weight, your beauty, etc., but if I am oriented slightly differently—like say, at an 135 degree angle—you might appear to be a fuzzy blur in

⁴² Note: I'll be using "orientation," "stance," and "attitude" interchangeably.

my peripheral vision (after all, my glasses don't cover that). The same is true practically speaking.

Thus, because he *loves* Lena Grove, Byron Bunch is oriented towards her in a way that makes fighting for her honor seem like a necessary thing to do, even though before he met her, he was perfectly mild-mannered.⁴³ Moreover this orientation Bunch takes towards Grove explains why it seems so natural for Bunch to follow her to Tennessee, while simultaneously explaining why it would be so odd—at least from his perspective—that he follow any other woman. In other words, it's because he is oriented towards her in this particularly loving way that she isn't fungible. Indeed, as long as Bunch is engaged with her in this way, she's literally irreplaceable from this practical perspective. And this suggests that our practical orientations shape how we perceive the practical landscape, through our evaluative judgments no doubt, but perhaps more fundamentally, through our emotional engagement with the world.

Moreover, Bunch's orientation towards Lena isn't explained by antecedent judgments about her loveliness or the choice-worthiness of loving her (in the novel, he simply comes to love her). Instead, such judgments are more plausibly thought of as issuing from his prior engagement with her as someone to be loved. This suggests that to be oriented towards something—say, a lover—isn't fundamentally a matter of having a particular set of beliefs or evaluative judgments. Rather it's to be related to that person or thing in specific ways that make such judgments possible. Just as how I am physically oriented with respect to you shapes and directs my perception of you, so too will my practical orientation shape and direct my concern for you.

⁴³ William Faulkner, *A Light in August* [1932].

To illustrate this point, suppose Jason sincerely believes (and frequently asserts) that he loves his children. Is it thereby true that he in fact loves them—that he is in fact oriented towards them and engaged with their welfare in a way that reflects his children’s significance? Unfortunately, it doesn’t. After all, like talk, beliefs are cheap, and we are often self-deceived. Though Jason might sincerely believe that his children are wonderful and that they deserve the best, it’s a sad fact about human psychology that such beliefs can be, and often are, held by unloving dead-beat dads. Thus, even if Jason sincerely believes that his children are wonderful, deserving children, plausibly, if he doesn’t pay child support, if he ignores his children when he has custody, if he tries to convince them that their mother is a horrible person, if he exposes them to dangerous and otherwise inappropriate things, etc., then he doesn’t really love them, and he isn’t oriented towards their welfare.⁴⁴

By contrast, now consider Jeff, who does love his children a great deal and who is concerned for their welfare. No doubt, it’s plausible that Jeff’s love for his children and their well-being will be accompanied by sincere beliefs and judgments about his love for them. Perhaps when he sees his children having fun or succeeding in their projects, he thinks to himself, “I sure do love my kids” and “my kids sure are the best!” But notice, such beliefs are not fundamental to his love for his children. Rather, beliefs of this sort *reflect* Jeff’s prior orientation towards his children and their welfare as individuals to be loved.⁴⁵

⁴⁴ All this also suggests that an individual’s evaluative judgments aren’t even a particularly reliable guide to how she is practically oriented. In my view, a better guide are her emotional dispositions.

⁴⁵ And indeed, even though beliefs and evaluative judgments can reflect our orientation towards an object, they can just as easily reflect out capacity for post-hoc rationalization. It seems to me that our emotions are much more revealing attitudes and consequently better at reflecting our concerns. I speculate that this is because our emotions actually play a much larger (perhaps even essential) role in structuring our concerns. I’ll say more about this in Chapter 7.

But Jeff's love for his children isn't merely reflected in his beliefs. Perhaps even more significantly, Jeff's love for his children is also reflected in his emotional engagement with them—he's happy at their successes, sad for their losses, and angry when they are wronged. Moreover, his love is reflected in his presumptive interpretations of their attitudes and actions—he's willing to give them the benefit of the doubt, willing to see things from their point of view, and willing to show deference when it's appropriate. We can also see his love reflected in the motivational states that it issues in—a desire that his children be happy, a hope that their futures are bright, and an evaluative judgment that they deserve the best. And finally, his love is reflected in his actions—he's there to pick them up even when it's inconvenient, he makes sure they eat first if there's not much food, and he reads them bedtime stories, even when this might cause him to miss the Big Game (you see, Jeff's an avid sports fan).

Simply put, Jeff's love for his children is an orientation he takes towards them and towards their welfare as individuals to be loved. And his love manifests itself in broad sets of conative dispositions—dispositions that motivate him to think about his children in particular ways (e.g., as loved and as valuable), to feel certain things about them (e.g., happiness, warmth, etc.), to engage with and react to them in particular ways (e.g., to praise, and to treat them in specific ways. Of course, if he messes up from time to time or if the “strains of involvement” become momentarily too burdensome (to borrow P. F. Strawson's [1974] phrase), it doesn't mean Jeff doesn't love his kids. After all, his practical orientation is reflected in his dispositions, his *proneness* to interact with his children in these particular ways. But if he lacks these sets of conative dispositions altogether—if he never interacts with his

children in these particular ways—then plausibly, he isn't practically oriented towards his children in a loving way.

Of course, love is just one way we can be practically oriented towards something or someone. It's a particularly intense and meaningful orientation, and so, when we love someone or something, our love for that person or thing will resonate in our conative dispositions to a great degree. It will affect us deeply, you might say. But we can also be oriented towards someone or something by caring about it or by being concerned for it. For example, Dan is concerned with jazz and with music more generally—i.e., he cares about it in a way that Joe does not. And while it is unlikely that Dan's concern for jazz is as deep as Jeff's love for his children,⁴⁶ structurally, Jeff's love and Dan's concern are relevantly similar. Just as Jeff's love orients him towards his children's well-being, Dan's concern for music orients him towards jazz. And again (as is also the case with Jeff), this orientation is manifested in a wide array of conative dispositions. In light of his orientation, Dan is prone to judge relatively obscure and hard-to-listen-to jazz as aesthetically pleasing; he's also disposed to feel disappointed when the record store doesn't have any new arrivals; he hopes for and wants to discover new, exciting jazz; and in virtue of his orientation towards music, he tends to spend a lot of money on jazz records. So Dan's orientation towards music, like Jeff's loving orientation towards his children's welfare, is a broad set of conative dispositions that plays some role in structuring his practical identity.

To sum up, the practical orientations that structure our practical identities can be understood as distinct stances or attitudes (in the aeronautical sense) we have towards our

⁴⁶ I certainly hope not!

physical, social, inter-, and intrapersonal environments that shape and direct our perception of and interaction with the world. Though our orientation towards a particular person or thing manifests itself in cognitive states (e.g., beliefs and evaluative judgments), our stance or attitude towards that person or thing is not reducible to or constituted by a set of beliefs or judgments.

Of course, this doesn't mean that we can never change our general orientations (though admittedly, doing so is often hard) or that some of them aren't reasons-responsive. But it does mean that we can only revise, alter, or abandon an orientation from the perspective of some other orientation. For example, suppose Dan's father thinks that not only is jazz a waste of time, but that people who like it are idiots. And further suppose that these judgments issue from Dan's father's general orientation towards music as being frivolous and wasteful. But now imagine that over time, after Dan's father discovers his son's love for jazz, he begins to soften his stance. For instance, when he hears jazz now, he doesn't immediately turn it off. And as he listens to it more, and as he talks to Dan about it more, he comes to appreciate it more. Soon he finds himself *wanting* to hear it. And through these small changes, eventually, he comes to be oriented towards jazz in an altogether different way. Indeed, the valence of his stance has shifted from negative to positive. But Dan's father doesn't change magically. Nor does his practical orientation change as a result of any particular decision or evaluative judgment that he makes (e.g., that it would be good to change his opinion). Instead, it changed because his stance on jazz conflicted, to a certain degree, with his orientation towards Dan, his son. Indeed, it's because he loved his son that he would give jazz a chance—that he could find it meaningful or significant at all.

But though many of our practical orientations are subject to change (through natural processes) or revision (through our own intentions), I'm doubtful that all of our orientations or stances can change so easily—or at all. For example, our orientation towards our own agency and towards the agency of others—an orientation that P. F. Strawson [1974] characterized as the “Participant Attitude”⁴⁷—is perhaps the most general practical orientation possible.⁴⁸ And if this orientation is the most general practical attitude possible—indeed, if it is necessary for having a practical orientation at all—then we cannot dissociate from it the way Dan's father dissociates from his unhappy orientation towards jazz. Without this orientation, we might think, we wouldn't have a practical identity at all. Such an orientation (if there is one) wouldn't just play an important or significant role in structuring our practical identities—it would play a constitutive role. But I bring this up here not as an argument in favor of thinking that, e.g., the Participant Attitude is an essential constituent of our practical identity. Rather, I introduce this possibility simply to suggest that it's not obvious that the ways in which we might be oriented towards the world are wholly up for grabs. Some orientations we find directing and shaping our engagement with the world might be basic, at least insofar as we are the kinds of agents we frequently take ourselves to be.

There's obviously more that I could say about the practical orientations or stances we take towards ourselves, others, and the world. But for present purposes, I think I've said enough to give us a good sense of what I mean when I say that for Structuralism, reasons-giving authority is tied to our practical orientations. So now that we've gotten a handle on

⁴⁷ Stephen Darwall [2006] similarly characterized this orientation as the “second-person standpoint.”

⁴⁸ I'll return to this possibility in Chapter 7 where I offer a Structuralist account of moral reasons.

what it is to be practically oriented towards someone or something, I want to say a bit more about the normative significance of being practically oriented towards the world in this way. Specifically, I want to explain why motivational states issuing from them have reasons-giving authority.

6.2.

It is probably already clear why I think those motivational states issuing from those practical orientations that structure our practical identities have reasons-giving authority. Because our practical orientations structure our practical identities—i.e., because who we are as agents is constituted by our particular *takes* or *attitudes* on the world and the concerns that are determined by these perspectives—when motivational states issue from our practical orientations, they speak for us in a way that, were those motivational states effective, we would be able to make sense of ourselves as agents, even if it is only, as the case would be instances of weakness of will, an ability to see ourselves as lazy or cowardly or otherwise irrational.

After all, to act for a reason is to act in a way that allows you to make sense of yourself.⁴⁹ Thus, because some motivational states, if effective, allow us to make sense of ourselves (given their connection to our practical identities), it's plausible that these motivational states would have reasons-giving authority. In other words, motivational states issuing from our practical orientations have reasons-giving authority precisely because these states are the sort that could explain reasons in a way that takes seriously the fact that when we act for reasons, we are thereby explicable to ourselves.

⁴⁹ Cf. David Velleman [1992].

§7.

So to conclude, Structuralism claims that an agent's reasons depend on her motivational states. But on Structuralism, not all motivational states are created equal. So, it's specifically those motivational states issuing from an agent's practical orientations that have reasons-giving authority. As I said earlier (in §3.1.), reasons are facts that bear on practical questions, and so we should understand Structuralism as a reductive analysis of *what it is* for some fact to bear on a practical question. Thus, on Structuralism, facts bear on practical questions if and only if an agent's acting on those considerations would promote the objects of those motivational states that issue from those practical orientations that play some role in structuring the agent's practical identity. And while Structuralism is indeed more complicated than other versions of HTR (e.g., Proportionalism), it seems to me that Structuralism is an elegant analysis of the *bearing on* relation that is constitutive of an agent's reasons.

Unfortunately theoretical elegance and truth sometimes come apart. And many have claimed that reductive analyses of reasons (e.g., Structuralism) are unlikely to be true. Such analyses are said to be ill-motivated, unable to explain the myriad ways in which reasons putatively bear on practical questions, extensionally inadequate, and unable to accommodate what many take to be indisputable—viz., that moral reasons are categorical and have practical priority. In the remainder of the dissertation I consider each of these objections in turn, and I hope to not only defend Structuralism but to positively argue that Structuralism isn't simply compatible with the categoricity and practical priority of moral reasons, it's

actually particularly well-suited to explain these putative properties of moral reasons. And if I'm right, then Structuralism won't just be another elegant but false theory—which are a dime a dozen anyway. It might possibly be true.

Chapter 3: Towards a Well-Motivated Humean Theory of Reasons

§1.

It's not enough to simply provide an elegant Humean theory of reasons. Such a theory must be well-motivated. That is, if the heart of the Humean theory is that an agent's normative reasons depend on non-normative psychological states of the agent, then similarly, we should expect that properties of an agent's reasons—e.g., the weights of those reasons—will also depend on non-normative psychological states of the agent. In other words, a Humean theory that explains the weights of reasons in wholly normative terms (i.e., a Humean theory that eschews an explanation of the weights of reasons in terms of properties of non-normative conative states) is not well-motivated. In this chapter, I consider a recent objection to Humean theories along these lines due to David Enoch, and I argue that Structuralism—as I developed it in Chapter 2—is able to answer Enoch's challenge in a well-motivated way. And this fact alone, I conclude, gives us reason to favor Structuralism to a brand of Humeanism developed by Mark Schroeder, Hypotheticalism.

David Enoch has recently argued that Humean theories of practical reasons should accept Proportionalism: a view of the weights of reasons according to which the *strength* of an agent's desires explain the *weights* of that agent's reasons.⁵⁰ This account of the connection between the desires that explain an agent's reasons and the weights of those reasons holds that the weight of *S*'s reason *R* to φ depends on (i) the strength of the desire *d* that explains *R* and (ii) how well *S* will promote the object of *d* by φ -ing.⁵¹ And at first blush,

⁵⁰ David Enoch, "On Mark Schroeder's Hypotheticalism," *The Philosophical Review* 120 [2011] 423-446.

⁵¹ Mark Schroeder, *Slaves of the Passions*, [2007a] and "Weighting for a Plausible Humean Theory of Reasons," *Nous* 41:1 [2007b] 138-160.

it is natural to think that Proportionalism and Humean theories go hand in hand, such that it would be, as Enoch suggests, embarrassingly ad hoc for the Humean to reject Proportionalism as an account of the weighting of reasons.⁵² After all, Humeanism is an attractive position precisely because of its explanatory power, a power it derives from a reductive analysis of reasons in terms of non-normative states of the agent. And rejecting Proportionalism appears to involve a rejection of a fully reductive account of the *bearing on* relation that is constitutive of practical reasons.

Unfortunately for the Humean, grounding the weights of reasons in the strength of desires is implausible for a variety of reasons. So blind acceptance of Proportionalism doesn't appear to be an attractive alternative for the Humean either. Indeed, as I suggested in Chapter 2, the fundamental problem with Proportionalism is its implausible commitment to "might makes right." Moreover, according to Schroeder, it's actually an unquestioned acceptance of Proportionalism that has led to many of the problems associated with Humean theories of practical reasons. For example, Humeanism is often charged with two-way extensional inadequacy. That is, Humeanism is often thought to entail that agents would have *too many* and *too few* practical reasons.⁵³ But Schroeder argues that these objections don't undermine Humeanism *qua* theory of practical reasons, and instead, these objections really only serve to highlight the problems with Proportionalism.

⁵² Indeed, as we'll see, Enoch [2011] motivates this natural thought.

⁵³ I won't rehearse these arguments here, but Schroeder [2007a, 2007b] and Enoch [2011] have given careful attention to these objections to Humeanism. I consider these objections to Structuralism in subsequent chapters.

After all, Schroeder argues, if most of the too many reasons I am said to have on Humeanism are of insignificant weight—i.e., such little weight that it would be a mistake for me to use them as premises in practical reasoning or mention them in conversation—then it is not a strike against Humean theories that they entail the existence of such reasons. But we can assign the intuitively correct weights to such reasons (i.e., insignificant) only if we reject Proportionalism since surely, some of the incredibly unweighty reasons in question—e.g., my reasons to eat my car—will depend on relatively strong desires—e.g., my desire to take care of the iron deficiency in my diet. Similarly, the thought goes, if we reject Proportionalism, then Humeanism can accommodate the idea that some particularly weighty reasons—e.g., moral reasons—are explained by relatively weak desires—e.g., a desire to make sense of myself or a desire to engage in meaningful reciprocal relationships. If this is correct, then plausibly Humean theories of practical reasons do not predict too many or too few reasons for action.⁵⁴ Thus, we can sidestep some of the worries concerning extensional adequacy that have plagued Humean theories simply by rejecting Proportionalism.

§2.

For all this, it looks like the Humean should simply reject Proportionalism (as I do in Chapter 2) and be done with it. Of course, as a consequence of rejecting Proportionalism as an adequate reductive analysis of the weighting of reasons, the Humean might conclude that although we can reductively analyze reasons in terms of an agent's desires, we cannot reductively analyze the *weights* of reasons in terms of natural properties of an agent's desires,

⁵⁴ Strictly speaking, a bit more needs to be said. Most notably, this is only correct if we accept certain theses concerning promotion. For more, see Schroeder [2007a] and Enoch [2011]. See also Chapters 5 and 6 of this dissertation.

viz., the *strength* of those desires.⁵⁵ And some Humeans, e.g., Schroeder [2007a, 2007b], are certainly happy to concede that this is the cost associated with offering a plausible account of the weighting of reasons. These theorists instead prefer to argue (or “settle” if Enoch is to be believed) for non-reductive accounts of the weights of reasons. This move itself can indeed seem strange. After all, if one of the commitments that motivates Humean theories is that what reasons an agent has will ultimately depend on who she is practically, then why would we expect an important property of her reasons—viz., their weight—to be independent of who she is practically? But this is precisely what Schroeder’s Hypotheticalism seems committed to. And while there is no outright contradiction in such a view, it does seem unmotivated.

Relatedly, Enoch provides some suggestive reasons to think that it isn’t so easy for a Humean theory of practical reasons to disentangle itself from reductive analyses like Proportionalism. Specifically, Enoch argues:

Desires come in different strengths. Reasons come in different weights. If you’re a Humean, you’re committed to a very close connection between one extreme on both these continua: no explaining desire, you are committed to saying, entails no reason. Furthermore, in that case you are committing to saying that there is no reason *because* there is no explaining desire. If you reject Proportionalism, though, you are committed to there *not* being a similar connection between other points on the continua. That the relevant, explaining desire is very weak, for instance, does not entail that the relevant, explained reason is very weak, and so forth. Why this discontinuity, though? The null-point on the continuum is just another point on the continuum. If we reject Proportionalism, isn’t it extremely surprising that we accept

⁵⁵ A note: in Chapter 2, I followed Schroeder [2007a] in understanding Humeanism [HTR] as a parity thesis. And as I said there, many forms of HTR simply accept that an agent’s desires determine her reasons. Obviously, it’s actually a great deal more complicated than this, and Structuralism reflects that fact. However, because Enoch develops his challenge against anti-Proportionalist versions of HTR as if desires were sufficient to explain an agent’s reasons, I will follow him for this chapter. Despite this, what I say here applies *mutatis mutandis* to a fully articulated Structuralism, which claims that an agent’s reasons are ultimately determined by those mental states that issue from those practical orientations that play some role in structuring the agent’s practical identity. So please forgive the gross simplification of Structuralism found in this chapter—it’s simply easier to deal with Enoch’s challenge in this way.

the particular instance of Proportionalism applied to the case of desire of strength zero, or no desire at all? ... if you want to dance, that there will be dancing at the party is a reason for you to go to the party, and if I don't want to dance, then it isn't a reason for me to go to the party. Just as plausibly, though, we can say that if you have a *very* strong desire for dancing, and I have a *vanishingly* weak desire for dancing, then that there will be dancing at the party is a much weightier reason for you to go than for me, [Enoch 2011, 439-40].

Thus, if Enoch is right, those Humean theories that reject Proportionalism are in trouble.

And although he (explicitly) isn't arguing that Humeanism *logically* entails Proportionalism in the sense that it would be inconsistent for a Humean to reject Proportionalism, Enoch *is* providing a challenge for Humeans who, like Schroeder, want to reject Proportionalism. And because I also think that all Humeans *should* reject Proportionalism, I think it's important that we have a convincing reply to Enoch's challenge. Moreover, because Structuralism is itself an alternative to Proportionalism, it must maintain a delicate balance between the motivations that lead us to Humeanism in the first place and a theoretically adequate account of the weighting of reasons. Only then will anti-Proportionalist theories (like Structuralism) have answered Enoch's challenge.

In other words, as I understand Enoch's challenge, it goes something like this: the Humean is *already* committed to thinking that if an agent lacks an explaining desire that will be promoted by φ -ing, then she also lacks a reason to φ . Indeed, this is precisely the commitment in virtue of which the view in question is recognizably "Humean." But if you accept this—i.e., if you accept that in the limiting case there seems to be a natural fit between having no desire (or perhaps a desire of zero strength) and having no reason—then once you abandon Proportionalism as a general account of the relation between strength of desires and the weighting of reasons, how can you explain why the necessary connection between strength of desires and weights of reasons is only applicable in the limiting case?

And more fundamentally, the challenge is simply captured with Enoch's thought(s) that "if you start off rejecting Proportionalism, it is not clear why you would accept Humeanism. And if you start off Humean, it is not clear why you would reject Proportionalism," [Enoch 2011, 440].

Now it seems to me that this is an important challenge to those Humean theories that, like Schroeder's Hypotheticalism and Structuralism as I developed it in Chapter 2, reject a connection between the strength of desires and the weights of reasons, i.e., those Humean theories that reject Proportionalism. And I think such theories need to motivate the rejection of Proportionalism in a way that appreciates the force of Enoch's challenge. Accordingly, in what follows I will sketch a reply to Enoch that involves a Structuralist account of the weighting of reasons that serves as an alternative to Proportionalism. Unlike Schroeder's [2007a] recursive account of the weighting of reasons, I will advance Structuralism as a reductive account of the weights of reasons. For reasons I suggested in Chapter 2, Structuralism appreciates a fundamentally *Humean* commitment to thinking that just as the existence of a reason is explained by an agent's desires, so too the weight of her reason is explained by natural, non-normative properties of her desires.⁵⁶

In §3 I will briefly present Structuralism. In §4 I sharpen Enoch's challenge by considering two distinctive ways of framing the challenge and argue that Structuralism has the resources to meet Enoch's challenge in either guise.

⁵⁶ Of course, like Schroeder's own account, my account of the weighting of reasons will be extensionally adequate (i.e., it will be consistent with Schroeder's replies on behalf of the Humean to the too many and too few reasons problems). Moreover, as I argue in Chapter 6, this account is also consistent with a conception of moral reasons as having practical priority.

§3.

Recall that, according to Proportionalism, the weight of S 's reason R to φ is proportional to (i) the strength of explaining desire d and (ii) how well d 's object is promoted by S 's φ -ing.⁵⁷ Thus, Proportionalism is a reductive analysis of the weighting of reasons because it explains the weight of R in terms of d 's natural non-normative properties, viz., d 's strength. Of course, as I have argued [Chapter 2], Proportionalism is committed to an implausible connection between a normatively irrelevant property of desires (viz., their causal strength) and the weight of reasons explained by those desires—a connection that mirrors the implausible “might makes right.” Further, as Mark Schroeder [2007a, 2007b] has argued, it's also this very feature of Proportionalism that renders it poorly suited to provide an extensionally adequate theory of reasons.

By contrast, according to Structuralism, the weight of S 's reason R to φ is proportional to (i) the role D plays in structuring S 's practical identity, where D is the broad set of conative dispositions that issues in d , and (ii) how well d 's object is promoted by S 's φ -ing. Although we might understand Structuralism itself as a form of Proportionalism, unlike Proportionalism as Schroeder defines it, Structuralism makes no mention of the *strength* of desires in its analysis of the weighting of reasons. In fact, on Structuralism, the comparative strengths of desires are wholly irrelevant to the weights of reasons. Despite this fundamental difference, Structuralism does follow Proportionalism in claiming that the weight of an agent's reason is *proportional* to some natural non-normative properties of the explaining

⁵⁷ As it turns out, there is often little discussion of this second feature of Proportionalism, and even though I will appropriate (ii) for Structuralism, I will continue the tradition of saying very little about it.

desires.⁵⁸ And instead of invoking *strength*, Structuralism invokes the notion of *the structural role that an explaining desire plays in an agent's practical identity*. Put simply, the idea behind Structuralism is that reasons which are explained by those desires that play a more significant role in structuring an individual's practical identity are comparatively weighty. And conversely, reasons which are explained by desires that play a less significant role in structuring an individual's practical identity are relatively unweighty.

To help see this, consider the following example. Say it's a typical Sunday afternoon in the fall. Like most Sunday afternoons in the fall, Jerry likes kicking back, putting his feet up, having a beer, and watching football. In fact, this is what he *really* likes doing on Sunday afternoons in the fall because he finds it a relaxing way to end the weekend.

But suppose that on Sunday afternoon Jerry's wife tells him that she's had a really rough week, and that she would like to spend some meaningful time with him—maybe by taking a nice walk. Now, when Jerry's wife tells him this, we can imagine that in this scenario, he has at best, only a weak desire to get off the couch and take a walk with her and that he retains his very strong desire to sit around watching football. Therefore, Proportionalism seemingly entails that those reasons explained by Jerry's desire to watch football are comparatively weightier than those reasons explained by his desire to take a walk with his wife. And if that's right, then there's at least *prima facie* reason to conclude that Jerry should ignore his wife and sit around watching football. After all, what we *ought* to do, all things considered, is a function of the weights of reasons. But this is obviously not the case since intuitively, he in fact has weightier reasons to take a walk with his wife.

⁵⁸ However, for the sake of clarity, I simply call this view "Structuralism" rather than "Proportionalism," which is identified specifically with thinking that it's the *strength* of the explaining desires that is relevant to the weighting of reasons.

Unlike Proportionalism, Structuralism can appreciate that Jerry should turn off the game and take a walk with his wife. And Structuralism would arrive at this conclusion by considering the practical orientations from which Jerry's desires to stay firmly fixed to the couch and to go on a walk with his wife respectively issue. Concerning the former, it's plausible to think that Jerry's desire to watch football issues from a practical orientation that structures some of Jerry's interests and hobbies. And no doubt, Jerry's particular interests and hobbies play a large role in structuring his practical identity. Much of what he discusses with friends and coworkers—the glue that keeps the fabric of many of these relationships together—centers on sports-related small talk. And this suggests that the particular desires that issue from this set of conative dispositions (e.g., his desire to watch football at that time) presumably give Jerry sufficiently weighty *pro tanto* reasons for taking it easy on Sunday afternoons in front of the television.

However, Jerry's relatively weak desire to go on a walk with his wife issues from the loving practical orientation that Jerry has towards his wife—an orientation that plays a much larger role in structuring Jerry's practical identity. Jerry's practical identity—i.e., his values, concerns, obligations, motivations, deliberative perspective, etc.—would be fundamentally different without his wife. Their joint activity of making a life together has played a significant role in molding Jerry into the person he is today. And while this practical orientation issues in a relatively weak desire to take a walk with his wife in this particular context, it's plausible to think that if Jerry were to dissociate himself from this practical orientation (i.e., if he were to abandon this attitude that shapes and directs his engagement with his wife, and come to regard it as somehow “alien”), he would do more serious violence to his practical identity than would be the case if he were to give up on the practical

orientation that leads him to enjoy football so much. And this reflects the fundamental fact (it's a fundamental fact for Jerry at least!) that the practical orientation that issues in his desire to take a walk with his wife play a more significant role in structuring Jerry's practical identity than does the practical orientation that issues in Jerry's desire to watch football. And as a result, on Structuralism, the reasons that are explained by these respective desires have different weights for Jerry, i.e., he has weightier *pro tanto* reason to take a walk with his wife than to watch football. And if this reason is undefeated, then plausibly Jerry should, all things considered, take a walk with his wife. Thus Structuralism is able to correctly predict and explain which of Jerry's reasons are weightiest in cases in which Proportionalism cannot.

Obviously, this restatement of Structuralism is hardly a full defense. After all there is a great deal more that could be said to develop and defend Structuralism as a plausible theory of the weighting of reasons. For example, as we saw in Chapter 2, unlike Proportionalism, Structuralism does not mistakenly attribute any right-making properties to the causal *strength* of an agent's desires. Moreover, as I argue in Chapters 5 and 6, Structuralism (but not Proportionalism) provides an extensionally adequate theory of reasons, predicting neither too many nor too few reasons. Relatedly, as I argue in Chapter 7, Structuralism is compatible with a conception of moral reasons as having *practical priority*.

Of course, these claims are simply a promissory note of things to come. However, meeting Enoch's challenge is something more preliminary—something that must be done before these further questions concerning Structuralism can even come up. After all, it is by meeting Enoch's challenge that we can see Structuralism to be not just an internally consistent Humean theory of reasons, but a *well-motivated* Humean theory of reasons as well. So on behalf of Structuralism, I want argue in §4 that it can meet Enoch's challenge.

§4.

To begin, it's worth noting that Enoch frames his challenge in two slightly different (but ultimately related) ways. First, Enoch's challenge targets a natural, often tacit, antecedent commitment to Proportionalism that seems implicit in Humean theories. And second Enoch's challenge leverages a general methodological point as an objection to Humean theories that, like Schroeder's [2007a] Hypoethicalism and Structuralism, reject Proportionalism as an adequate account of the weighting of reasons.

More exactly, Enoch first frames his objection as challenging anti-Proportionalist Humeans to explain why there appears to be a connection in the limiting case between strength of desire and weight of reason when according to anti-Proportionalism theories, this connection does not obtain in any other case. The thought here is that simply in virtue of accepting Humeanism, you are plausibly already committed to thinking that Proportionalism is the correct analysis of the weighting of reasons for some subset of an agent's reasons.

But second, Enoch's challenge can be understood as a more general (and perhaps implicit) methodological challenge concerning the theoretical motivation of anti-Proportionalist Humean theories—i.e., “if you start off a Humean, it is not clear why you would reject Proportionalism,” [Enoch 2011, 440]. This challenge rests on the idea that if, as the Humean claims, the existence of a reason ontologically depends on an agent's desires (along with some facts concerning the promotion relation), then it would be implausible to think that that a reason's properties, e.g., its weight, don't also depend on the desires in question in some way. And a natural way in which the weights of reasons could depend on

the desires in question would appeal to the strength of these desires, just as Proportionalism claims. Thus, according to Enoch, any form of Humeanism that rejects Proportionalism appears ad hoc and doesn't seem well motivated by *Humeanism's own lights*.

4.1

Regarding the first way in which Enoch frames his challenge, he claims that that “if we reject Proportionalism, [it would be] extremely surprising that we accept the particular instance of Proportionalism applied to the case of desire of strength zero, or no desire at all,” [Enoch 2011, 439]. The worry here is simply that if Humean theories are identified by an acceptance of the claim that desires of strength zero ground reasons with zero weight, then there is no principled reason for rejecting Proportionalism as a general account of the connection between the strength of desires and the weighting of reasons.

More exactly, what makes a theory of reasons distinctively *Humean* is its commitment to the following claim: if an agent lacks an explaining desire, then she does not have a corresponding reason. But according to Enoch, by accepting this commitment, the Humean *already* accepts a restricted form of Proportionalism. Thus, we can see Enoch as developing a parity argument—if Proportionalism explains the very case that identifies a theory as Humean, shouldn't the Humean think that Proportionalism will generalize? After all, it appears that the presence of this relationship (between strength of desires and weights of reasons) not only obtains in the limit case, but that it explains the uniquely Humean claim that reasons depend on desires generally.

In response to this way of framing Enoch's challenge, I want to distinguish between the following two claims.

- (1) If there is no explaining desire d , then agent S has no reason R to φ , and
- (2) If d has strength zero, then S 's reason R to φ has a weight of zero.

First, it's important to see that while Humean theories are committed to the truth of (1), they are not thereby committed to the truth of (2).⁵⁹ That is, on a plausible interpretation, (1) and (2) are nontrivially nonequivalent. According to (1), S 's reason R doesn't exist independently of d . By contrast, (2) is consistent with S having a reason R independently of d . Of course, according to (2) there is a tight relationship between d and the weight of R , such that if d is of strength zero, then R has no weight. But the weight of R is a property of R that presupposes R 's existence. Thus, while (1) makes claims concerning the ontology of reasons as such, (2) is a putative explanation of why some reasons have zero weight.

But the first way in which Enoch develops his challenge seemingly depends on (1) and (2) being equivalent. After all, Enoch's problem with Humean rejections of Proportionalism relies (in part) on thinking that Humeans accept (1), that (1) is equivalent to (2), and that it's unmotivated for a Humean to accept (2) without accepting Proportionalism as a general account of the relationship between the strength of desires and the weights of reasons (since Proportionalism is thought to explain the truth of (2)). Of course, the Humean can simply reject a conception of desires such that having no desire entails having a desire of zero strength. And if Humeans recognize the distinction here, then the Humean can reasonably reject the move from (1) to (2) that the first way of framing Enoch's challenge relies on.

⁵⁹ As I'm thinking of it, (1) corresponds to Enoch's claim with respect to the case of "no desire at all," and (2) corresponds to the case of "desire of strength zero."

Enoch himself notes this, but concludes that, at least in Schroeder's case, it's not clear that Schroeder's conception of desires allows for the move described above.⁶⁰ To this I would simply point out that whether or not Schroeder's own account of desires can make this distinction isn't germane to Enoch's deeper challenge to anti-Proportionalist Humean theories of reasons. After all, Enoch frames this challenge—the challenge of defending a well-motivated Humeanism while simultaneously rejecting Proportionalism—as a general problem for Humeans that reject Proportionalism, of which Schroeder is but one prominent example. Thus, I turn to Enoch's general challenge below.

4.2

So the first way we can understand Enoch as issuing his challenge fails because it commits the Humean to an entailment relation between (1) and (2) that is not required by Humeanism as such. But concerning the more general way of framing his challenge, Enoch writes:

The challenge ... is that *from a Humean perspective* there is something objectionably ad hoc about rejecting Proportionalism. To show that this is not so, it is just not enough to show that there are *other*, Humeanism-independent problems for Proportionalism. If Proportionalism is a very natural thesis for a Humean to accept, and if it is not sensible (because there are Humeanism-independent reasons to reject it), then this does not count in favor of a Humean view that nevertheless rejects Proportionalism. Rather, it counts against Humeanism. When a challenge has the form of (roughly speaking) a reduction ad absurdum of your view, you cannot just

⁶⁰ In fn 18 Enoch asks (and answers), "Can Schroeder perhaps utilize a distinction between no-desire and desire of strength zero? I am not sure, but I do not see how such a distinction could be defended given Schroeder's theory of desires (chapter 8)," [Enoch 2011, 439]. I have to confess that despite familiarity with Schroeder's account of desires, I'm not sure why Schroeder would be committed to thinking that there is no difference between no-desire and a desire of strength zero. But this interpretative point is beyond the scope of my reply to Enoch. Moreover, insofar as Enoch is making a point concerning problems with Humeanism more generally, then even if he is correct in thinking that Schroeder cannot distinguish between no-desire and desire of strength zero, such a distinction is possible given plausible assumptions concerning desires—assumptions that are well within the bounds of Humean framework.

retreat to being sensible, reject the absurdum (it is, after all quite absurd!), and then fix in your theory whatever needs fixing to avoid entailing the absurdum.... Rather, what is also needed is a convincing story of how the revision need to save the theory you started with from the absurd consequence can be motivated *from within your theory and its underlying motivations* [Enoch 2011, 438-39].

By building on this last point concerning philosophical methodology generally, we can sharpen Enoch's challenge in a way that makes problems for Humeanism as such. Specifically, on this way of understanding Enoch's challenge, if what motivates Humeanism is the attractiveness of a reductive account of normative reasons for action in terms of an agent's actual desires, then the Humean should be similarly concerned to reductively analyze an important property of reasons—viz., their weight—in terms of properties of desires as well. In other words, Enoch is offering the Humean a trilemma of sorts—either (i) accept Proportionalism and be saddled with its implausible consequences; or (ii) reject Proportionalism without offering a reductive alternative, thereby rendering the position embarrassingly ad hoc; or (iii) reject Proportionalism and in its place provide an alternative reductive analysis of the weighting of reasons that explains the distinctive Humean claim—i.e., (1) if there is no explaining desire d , then agent S has no reason R to φ —without appealing to extra-Humean motivations (i.e., non-reductive analyses). And it seems that for Enoch, the challenge will be met only by those Humean theories that (perhaps impossibly by his lights) satisfy (iii). So even if Schroeder's Humean theory of reasons—Hypotheticalism—cannot meet this challenge, if satisfying (iii) is in fact possible, the mere fact that Hypotheticalism putatively fails to do so shouldn't impugn Humean theories generally. Thus, we shouldn't understand Enoch as merely objecting to Schroeder's views, but as posing a challenge for all Humeans who want to reject Proportionalism. Accordingly,

I see Enoch's challenge as one Structuralism must meet if it is to be a well-motivated Humean theory of reasons.

Recall that Structuralism is a reductive analysis of the weighting of reasons. Because of this feature, in my view, Structuralism offers the Humean a distinctively *Humean* way of satisfying (iii). To see this, reconsider (1) and (2).⁶¹

(1) If there is no explaining desire d , then agent S has no reason R to φ , and

(2) If d has strength zero, then S 's reason R to φ has a weight of zero.

It seems to me that Structuralism can (for principled reasons) deny that the strength of an agent's desire explains the weight of an agent's reason—even in the limit case Enoch discusses. After all, (1) captures the distinctively Humean commitment to thinking that without desires, there are no reasons. However, the converse of (1) does not hold. That is, it's consistent with Humeanism that there are desires without corresponding reasons (e.g., Davidson's yen to touch the woman's elbow or Quinn's desire to turn radios on). Furthermore, because the converse of (1) does not obtain, it's consistent with (1) that an agent might have *strong* desires without corresponding reasons. So if the existence of a strong desire doesn't thereby entail the existence of a reason (let alone the existence of a weighty reason), then neither does the existence of a vanishingly weak desire (i.e., a desire of strength zero) entail the existence of a reason with zero weight. This is because the existence of a desire doesn't entail the existence of any reason at all—even Humeans must grant this!

And this would suggest that the Humean is not committed to the truth of (2), since the truth of (2) would wrongly entail that the existence of a desire of zero strength *entails* the

⁶¹ Although I have argued (convincingly in my view) that (1) and (2) are not equivalent, and that the first way of framing Enoch's challenge founders on this mistake, by reconsidering these claims, I think we can recognize a deeper point—viz., that strength of desire as such plays no role by the Humean's own lights.

existence of a corresponding reason with zero weight. Surely, even if the Humean accepts that desires of zero strength are possible, she does not mean to saddle herself to thinking that their existence entails the existence of corresponding reasons—even reasons of zero weight. That is, this point shows that the Humean isn't committed to thinking the strength of desires is tied to the existence of reasons, even in the limit case as Enoch suggests.

But the point I want to make here isn't simply that (1) and (2) aren't equivalent (though this is a lesson we can learn from this argument). Rather, it's that by the Humean's own lights, it's not the strength of an agent's desires that is explaining the existence of an agent's reasons. In thinking that we can reduce normative reasons to agents' desires, the Humean doesn't mean to commit herself to the absurd claim that any old desire entails a corresponding reason. And consequently, the Humean doesn't need to presuppose Proportionalism to be true in order to motivate or explain (1). If this is correct, then I think it takes some of the sting out of Enoch's challenge. But it doesn't suffice as a full answer to the challenge. After all, a full(er) answer doesn't merely provide a negative argument; it must deliver a positive motivation for the view in question. And this is where Structuralism comes in.

You see, Structuralism explains (1) by appealing to natural, non-normative properties of *S*'s desires. Thus, Structuralism, like Proportionalism, is a reductive analysis of the weighting of reasons in terms of natural properties of *S*'s desires. But unlike Proportionalism, Structuralism does not commit the Humean to an apparent connection between might and right, implausible conclusions concerning the extension of reasons, the practical priority of moral reasons, and so on. Instead, Structuralism offers the Humean an attractive, intuitive way of understanding the existence and weights of reasons.

Recall that according to Structuralism, the weight of an agent S 's reasons R to φ depends on the role O plays in structuring S 's practical identity, where O is the practical orientation that issues in the particular desire d whose object is promoted by S 's φ -ing. And (for the last time) recall (1):

(1) If there is no explaining desire d , then agent S has no reason R to φ .

As I've said, (1) doesn't entail that for any desire, there is a corresponding reason. But when conjoined with Structuralism, we *do* have a plausible account of when desires yield reasons. On Structuralism, a desire will serve as an explaining desire (i.e., it will generate a reason) if and only if that desire issues from some practical orientation that plays some role in structuring S 's practical identity. Thus desires-like states such as whims or flights of fancy do not generate reasons, even if such states are very strong, because these desire-like states do not issue from sets of conative dispositions that play a role in structuring our practical identities.

But this is just the intuition that leads many of us to Humeanism. That is, it's the thought that our reasons must come from our own motivational states. And of course, not just *any* motivational states an agent experiences are *her own* (in the relevant sense). It's the motivational states that are part of who we are *as agents*—i.e., those motivational states that play some role in shaping our own unique practical perspectives. In other words, the natural thought that our reasons are tied to our own motivational states is best understood in Structuralist terms—as picking up on the an analysis of reasons such that those desires which can generate an agent's reasons are the desires that issue from those practical orientations that constitute our practical identities.

And this thought explains why (1) is true (and notice, it does so without appealing to the strengths of desires). After all, if reasons depend on desires that issue from the practical orientations that play some role in structuring our practical identities, then without such states, there would be no reasons.

Once we see this, we can see how Structuralism is in a position to answer Enoch's challenge. Recall that Enoch challenges anti-Proportionalist Humeans to motivate their accounts of the weighting of reasons in distinctively Humean terms. This challenge constrains candidate accounts of the weighting of reasons to those accounts that are motivated by the same considerations that led us to Humeanism in the first place. And Structuralism, as an account of the weighting of reasons does just this. After all, it is Structuralist intuitions, which as I pointed out in Chapter 2, reflect a commitment to the thought that what reasons we have ultimately depends on who we are as agents, that led many of us to Humeanism in the first place, and these same instincts are present in the reductive analysis of the weighting of reasons that Structuralism offers. Contra Proportionalism, it isn't the strength of our desires that matters for the weights of our reasons; rather, it's the significance of those desires in our holistically-evaluated practical identities. *Of course* Jerry's relatively weak desire to spend time with his wife on a walk is weightier than his comparatively strong desire to ignore her, have another beer, and stay glued to the couch. After all, who Jerry is as an agent depends on the practical orientations that issues in the former desire to a much greater degree. And once you put it that way, it's not clear why Proportionalism would've ever seemed to be a good candidate as a reductive analysis of the weighting of reasons.

§5.

Of course, none of this entails that Structuralism is the correct kind of analysis of the weighting of reasons.⁶² Moreover, for everything I've said here, upon further review, we might realize that Structuralism founders on some of the same problems that ran Proportionalism aground. But even if this is correct (and I hope not!), Structuralism is significant because it answers Enoch's challenge. That is, Structuralism offers us a well-motivated, principled way to reject Proportionalism while maintaining a version of Humeanism that is true to its basic commitment—that the normative realm is to be explained in terms of non-normative conative states of the agent.

⁶² Indeed, if Schroeder is correct then reductive analyses of the weighting of reasons like Structuralism aren't even the right *kind* of analyses.

Chapter 4: The Extensional Adequacy of Structuralism Part I:

Too Many Reasons

§1.

Structuralism, like all versions of HTR, begins with the idea that an agent's reasons ontologically depend on her motivational states for their existence. Yet Structuralism respects the intuition that not all motivational states can plausibly serve as the basis for reasons. For example, take Sean. Sean has decided to paint his bedroom, but as he stirs the paint, an odd sensation pops unbidden into his head: "What if instead of painting the walls with this paint, I ate it?" Ignoring it, Sean sets about painting. But as he's painting, Sean can't help but to wonder about what it would be like to eat the paint. And as he's wondering about what it would be like to eat paint, he finds himself wanting to taste it, and perhaps more shockingly, wanting to devour it. Thus, Sean has motivational states that would be promoted by his eating the paint.⁶³ And this would seemingly suggest—at least according to a bare statement of HTR—that the facts that the paint is readily accessible and that eating it would satisfy his desire are reasons for Sean to eat the paint.

But this putative implication of HTR is quite understandably, thought to be crazy. Suppose you were the attending physician who revived Sean after he ate the paint. You'd almost certainly ask something like, "What the hell were you thinking?!" And if Sean replied by saying that he had a reason to eat the paint and had gone on to explain this claim by appealing to something like HTR, you'd likely be even more incensed: "That's not a reason

⁶³ It's important that in the case, the motivational states in question are not motivated by instrumental concerns (say, Sean has made a substantial wager that he could drink a thimble of paint). For more, see Davidson [2001] and also, Quinn [1995].

to eat paint at all!”⁶⁴ And any theory of normative reasons that predicts that Sean’s yen to taste the paint is a motivational state that has reasons-giving authority for Sean, erroneously posits *too many reasons*, and is thereby extensionally inadequate as a theory of practical reasons.

Notice however, that Structuralism does not entail that Sean has a reason to eat the paint. On Structuralism, the existence of an agent *S*’s reason *R* to φ depends on a motivational state (or a set of motivational states) *m*, where *m* is a state issuing from one of the practical orientations that plays some role in structuring *S*’s practical identity. But despite the ease with which Structuralism skirts this worry, related concerns arise. Thus, in what follows I will consider two related problems concerning the extensional adequacy of Humean theories of reasons like Structuralism.

§2.

Of course, Sean’s is an easy case. It’s relatively clear that for Sean—like for most of us—yens to eat paint, even if we find ourselves confronted with such things from time to time, do not issue from practical orientations that play some role in structuring our practical identities. No one would say to Sean *sams* the desire to eat paint, “I don’t even know who you are any more,” the way someone might say that to Levi when he dissociates himself from his concerns for his family to run off with a younger woman. And plausibly, this is because unlike Levi’s concern for his family issues from a practical orientation that plays a large role in structuring his practical identity, Sean’s desire to eat the paint is merely a passing

⁶⁴ And if you were not only Sean’s physician but also his metaphysician, you might tell him that he needs to reject his theory of the metaphysics of reasons.

whim. And passing whims are just that—*passing*—and accordingly, they disconnected from the practical attitudes that structure our agency. And as a result, on Structuralism, they don't enjoy reasons-giving authority. So Sean really does have no reason to eat the paint.

But what can Structuralism say to a more difficult case—a case in which Structuralism posits that an agent has a reason, but intuitively, the agent has *no reason* at all? Because Sean's yen to drink the paint wasn't appropriately tied to his practical identity (i.e., it didn't issue from one of the practical orientations that played some role in structuring his practical identity), it lacks reasons-giving authority. However, unlike his yen to eat paint, many of Sean's motivational states *do* issue from practical orientations that play a role in structuring his identity. It would seem then, that on Structuralism, these motivational states do have reasons-giving authority. But this isn't obvious since it can plausibly seem as if there are possible cases in which Structuralism entails that an agent has a reason to φ , but that intuitively, in such cases the agent does not have a reason to φ .

And it's precisely *these* sorts of cases that generate the force of the *Too Many Reasons* objection. In what follows I want to consider what I take to be two of the strongest versions on this objection, the first of which I'll call the "Way Too Many Reasons" problem and the second of which I'll call the "Still Too Many Reasons" problem. I turn to the Way Too Many Reasons objection to Structuralism below in §3.

§3.

According to the Way Too Many Reasons problem, even granting Structuralism its circumscription of motivational states with reasons-giving authority (such that Quinn's functional state and Davidson's yen lack reasons-giving authority), its truth would

nevertheless entail the existence of too many reasons. To see this, consider Tracy. Tracy loves the outdoors. She especially enjoys bird-watching and hiking. Moreover, Tracy's concern for the outdoors issues from her practical orientation on the world. She is directed towards and engaged with nature as a thing to be not only enjoyed but also valued, to be conserved and never exploited. In light of this general attitude Tracy has towards nature, she values it in specific ways, and she is disposed to desire to spend time in it, to become angry when she sees it unnecessarily spoiled, to judge it worthy of Federal protection, etc. Among other things, Tracy finds special enjoyment in the grandeur of nature. Thus, given how much Tracy loves spending time in nature, it's plausible that on Structuralism the fact that a hiking trail offers spectacular views gives Tracy a reason to take a hike. After all, if Tracy were to take a hike on that trail on the basis of that fact, she would be promoting the object of a desire of hers (e.g., that she sees beautiful scenery) that issues from a practical orientation that plays some role in structuring her practical identity (viz., her orientation towards nature).

But suppose that one night, as she's out with friends, in transit to a bar, Tracy's friend Marvin has a heart attack. Tracy immediately changes directions and heads quickly towards the hospital. But further suppose that en route, she passes the trailhead of one of her favorite places to hike. If Structuralism is true, then it would seem that at that moment, Tracy has a reason to stop and take a hike. After all, by hiking the trail, she would promote the object of a desire that issues from one of the practical orientations that structures her practical identity. But intuitively, given Marvin's immediate needs, Tracy has *no reason at all*

to stop for a hike. Thus like other forms of HTR, Structuralism seems to be extensionally inadequate.⁶⁵

We can restate the problem a bit differently. You might think that (among other things) a practical reason is something that can serve as a premise in an agent's practical reasoning.⁶⁶ But intuitively, the fact that the trail in question would provide spectacular views should not play a role in Tracy's practical reasoning in the current situation. That is, in deciding what to do, Tracy should not consider putting off taking her friend to a hospital so that she might have a chance to take a hike. Thus, the fact that the hike would provide potentially spectacular views is not a reason for Tracy to stop in this case. And again, this points to the extensional inadequacy of Structuralism, since on Structuralism, this fact is a reason for Tracy.

§4.

In response to the Way Too Many Reasons problem under consideration, I first consider a specific reason for thinking that in the case described above (and in relevantly similar cases), Tracy does in fact have a reason to take a hike.⁶⁷ I then appeal to Structuralism's account of the *weighting of reasons* to assuage any further worries you might have concerning whether these sorts of cases pose challenges for the extensional adequacy of Structuralism.

⁶⁵ This example is due to Gary Watson [personal conversation].

⁶⁶ Watson [personal conversation] suggested this version of the objection to me, which is inspired by the account of reasons in Richard Kraut's *What is Good and Why: The Ethics of Well-Being* [2007].

⁶⁷ This response is due to Mark Schroeder [2007a].

In defense of Humean theories of reasons more generally, Mark Schroeder [2007a, 2007b, and 2007c] has argued that judgments concerning negative reasons existentials (e.g., “Tracy has no reason to take a hike”) are unreliable.⁶⁸ According to Schroeder, we often initially judge that an agent has no reason to φ even though, upon reflection, it is clear from the case that the agent does in fact have a reason to φ . Call this “Schroeder’s point.”

To illustrate Schroeder’s point, imagine that as Tracy is turning the car around to head towards the hospital another of her friends suffers a heart attack. In this case, with two people’s lives on the line (imagine that neither has died of complications yet), plausibly, when she passes the trail that affords spectacular views, Tracy has *even less reason* to stop at the trailhead than she would if only one friend’s life was on the line. But if she has *less* reason in this case, then it would follow that she had *some* reason to stop in the previous case (i.e., the case in which only Marvin’s life was on the line).⁶⁹ In other words, because we think that Tracy would be acting more objectionably and less rationally in the second case, we can conclude that she has some reason, albeit small and easily outweighed, for stopping at the trailhead in the first case.

Of course, Tracy shouldn’t use her reason to stop at the trailhead as a premise in practical reasoning given the features of her current context. And because in this context her reason can’t serve this role that we often associate with practical reasons, we’re prone to make negative existential judgments concerning Tracy’s reason(s) to stop at the trailhead in the original case. But as Schroeder’s point teaches us, the connection between the role

⁶⁸ Schroeder argues for this by appealing to a case involving epistemic reasons. Although the case I offer below involves practical reasons, it is structurally parallel to Schroeder’s case.

⁶⁹ While colloquial, the locution “some reason” is imprecise. More precisely, to say that an agent has some reason φ is to say that she has a reason to φ of some non-zero weight.

reasons play in practical reasoning and their existence is too tenuous to warrant judgments of this sort. After all, just as it's not always appropriate to assert that p in conversation, it seems similarly plausible that it's not always appropriate to take some consideration R to be a premise in practical reasoning. But the impropriety of asserting that p in particular contexts doesn't affect whether p is true (i.e., in some cases, that p may be true, even though it's infelicitous to assert that p). Similarly, the impropriety of S taking R to be a premise in practical reasoning doesn't affect whether R is a reason for S .⁷⁰

Schroeder's point then, gives us a plausible basis for rejecting the intuition that motivates the Way Too Many Reasons Problem—viz., that in many cases, our judgments about an agent's reasons are mistaken. In §§4.1-4.2, I'll argue that Structuralism itself offers us further reasons for thinking that Humean theories can be extensionally adequate.

4.1

Building on Schroeder's point, Structuralism can explain why, in such cases, agents like Tracy should not take the spectacular views afforded by the trail to be deliberatively significant. Recall that on Structuralism, the *weight* that an agent S should place on a particular reason R to φ in deliberation will depend on:

- (1) The degree O plays a role in structuring S 's practical identity, where O is one of the practical orientations that play some role in structuring S 's practical identity, and
- (2) *How well* φ -ing will promote m , where m is the motivational state that issues from O .

⁷⁰ I follow Schroeder [2007a, 2007c] on this point.

I begin by focusing on (1) in §4.1, before turning to (2) in §4.2.

Consider that in Tracy's case, her reasons to stop at the trailhead stem from her love of the outdoors and her more general practical engagement with nature; she is engaged with and oriented towards it in a way that leads to her concern. But her reasons to go to the hospital without delay stem from her concern for her friendships and her particular love for her friends, and from the more general *participant orientation* that structures Tracy's love for her friends.⁷¹

And while Tracy's concern-based orientations towards the outdoors is an important feature of her practical identity, plausibly, it plays a significantly less important role in structuring her practical identity than does her general concern for meaningful social interactions or her specific concern for and love of her friends. So on Structuralism, any reasons that depend on motivational states issuing from practical orientations that play a less significant role in structuring her practical identity, such as her love of hiking, will be comparatively less weighty. And those reasons that depend on motivational states (e.g., her concern for friendship and her love of particular individuals as friends) that issue from a practical orientation that plays a significant role in structuring her practical identity will be comparatively weightier. Thus, Tracy plausibly has weightiest reasons to refrain from hiking.

4.2

Moreover, on Structuralism, the weight of a reason isn't simply a function of the role a relevant practical orientation plays in structuring an agent's practical identity. Recall that

⁷¹ I'll discuss the "participant orientation" or "participant *attitude*," in greater detail in Chapters 7-8. See also, Strawson [1974] and Darwall [2006].

Structuralism also holds that the weightiness of S 's reason R to φ will depend on: (2) *How well* φ -ing promotes m , where m is the motivational state that issues from O . Notice that according to (2), the details of the *promotion* relation that obtains between the object of a motivational state with reasons-giving authority and a particular course of action (e.g., φ -ing) are relevant for determining the weight of a reason to φ . Indeed, the details picked out by (2) are highly context sensitive, and this puts Structuralism in a position to better explain why Tracy's reason to stop at the trailhead is insignificant.

First, suppose that Tracy and her friends, like most folks who go to bars, are out at night. Though the trail in question offers spectacular views, when Tracy is on her way to the hospital, those views are mostly obscured by darkness. Moreover, if the trail is strenuous, it's not obvious that Tracy could safely get to the beautiful vistas in the first place. So even if Tracy acted on her concern for nature and stopped to go hiking, the fact that the trail offered beautiful views would do little to promote this concern in the context, since it is dark and it is unlikely that she would be able to enjoy the spectacular views. Thus, if Tracy acted on the fact that the trail offers spectacular views in the context, she would do very little to promote the object of her concern to enjoy nature. And this would further support Structuralism's claim that Tracy's reason is of little weight.

Of course, this point might seem like a Band-Aid, since it is merely a contingent fact about the case that it was dark. Wouldn't the object of Tracy's love of spectacular natural views be well promoted by stopping if it wasn't dark? And whether Tracy has Way Too Many Reasons or not shouldn't depend on the time of day Marvin has a heart attack.

Structuralism can further extend the above point by recalling that Tracy's practical orientation towards her friend Marvin plays a more significant role in structuring her

practical identity. Because our practical orientations interact with one another holistically, exerting a kind of internal pressure of integrity, Tracy's practical orientation toward her friends affects the way that her orientation towards nature shapes and directs her concerns. Indeed, the reasons-giving authority of any particular motivational state is not invariant; because the role that particular practical orientations play in structuring our practical identities is a function of how they holistically interact with other practical orientations, the motivational states issuing from any given practical orientation will not always have the same reasons-giving authority.⁷² So in light of Tracy's orientation towards Marvin, we shouldn't expect those desires and concerns that would be promoted by a late night hike to have significant reasons-giving authority for Tracy in the circumstances. After all, how well could she enjoy the spectacular views, even in the middle of a clear, beautiful day, if the whole time she was worried about Marvin and feeling guilty by abandoning him in his time of need? Presumably, not very well.

Taken together, (1) and (2) provide an explanation of why Tracy shouldn't use the fact that the trail affords spectacular views as a reason in her deliberation about what to do. Simply put: because this fact, if acted on, would promote the object of a motivational state with comparatively less reasons-giving authority, and because this fact, if acted on in the circumstances, would at best poorly promote the object of the motivational state in question, this fact is a reason for Tracy of relatively insignificant weight. Moreover, given the

⁷² This suggests that our practical orientations—at least many of them, but perhaps all of them—structure our practical identities in a relatively fluid way. This is perhaps reminiscent of Nietzsche's claim that the soul or person is the social structure of her drives [BGE 12, 19]. This view would suggest that fundamentally, all of an agent's practical orientations are on a par, and that it is merely a contingent fact that at some or other time, this or that practical orientation plays a larger role in structuring an agent's practical identity. However, in Part II I will explore the possibility that at least one of our practical orientations—what P.F. Strawson called the "Participant Attitude"—is actually fundamental in the sense that it is internal to the agent and no merely contingently related to her practical identity.

circumstances of her deliberation, this fact is irrelevant to Tracy, who in asking herself “what should I do?” has already bracketed a range of possible courses of action. After all, while the aim of deliberation as such might be something general like determining what an agent should do, this isn’t the case when we *actually* deliberate. When we actually deliberate, our aim is more particular; it is to determine what we should do *in this particular context*. And once we recognize the context-sensitivity of the deliberative question itself, we can get a better sense of why some facts aren’t admissible, even though they are reasons. Because Tracy’s reason to stop and hike bears on her particular practical question to such a small degree, she can (and should) ignore it in her deliberations—just as the irrelevance of the date of your mother’s birthday to the conversation at hand means you can (and should) leave it unsaid.⁷³

Thus, for Tracy, while she does have some reason to stop at the trailhead, it is decisively outweighed by her reasons to go to the hospital without delay. And given that it is outweighed, and more importantly, that it is so decisively outweighed, it would be downright offensive for Tracy to consider the spectacular views offered by that trail to be a reason for stopping in this case. So while Structuralism does predict that Tracy has a reason to stop in the case described above, it does so correctly (recall Schroeder’s point that our judgments concerning negative reasons existentials are particularly unreliable). Moreover, Structuralism can explain why, despite there being a reason for Tracy to stop, she should not consider that reason as a premise in her practical reasoning. Her reason to stop is decisively outweighed and rendered irrelevant given the context-dependent aims of deliberation. And because this

⁷³ The thought that the pragmatics of deliberation are relevantly similar to the pragmatics of conversation in this respect is due to Mark Schroeder [2007a]. However, it seems like a plausible claim independently of whether some versions of Humeanism is true. After all, it could explain why, in Williams’ case, deliberating about what the Principle of Utility or Categorical Imperative would prescribe is “one thought too many” [Williams 1981].

reason is so grossly outweighed by reasons stemming from Tracy's concern for Marvin and because her reason to stop and hike is bracketed in the particular context, it would be inappropriate for Tracy to consider it in deliberation. But as we know from Schroeder's point, this doesn't mean that such reasons don't exist. Thus, Structuralism does not erroneously predict Way Too Many Reasons for Tracy, and so it is not extensionally inadequate in this way. And furthermore, by invoking Schroeder's point about negative reasons existentials, Structuralism respects the intuition that leads many to think Humeans posit Way Too Many Reasons—viz., the intuition that if some form of Humeanism is true, then an agent will have lots of reasons that should not play a role in deliberation.

But even if Structuralism doesn't posit *Way* Too Many Reasons, if Structuralism is true, we might still worry that agents have *Still* Too Many Reasons.

§5.

Suppose we embrace Schroeder's point about negative reasons existentials. And further suppose that we embrace Structuralism's claim that reasons relevantly similar to Tracy's reason to stop and hike should play no role in deliberation because they are of insufficient weight and are bracketed by the context-dependent aims of deliberation. At this point, it seems that Structuralism has a compelling answer to the Way Too Many Reasons problem. Because Structuralism neatly explains the comparative weights of reasons, agents do not have way too many reasons.

But there is a related problem particular to Structuralism lurking nearby—call it the *Still Too Many Reasons* problem. According to the Still Too Many Reasons problem, although Structuralism can explain why Tracy's reasons are of such insignificant weight that they

should play no role in deliberation, we don't yet have an explanation of why, on Humean theories like Structuralism, an agent won't nevertheless have too many reasons that are significantly weighty. Thus, the Still Too Many Reasons problem is a problem for Structuralism's extensional adequacy qua theory of practical reasons. But it does challenge the extension of weighty reasons that Structuralism putatively entails, and because agents have too many weighty reasons—reasons which intuitively aren't weighty at all—Structuralism is inadequate as a theory of practical reasons.

To illustrate the Still Too Many Reasons problem, consider the following vignette from Dostoevsky's *The Brothers Karamazov*:

It was in the darkest days of serfdom, back at the beginning of the century—and long live the liberator of the people! There was a general at the beginning of the century, a general with high connections and very wealthy landowner, the sort of man (indeed, even then they seem to have been very few) who, on retiring from the army, feels all but certain that his service has earned him the power of life and death over his subjects. There were such men in those days. So this general settled on his estate of two thousand souls, swaggered around, treated his lesser neighbors as his spongers and buffoons. He had hundreds of dogs in his kennels and nearly a hundred handlers, all in livery, all on horseback. And so one day a house-serf, a little boy, only eight years old, threw a stone while he was playing and hurt the paw of the general's favorite hound. 'Why is my favorite dog limping?' It was reported to him that this boy had thrown a stone at her and hurt her paw. 'So it was you,' the general looked the boy up and down. 'Take him!' They took him, took him from his mother, and locked him up for the night. In the morning, at dawn, the general rode out in full dress for the hunt, mounted on his horse, surrounded by spongers, dogs, handlers, huntsmen, all on horseback. The house-serfs are gathered for their edification, the guilty boy's mother in front of them all. The boy is led out of the lockup. A gloomy, cold, misty autumn day, a great day for hunting. The general orders them to undress the boy; the child is stripped naked, he shivers, he's crazy with fear, he doesn't dare make a peep ... 'Drive him!' the general commands. The huntsmen should, 'Run, run!' The boy runs ... 'Sic him!' screams the general and looses the whole pack of wolfhounds on him. He hunted him down before his mother's eyes, and the dogs tore the child to pieces ... ! [Dostoevsky 2002, 242-43].

As usual, Dostoevsky is able to eloquently put his finger on a feature of human psychology that is as repulsive as it is familiar: not only are we capable of performing

horrible, repugnant evil, but we are ourselves very often horribly and repugnantly evil. Moreover, the General's treatment of the young boy isn't an evil that issues from some alien desire. That is, it's no Davidsonian yen or Quinean functional state that motivates the General to hunt the young boy down like an animal. Rather, his callous disregard for the boy issues from his orientation towards serfs as *things* or as *objects*. And unfortunately, because of his status, the culture in which he finds himself, and no small amount of self-deception, this practical orientation plays a very important role in structuring the General's practical identity. Dostoevsky himself hints at this by noting that the General "feels all but certain that his service has *earned him the power of life and death over his subjects*" [Dostoevsky 2002, 242; emphasis added]. And such a feeling undoubtedly issues from a more general, objective orientation towards serfs as creatures to be ordered about and not as persons to be engaged with reciprocally.⁷⁴

In the circumstances Ivan Karamazov describes above, the General's practical orientation towards his serfs as disposable manifests itself in any number of motivational states whose objects are promoted by the boy's death (e.g., his desire that the young boy be made an example of). And on Structuralism, this entails that the General has reason(s) to kill

⁷⁴ Presumably the General relates to his serfs from the Objective Attitude. According to P. F. Strawson, the Objective Attitude is an attitude or practical orientation we take towards things. And when we take the Objective Attitude towards persons, we are not regarding them as such. Concerning this practical orientation, Strawson writes:

To adopt the objective attitude to another human being is to see him, perhaps, as an object of social policy; as a subject for what, in a wide range of sense, might be called treatment; as something certainly to be taken account; perhaps precautionary account, of; to be managed or handled or cured or trained; perhaps simply to be avoided ... If your attitude towards someone is wholly objective, then though you may fight him, you cannot quarrel with him, and though you may talk to him, even negotiate with him, you cannot reason with him. You can at most pretend to quarrel, or to reason, with him [Strawson 1974; 10].

The General certainly takes it that the young boy is to be "managed or handled" for his "crime" rather than reprimanded or even taught, as one might teach a child they love with an eye toward their future standing as one with whom you can engage with reciprocally.

the boy. And because the General's orientation towards serfs in particular—indeed, his orientation towards all those he took to be his inferiors—plays such an important role in structuring his practical identity, it would follow that any reasons that the General has which are explained by motivational states issuing from this practical orientation are quite weighty.

So while Humean theories can reply to the Way Too Many Reasons problem by appropriating Schroeder's point and Structuralism's account of the weighting of reasons, it seems that no such response is available to the Humean in the face of the Still Too Many Problem. Put differently, the General poses a problem for Structuralism because he seems to present us with an agent who, because of his depravity, has too many weighty reasons to act in ways that strike us as irrational. In this case, he has too many weighty reasons because, intuitively, he has either no reason to kill the boy or (more likely given Schroeder's point) any reasons he does have to kill the boy are of insignificant weight. But it seems that on Structuralism, he has particularly weighty reasons to kill the boy. Hence, because such reasons are sufficiently weighty, the General still has too many reasons.

And notice, the General will still have too many reasons even if he has a *moral* reason to refrain from killing the boy—the problem of specifically moral reasons on a Humean view like Structuralism is distinct from the problem I'm currently considering (though as we'll see in Chapters 5-8, there is a unified solution to these problems). The problem under consideration here is simply that it looks like on Structuralism, agents can have extremely weighty reasons to act in evil, irrational ways—reasons that should play an important deliberative role for those agents. And to many, this seems to be a problem.

So, what can Structuralism say concerning horrible, repugnant individuals who have oriented themselves similarly to the General? Do they have very weighty reasons to act

wrongly? Do they even have *most* reason to act wrongly? Can Structuralism itself help us solve this problem?

§6.

There are two ways that Structuralism might respond to the Still Too Many Reasons problem. The first of these, which I consider in the remainder of Chapter 4, is modest. The second of these is a great deal more ambitious and I save my discussion of this ambitious reply for Part II: Moral Reasons.

The modest Humean response, traditionally associated with Hume himself, but taken up by contemporary moral theorists like Harry Frankfurt [2004] and Susan Wolf [1982] alike, rejects “the pan-rationalist fantasy of demonstrating—from the ground up—how we have most reason to live” [Frankfurt 2004, 28]. In so doing, this modest approach admits that although what the General did was ultimately wrong, because of who he was practically speaking, he plausibly had weighty reasons for doing it. This is undoubtedly unfortunate; the modest Humean need not deny this. And indeed, in light of how he acted, the modest Humean can say that we can *morally blame* the General in any number of ways: we might, as Alyosha Karamazov thought appropriate, kill him; we might, as it happens in Ivan’s story, remove him from his seat of power; we might hate him and allow our contempt to manifest itself in our treatment of him; we might set our wills decisively against the values and motives that led the General to act [cf. Tognazzini MS] or mark an impairment in our

relationship with him in light of the meaning of his action [cf. Scanlon 2008] or protest his wrongdoing in some other way [cf. Smith forthcoming].⁷⁵

But though we can lament his crime and blame him for it, we can't truthfully say that the General acted for no reason. Indeed, he did act for a reason—a reason that, given his practical orientations towards serfs, along with his loving orientation towards his hounds, was, according to Structuralism, quite weighty. The modest Humean then, admits that yes, in a sense, Structuralism posits Still Too Many Reasons. But she is quick to point out that it is only by rationalism's lights that it does so, and since rationalism is an Enlightenment fantasy, it's no strike against Structuralism that it fail to accommodate its aims.

§7.

Of course, even the modest Humean might deny that in Dostoevsky's case the General really does have *all things considered* most reason for murdering the young boy. After all, if some practical orientation that plays a more significant role in structuring the General's practical identity issues in motivational states whose objects would be promoted by not murdering the boy, then on Structuralism, those reasons will be weightier. And so while the General can consider reasons for killing the young boy in deliberation,⁷⁶ he should, all things considered act on those reasons which are weightiest. Of course, this may nevertheless leave the General with Still Too Many (weighty) Reasons. But, we needn't worry, says the modest

⁷⁵ Of course, you might worry that blame—specifically moral blame— would be inappropriate in such a case since, by hypothesis, although the General acted wrongly, he can reasonably defend himself by pointing out that he didn't act irrationally since he acted on his weightiest (set) of reason(s).

⁷⁶ Unlike Tracy's deliberation, the context-dependent aim of the General's deliberative question—“what should I do in light of the harm that this young boy has done to my favorite hound”—doesn't bracket his reasons for murdering the young boy.

Humean, since Structuralism can explain why, in this case, it would be irrational for the General to kill the young boy. To demonstrate the General's practical irrationality, the modest Humean must point to some practical orientation of the General's that (i) plays a more significant role in structuring his practical identity than that of his practical orientation towards serfs and (ii) issues in motivational states whose objects will be well-promoted by refraining from killing the boy. If this can be done, then the modest Humean could plausibly maintain that the General acted irrationally and that, though he still had weighty reason to murder the boy, he shouldn't have.⁷⁷

But this modest Humean strategy is perhaps easier said than done. We know from Ivan's telling of the tale that the General is especially sure of his place in society and by all appearances, wholly unconcerned with those who were beneath him. Indeed, it's plausible that (with the aid of a great deal of self-deception) the General himself would admit to this and simply laugh at any suggestion that this particular perspective to take on the world was objectionable. However, the modest Humean would never deny that the General is practically oriented (in a really fundamental way) towards his serfs as insignificant peons. But she might follow Hume in suggesting that plausibly, the General has even more fundamental practical orientations—ones that would issue in motivational states whose objects would be well-promoted by his refraining from murdering the young boy: This form of modest Humeanism can be traced to Hume's discussion of human nature in *A Treatise of Human Nature*:

⁷⁷ Again, notice that this modest Humean strategy, even if successful, concedes that an agent can have Still Too Many Reasons. So the aim of this strategy isn't to reply to the Still Too Many Reasons problems on its own terms. Rather it's to take much of the teeth out of the problem.

Let all the powers and elements of nature conspire to serve and obey one man: Let the sun rise and set at his command: The sea and rivers roll as he pleases, and the earth furnish spontaneously whatever may be useful or agreeable to him: He will still be miserable, till you give him some one person at least, with whom he may share his happiness, and whose esteem and friendship he may enjoy [B2.2.5].

Here Hume suggests that human persons—even evil ones like the General—have a very deep concern to be engaged with others in meaningful ways: to be loved and to be treated with respect; to be enjoyed and to be taken seriously; to be friends and confidants. For without such things, Hume suggests, an individual would not find life worth living. Indeed, it's plausible that much of the General's desire to be held in esteem, to be respected and taken seriously, and to be feared issues from this more general concern. And moreover, this concern to be held in esteem and enjoy mutual respect arises because we are oriented towards other agents not simply as things we encounter in the world, but as persons to be engaged with in meaningful, reciprocal ways.⁷⁸

If this is correct, then the modest Humean can plausibly claim that although the General had weighty reasons to murder the young boy, he had even weightier reasons to refrain. After all, because he is, by hypothesis, oriented towards others in the way that Hume described, at least some of the desires or concerns that issue from this more general orientation will undoubtedly be well-promoted by refraining from murdering the young boy. The General's orientation towards other persons as necessary participants in a meaningful life, will issue in motivational states whose objects will be well-promoted by attitudes and actions that *facilitate* meaningful personal engagement like, at a minimum, actions that display a sufficient respect for persons. And on Structuralism, this would entail that facts like 'the

⁷⁸ This point recalls Martin Buber's [1970] "I-Thou" perspective (or orientation) that we can take on the world, Peter Strawson's [1974] "Participant Attitude," or more recently, Stephen Darwall's [2006] "Second-Person Standpoint."

young boy is a person,' 'he didn't mean to hurt the hound,' and 'his mother loves him very much,' are reasons for the General to refrain. And because their status as reasons is explained by motivational states issuing from a practical orientation that plays a really fundamental role in structuring his identity, they are especially weighty reasons for the General.

Moreover, it's plausible that the practical orientation that Hume describes plays a larger role in structuring the General's practical identity than does his practical orientation towards serfs as objects. Some evidence for this comes from the fact that intuitively, even really evil people like the General would do more damage to their practical identities by dissociating from the naturally personal practical orientation described by Hume than they would by dissociating themselves from those orientations that lead them to view and consequently engage other persons as normatively insignificant. After all, if the General were to abandon the latter orientation, he would only cease to be especially vile, but if he were to abandon the former orientation, then by Hume's lights, he would cease to be recognizably human. Accordingly, the former orientation plays a more significant role in structuring the General's practical identity, and so, reasons that are explained by motivational states issuing from this orientation are weightier. Thus, on thoroughly Humean grounds, we can modestly conclude that the General is irrational when he murders the boy since by hypothesis he has weightier reasons to refrain from murdering the boy.

Moreover, we could extend this modest Humean reply. Perhaps any creature that is recognizable to us as a rational agent will also have the naturally personal practical orientation that Hume describes. And perhaps this practical orientation will play a significant role in structuring any agent's practical identity—thus ensuring that it

motivational states issuing from it will have weighty reasons-giving authority. If this is correct then the modest Humean can plausibly claim that for any agent for whom meaningful reciprocal interpersonal relationships are of special practical import, that agent will have especially weighty reasons to act in ways that will facilitate such relationships. And these ways of action are familiar from much of moral philosophy: actions that take seriously other agents as ends in themselves and not merely as means and actions that can be justified to others who are similarly motivated to participate in meaningful interpersonal relationships.⁷⁹ And to the modest Humean's credit, this seems correct as far as it goes.⁸⁰

But despite the plausibility of such a reply, the modest Humean also has to admit that perhaps there really are rational agents who, like the General have weighty reasons to act in horribly repugnant ways, but who, unlike the General, don't have weightier reasons that make evil activity irrational. Of course, she might protest that this is no strike against modest Humeanism since such an agent is not recognizably human. Moreover, as Bernard Williams [1972] points out, the bare possibility of such an agent is problematic only if that agent presents us with a picture of a life that we might be tempted to choose from the practical point of view. But who among us would find a life that rejects the importance of reciprocal interpersonal relationships? I certainly wouldn't. And I'm willing to bet that you wouldn't either.

But this raises a new set of issues since we might worry that this is too thin of a basis on which to ground reasons to treat others as ends or reasons to take seriously others' standing as participants in relations of mutual recognition. In short, the problem is less

⁷⁹ Cf. Kant [4: 4??]; Scanlon [1998].

⁸⁰ Indeed this discussion of a possible modest Humean reply is structurally similar to Raz's [2002] discussion of the amoralist.

about extensional issues—perhaps there really are no agents who lack especially weighty moral reason. Rather, the problem seems to be foundational. How can practical reasons, and particularly the class of moral reasons, be grounded in something that is contingent?

§8.

To sum up. The Still Too Many Reasons problem accuses Structuralism of entailing that agents have too many weighty reasons, which in certain contexts might be part of a weightiest set of reasons that dictates what they ought to do all things considered. The modest Humean reply to this objection concedes that on Structuralism, we'll sometimes find that agents (like the General) have weighty reasons to act in ways that are evil and perhaps irrational. But the modest Humean continues: if the agent is recognizably human, then she will be naturally personally orientated towards others in the ways that Hume himself described. And moreover, this particular practical orientation will play a very significant role in structuring the agent's practical identity. So accordingly, any reasons explained by motivational states issuing from this practical orientation will be very weighty.

But here we might worry that the modest Humean's solution to this extensional problem opens her up to more serious, foundational worries—particularly in the case of moral reasons. Of course, if we conceive of moral reasons as being categorical and as having practical priority, then these worries will dissolve. After all, if moral reasons are categorical, then whether an agent has them is not a contingent matter. And moreover, if moral reasons have practical priority, then they will necessarily outweigh non-moral reasons. Given this necessity, we can conclude that whatever would ground moral reasons as categorical and as

having practical priority could not be thin in the way that the modest Humean's basis for moral reasons seems too thin.

However, this opens up a whole new set of issues, most notably: how can a Humean theory of practical reasons like Structuralism ground the categoricity and practical priority of moral reasons? On the face of it, Humean theories generally seem to be incompatible with a conception of moral reasons as having these two properties. But on this point, the ambitious Humean tells us to hold our horses: Structuralism, together with some plausible assumptions about rational agency, is indeed consistent with a theory of moral reasons as being categorical and as having practical priority. If this bold suggestion is correct, the Structuralism is not only extensionally adequate but is *necessarily* extensionally adequate.

The remainder of the dissertation then, is an attempt to develop this ambitious Humean account of moral reasons as being categorical and as having practical priority. In Chapters 5 and 6, I consider potential roadblocks for this ambitious Humean reply, and I suggest a possible Humean reply. Then in Chapter 7 I argue that the Humean reply sketched at the end of Chapter 6 is not just possible, but also plausible. I tentatively conclude then that Structuralism can account for the putative categoricity and practical priority of moral reasons, thereby providing us with a well-motivated, extensionally adequate theory of reasons.

Part II

Moral Reasons

“The great truth that is too often forgotten is that it is in the nature of people to do good to one another.”

Marilynne Robinson [2012, 33]

Chapter 5: The Extensional Adequacy of Structuralism Part II: *Too Few Reasons*

§1.

In Part I: Practical Reasons, I introduced Structuralism as a theory of practical reasons. I argued that it offers the Humean a plausible analysis of the *bearing on* relation that is constitutive of some fact being a reason to φ —an analysis that, unlike Proportionalism, is the *right kind* of analysis. I then argued that Structuralism was well motivated on thoroughly Humean grounds; accordingly, it fares better by Humean lights than does Mark Schroeder’s alternative Humean theory of reasons, Hypotheticalism, according to which the weight of a reason is an irreducibly normative property of that practical reason. Most recently, I have considered two versions of the Too Many Reasons objection to Humean theories: the *Way Too Many Reasons* problem and the *Still Too Many Reasons* problem. There I argued that Structuralism clearly has the resources to deal with the Way Too Many Reasons problem.⁸¹ But I admitted that it’s not at all obvious that Structuralism can solve the Still Too Many Reasons problem—a problem for Structuralism that arises because Structuralism putatively entails that because of their actual practical orientations, some agents might have especially weighty reasons which intuitively, they actually do not have.⁸² After considering a modest Humean reply (i.e., an anti-rationalist reply) to this argument, I suggested that if Structuralism is consistent with the categoricity and practical priority of moral reasons, then it has the resources to solve the Still Too Many Reasons Problem.

⁸¹ Though, as I say in Chapter 4, it is only in conjunction with Schroeder’s Point that Structuralism can plausibly answer the Way Too Many Reasons problem.

⁸² Or better: they have weighty reasons which intuitively, are actually so insufficiently weighty that those reasons are silenced (i.e., it would not be appropriate for the agent to use those reasons in practical reason).

In Part II: Moral Reasons, I hope to show that Structuralism is consistent with the categoricity and practical priority of moral reasons. I begin in Chapters 5 and 6 by discussing two problems that are related to (but ultimately distinguishable from) the Still Too Many Reasons problem: the Too *Few* Reasons Problem (which I consider in this chapter) and the Problem of Practical Priority (which I consider in Chapter 6). As we'll see, these are problems that arise once we acknowledge that to resolve the Still Too Many Reasons problem, moral reasons must be categorical (i.e., reasons for any agent) and those reasons must be practically necessary (i.e., overriding). Of course, if we simply accept, like the modest Humean, that some agents can have very weighty—perhaps decisive—reasons for acting immorally, thereby severing the link between acting rationally and acting morally, then these problems do not arise. And if we accept this view, then it's simply no strike against a theory like Structuralism that it cannot account for the putative categoricity and practical priority of moral reasons since morality and rationality are not constitutively linked.

Though this modest form of Humeanism might be correct, I and many others have our doubts: moral reasons do seem to be tightly tied to rationality.⁸³ Thus, in Part II, after introducing the Too Few Reasons problem and the Problem of Practical Priority, I will attempt to argue that Structuralism can solve these problems, as well as the Still Too Many Reasons problem, and that it can do so without severing the link between morality and rationality. But first, I turn to the Too Few Reasons Problem.

⁸³ Most notably, Kant and those in the Kantian tradition have taken moral reasons to have the properties of categoricity and practical priority.

§2.

Recall that the Still Too Many Reasons problem presents us with a reason to doubt the extensional adequacy of Structuralism. On the Still Too Many Reason problem, evil agents like the General have too many *weighty* reasons for countenancing immoral actions. Indeed, in some cases, it might seem as if evil agents like the General have sufficiently weighty reasons for countenancing immoral actions, such that it would not be irrational for them to act in immoral ways. The problem then, is that Structuralism is overly capacious as an account of practical reasons. In other words, reasons—particularly weighty reasons—are simply too cheap on Structuralism. Consequently, it seems that on Structuralism (given that the all things considered rational thing to do is a function of the weights of an agent's reasons for and against a particular course of action), the General might be acting rationally by murdering the little boy.

But bracketing modest Humean replies and speaking to those moral theorists who are worried by this conclusion, I suggested that we can avoid this anti-rationalist conclusion if we take moral reasons to have the properties of categoricity and practical priority. Unfortunately, this leads to new problems, since Humean theories of reasons are infamous for their rejection of categorical moral reasons—a rejection that is in no small part motivated by many Humeans' anti-rationalism.⁸⁴ In other words, a further problem for Humean theories of reasons alleges that such theories cannot account for moral reasons for all agents. Therefore, Humean theories seem to predict *too few* reasons for agents, and are thus extensionally inadequate.

⁸⁴ Cf. Frankfurt [2004] and Goldman [2009].

To help see this, consider *Categoricity*, the alleged property of moral reasons that putatively leads to the problem.

Categoricity: R is a categorical reason to φ if and only if necessarily, for any agent S , if S is a rational agent, R is a reason for S to φ .

According to *Categoricity*, categorical reasons are reasons that reasons for all rational agents. In other words, for the class of agents for whom practical reasons have significance, if some particular reason is a categorical reason, then it will be a reason for all agents in this class. And if any reasons have this property, the thought goes, it's moral reasons.

But it's hard to see how moral reasons could have this property if we accept a relatively straightforward Humean theory of reasons. Consider a basic version of HTR.

HTR_{BASIC} : f is a reason R for S to φ if and only if S can, by φ -ing on the basis of f , promote the object of (at least) one of her desires.

If we accept HTR_{BASIC} (and many forms of Humeanism do seem to accept something like HTR_{BASIC}), then we'll apparently run into problems from the get-go. For example, it's plausible that when Colin sees someone suffering in a way that could be alleviated by a relatively small charitable donation that he has a reason to make such a donation.⁸⁵ And if such a reason has the property of *Categoricity*, then on HTR_{BASIC} , this would seem to require that necessarily, for any agent, that agent will have some desire whose object will be promoted by making a charitable donation.

But it's not clear that all agents who have this reason will necessarily have any desires that are promoted by donating money, since no such desire (and indeed, no particular desire

⁸⁵ Of course, it's possible that such a reason is outweighed by other reasons. Though, I suspect that it is outweighed in fewer cases than most of us care to admit.

at all) seems to be necessary for an agent qua agent.⁸⁶ Indeed some agents—perhaps those slavishly devoted to the Objectivist “philosophy” of Ayn Rand—will have regulated their desires so well that they no longer have any desire to help those less fortunate than themselves.⁸⁷ Apparently then, on HTR_{BASIC} , these John Galt characters will have no reason to help the less fortunate. Thus, we conclude, by accepting HTR_{BASIC} (or some relevantly similar form of HTR), we must thereby reject Categoricity for moral reasons and indeed, for any class of practical reasons.

This problem is even more apparent when we consider Structuralism. After all, for reasons familiar from modest Humeanism, it seems implausible given what we know about actual agents, that an agent could really have no desire at all whose object would be promoted by making a relatively small donation. But recall that Structuralism, unlike HTR_{BASIC} , doesn’t grant reasons-giving authority to just any old desires. Rather, on Structuralism only those motivational states that issue from a practical orientation that plays some role in structuring the agent’s practical identity have any reasons-giving authority. And it certainly seems coherent that for some possible agent, none of those practical orientations that play a role in structuring her practical identity issue in motivational states whose objects will be promoted by making charitable donations. Of course, perhaps those agents sometimes *want* to help or *feel concerned* that they aren’t helping. But if these motivational states float free of the practical orientations that structure their practical identities, they lack

⁸⁶ J. David Velleman [1992] denies this when he claims that the desire to make sense of oneself is constitutive of rational agency as such. As we’ll see in Chapter 7, although Structuralism does not claim that any particular desire is necessary, it will claim that a particular practical orientation—specifically the Participant Attitude—is constitutive of agency.

⁸⁷ If it seems like I’m picking on Ayn Rand and her devotees, it’s because I am.

reasons-giving authority, so it's not clear that on Structuralism all actual agents will have motivational states with reasons-giving authority whose objects will be promoted by making charitable contributions. Notice that this worryingly suggests that Structuralism cannot account for Categoricality in the case of moral reasons. Unfortunately, if this is true then on Structuralism agents have too few reasons, since intuitively moral reasons are reasons for any agent.⁸⁸ And if this is correct, then prospects for an ambitious Humeanism aren't just bleak, they're non-existent.

§3.

However, Mark Schroeder [2007a] has argued that this is too quick. Neither HTR_{BASIC} , nor any plausible version of HTR (including, as we'll see, Structuralism) entails that an agent has a reason only if she has a desire whose object will be directly promoted in the way that say, the object of a desire that the homeless be taken care of will be directly promoted by taking care of the homeless. In fact, given a plausible account of the promotion relation, Humean theories can allow for a great deal more flexibility in grounding agents' reasons. After all, as Schroeder rightly points out, on HTR, reasons can be overdetermined.⁸⁹

To illustrate this, consider the fact that Maury needs financial help to make ends meet. This fact, I take it, is a reason for us to give Maury some rent money. Now obviously, it might not be a particularly weighty reason for all agents (e.g., it is a much weightier reason

⁸⁸ Perhaps you think that moral reasons are not simply reasons for any agent but that moral reasons are reason for any agent *independently* of that agent's practical orientations. Talk about Schroeder on strong and weak modal force [Schroeder 2007a] here, and admit that Structuralism will not be able to capture strong modal force. Of course this is no problem since weak modal force captures the heart of categoricity.

⁸⁹ The case I develop below parallels Schroeder's discussion of Susan's overdetermined reason to go into the lounge [2007a, 108-9].

for Maury's father than it is for a stranger living in Timbuktu), but plausibly it is nevertheless a reason for all agents to give Maury some rent money since moral reasons putatively instantiate Categoricality. And Humean theories explain how this fact could be a reason for all agents by appealing to the general way in which desires give rise to an agent's reasons. After all, on HTR this fact could be a reason for me in virtue of any number of different motivational states, each of whose object would be promoted if I were to act on the basis of this fact. This is especially clear when we consider HTR_{BASIC} which simply holds that an agent will have a reason to φ if she will promote the object of *at least one* of her desires.⁹⁰ But this leaves open the possibility that in some cases, φ -ing will promote the objects of several or perhaps most of her desires. In such cases, the existence of R as a reason for S to φ will be overdetermined.⁹¹

So, for example, suppose Maury is my brother and that out of fraternal concern, I genuinely desire that things go well for Maury. In light of this desire, it's plausible that the fact that he needs help is a reason for me to loan him some rent money since, by giving him rent money, I promote the object of my desire that things go well for Maury. Alternatively, I might simply have a desire to always do the right thing. If I have such a desire, then the fact that Maury needs help will be a reason for me to help him, since by helping him, I will promote the object of my desire to always do the right thing. Or suppose that even if I'm no Dudley Do-Right, I do have a desire to impress my date. And further suppose that in such a case, the fact that Maury needs help will be a reason for me to offer him some rent money because, given my date's ideas about social justice, I will promote the object of my

⁹⁰ The same can be said *mutatis mutandis* for Structuralism.

⁹¹ Again, see Schroeder [2007a, 108-9].

desire to impress my date by giving him money. Of course, even if I'm not Maury's brother, nor explicitly concerned to be moral, nor on a date, I might still have some desire whose object will be promoted by giving Maury some rent money.

To see this, note that even if I was a kind of Randian hero who desired to live in a society without mooches like Maury, it's plausible that on HTR, I would nevertheless have a reason to give Maury money. After all, even in *this* case, if I were to give Maury money because of the fact that Maury needs money, I would thereby promote the object of my desire, since by giving Maury some money, I make it somewhat less likely that he'll need to mooch off others again.

Of course, the Randian might try to deny this by arguing that I would better promote the object of my desire by refraining from giving him the money, since that would either contribute to Maury making changes to his "parasitic" lifestyle or perhaps lead to his death. In either case, the Randian might say, I'd better promote the existence of a society without moochers by refusing to ever give charity to folks like Maury. This is perhaps correct—though I have some serious doubts—but it is irrelevant to HTR's claim that I have a reason in such a case. On Humean theories, even thoroughly reductive analyses like Structuralism, the existence of a reason depends on whether the object of an agent's motivational state(s)—whichever particular motivational state(s) the version of HTR in question takes to have reasons-giving authority—is promoted to some extent. Even *minimally* promoting the object of some motivational state is sufficient on Humeanism to generate a reason, and since in the case envisioned, I would minimally promote the object of my Objectivist desire to some extent if I were to give Maury money, I have a reason to give Maury money. Thus it looks like a single fact can be a reason for me given the existence of many different and in most

cases, wholly unrelated motivational states—even motivational states which we might have antecedently thought wouldn't have generated the particular reasons that they do in fact generate.

In light of the fact that an agent's reasons can be overdetermined (in just those ways that my reason to give Maury some money can be overdetermined), Schroeder further claims:

For all [the Humean theory of practical reasons] says, there might be reasons that can be explained by *any* possible desire.

[An ambitious Humean's] proposal for how there could be genuinely [categorical] reasons is therefore that genuinely [categorical] reasons are *massively* overdetermined. They are reasons for anyone, no matter what she desires, simply because they can be explained by any (or virtually any) possible desire. Suppose that some reason is like this—it can be explained by any possible desires. Then anyone who has any desires whatsoever will have this reason. So if to count as an agent, you must have *some* desire or other ... then any possible agent would have this reason. And there would be no desire such that if you did not have that desire, then it would not be a reason for you [Schroeder 2007a, 109].

Here Schroeder is suggesting that the Humean can explain the categoricity of moral reasons by exploiting the simple point that nothing about HTR *as such* will generate the Too Few Reasons problem since HTR simply claims that the existence of an agent's reasons will depend on some desire (or on some relevant motivational state) whose object is promoted by a particular course of action. And if the promotion relation that Humeans point to is sufficiently weak, then it's possible that some reasons will be generated by *any desire at all*.⁹²

If this is correct, then there is no *logical* problem for Humean theories that arises from worries about Categoricity since nothing about HTR as such rules out the possibility that some reason R can be generated by any desire at all. Of course, even if there is no logical problem, it still seems as if there is *some* problem in the neighborhood. After all,

⁹² Schroeder [2007a] defends this point.

unless we can say more about the class of categorical reasons that putatively can be explained by any desire at all, it's not clear that we have any justification in thinking that there are such reasons. Indeed, on this point, Schroeder has only provided a Humean *defense* of the possibility of there being reasons that instantiate the property Categoricality. But what we really need a Humean *theodicy* (to borrow an image from another philosophical problem). In other words, what we really need is a plausible account of *how* some reasons will be reasons for any agent given any desires at all.

To generate such an account however, we'll need a better sense of the promotion relationship that is implicated in Humean theories. That is, to know whether there is some reason to suspect that some reasons can be generated by any desires at all, we'll need to know what is involved in *promoting* the objects of an agent's desires. Accordingly, I turn to that task below.⁹³

§4.

According to Humean theories like HTR_{BASIC}, Proportionalism, Hypotheticalism, and Structuralism, the existence of an agent's reason R to φ will ultimately depend on whether φ -ing will *promote* the object of some particular motivational state (be it desires—strong or weak—or motivational states issuing from those practical orientations that play some role in structuring an agent's practical identity). But what is it to promote the object of one of these motivational states? Plausibly, it has something to do with making the state of affairs picked out by the object of a particular motivational state *more likely*. Of course, this is too

⁹³ The discussion of the promotion relation below is something of a technical detour. Feel free to skip ahead to §5.

imprecise. So to clarify the promotion relationship, I will consider a recent critique of two probabilistic analyses of the promotion relationship. Specifically, Jeff Behrends and Joshua DiPaulo have argued that two influential accounts of *promotion* due to Stephen Finlay [2006] and Mark Schroeder [2007a] are inadequate.⁹⁴ After considering Behrends and DiPaulo's criticisms of Finlay's and Schroeder's reductive probabilistic analyses of promotion, I go on to suggest an alternative probabilistic analysis of the promotion relationship according to which an action will promote the object of a motivational state if and only if by performing that action, the agent will render the object of the motivational state more likely to obtain, relative to its antecedent intrinsic likelihood of obtaining.

4.1.

To begin, consider Stephen Finlay's analysis of the promotion relationship:

Promotion₁: For some agent X, desire D, and action A, A *promotes* p—the object of D—iff X's doing A renders p more likely than it would have been had X not done A.⁹⁵

As a first pass, this seems plausible. Intuitively, I promote the object of my desire to have a drink by bringing the cup to my lips because by bring the cup to my lips, I make it more like that the object of my desire will obtain than it would had I not brought the cup to my lips.

However, Behrends and DiPaulo offer a counterexample to Promotion₁. Consider *Buttons*.

Debbie has some desire. There are three buttons in front of her. If she pushes either Button A or Button B, her desire is guaranteed to be fulfilled. If she pushes Button C, her desire will not be fulfilled. Debbie in fact pushes A. Had she not pushed A, though, she would have pushed B instead [Behrends and DiPaulo 2011, 2].

⁹⁴ Jeff Behrends and Joshua DiPaulo [2011].

⁹⁵ Stephen Finlay [2006]; As Behrends and DiPaulo note, though Finlay doesn't use the language of "promotion," he is clearly interested in analyzing this relationship.

In this case, on Promotion₁, Debbie does not promote her desire by pressing Button A, since, had she refrained from pressing Button A, she would have pressed Button B, which would have made it no less likely that the object of her desire obtain. But this seems incredible. How could pressing Button A not count as promoting her desire? After all, by pressing Button A, Debbie guarantees the satisfaction of her desire, and if guaranteeing the satisfaction of a desire doesn't count as promoting that desire, what would?

Of course, Behrends and DiPaulo admit that it's not wholly unmotivated for defenders of Promotion₁ to simply bite the bullet and accept that in some cases, guaranteeing the satisfaction of a desire won't count as promoting that desire. In light of this, Behrends and DiPaulo offer a further reason for thinking that Finlay's Promotion₁ fails as an analysis of the promotion relationship relevant to practical rationality. Specifically they argue that Promotion₁ is inconsistent with a plausible account of instrumental rationality. According to this account—*Instrumental Reason*—“for X to have an instrumental reason to A is for there to be some p such that X has a desire the object of which is p, and for there to be some fact that is part of what explains why X's doing A promotes p” [Behrends and DiPaulo 2011, 3].⁹⁶ Though this is a natural picture of what it is for an agent to have an instrumental reason to A, because Promotion₁ entails that Debbie does not promote her desire by selecting Button A, it would entail that she had no reason in *Buttons* to press Button A. But surely this is false, since given her desire she clearly has more reason to press Button A than to press Button C, and if she has *more* reason, then that entails she has some reason

⁹⁶ Behrends and DiPaulo's *Instrumental Reason* is modeled on Schroeder's [2007a] account of reasons.

to press Button A. Thus, Promotion₁ fails as an account of the promotion relationship since it is inconsistent with *Instrumental Reason*.

4.2.

Having offered us some reason to reject Finlay's Promotion₁, Behrends and DiPaulo turn their attention to Mark Schroeder's analysis of the promotion relation. Consider Schroeder's analysis.

Promotion₂: For some agent X, desire D, and action A, A *promotes* p—the object of D—iff X's doing A renders p more likely than it would have been had X done nothing.⁹⁷

As was the case with Promotion₁, Schroeder's Promotion₂ seems initially plausible.

Promotion₂ seems to correctly predict that I promote the object of my desire to watch the game by turning on the TV. And it seems to offer an intuitive explanation of why turning on the TV promotes the object of my desire: by turning on the TV, I make it more likely that I'll get to see the game than it would have been had I done nothing.

But again, Behrends and DiPaulo offer us two reasons to doubt its adequacy. First, they offer a variant of *Buttons*. Consider *Buttons 2*.

Julie has some desire. There is one button in front of her. She knows that if she pushes the button, her desire is guaranteed to be fulfilled. However, unbeknownst to Julie, if she does not push the button, Black will ensure that her desire is fulfilled [Behrends and DiPaulo 2001, 4].

Once we consider *Buttons 2*, we'll see that if Promotion₂ is true, then Julie will not promote her desire by pressing the button. But this is surprising since by pressing the button she thereby *guarantees* the satisfaction of the desire. Of course (as was the case with *Buttons* as a

⁹⁷ Mark Schroeder [2007]. As was the case with Finlay's Promotion₁, I borrow Behrends and DiPaulo's statement of Schroeder's analysis.

counterexample to Promotion₁), Behrens and DiPaulo again admit the defender of Promotion₂ can bite the bullet about such cases and stand firmly behind her analysis. Accordingly, as they did with Promotion₁, Behrens and DiPaulo point to a deeper problem for Promotion₂.

Consider *Do Nothing*.

At t_1 Austin forms the desire that p be the case at t_3 . Black has arranged things such that if Austin does nothing at t_2 , p will be the case at t_3 ; Black has further arranged things such that any other behavior at t_2 on Austin's part will result in not- p [Behrens and DiPaulo 2011, 4].

In this case, Behrens and DiPaulo rightly note that it seems as if Austin has an instrumental reason to do nothing at t_2 . But if Promotion₂ is true, it's unclear how any agent could have a reason to do nothing since doing nothing can never render a state of affairs more likely to obtain than if the agent were to do nothing. After all, doing nothing will make it exactly as likely that the object of some motivational state will obtain as will doing nothing, since they are the same activity (or omission). So, doing nothing can never render p *more likely* than doing nothing. Therefore, it seems that on Promotion₂, doing nothing can never promote a desire. But surely doing nothing can promote desires, since the best explanation of why Austin has an instrumental reason to do nothing at t_2 is that by doing nothing, he'll promote the object of his desire that p . Therefore, we should reject Schroeder's Promotion₂ as an adequate analysis of the promotion relationship.

4.3.

Now that we've rejected Finlay's Promotion₁ and Schroeder's Promotion₂, where should we look for an adequate account of the promotion relationship? Behrends and DiPaulo conclude their discussion with some suggestions. Specifically they claim that:

[These counterexamples] may even suggest that promotion is best thought of non-probabilistically. For, one might think, the central motivation for offering a probabilistic account of promotion is that, *prima facie*, promotion seems to require an increase in probability. Once serious doubt has been cast upon that position, though, probabilistic accounts might turn out to be unmotivated [Behrends and DiPaulo 2011, 5].

But I think this suggestion is too quick. Plausibly, Behrends and DiPaulo have given us reason to doubt the adequacy of Promotion₁ and Promotion₂. However, the problem, as I see it, isn't that these analyses of the promotion relationship are probabilistic; rather it's that they are counterfactual analyses. Note that in both Promotion₁ and Promotion₂ the baseline relative to which a particular course of action is compared is what *would* happen if the agent did otherwise (in the case of Finlay's Promotion₁) or what *would* happen if the agent *would* have done nothing (in the case of Schroeder's Promotion₂).

But plausibly, whether A-ing promotes p—the object of D—is not a matter of alternative sequences. Rather, whether A-ing promotes p seems to be a matter of the actual-sequence. Indeed the very case that Behrends and DiPaulo use to undermine Schroeder's Promotion₂ is one that is relevantly similar to Frankfurt style cases.⁹⁸ And the 'moral' of the Frankfurt cases is that the alternative-sequence—i.e., what the agent *would have done*—is irrelevant to the question of whether she is morally responsible for her action.⁹⁹ Of course,

⁹⁸ Frankfurt [1969]. Behrends and DiPaulo themselves note the structural similarities between *Buttons 2* and Frankfurt style cases.

⁹⁹ For more on the 'moral' of the Frankfurt cases, see Fischer and Ravizza [1998] and Fischer [2010].

whether Frankfurt cases succeed in showing that the alternative-sequence is irrelevant to moral responsibility is quite controversial.¹⁰⁰ But it seems to me that Behrends and DiPaulo's *Do Nothing* actually gives us good reason to doubt that alternative-sequence accounts of promotion are inadequate.¹⁰¹ However, this doesn't mean that probabilistic accounts of promotion need to be put out to pasture, only that alternative-sequence probabilistic accounts should be rejected.

With this in mind, I offer a third probabilistic account of the promotion relation. Because it's the account of promotion that I marry to Structuralism, I'll simply call it "Promotion."¹⁰²

Promotion: For some agent X, desire D, and action A, A *promotes* p—the object of D—iff X's doing A renders p more likely relative to the antecedent intrinsic likelihood of p obtaining.

On Promotion, it's not what *would* happen that determines whether a particular course of action promotes the object of an agent's motivational state. Instead, it's whether the likelihood of p obtaining is greater *after* the agent acts than it was before she acted. To illustrate, I'll begin with a simple case. Take my desire to finish my dissertation—a desire that, according to Structuralism, has reasons giving authority since it issues from a practical orientation that plays some role in structuring my practical identity. If I work on my dissertation for eight hours a day, then I promote the object of my desire—viz., that I be finished with my dissertation. And Promotion gives us a plausible explanation of this. After

¹⁰⁰ Cf. Pretty much all non-Pereboom incompatibilists.

¹⁰¹ Notice, *Do Nothing* isn't just a problem for Schroeder's Promotion₂ as Behrends and DiPaulo take it to be. It's also a problem for Finlay's Promotion₁ since doing nothing instead of A-ing is a way of doing something other than A.

¹⁰² I confess that it is only because Promotion is my account that it isn't sullied with an inelegant subscript.

all, the antecedent intrinsic likelihood of my dissertation being finished without eight hours per day of work is much lower than is the likelihood of my dissertation being finished with eight hours per day of work.

Of course, even Finlay's Promotion₁ and Schroeder's Promotion₂ can explain this. But Promotion is better than these alternative-sequence analyses because it can also explain why Debbie and Julie promote their respective desires in *Buttons* and *Buttons 2*. Consequently, Promotion can explain why Debbie and Austin have instrumental reasons in *Instrumental Reason* and *Do Nothing* respectively.

Notice that in *Buttons*, because it is possible (however improbable) that Debbie push C, the fact that she actually pushes A—an action that guarantees that the object of her desire is satisfied—will promote the object of her desire if Promotion is true. Of course, it is true that she *would have* pushed B if she hadn't pushed A and that B would have also satisfied the object of her desire. But this fact (that is made true in virtue of the nearest set of possible worlds) is irrelevant to whether the probability of the object of her desire being satisfied is made more likely by her pressing A relative to its antecedent intrinsic likelihood. After all, though it is true that in the nearest set of worlds in which Debbie doesn't press A she does press B, it doesn't follow from this that there are no worlds in which she presses C. Moreover, the antecedent intrinsic likelihood that Debbie will satisfy the object of her desire must be less than 1, since nothing about state of affairs before Debbie's decision to push A entails that Debbie will choose anything, since it's possible that the world end. Because *Buttons* doesn't show Promotion to entail that Debbie doesn't promote the object of her desire even though she guarantees its satisfaction, Promotion can subsequently explain why

Debbie has an instrumental reason to press A. And this shows Promotion to be superior as an analysis of the promotion relationship than is Finlay's Promotion₁.

Likewise, with the help of Behrends and DiPaulo's cases, we can see why Promotion is better than Schroeder's Promotion₂. Unlike Promotion₂, Promotion would predict that Julie promotes her desire in *Buttons 2* and that Austin has an instrumental reason to do nothing in *Do Nothing*, which requires that he can promote the object of his desire by doing nothing. To see this, first recall *Buttons 2*.

Julie has some desire. There is one button in front of her. She knows that if she pushes the button, her desire is guaranteed to be fulfilled. However, unbeknownst to Julie, if she does not push the button, Black will ensure that her desire is fulfilled [Behrends and DiPaulo 2001, 4].

In this case, suppose that Julie pushes the button. Intuitively, she thereby promotes her desire. But on Schroeder's Promotion₂ she does not, since by pushing the button she does not increase the likelihood of her desire being fulfilled relative to her doing nothing.

However, while this is plausibly a counterexample to Promotion₂, it gives us no reason to doubt Promotion as I articulated it above. Consider that Black will ensure that Julie's desire is fulfilled only if Julie refrains from pushing the button. And because we're imagining that Julie does in fact push the button, it follows that Black plays no role in the actual-sequence; he is merely a counterfactual promoter.¹⁰³ But on Promotion, alternative sequences are irrelevant to whether or not a particular course of action will promote the object of an agent's desire. Rather, it's simply a matter of whether the comparative likelihood that the agent will satisfy the object of her desire relative to the antecedent intrinsic likelihood that determines whether a particular course of action *promotes* an agent's

¹⁰³ Compare this to Black's role as a counterfactual intervener in Frankfurt-style cases.

desire in the relevant sense. Of course, the presence of Black, together with the fact that he is poised to ensure that Julie's desire is satisfied tells us that he *would* be the explanation of how the object of her desire was promoted in the alternative sequence. But in the actual sequence, what raises the likelihood that the object of her desire will obtain is Julie's act of pushing the button.

Now, you might worry that even here, given Black's presence, the antecedent probability that the object of her desire is satisfied is 1. After all, if Julie presses the button then she guarantees that the object to her desire will be satisfied, and if she doesn't push the button then Black ensures that the object of her desire will be satisfied. Since these seem to be the only options, it might seem that the probability that her desire is satisfied is 1 and that nothing she can do can increase the likelihood of something that has a probability of obtaining of 1.

But this isn't quite right, since probabilities depend on more factors than what happens in nearby possible worlds. Specifically, *intrinsic* probabilities of the sort that are relevant to Promotion depend on more than what would happen in light of extrinsic features of the actual-sequence. To see this, consider a fair six-sided die that is rolled. Making some standard assumptions, the antecedent intrinsic probability that a 3 will come up is roughly .166. But this is true even if God (or some Jedi as the case may be¹⁰⁴) would keep the die rolling until it came up 3 if it would've been the case that the die was about to come up 1, 2, 4, 5, or 6. No doubt, God's presence ensures that the die will come up 3, such that if we rolled it n times, it would come up 3 n times. But again, this doesn't entail that the

¹⁰⁴ Indeed, this was a terrible plot device in the widely panned *Star Wars I: The Phantom Menace*, which is quite possibly one of the worst, most incoherent movies ever.

antecedent, intrinsic probability of the die coming up 3 is 1, since the intrinsic probability of some outcome, is presumably determined intrinsic properties of the event itself.

The same will be true for Julie and the antecedent intrinsic likelihood that the object of her desire will be satisfied. Although in the case she knows that pressing the button will guarantee that her desire will be satisfied, this by itself does not entail that her desire will be satisfied since it does not entail that she will press the button. For example, if the desire in question is relatively insignificant, it would not be incredible that she would forgo pushing the button. Or if Julie thinks that a particularly bad desire of hers will be satisfied by pressing the button, she might similarly refrain. In any case, this shows that given the range of options open to her, the antecedent intrinsic probability that she will satisfy the object of her desire is something less than 1. So, in those cases in which she does push the button, she raises the likelihood of her desire being satisfied relative to this baseline (i.e., the antecedent intrinsic probability). Consequently, in such cases she promotes her desire and this shows Promotion to be adequate in a way that Schroeder's Promotion₂ is not.¹⁰⁵

4.4.

If the above is correct, then we have a plausible probabilistic account of promotion—viz., Promotion. On Promotion, the baseline relative to which a particular course of action

¹⁰⁵ Of course, I take it that when Schroeder identifies an agent's doing nothing with the baseline relative to which a course of action must increase the likelihood of the object of an agent's desire obtaining, he was thinking that this baseline is related to the antecedent intrinsic probability of the object of an agent's desire obtaining. After all, I suspect these two baselines are highly correlated; it's just that as *Buttons 2* shows us, they're not identical.

Likewise, although *Do Nothing* is bad for Promotion₂, it is fairly clear that *Do Nothing* poses no problems for Promotion since *Do Nothing* exploits Schroeder's claim that it is specifically relative to the baseline of doing nothing that a particular course of action might promote the object of an agent's desire. Since Promotion does not implicate doing nothing as the relevant baseline, it is not undermined by the case.

counts as promoting the object of an agent's desire is the antecedent intrinsic likelihood of the object of that desire being satisfied. And with this in mind, we can return to the important question of whether Humean theories of practical reasons (and specifically Structuralism) fail to accommodate Categoricality as a property of moral reasons.

§5.

Recall that following Mark Schroeder [2007a], in §3. I noted that there was nothing about Humean theories of reasons *as such* that entailed Too Few Reasons since Humeanism is consistent with it being the case that some reasons can be explained by *any desire at all*. To see this, we needed to understand when the object of a motivational state with reasons-giving authority is promoted by a particular course of action. After all, on Humean theories, motivational states explain agents' reasons in the sense that when, in light of some fact, particular courses of action promote the object of those motivational states. And as we saw in §4., a course of action will promote the object of some motivational state just in case that course of action will raise the probability that the object of motivational state in question relative to its antecedent intrinsic probability.

But this logical point, which is sufficient to show that Humean theories as such aren't incompatible with Categoricality, has done nothing to explain how moral reasons could be explained by *any motivational state at all*.¹⁰⁶ What we need is a fuller account—a *theodicy* to return to an analogy I made earlier—of how moral reasons can be explained by any

¹⁰⁶ More precisely: making this logical point has done nothing to explain how moral reasons could be explained by any motivational state with reasons-giving authority.

motivational state with reasons-giving authority in a way that is compatible with the content and distinctively Humean motivations of Structuralism.

It turns out that such an account is relatively straightforward. Consider that moral reasons—reasons that putatively enjoy the property of Categoricality—are such that, if we act on the basis of these reasons, we act in ways that promote a *kingdom of ends*. Now of course, there is some controversy as to exactly what Kant means when he claims that “every rational being must act as if he were by his maxims at all times a lawgiving member of the universal kingdom of ends,” [G 4:438]. But leaving aside delicate interpretive issues surrounding Kant’s invocation of a “kingdom of ends,” I think we can (uncontroversially) say is that the kingdom of ends is an ideal of morality—an ideal taken up by contractualists as well as dyed-in-the-wool Kantians alike.¹⁰⁷

The ideal of a kingdom of ends is grounded in a picture of moral reasons that presupposes that by acting on moral reasons, we are able to facilitate and underwrite a community of individuals related to one another in meaningful, reciprocal ways. Consequently, if what I do when I act for a moral reason promotes this ideal, then it will turn out that by acting for moral reasons I will promote the object of any of my motivational states since the object of any motivational state at all is more likely to obtain relative to its antecedent intrinsic likelihood of obtaining if I act in a way that promotes a community of equals in which I am given latitude to direct my own affairs. Indeed, a condition that strongly promotes the object of any desire obtaining is a relatively stable environment, and insofar as I act for moral reasons, I work to facilitate such an environment, thereby promoting the objects of any of my motivational states with reasons-giving authority.

¹⁰⁷ Cf. Scanlon [1998].

Of course, if this is correct, then it will give Structuralism an account of how it is that the object of any motivational state at all will be promoted by acting for moral reasons, and thus, it will explain how Structuralism is compatible with the Categoricality of moral reasons. But it will only do so by committing Structuralism to controversial claims concerning first-order normative theory—viz., that some form of Kantian deontology, contractualism, or perhaps certain forms of normative pluralism must be true.¹⁰⁸ Accordingly, this means that if Structuralism is to capture the Categoricality of moral reasons, it rules out familiar forms of pure consequentialism.¹⁰⁹

I admit that this is a cost of Structuralism since ideally we'd want an inclusive account of practical reasons, one that doesn't presuppose any particular first-order normative theory. But because it's the putative Categoricality of moral reasons that leads us to this anti-consequentialist position, I think the cost of accepting Structuralism is quite low. After all, for those of us committed to (or at least interested in exploring) Categoricality, we might think that our meta-ethical theories shouldn't be wholly unresponsive to our first-order normative theories. As a result, I think that the fact that the conjunction of Structuralism and Categoricality precludes pure consequentialist theories isn't too problematic, since after all, there are good independent grounds for rejecting pure consequentialism.

It seems then that Structuralism has the resources for dealing with the Too Few Reasons problem since, on Structuralism, an agent's reasons can be overdetermined and it

¹⁰⁸ A pluralism that accepted deontological side constraints (along with consequentialist considerations) might also embrace something like the ideal of the "kingdom of ends," and so Structuralism could generate similar claims using this kingdom of ends-like ideal.

¹⁰⁹ Interestingly, if a consequentialist claimed that the only consequences with normative significance were those that are related to the facilitation and maintenance of a community of equal persons, then they too would be able to appeal to this sort of view. However, I know of no such consequentialist.

will turn out that moral reasons will be generated by any motivational state with reasons-giving authority at all. Moreover, note that because moral reasons can be explained by any motivational state at all (so long as it is a motivational state with reasons-giving authority), an agent's moral reasons will be abundantly overdetermined, since they will be generated by *all* of her motivational states with reasons-giving authority. Accordingly, Structuralism (and HTR more generally) can not only defend itself against the Too Few Reasons problem; it can also explain why it is that moral reasons can be explained by any motivational state whatsoever.

§6.

But even if Structuralism is successful on this count, it hasn't thereby solved the Still Too Many Reasons problem. After all, nothing I've said here requires that moral reasons—which, on Structuralism can be reasons for any rational agent—are particularly weighty. And indeed, we might worry that on Structuralism, such reasons will be of comparatively little weight, since Structuralism claims that the weight of a reason is in part a function of *how well* it promotes the object of an agent's motivational state. And many of the motivational states that will generate agents' moral reasons will only be *minimally* promoted by acting for moral reasons. And this would suggest that for many agents, moral reasons could be of comparatively little weight. Thus, it might be possible that for many agents (like the General discussed in Chapter 4), it would be rational for them to act immorally. Consequently, the Still Too Many Reasons problem is still a problem for Structuralism.

To solve the Still Too Many Reasons problem, Structuralism will have to argue that moral reasons not only enjoy Categoricity, but that they have Practical Priority. If moral

reasons have practical priority, then they will be comparatively weighty and perhaps, if ambitious Humeans are correct, they will entail an essential connection between rational action and moral action. Of course, as was the case with Categoricality, it's not at all obvious that Structuralism (or other forms of HTR) is compatible with Practical Priority. Accordingly, I turn my attention to the problem of Practical Priority as it arises for Structuralism below in Chapter 6.

Chapter 6: Structuralism and the Practical Priority of Moral Reasons

§1.

Consider Davis who is very angry at his roommate Marcus (Marcus forgot to pay the electricity bill—his responsibility—and the power was cut off). Because of his anger, Davis is considering telling Marcus's fiancé that he cheated on her several years ago. Davis knows that Marcus has changed a lot since then and that he has never done anything like that since—Marcus has recognized his mistake, learned his lesson, and has become a model partner. He also knows that even though Marcus has changed, his fiancé would likely end the relationship if it was disclosed to her by someone other than Marcus. But Davis is angry and frustrated and above all, he wants to hurt Marcus; so destroying his relationship seems like a good thing to Davis at the moment. Even agreeing that infidelity is very bad, I think it's plausible to think that all things considered, Davis shouldn't disclose Marcus's long-ago affair to his fiancé. For starters, it is simply none of Davis's business. But moreover, from the moral point of view, keeping quiet (at least in this context) is the right thing to do, since vindictive and disproportional retributive anger (and its expression) are unjustified responses to wrongdoing. Rather than facilitate reconciliation and moral repair, such responses undermine the fabric of moral relations.

Of course, as I argued in Chapter 5, Structuralism can account for Davis's moral reasons to refrain from disclosing Marcus's long-ago infidelity, regardless of Davis's actual motivational states. And indeed, such reasons will be overdetermined: that Marcus is now completely devoted to his fiancé, that Marcus doesn't deserve to have his relationship destroyed simply because he forgot to pay the electricity bill, that Davis has to live with

Marcus, etc. After all, if, in light of one of these facts, Davis acts in a way that takes seriously Marcus's equal standing (say, by acting in a way that Marcus couldn't reasonably reject), then he will promote the object of at least one (but in reality, a great many) of the motivational states that have reasons-giving authority for him. Because moral reasons are so cheap, any agent who has motivational states with reasons-giving authority (i.e., any agent who has motivational states issuing from a practical orientation that plays some role in structuring her practical identity) will have moral reasons. Moral reasons therefore, instantiate the property of Categoricity.

However, even though it is true that on Structuralism, it's possible that all agents have moral reasons, we haven't thereby solved the Still Too Many Reasons problem. Recall from Chapter 4 that the Still Too Many Reasons problem is a challenge to Structuralism because on Structuralism, it seems possible that some agents will have too many *weighty* reasons—reasons which are sufficiently weighty to make it rational (at least in some cases) for agents to perform heinous immoral actions. But intuitively, evil agents like the General don't have such weighty reasons, and they are not rational when they act immorally. And simply pointing out that moral reasons are categorical won't solve *this* problem since what is all things considered rational for an agent to do is a function of the weightiest set of reasons. But it is not clear that given Structuralism's reply to the Too Few Reasons problem, moral reasons will be sufficiently weighty to make it irrational for any agent to refrain from acting immorally. At least, this is the worry for Structuralism.

Obviously, the problem is solved if moral reasons have practical priority—i.e., if moral reasons decisively outweigh non-moral reasons. But just as it wasn't initially apparent that Structuralism was compatible with Categoricity, neither is it apparent that Structuralism

is compatible with Practical Priority—the property a reason has when it is decisively weighty. After all, on Structuralism, the weight of a reason depends (in part) on the role that the practical orientation issuing in the relevant motivational state plays in structuring the agent’s practical identity, with reasons explained by motivational states issuing from practical orientations that play a significant role in structuring being weightier. But it is not at all obvious that for any agent, the practical orientations from which the relevant motivational states issue will necessarily play a more significant role in structuring her identity, such that her moral reasons will be of particular weight. In response to this important challenge, in what follows I hope to argue that Structuralism is compatible with Practical Priority, and that it is consequently well-positioned to serve as a theory of moral reasons of the sort required by ambitious Humeanism.

§2.

Recall from Chapter 5 the various motivational states that might ground my reason to help Maury: my fraternal love, my desire to do the right thing, my hope to impress my date, my concern to live in a society with no mooches, etc. If it turns out that at any one time, I have all these motivational states, then my moral reason to help Maury will be overdetermined. And indeed this will be true even if I lack any one of these motivational states—e.g., the desire to do the right thing.

But suppose I do lack that particular desire. Or suppose that the particular desire issues from a practical orientation that plays a relatively insignificant role in structuring my practical identity. In such a case, what can we say about the weightiness of my moral reason to help Maury? Ultimately, on Structuralism, its weight will depend on the role that the

practical orientations issuing in these other motivational states play in structuring my practical identity. If I'm like a relatively normal person, then I am oriented towards my family in a particularly intimate way. Consequently, it's likely that the practical orientation that shapes and directs my concern for my brother plays a more significant role in structuring my practical identity. And on Structuralism this would entail that my moral reason to help Maury is particularly weighty, since at least one of the motivational states that overdetermine its existence issues from a practical orientation that plays a significant role in structuring my practical identity.

But suppose my identity was not structured in this way. Would moral reasons still be comparatively weighty? Necessarily?

In *Rammer Jammer Yellow Hammer*, Warren St. John details the story of Freeman and Betty Reese, a pair of committed University of Alabama Crimson Tide football fans.¹¹⁰ Talking to a reporter before a game, Freeman Reese was once asked what he and his wife had given up because of their commitment to attending every Alabama football game—something they had done (at that point) for over fifteen years. St. John documents the exchange:

“Let’s see,” [Reese] says in a soft Southern drawl. “We missed our daughter’s wedding.”

“You what?”

“We told her, just don’t get married on a game day and we’ll be there, hundred percent, and she went off and picked the third Saturday in October which everybody knows is when Alabama plays Tennessee, so we told her, hey, we got a ball game to go to. We made the reception—went there soon as the game was over” [St. John 2004, 10].

¹¹⁰ Warren St. John [2004].

Unlike my concern for Maury, this case seems to provide a counterexample to the old adage that blood is thicker than water. St. John goes on to document others' reaction upon hearing of the Reese's commitment to Alabama football games.

For over a decade, whenever the subject of fans came up, I deployed the story of the wedding skippers as a kind of archetypal example of how far people would go. There was something interesting about the way people reacted. Most seemed to think the couple aboard the *Crimson Express* [the \$300,000 motor home that the Reese's used to follow the Crimson Tide] were nutjobs or worse. Some used terms like "child abuse," or compared the couple's passion for football to alcoholism. Skipping a child's wedding was simply *beyond*. To people in this camp the most disturbing part of the story was that the couple had made the reception. They weren't not talking to their daughter; the family hadn't experienced some irreparable rift. They just weren't missing football games for her.

A not insignificant few, however—an unscientific estimate would be 25 percent or so—had a different view: they blamed the daughter. These people wanted to know what kind of person would force someone they loved to make such a choice. Those who blamed the daughter usually shared something in common: they were fans too. If they knew anything about Alabama football, they were likely to condemn the daughter even more harshly for the simple reason that she had scheduled her wedding for the day of the Tennessee game. Tennessee is a huge Alabama rival. You don't devote your life to Alabama football and miss the Tennessee game. Interesting too was that while those who thought the parents were bonkers usually reacted with a chuckle and a shrug, those who blamed the daughter were more likely to get *angry* over the story [St. John 2004, 12-3].

So while the Reese's might have been singular in their commitment to Alabama football, they weren't alone in thinking that their commitment to the team might give them reason to do some pretty drastic things. But do the existence of such individuals (i.e., the group of folks unscientifically estimated to be 25 percent of St. John's interlocutors) pose a challenge to Structuralism?

I think they might.

The problem arises from the fact that these agents seem to have a disordered or otherwise confused set of concerns (they *are* nutjobs). But these concerns undoubtedly issue from a practical orientation that plays a major role in structuring their practical identities—"I

just love Alabama football, is all I can think of” [St. John 2004, 15], Reese said in reply to a question about his motives for missing his daughter’s wedding. And presumably, Reese and his wife didn’t make their decision lightly, and it is likely that they—like many of the other *fans* that St. John interviewed—took themselves to be justified in going to the Alabama-Tennessee game. But this isn’t obvious, and like the majority of those St. John talked to, I do catch the whiff of practical irrationality when I reflect on the Reese’s story.¹¹¹ But can a Humean, and specifically, can a Structuralist account for such irrationality? And more importantly, can Structuralism account for the putative irrationality that would arise if, instead of merely skipping their daughter’s wedding (which may or may not be a moral offense), a family of superfans like the Reeses—who I’ll call the Houndstooths so as not to unfairly presume anything about the actual Reeses’ moral characters—would regularly avoid moral obligations to attend Alabama football games. The question then, is whether Structuralism can explain the Houndstooths’ practical irrationality given the fact that their love of Alabama football issues from a practical orientation that plays such an important role in structuring their practical identities? In other words, on Structuralism, is it possible that the Houndstooths’ moral reasons have practical priority?

§3.

To determine whether the Houndstooths’ moral reasons have practical priority, we must first recall Structuralism’s account of the weighting of reasons. On Structuralism, the weightiness of S ’s reason R to φ will depend on:

¹¹¹ I say this as a college football fan, one who often finds Alabama, with their smothering defense, entertaining to watch.

- (1) the degree O plays a role in structuring S 's practical identity, where O is one of the practical orientations that play some role in structuring S 's practical identity,
- (2) *how well* φ -ing will promote m , where m is the motivational state that issues from O .

So just how weighty are the Houndstooths' reasons for attending the game? On Structuralism, probably very weighty. After all, they are oriented towards Alabama football as a particularly meaningful use of their time and their money; it structures both how they spend their time, and it also significantly affects how they spend their money (imagine that like the actual Reeses, the fictional Houndstooths also spent \$300,000 on a motor home). Indeed, the amount of time and money that the Houndstooths poured into Alabama football clearly reflects how significantly this orientation structures their values. Moreover, this practical orientation undoubtedly issues in any number of motivational states that will be especially well-promoted by their attending the Alabama-Tennessee game. Thus, because their reason(s) to attend the game will be massively overdetermined by motivational states issuing from a practical orientation that plays a major role in structuring their practical identity, the Houndstooths' reasons will be especially weighty.

But if we follow Schroeder by explaining the Houndstooths' *moral* reasons as overdetermined by almost any motivational state with reasons-giving authority, will we be able to correctly (by the ambitious Humean's lights) account for the fact that their moral reasons have practical priority over non-moral reasons—even particularly weighty non-moral reasons like their reasons to attend the game? This isn't clear, since on the approach to explaining the existence of moral reasons for all agents developed in Chapter 5 (an approach that, in broad outline, follows Schroeder [2007a]), an agent's moral reason to φ will exist because some moral course of action will promote the object of at least one of her

motivational states. But the degree to which φ -ing will promote the object of any old motivational state is quite minimal, where φ -ing involves e.g., acting on a principle that no one (who is suitably motivated) can reasonably reject. Indeed, unless the Houndstooshs happen to have motivational states with reasons-giving authority that are directly well-promoted by φ -ing, then their moral reason to φ will likely be of little comparative weight since, in light of (2), it doesn't *well* promote any of the relevant motivational states. Thus, by simply sticking to the strategy outlined in Chapter 5, it's not clear that Structuralism will be able to explain Practical Priority.

Therefore, Structuralism can't simply bootstrap Practical Priority from Categoricity. We'll need to develop an alternative strategy—one that will explain how, in light of Structuralism's commitment to (1) and (2), moral reasons (which are reasons for all rational agents) will have practical priority. Accordingly, I want to turn my attention to (1), and specifically, to the question of how particular practical orientations can play a role in structuring our practical identities.¹¹²

§4.

To help understand the claim that the weighting of reasons depends (in part) on the degree to which O plays a role in structuring S 's practical identity, where O is one of the practical orientations that play some role in structuring S 's practical identity, I borrow an idea from Harry Frankfurt's discussion of internal states. To make sense of this idea in a theory-

¹¹² You might be wondering why I'm not also considering (2), but I'm not sure there's much to say beyond the analysis of the promotion relationship I offered in Chapter 5. On Promotion, it'll turn out that some motivational state is *well-promoted* in the way described by (2) just in case S considerably raises the likelihood of the object of that motivational state obtains, relative to its antecedent intrinsic probability of obtaining.

neutral way (I return to what internality looks like for Structuralism below), we can say that for Frankfurt, a psychological state or attitude (and for our purposes, a motivational state with reasons-giving authority) is *internal* to an agent if and only if the agent cannot dissociate herself from that state as if it were merely an alien cause that produces her behavior independently of her will. In other words, if an agent cannot dissociate herself from a particular motivational state, that motivational state is internal to the agent. Put differently, such motivational states are not merely causal forces, pushing the agent to and fro; rather, these motivational states speak for the agent—they are *her own*. Agnieszka Jaworska helpfully glosses this idea when she claims that a motivational state is internal to agent when “[it] cannot be construed as merely popping up in the agent’s psychology” [Jaworska 2007, 531]. In such a case, the agent is *identified* with the motivational state in the sense that it belongs to him.¹¹³ Again, the thought here is that an internal motivational state is one with which we can identify the agent.

Of course, I can want coffee now and not want it later. I can want to work on a paper now and not want to work on it later. And so on. But I endure even as these particular motivational states (that I’m simply assuming have reasons-giving authority) come and go. So if some motivational state *m* is to be identified with me, as internal motivational states are identified with me, it must be the case that *m* similarly endures. Consequently, internal motivational states must be necessary constituents of an agent’s practical identity, such that were an agent to dissociate herself from an internal motivational state, she would destroy her practical identity as it is currently constituted. In other words, if *S* has an internal motivational state *m*, then if *S* were to dissociate herself from *m*, she would cease to

¹¹³ For more, see Jaworska [2007].

be who she *is* as an agent. And possibly, some internal states are such that if, per impossible, *S* were to dissociate herself from such a state that she would cease to be an agent at all. That is, some internal motivational states might be, quite literally, constitutive of our agency.

Now, on Structuralism, motivational states have reasons-giving authority only because they issue from a particular class of practical orientations, which aren't atomic states or attitudes. Rather they are directed volitional structures that shape our evaluative, cognitive, and conative engagement with the world.¹¹⁴ But the notion of internality can be applied equally well to an agent's practical orientations. On this view, we can say that a practical orientation *O* is internal just in case if *S* were to dissociate herself from *O*, then she would cease to be who she *is* as an agent (or perhaps more strongly, that she would cease to be an agent at all). And a consequence of *O*'s internality would be that motivational states issuing from *O* would also be internal to the agent, since by dissociating herself from such a motivational state, presumably the agent would thereby be dissociating herself from *O*. And given the damage that agents would do by dissociating themselves from internal orientations (or from the motivation states that issue from such orientations), we can conclude that practical orientations that are internal to the agent play a necessary role in constituting (which is a particularly robust form of structuring) the agent's practical identity.

Thus, because of Structuralism's commitment to (1), any reasons explained by motivational states issuing from such orientations would be of considerable weight, since they are reasons that the agent will necessarily have given her practical identity. To help see this point, consider a dilemma in which *S* has a reason R_1 to φ and a reason R_2 to refrain from φ -ing. If R_1 is explained by a desire m_1 issuing from a practical orientation that is not

¹¹⁴ For more, see Chapter 2. See also Frankfurt [2004].

internal to S , and R_2 is explained by a motivational state m_2 issuing from a practical orientation that is internal to S , then on Structuralism R_2 is a weightier reason for S *because*, given the internality of the practical orientation from which m_2 issues, m_2 issues from a practical orientation that plays a more significant—indeed essential—role in structuring S 's practical identity. And this is just what we should expect: given Reasons Internalism, reasons will depend on who the agent is, and so plausibly, so too will the weight of an agent's reason, such that, if the reason is grounded in a practical orientation that is necessarily part of the agent, that reason will be weighty.

Moreover, given that our internal practical orientations (and the derivatively internal motivational states that issue from them) play a constitutive role in structuring an agent's practical identity, it is plausible to think that a reason (or a set of reasons) that is grounded in a practical orientation that is not internal to the agent—for our purposes, an *external* practical orientation—can never outweigh reasons grounded in and explained by the agent's internal practical orientations. After all, if I were to dissociate myself from *all* of my external practical orientations, I would do significantly less violence to my practical identity than I would do if per impossible, I were to dissociate myself from those internal practical orientations that play a constitutive role in structuring my practical identity. This suggests that any reasons that is grounded in and explained by an agent's internal practical orientations will be decisively and necessarily weightier than reasons that are grounded in and explained by the agent's external practical orientations. Thus, by introducing Frankfurt's distinction between internal and external practical orientations, Structuralism has the resources to explain how, in principle, some reasons decisively and necessarily outweigh

other classes of reasons. In other words, Structuralism has the resources to explain how moral reasons might have the property of Practical Priority.

Of course, most of our practical orientations—perhaps all of our practical orientations (as well as the motivational states issuing from any of practical orientations)—lack the property of internality. Right now, off the top of my head, I would *like* to be a couple of inches taller, I *judge* that having a 21 year single malt Scotch would be a good thing, and I *love* bacon. But despite the fact that I am motivated in these ways, and that such motivations issue from practical orientations that do play some role in structuring my practical identity, I can dissociate myself from any of these motivations (and if I am sufficiently transparent, I could also dissociate myself from any of the practical orientations that issue in these motivational states. Indeed, I have tried to dissociate myself from my orientation towards meat, even delicious meat like bacon, as to be eaten.¹¹⁵ Accordingly, none of these motivational states or the practical orientations that issue in these states are internal to me. But notice, most of our motivations, as well as the practical orientations from which these states issue, are relevantly similar to the ones I've listed. And since these external practical orientations issue in motivational states with genuine reasons-giving authority, they will generate most of the reasons that I do have. So Structuralism will also need an account of cases in which I have reasons both for and against φ -ing and that each set of reasons is exclusively grounded in and explained by motivational states issuing from external practical orientations.

To see how this would go, first consider an analogy. Unlike my brain, the rest of my central nervous system, my circulatory system, etc., my hair plays no necessary role in

¹¹⁵ I've succeeded in the dissociative aspect of vegetarianism better than the avoiding all meat aspect.

structuring my living body. If I lost my hair, I might be sad, but I wouldn't die. Thus, with respect to my living body, we could say that my hair is external (in the relevant sense). Similarly, we could say that left leg is external to my living body. After all, if my leg was amputated, I wouldn't die. Thus, my leg does not play a *necessary* role in structuring my living body. Nevertheless, I think that we can all agree that my left leg plays a much more important role in structuring my living body than does my hair, even though my left leg is not internal to my living body in the way that we might describe my brain or heart as internal to my living body. Some evidence for this is that I would do a great deal more violence to my bodily integrity by cutting off my leg than I would be getting a haircut. So although my leg doesn't thereby reach the threshold for internal (in the relevant, analogous sense), it would be more like something internal than is my hair. And what we can say about non-essential body parts can also be said about various external practical orientations—some play much more important (though still inessential) roles in structuring our practical identities.¹¹⁶

Thus Structuralism claims that in a case in which all of an agent's reasons are exclusively grounded in and explained by external practical orientations, the weightier reasons will depend on those external practical orientations that play a more significant role in structuring an agent's practical identity—i.e., those practical orientations that, were the agent to dissociate herself from them, she would do significant, but perhaps not irreparable violence to her practical identity.

This is, I take it, what actually happens when the Houndstooths are trying to decide whether to spend \$25,000 fixing up their motor home or to go on an exciting deep sea fishing adventure in the Gulf of Mexico. The reasons to spend the money on RV repairs are

¹¹⁶ I recognize that this example is somewhat imprecise, but I think it nevertheless illustrates my point.

grounded a practical orientation that while short of internal—the Houndstooths wouldn't *literally* be new agents if they decided to stop following Alabama football—nevertheless plays a very, very important role in structuring their practical identities. By contrast, while the reasons to spend the money on the fishing trip are grounded in a deep enjoyment for fishing, it is an enjoyment that they could easily do without, since it issues from a relatively peripheral practical orientation. But while the reasons for and against spending the money on RV repairs do not issue from internal practical orientations, it's clear from what we know about the Houndstooths—they'd miss their daughter's wedding or act immorally just to ensure that they were at every Alabama game—that the practical orientations that shapes and directs their love of Alabama football plays a much bigger role in structuring their practical identities than does the practical orientations that grounds their enjoyment in deep sea fishing. Thus, the Houndstooths' have weightier reasons to spend the money on the motor home.

This suggests that not all external practical orientations are created equally. In some cases, I can dissociate from an external practical orientation without risking violence to my practical identity, since the practical orientation in question plays a very minor role in structuring my practical identity. But in other cases, dissociation—while possible—would involve a radical restructuring of my practical identity. And in the latter case, reasons grounded in and explained by these practical orientations are comparatively weighty.¹¹⁷

¹¹⁷ It is here that we can see a further reason for preferring Structuralism to Proportionalism. In ordinary cases Proportionalism predicts that because an agent has a strong desire that is well-promoted by φ -ing that she has a correspondingly weighty reason to φ . But suppose that on the basis of reflection, the agent dissociates herself from this desire without significant structural damage to her practical identity. Given the possibility of dissociation without structural damage, why should her desire, which is admittedly still stronger, have decisive reasons-giving authority for her? Proportionalism can only say, "Well, it's just because the desire

* * *

By better understand what is involved in an agent's practical orientation playing some role in structuring her practical identity, we have seen that Structuralism is consistent with some reasons having practical priority. Unlike the Houndstooths' reasons for spending another \$25,000 on their motor home, reasons with genuine practical priority would have to issue from practical orientations that are internal to an agent's practical identity. But are there any such reasons? And do moral reasons specifically issue from such a practical orientation?

In the remainder of Chapter 6 then, I want to sketch answers to these questions. I will first suggest that plausibly, there is at least one practical orientation that is internal to agents' practical identities. Accordingly, the class of reasons grounded in and explained by motivational states issuing from this practical orientation will have the property of Practical Priority—i.e., they will be necessarily decisive. I will then suggest that moral reasons fall within the class of reasons grounded in and explained by motivational states issuing from this practical orientations. Thus I will claim that on plausible assumptions, Structuralism is not only consistent with, but can explain why, moral reasons have practical priority. I will further develop this claim in Chapter 7 where I will explore Structuralism's commitments concerning moral reasons in more detail.

is stronger ... so even if an agent dissociates herself from it without doing violence to her practical identity, it still has reasons-giving authority for her." And this is unsatisfying.

By contrast, Structuralism is sensitive to these considerations. In fact, Structuralism is in part motivated by the thought that the possibility of dissociation without structural damage can obviously affect the reasons-giving authority of some practical orientation and the set of motivational states issuing from that orientation.

§5.

Freeman and Betty Reese really did decide to go to the Alabama-Tennessee game rather than their daughter's wedding ceremony. But suppose that the fictional Houndstooths' one up that. Suppose that knowing that the University of Alabama has returned to its repugnant practice of not integrating its team, the Houndstooths nevertheless go to the Alabama-Tennessee game.¹¹⁸ I take it that supporting a state institution that systematically discriminates in blatant and harmful ways is morally objectionable. And to the extent that we know about such discrimination, we should not support it, and we should work to eradicate it. But the Houndstooths simply don't care. They just want to see Alabama beat Tennessee.

If the argument of Chapter 5 is correct, then the Houndstooths certainly have moral reasons to refrain from going to any more Alabama games. But these reasons are explained by the fact that almost any motivational state with reasons-giving authority can generate reasons to take others seriously in the ways required by morality. Of course, acting in these ways towards others, which in this case involves boycotting Alabama football, only minimally promotes the objects of the explaining motivational states. Thus, given Structuralism's commitment to (2), which holds that the weight of an agent's reason to φ is (in part) a function of how well φ -ing promotes the object of the relevant motivational state, the Houndstooths' reasons to boycott the racist, discriminatory policy are not particularly weighty. And given that their reasons to attend the game are explained by motivational

¹¹⁸ The Crimson Tide only integrated its football team in the early 1970s, years after the Civil Rights Amendment passed—'Bear' Bryant might have been a good football coach, but he certainly wasn't a good man. Indeed, Bryant was once asked about the changed, and he explained that it was only because integrated teams like USC were trouncing Alabama that he would even consider it. Of course, this is something that is often left out of the panegyrics that pass for Bryant biographies.

states issuing from a practical orientation that, while external, nevertheless plays a major role in structuring their practical orientation, and that attending the game would promote the object of these motivational states especially well, on Structuralism, their reasons to attend the game will be especially weighty. So weighty in fact, that plausibly, the Houndstooths should go to the game and support Alabama even though by doing so, they are outwardly supporting racist policies unconstitutionally funded by a state institution.

Of course, if moral reasons have the property of Practical Priority, then it is irrational for the Houndstooths to go to the game since their moral reasons will necessarily and decisively outweigh their non-moral reasons. But it appears that on Structuralism, when we consider what is going on in this case, the weight of reasons favors going to the game, and this would entail that the Houndstooths would be irrational if they refrained from attending. But recall that Structuralism is compatible with some reasons having Practical Priority, as long as those reasons are grounded in and explained by motivational states issuing from internal practical orientations of the agent. Thus, if the Houndstooths' moral reasons can be grounded in and explained by some motivational state issuing from an *internal* practical orientation, then they will necessarily have decisive reason to boycott Alabama's football games.

Of course, this leads us to the question of whether it is plausible to think that there might be any practical orientation that is (i) internal to all agents and (ii) that can serve to ground and explain all of an agent's moral reasons. Concerning (i), Frankfurt himself claimed that it is possible for an agent to dissociate herself from any of her first-order

desires—i.e., desires that take some state of affairs involving an action as their object.¹¹⁹ Building on this, Gary Watson has argued that plausibly, there are no normative differences between first-order and higher-order desires.¹²⁰ And this suggests that by Frankfurt's own lights, it is possible for an agent to dissociate herself from any of her desires whatsoever—thereby showing that none of her desires could have issued from an internal practical orientation since, had they done so, they too would have been internal. Thus, if Watson is right, no desires have the property of internality. Similar arguments can, and indeed have been made about other motivational states, such that things look pretty bleak for the Humean who wants to maintain that moral reasons instantiate the property of Practical Priority¹²¹

But perhaps this is too quick. Perhaps *some* motivational states with reasons-giving authority do issue from practical orientations that play a necessary role in structuring our practical identities. Such practical orientations, if they were to be necessary constituents of our practical identities must be *constitutive* of our not just our particular practical identities, but of our form of agency itself. Of course, this class of practical orientations that are putative constitutive of our form agency will, in many ways, be different from other practical orientations and the motivational states that issue from our practical orientations. Indeed, the practical orientations I have in mind may not issue in particularly *strong* motivational states (in the sense at stake in Proportionalism). Moreover, the practical orientations I have in mind might be very difficult to access introspectively, such that, even if we were

¹¹⁹ Harry Frankfurt [1971].

¹²⁰ Gary Watson [1975].

¹²¹ For an example of such arguments, extended to other motivational states, see Velleman [1992].

attentively reflecting on our reasons, we might very well not notice that we have some motivations and that those motivations issue from a particular kind of practical orientation. Finally, it is almost certainly true that such practical orientations, as well as the motivational states issuing from them, are not transparent to us the way that say, our concern for a loved one is often readily available to us. But, the practical orientations I have in mind would be *essential* for agents of the sort we take ourselves to be, i.e., rational agents.

With this in mind, consider J. David Velleman's claim that a particular desire—viz., the desire to make sense of ourselves—is constitutive of our status as rational agents.¹²² According to Velleman, the fact that we are rational agents is grounded in and explained by a desire to make sense of ourselves. This desire, unlike say, a desire for another slice of pizza, is a desire that structures and organizes our agency. Of course, on Structuralism, no individual desire, even a desire to make sense oneself, can ground and explain our status as rational agents, since our reasons depend on how we are practically oriented and how those particular orientations structure our identity. Despite this difference, Velleman is nevertheless instructive on this point.

So suppose that some practical orientation that issues in motivational states that can explain moral reasons does play an essential constitutive role in structuring our practical agency as such. Suppose that we are, as P. F. Strawson has claimed, practically oriented towards ourselves and others not as mere objects in the world, but as *participants* in a shared way of life, made rich by meaningful reciprocal interpersonal engagement. If this practical orientation, what Strawson has called “the Participant Attitude,” played an essential

¹²² Velleman [1992].

constitutive role in structuring our practical identities as agents, then moral reasons explained by motivational states issuing from this practical orientation would have practical priority.

Though I will go into this in more detail in Chapter 7, I want to briefly gesture at how identifying Strawson's Participant Attitude as playing an essential role in structuring our practical identities—as being *internal* to our practical identities—might aid Structuralism as an account of moral reasons. Suppose that because of this orientation, I engage with others as apt targets of the reactive emotions (e.g., resentment and gratitude). In virtue of this form of engagement, I come to desire to relate to others in meaningful ways.¹²³ And it seems clear that a desire to relate to others in meaningful ways would (among other things) be well-promoted by treating others as ends and never as means. After all, by treating others as such, we are thereby able to fully engage in meaningful reciprocal relations of mutual regard. And so this desire could explain why I have moral reasons not to murder or coerce you, why I have moral reason to help you in need, and why I have moral reason to generally take you seriously in a way that facilitates meaningful reciprocal relations. Furthermore, if, as I claim, this desire issues from a practical orientation that really does play an essential role in structuring our identities, then it would be an internal desire. And if so, then the reasons it explains will have practical priority.

§6.

This all suggests that Structuralism is compatible with practical priority and that an ambitious Humean reply to the Still Too Many Reasons problem might be possible. But more needs to

¹²³I take it that such a desire wouldn't be surprising in light of my emotional engagement since, by seeing myself and others as apt targets of the reactive emotions, I see myself and others as potential friends, intimates, colleagues, etc

be said. After all, in showing Structuralism to be compatible with the practical priority of moral reasons, I haven't thereby shown that on Structuralism, moral reasons *do* have practical priority. Therefore, in the next chapter I will discuss Strawson's Participant Attitude, and I will argue that plausibly, it is a practical orientation that is internal to the practical identity of any rational agent. And if this is correct, then perhaps the ambitious Humean reply to the Still Too Many Reasons problem isn't just a bare possibility, but a live philosophical option.

Chapter 7: The Participant Attitude and Moral Reasons

§1.

As we have seen, Structuralism can explain the existence and weight of ordinary practical reasons—i.e., the reasons we have to enjoy an ice cream cone on a hot summer day, the reasons we have to choose this dish from the menu rather than that dish, the reasons we have to spend our time working rather than goofing off, and the reasons we have to buy an Art Tatum record instead of some mindless adult contemporary album of saxophone covers.

So too, Structuralism can explain our moral reasons—e.g., our reasons to help Marty when he is in need. But this is hardly an impressive feat, since *any* theory of practical reasons can explain the existence of some moral reasons. The real question then, is whether Structuralism can explain a class of moral reasons that instantiate the properties of Categoricity and Practical Priority. As a first step in answering this question, in Chapters 5 and 6, I argue that Structuralism is *compatible* with the existence of moral reasons that instantiate these properties.

But though this might be a surprising claim—many have simply *assumed* that no Humean theory could be compatible with a class of moral reasons that instantiate these properties—it is a relatively weak claim. Logically, it's on par with showing that God's existence is compatible with the existence of evil, which of course, is something in itself. But it's also not what most people worry about when they're worrying about the existence of evil. When Elie Wiesel curses God in *Night*, it's not because he's thinking that maybe given the existence of so much gratuitous suffering, it's logically impossible that God exists.¹²⁴

¹²⁴ Wiesel [1960].

Rather, he's cursing God because he can't see *how* such suffering could be allowed by a putatively good god. What Wiesel needed (and what wasn't forthcoming) wasn't a *defense*, it was a *theodicy*.

Now, concerns about God's existence no doubt have a greater existential pull than concerns about our best theories of practical reasons. But there is a similar distinction to be made. Whereas in Chapters 5 and 6 I offered a defense of Structuralism, in this chapter, I hope to do something a bit more ambitious—something that is a bit more on par with a theodicy. After all, it's not enough to simply argue that Structuralism is compatible with a class of moral reasons instantiating the properties of Categoricity and Practical Priority, if this nevertheless leaves open the further practical question: “yes, I see that on Structuralism it is possible that moral reasons are practical authoritative, but are they practically authoritative reasons *for me*?”

Accordingly, in this chapter I want to close this question.

To do so, I will sketch an account of *how* it might be the case that some, and perhaps all, of *our* moral reasons (i.e., moral reasons that *we* actually have) instantiate the properties of Categoricity and Practical Priority. Of course, this sketch of how Structuralism explains the actual existence of moral reasons that instantiate the properties of Categoricity and Practical Priority is rather complicated. Therefore I want to perspicuously (though some might say “tediously”) foreshadow the argumentative strategy of this chapter.

First, before getting to the meat of the chapter,¹²⁵ I'll briefly remind us of the conditions under which Structuralism is compatible with the existence of categorical and practically authoritative moral reasons. If you've just finished Chapter 6, then this might be

¹²⁵ For the vegetarians and vegans reading this, the “hearty Portobello mushroom of the chapter.”

old hat. But if you're having some trouble keeping straight the finer points of Structuralism, this will be a welcome refresher.

Second, I will discuss Peter Strawson's "Freedom and Resentment." I begin my discussion with a broad, synoptic overview of Strawson's argument, which is primarily an attempt to argue that the "facts as we know them" provide an adequate justificatory basis for our responsibility practices. I will then consider an important objection to Strawson due to Paul Russell (but echoed in the work of many others). And I will go on to argue that Russell's objection rests on a mistake. More precisely, I will argue *pace* Russell, that Strawson's argumentative strategy is not naturalistic at all. Rather, it is a transcendental argument (which isn't surprising given Strawson's place in the Kantian tradition). Therefore, according to my interpretation of Strawson, our responsibility practices are justified because *from the practical point of view*, it is impossible for us to be oriented towards others in those ways required for meaningful relationships if we do not antecedently presuppose that we (both us and others) are apt targets of the reactive emotions.

Third, I will explore the possibility of extending Strawson's argument beyond the realm of moral responsibility. Just as Strawson claims that moral responsibility is grounded in the practical orientation that structures our status as relational agents—the so-called "Participant Attitude—I want to claim (for reasons very similar to those of Strawson) that the very same practical orientation grounds a class of moral reasons that instantiate Categoricality and Practical Priority. After all, according to Structuralism, if a practical orientation is inescapable *from the practical point of view* then it will be *internal* (in the sense introduced in Chapter 6 and to be reintroduced below) to our practical identities as rational agents. So if moral reasons could be explained by motivational states issuing from this

internal practical orientation, then they would be reasons of decisive weight for anyone with such a practical orientation—which would be all rational agents. Thus, this class of moral reasons would instantiate Categoricality and Practical Priority. I conclude by arguing that the same Participant Attitude that Strawson appeals to as the basis of our responsibility practices, does in fact serve to explain moral reasons. So on Structuralism, moral reasons are not just compatible with Categoricality and Practical Priority; they actually instantiate these properties.

So, without further ado, let's get to the details.

§2.

Structuralism isn't just a reductive analysis of the existence of an agent's reasons. It's also a reductive analysis of the weights of an agent's reasons. And on Structuralism, the weightiness of any particular reason R to φ depends (in part) on whether the motivational state m that explains R issues from a practical orientation that plays an important role in structuring the agent's practical identity. If so, then, R is a weighty reason for the agent to φ . If not, then R is of comparatively little weight.

Of course, not all roles that a practical orientation might play in structuring an agent's practical identity are equal. Not even all *important* roles are equal. That is, while some practical orientations might play an important, but non-essential role in structuring an agent's practical identity, other practical orientations might play an essential role in structuring an agent's practical identity. And if some practical orientation plays an essential role in structuring an agent's practical identity, then because she cannot dissociate herself from it—at least not without destroying her practical identity—it is *internal* to her practical

identity. Now, as I said in the previous chapter, a practical orientation might be internal to an agent in one of two ways: (i) it might play an essential role in structuring her particular, contingent practical identity, and (ii) it might play an essential role in structuring her practical identity *as a rational agent*. Now, for reasons discussed in the previous chapter, I doubt that there is any actual practical orientation that is internal to an agent because it satisfies (i). But perhaps there is a practical orientation that we must have if we are to be rational agents at all, as is required by (ii). If there is a practical orientation that is internal to our practical identities as rational agents, then it would serve as the basis of a class of moral reasons that instantiate Categoricity and Practical Priority.

The question then, is whether there is any *actual* practical orientation that is internal to our practical identities as rational agents? I think there might be. And I think it might be what Peter Strawson has called the “Participant Attitude”—the orientation from which we are able to relate to others reciprocally. Accordingly, I want to explore Strawson’s account of the Participant Attitude, beginning with its role in justifying our responsibility practices. Of course, this will require that we appreciate the general outline of Strawson’s project. So get ready for some heavy duty interpretative work.

§3.

Before I begin the real work of explicating Strawson, I want to assuage any worries that a trip through Strawson’s theory of responsibility is an unnecessary detour by assuring you that it’s not. As we’ll see in §4, the role that the Participant Attitude plays in grounding our responsibility practices will also serve as a basis for moral reasons that instantiate Categoricity and Practical Priority. So understanding Strawson’s project in “Freedom and

Resentment” will be invaluable for understanding how the Participant Attitude can serve as the basis of moral reasons.

3.1.

Strawson begins by distinguishing between two positions—“optimism” and “pessimism”—that you could adopt as a way of spelling out the relationship between causal determinism and the justification of our responsibility practices. On the one hand, optimists claim that determinism is no threat to the justification of our responsibility practices (and so, optimists correspond roughly to contemporary compatibilists). And on the other hand, pessimists worry that determinism would threaten the justification of our responsibility practices (so, pessimists correspond roughly to contemporary incompatibilists).¹²⁶

But Strawson doesn’t just come out and defend (or reject) one of these positions outright. Rather, Strawson interestingly takes his project to be that of reconciliation. In other words, in “Freedom and Resentment” he’s looking to marry the virtues of each of these views. With this task in mind, Strawson asks about the possibility of reconciling optimism and pessimism—i.e., *what would reconciliation between optimist and pessimist look like?* In answering this question, Strawson writes:

Well, there might be a formal withdrawal on one side in return for a substantial concession on the other. Thus, suppose the optimist’s position were put like this: (1) the facts as we know them do not show determinism to be false; (2) the fact as we know them supply an adequate basis for the concepts and practices which the pessimist feels to be imperiled by the possibility of determinism’s truth. Now it might be that the optimist is right in this, but is apt to give an inadequate account of

¹²⁶ I say “roughly” because unlike optimists as Strawson describes them, many contemporary compatibilists are not concerned to ground our responsibility practices in “the facts as we know them to be.” Likewise, many contemporary incompatibilists are concerned with the facts as we know them to be. So in comparing optimists and pessimists to contemporary compatibilists and incompatibilists, you might find that your mileage varies.

the facts as we know them, and of how they constitute an adequate basis for the problematic concepts and practices; that the reasons he gives for the adequacy of the basis are themselves inadequate and leave something out vital. It might be that the pessimist is rightly anxious to get this vital thing back and, in the grip of his anxiety, feels he has to go beyond the facts as we know them; feels that the vital thing can be secure only if, beyond the facts as we know them, there is the further fact that determinism is false. Might *he* not be brought to make a formal withdrawal in return for a vital concession [Strawson 1974, 1-2]?

As we can see, Strawson begins his task of reconciling optimism and pessimism by first diagnosing their problems.

According to Strawson, the problem with optimism is that the optimist reduces our responsibility practices to mere instrumental relations—relations that “*exploit* our natures” as Strawson will go on to say. That is, the optimist seeks to justify our responsibility practices on consequentialist grounds, by appealing to their *usefulness* in securing a well-ordered society—e.g., “optimists about determinism point to the efficacy of the practices of punishment, and of moral condemnation and approval, in regulating behavior in socially desirable ways,” [Strawson 1974, 2].

By contrast, the pessimist rejects the optimist’s putative justification for our responsibility practices. Indeed, she “recoils” at the notion that our responsibility practices are mere instruments for social regulation, to be justified only if they efficiently bring about certain behavior. Rather, by the pessimist’s lights, our responsibility practices are justified because agents who act in morally significant ways *deserve* to be praised and blamed, rewarded and punished. But this too leads to a problem since, by the pessimist’s lights, desert requires freedom and freedom requires the falsity of causal determinism. Unfortunately, the falsity of causal determinism is not among the facts as we know them. So for Strawson, the problem with pessimism is that the pessimist fails to take seriously the idea that the facts about our

own humanity as we currently know them are sufficient to ground our responsibility practices.

But despite this deep disagreement between optimist and pessimist, Strawson thinks that each of these views that must be taken seriously, since each of these view does get *something* right. On the one hand, the optimists are right to think that the facts as we know them provide an adequate justification for our responsibility practices.¹²⁷ But, because they appeal to the efficacy of our responsibility practices in regulating behavior as a justification for the practices, Strawson takes the optimists to be missing something vital; as Strawson says, “[utility] is not a sufficient basis, it is not even the right *sort* of basis, for these practices as we understand them,” [Strawson 1974, 4].¹²⁸ In other words, the optimist undermines her own insight by appealing to the wrong set of facts.

On the other hand, the pessimist is also right about something important since she clearly recognizes that any attempt to justify our practices in terms of their social efficacy in regulating behavior fails to appreciate the significance of these practices. For Strawson then, the pessimist position rightly rejects consequentialist attempts to justify our responsibility practices. But it is here that the pessimist overreaches since by her lights, only freedom of the will and a metaphysical notion of desert, where this would entail the falsity of causal determinism can justify our practices. However, the falsity of causal determinism is not among the facts as we know them, and so the pessimist takes the facts as we know them to

¹²⁷ Of particular relevance, Strawson claims that the truth of casual determinism—whatever that thesis might be—is not among the facts as we know them. So when he commends the optimists for respecting the facts as we know them, Strawson is simply claiming that the optimists rightly take our current understanding of ourselves and the world to be an adequate basis for justifying our responsibility practices, it’s just that they do not correctly identify this adequate basis.

¹²⁸ For more on this point—what Stephen Darwall has called “Strawson’s point”—see Darwall [2006].

be an inadequate basis for our practices. But Strawson worries that by seeking to ground our responsibility practices outside of the facts as we know them, the pessimist fails to recognize and appreciate her “own humanity” [Strawson 1974, 24], which should be a sufficient basis upon which to ground our responsibility practices.

So understood, Strawson accuses the optimist of *over-intellectualizing* our responsibility practices, since she attempts to ground our responsibility practices in a calculating appeal to their instrumental value in bringing about a well-regulated social order. As a result, the optimist fails to appreciate the human significance of these practices because she fails to take into account *all* of the facts as we know them—viz., she ignores the fact “our practices do not merely exploit our natures [for the good of social regulation], *but express them*,” [Strawson 1974, 25, emphasis added].

But as it turns out, the pessimist also over-intellectualizes our responsibility practices, since she seeks to explain their significance by going *beyond* the facts as we know them to be. Thus, according to the pessimist, the vital thing that optimist accounts ignore, must be grounded not in our own humanity and in the attitudes that are constitutive of what it is to be a person, but in some further metaphysical fact—which unfortunately turns out to be nothing more than “a pitiful intellectualist trinket” [Strawson 1974, 24].

So whereas “the optimist’s style of over-intellectualizing the facts is that of a characteristically incomplete empiricism, a one-eyed utilitarianism” [Strawson 1974, 23], the pessimist’s style of over-intellectualizing the facts is that of (to borrow from Strawson’s [1966] description of Kant’s objection to Rationalism) a “dogmatic rationalism [that] exceeds

the upper bounds of sense” [Strawson 1966, 12].¹²⁹ This latter mistake is a failure to allow the conditions that facilitate the possibility of interpersonal relations serve as the final arbiters of the conditions of morally responsible agency.¹³⁰ If this is correct, then we can understand Strawson’s diagnosis of the problem of moral responsibility to be roughly analogous to Kant’s critique of Empiricists and Rationalists in his resolution to the Third Antinomy [CPR A444/B472 – A451/B479].

There Kant argues that the Empiricists, who defend the “Antithesis”—the view that “there is no freedom; everything in the world takes place solely in accordance with laws of nature” [CPR A445/B473]—fail to take seriously our humanity because on their view cannot accommodate our practical perspectives, to which Kant takes us to be committed. Likewise, Kant criticizes the Rationalists, who defend the “Thesis”—the view that “causality in accordance with laws of nature is not the only causality... [and that] it is necessary to assume that there is also another causality, that of freedom” [CPR A444/B472]. The Rationalists also fail to take seriously our humanity, since they fail to appreciate our place in the natural world, instead resorting to “speculative reason” and metaphysics (in the pejorative—at least for Kant) to ground our agency.¹³¹ And while Kant’s diagnoses of the mistakes of

¹²⁹ “Rationalism” in this context refers to the *metaphysical* rationalism of Plato, Leibniz, and Wolff which is contrasted with the empiricism of Epicurus and Hume rather than the *metaethical* view which I discuss elsewhere in the dissertation according to which it is never rational to act in ways contrary to morality.

¹³⁰ In *The Bounds of Sense* [1966], Strawson is attempting to reconstruct Kant’s project as he develops it in the First Critique. So when Strawson claims that the pessimists are mistaken in their search for some further metaphysical basis for the justification these relationships that abstracts away from the facts as we know about the relationships themselves—a basis that seems irrelevant to any possible experience we might have as participants in reciprocal interpersonal relationships, he is clearly drawing some analogy to Kant’s critique of Rationalists as it is developed in the Dialectic. And indeed, Strawson points to a mistake that exactly parallels that of the dogmatic Rationalists who seek to ground experience in a metaphysics which appeals to facts that are independent of any possible experience.

Empiricism and Rationalism do not map perfectly onto the mistakes of optimists and pessimists, they are nevertheless instructive. Strawson, like Kant, takes each side to be failing to capture who we are as human persons.

Of course, although it is an important first step, diagnosing the mistakes of optimism and pessimism is not Strawson's goal in "Freedom and Resentment." As I've already suggested, he primarily aim to offers us a way of *reconciling* (i) the optimist's commitment to the facts as we know them with (ii) the pessimist's commitment to what's vital about our standing as morally responsible agents. In his positive project, Strawson introduces the idea of the *Participant Attitude*, which will turn out to be the key to reconciling (i) with (ii). Accordingly, I turn to Strawson's positive project below.

3.2.

Strawson begins his reconciliation project by discussing a "commonplace" of human life, something he takes to be a basic feature of our humanity: *we participate in meaningful forms of reciprocal relations*. This fact about us, Strawson thinks, is quite important. After all, it is through these relationships that we, "as sharers of a common interest; as members of the same family; as colleagues; as friends; as lovers; as chances parties to an enormous range of transactions and encounters" [Strawson 1974, 6], are morally responsible to one another. In

¹³¹ Of course, I do not mean to suggest that Kant denies that the free will is its own causality. But it does seem to be a theme of the Dialectic that the Rationalist makes the same mistake as the Empiricist, which manifests itself for Rationalists as a willingness to disregard who we are as human cognizers by indulging in speculative flights of metaphysical fancy.

other words, for Strawson, the property of *being* morally responsible is a property that agents instantiate because they participate in these sorts of relationships.¹³²

Now, according to Strawson, we are to understand these relationships in terms of a general concern for good will. Furthermore, this concern is manifested in the “non-detached” emotions and attitudes that play a role in structuring interpersonal engagement.¹³³ These emotions, which Strawson calls the reactive attitudes, include gratitude and resentment, esteem and indignation, and many forms of pride and guilt.¹³⁴ So being engaged with others in a way that presupposes the importance of mutual respect is part of *what it is* to be involved in reciprocal interpersonal relationships at all. And such engagement is undeniably mediated by the reactive emotions, the proneness to which just is, according to Strawson, “the making of the demand [for good will]” [Strawson 1974, 22]. Thus, for Strawson the significance of respect is *presupposed* by the commonplace that we participate in reciprocal relations

Strawson certainly seems to be onto something important here. For example, it’s not clear how we could enjoy say, a friendship, if we weren’t engaged with one another in ways that involve an expectation for good will—an expectation that disposes us to feeling grateful when our friend shows us kindness, resentful when she is indifferent or cruel, and guilty

¹³² This important insight regarding the nature of moral responsibility is frequently forgotten in the contemporary debate, which too often treats the property of *being morally responsible* as a sui generis moral property, whose instantiation conditions obtain wholly independently of human life and practice. Wallace [1994], Fischer and Ravizza [1998], and Darwall [2006] are notable exceptions.

¹³³ I will often follow Strawson in labeling the object of the reactive emotions as “good” and “ill will,” but sometimes I will also talk of reasonable regard and unreasonable disregard (following Watson [1987]) or of respect and disrespect. As I mean them, they are synonymous.

¹³⁴ Strawson also includes some harder to characterize emotions such as love, forgiveness, and hurt feelings. Although these are of much interest, because they are not Strawson’s focus, I will have little to say about them here.

when we hurt her. Quite simply, this form of emotional engagement is necessary for mature, intimate friendships. And if another person was not disposed to this range of emotions, then it is plausible that she's not concerned with showing good will towards others or being shown good will by others.¹³⁵ But how could we relate to such a person? It is difficult to see how reciprocal relations would be possible, since good will and respect is not even on her radar *as something to be done*. Thus, because she lacks the relevant form of emotional engagement, she plausibly stands outside of the network of normal reciprocal interpersonal relationships. This all seems correct, and so it is fodder for Strawson's claim that normal human participation presupposes the significance of respect.

Of course, we could still interact with such a person. But such interactions would be "cold." We would naturally come to see interactions with such a person much more calculatingly, as do φ when around S in order to get ψ . And indeed we probably *should* see our interactions with S in this way, since any engagement with S that involved the full panoply of reactive emotions would be unreciprocated and leave us frustrated and render S the unfair target of our ire. Thus, we would, as Strawson puts it, begin to develop a kind of *objectivity* of attitude towards this individual. For Strawson, to take this objective attitude towards another involves *reorienting* yourself with respect to them, adopting new presumptive interpretations of their behavior, and developing strategies for "dealing" with them.

About this practical orientation we take towards those with whom normal interpersonal relationships are sufficiently impaired or impossible, Strawson writes:

To adopt the objective attitude to another human being is to see him, perhaps, as an object of social policy; as a subject for what, in a wide range of sense, might be called

¹³⁵ This is something like a restatement of Strawson's claim that it is the proneness to the reactive emotions that constitutes the basic demand for reasonable regard.

treatment; as something certainly to be taken account, perhaps precautionary account, of; to be managed or handled or cured or trained; perhaps simply to be avoided [Strawson 1974, 9].

When we are practically oriented towards another through the Objective Attitude, the ways in which we are emotionally engaged with her are significantly altered. Although we are still prone to certain forms of emotional engagement, such engagement is *toned* or *muted*, and it reflects the unfortunate fact that something has gone wrong. Accordingly Strawson continues:

The objective attitude may be emotionally toned in many ways, but not in all ways: it may include repulsion or fear, it may include pity or even love, though not all kinds of love. **But it cannot include the range of reactive feelings and attitudes which belong to involvement or participation with others in inter-personal human relationships;** it cannot include resentment, gratitude, forgiveness, anger, or the sort of love which two adults can sometimes be said to feel reciprocally, for each other. If your attitude towards someone is wholly objective, then though you may fight him, you cannot quarrel with him, and though you may talk to him, even negotiate with him, you cannot reason with him. You can at most pretend to quarrel, or to reason, with him [Strawson 1974, 9; emphasis added].

The Objective Attitude then, is not devoid of emotion, but it is devoid of the forms of emotion and personal engagement that characterizes human life as we know it to be. So when we engage with another person from the Objective Attitude, our engagement is comparatively thin. It is less satisfying and meaningful than full engagement with other persons.

Thus, the Objective Attitude can be contrasted with an attitude of participation, which is a practical orientation we *presumptively* take towards our fellow persons and which manifests itself in the full suite of reactive emotions. In other words, because the Objective Attitude is a practical orientation that precludes certain forms of meaningful reciprocal interpersonal relationships, it's plausible to conclude that there is another practical

orientation—an attitude of participation—that underwrites and facilitates our participation in such relationships. Accordingly, our engagement with others through responsibility-entailing attitudes like resentment must be grounded in an antecedent orientation we take towards others as *persons to be fully engaged with*.¹³⁶ For Strawson then, our responsibility practices are grounded in our orientation towards others via this particular attitude, the Participant Attitude.

Notice that Strawson’s claims about the Participant Attitude stand in stark contrast to the view defended by the optimist, who regards our responsibility practices with an attitude infused with objectivity. But despite this fundamental point of disagreement, Strawson does seem to be respecting the optimist’s commitment to the facts as we know them. Thus, he ultimately sides with a “radically” altered version of the optimism because he thinks that the facts as we know them, which include our awareness of the Participant Attitude, are sufficient to justify our responsibility practices.

But this is simply an assertion, rather than an argument. After all, this doesn’t answer the lingering question: how does Strawson move from the fact that we are practically oriented towards other through the Participant Attitude to the conclusion that our responsibility practices are justified? As an answer to this question, Strawson claims that our responsibility practices are grounded in and justified by

[our] human commitment to participation in ordinary inter-personal relationships [which] is ... too thoroughgoing and deeply rooted for us to take seriously the thought that a general theoretical conviction [like the truth of causal determinism] might so change our world that, in it, there were no longer any such things as inter-personal relationships as we normally understand them [Strawson 1974, 11].

¹³⁶ Here I’m contrasting “full engagement” to the kind of emotional engagement that we might have with someone when we regard him or her from the Objective Attitude.

In other words, our practices are justified by our *inescapable* commitment to reciprocal interpersonal relationships, which itself depends on our antecedent engagement with others from the Participant (rather than the Objective) Attitude.

Indeed, for Strawson it is “practically inconceivable” [Strawson 1974, 11] that we give up the Participant Attitude—a practical orientation that directs us to regard ourselves and others as possible participants in meaningful relations. And as Strawson points out, a theoretical belief in determinism seems irrelevant to this sort of engagement. After all, even if causal determinism is true, it couldn’t affect the status of our interpersonal relationships as meaningful. We already know them to be!¹³⁷ Consequently, causal determinism cannot undermine the status the responsibility-entailing reactive emotions that reflect the Participant Attitude and *presuppose* our engagement with others through this emotion.

In other words, because it is “practically inconceivable” that a theoretical commitment to the truth of causal determinism could lead us to abandon the Participant Attitude, Strawson claims that the truth of causal determinism is not a threat to the legitimacy of our responsibility practices. And this, Strawson concludes, allows us to reconcile the optimist’s insights with those of the pessimist—“to make a formal withdrawal in return for a vital concession,” [Strawson 1974, 2]. After all, by grounding our responsibility practices in our practically inescapable commitment to the practical orientation that grounds certain forms of reciprocal interpersonal relationships, Strawson is able to respect the optimist’s constraint that we stick to the facts as we know them and take seriously the pessimist’s commitment to the non-instrumental conception of our practices.

¹³⁷ I take it that for Strawson one of “the facts as we know them” is that our reciprocal interpersonal relationships as we currently understand them are in fact meaningful.

Strawson's idea then, is simply that our commitment to the Participant Attitude is *inescapable* for us. It is practically inconceivable that we presumptively and necessarily engage others from the Objective Attitude, as if that was the normal or only rational form of human interaction. Likewise, we could not abandon the meaningful relationships made possible by our orientation towards others through the Participant Attitude if we simply discovered the truth of causal determinism. Therefore, according to Strawson, the justification for our responsibility practices is found in the inescapability of the Participant Attitude.

3.3.

However, on this point, many commentators—even contemporary compatibilists who agree that causal determinism *per se* does not rule out moral responsibility—have thought Strawson to be over-reaching at best or wholly misguided at worst.¹³⁸ After all, how could the *mere* fact that it would be practically inescapable or that it would be *very hard* to conceive of our lives independently of the Participant Attitude, rationalize our responsibility practices? Put differently, why should we think that the natural human commitment to interpersonal relationships (and to the reactive emotions internal to such relationships)—thoroughgoing though it may be—serves as a legitimate rational or justificatory basis for our responsibility practices? With no obvious answer forthcoming, it would seem that Strawson's biological, psychological, and sociological ruminations about our human commitment to the Participant Attitude, no matter how accurate they might be as descriptions of human nature, cannot credibly rationalize our responsibility practices. And this suggests that Strawson's attempt to

¹³⁸ Cf. Russell [1992], Sommers [2007, 2011], and Vargas [2004, forthcoming]. Bennett Helm has also suggested this interpretation of Strawson to me in conversation.

ground the legitimacy of our responsibility practices in facts about human nature, while avoiding the mistakes that plagued the optimists and the pessimists, is itself inadequate.

Paul Russell develops this line of thought into a full-fledged objection to Strawson's "naturalistic strategy" for justifying our responsibility practices.¹³⁹ For Russell, Strawson's naturalistic strategy relies on two claims.

- (1) Given certain features of human nature, it is psychologically impossible for us to abandon the Participant Attitude (or its attendant emotions).
- (2) Given this psychological impossibility, a theoretical belief that causal determinism obtains should not undercut our responsibility practices.

And according to Strawson, it follows from (1) and (2) that causal determinism is no threat to the legitimacy of our responsibility practices.

But Russell claims that replying to the pessimist in this way completely misses her point. To see this, distinguish between two forms of naturalism that might be underwriting the move in Strawson's argument from (1) to (2): type- and token-naturalism.

Type-Naturalism: Our liability to adopt the Participant Attitude and to engage with others via the reactive emotions "is natural to humans and requires no general justification of any sort" [Russell 1992, 294].

Token-Naturalism: Our liability to adopt the Participant Attitude and to engage with others via the reactive emotions is so deeply rooted in human nature that "we will inescapably or inevitably continue to entertain or feel [these emotions], whatever reason suggests to us" [Russell 1992, 295].

After he distinguishes between these type- and token-naturalism, Russell then argues that while type-naturalism might be plausible, it doesn't rule out the pessimist's view. After all, even the pessimist can maintain the relatively weak claim that natural facts about humans

¹³⁹ Russell [1992].

don't require a general justification—simply that they require a justification in specific contexts, and that causal determinism would offer us such a context. As Russell says:

The Pessimist does not (or need not) claim that we are capable of suspending our disposition or liability to the reactive attitudes—much less that the thesis of determinism requires us to do so. ... Rather, the Pessimist claims only that we can and must cease to entertain reactive attitudes insofar as we have reason to believe that *everyone* is incapacitated in the relevant ways. If the thesis of determinism is true, the Pessimist argues, then we are, indeed, all morally incapacitated [Russell 1992, 296].

In other words, the pessimist can accept type-naturalism since she can say that of course we would be thoroughly disposed to engage with others via the reactive emotions even if determinism were true without committing herself to the further claim that any particular instance of resentment would be fitting in that case. And if this is correct, then Strawson cannot appeal to the truth of type-naturalism to undermine the pessimist's position.

Thus Russell claims that if Strawson is to undercut the pessimist's position, he must do so by embracing token-naturalism. And although type-naturalism might be plausible, token-naturalism is a much stronger claim (than type-naturalism) since on token-naturalism (but not type-naturalism) the natural human proclivity to adopt the Participant Attitude is one that is not sensitive to reason. That is, even if we see that there are undeniably weighty reasons to refrain from regarding others from this perspective—reasons that will undercut the legitimacy of our emotional engagement with others via the reactive emotions—according to token-naturalism, the commitment to the Participant Attitude is wholly unshakeable. And so, it can't be any more rationally objectionable than other given facts of human nature—e.g., that we have hair or that we are bipedal. Thus by Russell's lights, Strawson must assume token-naturalism to move from (1) to (2).

Unfortunately for Strawson however, while token-naturalism can answer the pessimist's challenge and perhaps explain the move from (1) to (2), it is incredibly implausible. Russell makes this point forcefully:

...it seems clear why the Pessimist finds Strawson's [token-] naturalistic reply both misguided and disturbing. What is particularly disturbing about Strawson's naturalistic strategy, expressed in more general terms, is that it casts doubt on our ability or capacity to curb or control our emotional life according to the dictates of reason. More specifically, it seems clear that, despite disclaimers to the contrary, Strawson's naturalistic strategy invites us to accept or reconcile ourselves to reactive attitudes (and their associated retributive practices) even in circumstances when we have reason to repudiate them. Given this, it seems evident that we have good reason to reject Strawson's suggestion that we dismiss the Pessimist and refuse to take his arguments seriously. We have, on the contrary, every reason to take the Pessimist seriously [Russell 1992, 297-298].

Because token-naturalism commits us to an implausible conception of human nature—one in which we genuinely are slaves to destructive forces of our reactive emotions—we have good reason to doubt its truth. On this point, I hardly agree with Russell. And this suggests that if Strawson is relying (even implicitly) on token-naturalism as part of his objection to pessimism or as part of what explains the move from (1) to (2), then his argument is unsound.

Luckily for Strawson, I do not think this is a problem for his reply to the pessimist, since I do not think his reply is a naturalist one (and even if it is a kind of naturalist reply, it is not one that commits Strawson to token-naturalism). As I understand him, Strawson's appeals to human "nature" are like Kant's appeals to "pure psychology" in the First Critique: such appeals are neither naturalistic nor empiricist; they are transcendental.¹⁴⁰ Indeed, if we

¹⁴⁰ The relevant notion of psychology here is specifically transcendental psychology rather than rational psychology, which Kant criticizes in the Dialectic.

understand Strawson as sketching the transcendental basis of our practices (rather than a naturalistic one), then Russell's worries will be vitiated.

3.4.

Of course, in favor of Russell's interpretation, Strawson does argue that we can ignore the pessimist's concerns because of our thoroughgoing commitment to the Participant Attitude. And, Strawson undeniably uses language that makes Russell's interpretation seem correct. For example, Strawson does put our commitment in terms of a "*human* commitment" (emphasis mine), and this certainly suggests that this commitment is merely a contingent (though very deep) feature of human beings qua human beings (rather than say, human beings *qua* rational agents). This explains why Russell and others (Vargas [2004, forthcoming] and Sommers [2007, 2011]) interpret Strawson's claim that it is "practical[ly] inconceivable" that we abandon this commitment as a biological, psychological, or sociological claim—viz., that given our human psychologies, it would be very hard, if not impossible, for us to abandon the Participant Attitude. And this, they (rightly) claim, is an insufficient justification of our responsibility practices.

But I think this interpretation is mistaken, particularly on what it means to say that it is "practically inconceivable" that a theoretical belief in determinism could lead us to no longer regard ourselves and others from the Participant Attitude. As I understand him, when Strawson claims that it is "practically inconceivable" that we could abandon our commitment to the Participant Attitude, to the reactive emotions, and to interpersonal relationships, he is not making a psychological (or biological or sociological) claim at all. Nor is he using "practically" as a way of modifying the inevitability of our commitment.

Rather, to say that it is *practically* inconceivable that we could abandon the Participant Attitude is to say that it is inconceivable or inescapable from the *practical point of view of agency* that we abandon this orientation. In other words, Strawson is claiming that from the perspective we take on the world as agents, we cannot fail to regard engage with others from the Participant Attitude.

Thus, it isn't that this orientation is psychologically or even epistemically inescapable (Strawson admits this when he says that "it does not seem to be self-contradictory to suppose that this might happen. So I suppose we should say that it is not *absolutely* inconceivable that it should happen" [Strawson 1974, 11]); it's that it is practically inescapable for us insofar as we are agents engaged in interpersonal relationships. And regarding myself as a participant in interpersonal relationships has implications not only for how I am oriented towards you, but also how I understand my own agency. As a result, if we regard ourselves as rational agents (i.e., agents for whom the activity of asking for and giving reasons has significance), we must first regard ourselves (and others) from the Participant Attitude.

Now if we read Strawson in this way, then I think Russell's argument loses much of its force. For then Strawson surely isn't claiming that our engagement with others via the reactive emotions is *irrational* or wholly immune to rational revision.¹⁴¹ If, as the case may

¹⁴¹ In a footnote Strawson does compare the justification of our responsibility practices with the justification of inductive belief-formation, calling the latter "original, natural, non-rational (not *irrational*)" [Strawson 1974, 23]. He then approvingly cites Hume as giving making clear the facts about the external basis of induction. This suggests—more strongly than anything else—that Russell's interpretation of Strawson is a bit closer to Strawson's project than the interpretation I offer. Of course, I think there a host of reasons for thinking that this footnote is anomalous, but I would just say that even if Strawson is implicitly trying to ground moral responsibility through token-naturalist commitments, there is a position very close to the one that Strawson articulates—one that he hints at—and it is *this* position, one closer to Kant than to Hume, that I want to consider.

be, reason dictates that we prune certain retributive features of the reactive emotions, such that we disengage from the more destructive aspects of the strains of involvement, then we should do so.¹⁴² And indeed, I suspect we should do so by Strawson's own lights. After all, in discussing two kinds of pleas, which correspond to excusing and exempting conditions, Strawson freely admits that we must be sensitive to facts that affect an agent's capacity to participate fully in reciprocal interpersonal relations.¹⁴³ That an agent wasn't himself, that he was pushed, that he was coerced, that he was under significant duress, etc., are all facts that are relevant to the fittingness of the reactive emotions. Indeed, as Strawson says:

Inside the general structure or web of human attitudes and feelings of which I have been speaking, there is endless room for modification, redirection, criticism, and justification. But questions of justification are internal to the structure or relate to modifications internal to it. The existence of the general framework of attitudes itself is something we are given with the fact of human society [Strawson 1974, 23].

Here Strawson is simply reiterating the point that these attitudes, and we might assume, the more general Participant Attitude that structures them, are constitutive of human society, of our particular way of life. So Strawson does think that within our shared way of life, there is room for rational revision of the sort Russell thinks is impossible given the truth of token-naturalism.

But how does this help? That is, how does it help Strawson's argument that we read the notion "practical inconceivability" in a more Kantian way, as part of a transcendental argument for moral responsibility, one that would demonstrate that the justification for our responsibility practices doesn't hang on the truth or falsity of causal determinism? After all,

¹⁴² Perhaps for reasons of the sort Gary Watson considers in "Responsibility and the Limits of Evil," [1987], we should prune the retributive features of the reactive emotions. Of course, I do not think this entails that we should wholly abandon these emotions or the Participant Attitude that underwrites their communicative and expressive significance.

¹⁴³ For a helpful discussion of Strawson on excusing and exempting conditions, see Watson [1987].

it is possible that a belief in the truth of causal determinism would be devastating our conception of the meaningfulness of our interpersonal relationships.¹⁴⁴ Perhaps if we really believed—I mean *really* believed—that determinism was true, we *would* give up our concern for human relations as they are currently constituted.

Or alternatively, as Derk Pereboom has argued, it might be that a belief in causal determinism would not undermine the meaningfulness of our reciprocal interpersonal relationships. But as Pereboom also suggests, this fact—that interpersonal relationships would be truly meaningful even if we believed causal determinism to be true—does not show, as Strawson takes it to show, that we are thereby morally responsible in a basic-desert entailing sense. Rather, it simply shows that meaningful interpersonal relationships do not require moral responsibility.¹⁴⁵ And this suggests that in trying to fill the gap left by optimist and pessimist accounts, there is a lacuna in Strawson's own account of the justification of our responsibility practices.

However, these worries fail to undermine Strawson's point, which is just that our responsibility practices are grounded in the regard we have for person from a distinctively practical point of view, viz., the Participant Attitude. Indeed, insofar as we regard ourselves as rational agents (i.e., as participants in a way of life that involves among other things, the practices of asking for and giving reasons), we are rationally committed to our status as morally responsible agents. After all, how could I see (consistently) see you as a friend if I don't regard you as owing me respect *and* as being owed respect by me in return (after all,

¹⁴⁴ Compare Ivan Karamazov's claim that if God doesn't exist then everything is permissible. As a matter of sociology, it certainly seems like many people accept this. Perhaps too, if people came to believe in the truth of determinism, they would similarly revise many of their deeply held beliefs.

¹⁴⁵ Pereboom [2001, 2012].

friendship isn't a one way street)? And if I see you as owing me respect, then I must see you as capable of showing me respect. But you can be capable of respect only if your actions can convey or express interpersonally significant forms of respect. And your actions can convey or express interpersonally significant forms of respect only if you are sufficiently sensitive to and responsive to reasons—i.e., only if you control your actions. Of course, to be in control of your actions *is* to be, in some sense, morally responsible for your actions. Since causal determinism doesn't threaten your sensitivity or responsiveness to reasons, then it can't, from the practical point of view, threaten your responsibility, at least understood as this minimal form of responsible agency that is internal to our lives as agents involved in reciprocal interpersonal relationships of mutual respect. Responsibility of this sort is presupposed by our orientation towards ourselves and others as participants in meaningful interpersonal relationships.¹⁴⁶

If this is correct then Strawson provides us with a strategy for responding to Russell's restatement of the pessimist's worries. It's not that we're unfortunately locked into the Participant Attitude the way some of us find ourselves unfortunately locked into our heights, attractiveness, or senses of humor. Rather, we're locked into the Participant Attitude because we are rational agents engaged in reciprocal interpersonal relationships. And this distinctive class of relationships requires the capacities to act for reasons and to evaluate the reasons that we and others act on. Moreover, part of what it is for us to evaluate the reasons that we and others act on is for us to hold ourselves and others to

¹⁴⁶ I admit that this is just a sketch of what a fuller argument would look like since as it stands, this argument doesn't close off all skeptical avenues. Pereboom [2001] is especially insightful in his criticisms of Strawson on this point. However, I do advance the argument here because I am optimistic that, with some changes, the skeptical avenues could be closed off. Moreover, it can serve as an analog to an argument that is of particular interest to me in this dissertation—one that concludes that moral reasons that instantiate Categoricality and Practical Priority might plausibly depend on the Participant Attitude.

certain standards (e.g., that they act for reasons that we cannot reasonably reject). In other words, in virtue of our Participant orientation towards others, we *hold ourselves and others responsible*. Consequently, we must hold ourselves and others responsible if we are to engage in a form of life which prizes reasons, which is of course, the only life we know. Thus, to give up on the game of moral responsibility is to give up on the game of asking for and offering reasons, which from the practical point of view of rational agency, we cannot do since there cannot be a reason to act for no reasons.¹⁴⁷

* * *

Of course, my purposes, it might not matter if it turned out that no adequate transcendental defense of Strawson's claims about the basis of our responsibility practices can be offered. For my purposes, what's important about Strawson's discussion of the Participant Attitude is his claim that this orientation is inescapable from the practical perspective. After all, this would suggest that the Participant Attitude might be a practical orientation that is internal to rational agency as such. That is, the Participant Attitude might play an essential constitutive role in structuring our practical identities because it might play an essential constitutive role in structuring our status as rational agents.

And if this is correct, then according to Structuralism, any reasons that are explained by motivational states issuing from the Participant Attitude will have the properties of Categoricity and Practical Priority. Perhaps then, the ambitious Humean, who wants to say

¹⁴⁷ And even if per impossible we could give up on the game of asking for and giving reason, to give up on *this* game, is to abandon the interpersonal relationships that structure our lives and give them meaning.

that evil agents like the General don't have such weighty reasons to act immorally that it might sometimes be rational for them to do so, can appeal to the resources introduced by Strawson in his discussion of moral responsibility as a way of defending the extensional adequacy of Humean theories of reasons. In the remainder of this chapter then, I consider what such a position would look like.

§4.

At this point, I want to extend Strawson's argument. I want to use Strawson's claim that the Participant Attitude is a practical orientation from which there is no escape and Structuralism's claim that reasons explained by motivational states issuing from an internal practical orientations instantiate Practical Priority. These two claims are the ingredients we need explain the putative Categoricity and Practical Priority of moral reasons.

Therefore in the remainder of this chapter, I will consider some reasons for thinking that the Participant Attitude is inescapable (some of which were anticipated by the discussion in §2). I will then argue why this fact, together with the truth of Structuralism will entail that a class of moral reasons instantiates Categoricity and Practical Priority.

4.1.

The first ingredient in securing moral reasons instantiating Categoricity and Practical Priority is *practical inescapability*. But what does it mean to say that the Participant Attitude is *practically* inescapable? How does this differ from other forms of inescapability—like the inescapability of Alcatraz? And why might inescapability of any sort matter, normatively speaking?

These are the questions I will answer in §§4.1. – 4.2., and hopefully, in so doing, I'll offer some good reasons to think that the Participant Attitude, as described by Strawson, really is practically inescapable and that this can ground moral reasons that instantiate Categoricality and Practical Priority.

A first pass at a definition: an activity or practical orientation is *practically* inescapable if and only if the conditions that would allow you to abandon that activity or practical orientation only arise from *within* that activity or practical orientation itself. This is somewhat opaque, so let's consider an argument for why activity itself might be practically inescapable. Concerning activity itself, Christine Korsgaard reasons:

Human beings are *condemned* to choice and action. Maybe you think you can avoid it, by resolutely standing still, refusing to act, refusing to move. But it's no use, for that will be something you have chosen to do, and then you will have acted after all. Choosing not to act makes not acting a kind of action, makes it something that you do.... You can't *undertake* [the activity of failing to act]—if you did, you'd be faking, and what's more, you'd be acting, in a wonderfully double sense of the word. So as long as you're in charge, so long as nothing happens to derail you, you must act. You have no choice but to choose, and to act on your choice [Korsgaard 2009, 1].¹⁴⁸

Korsgaard's idea is simply that we necessarily act, at least insofar as we are agents. In other words, action is inescapable for us from the practical point of view of agency.

Sure, we can kill ourselves or severely injure ourselves, but then we only find respite from the necessity of acting by giving up our practical identities as agents. As long as we have a practical identity at all, we must act. Action is thus practically inescapable because it is only escapable by abandoning the practical point of view altogether, which by the way, you can only do by acting.

¹⁴⁸ Korsgaard [2009].

Of course, while I find Korsgaard helpful on this point, I think that there is something more to be said. After all, you might think that we act only because we are antecedently oriented towards the world in a particular kind of way. The idea here is simply that action seems to presuppose that we have some practical orientation. To see this, consider this simple argument—the *Action Requires Orientation* (ARO) argument.

- (1) Ration action is action done for reasons.
- (2) On Structuralism, an agent's reasons are explained by motivational states issuing from practical orientations that play a role in structuring the agent's practical identity.
- (3) So, an agent can act rationally only if she has motivational states issuing from practical orientations that play a role in structuring her practical identity.
- (4) So, an agent can act rationally only if she has some practical orientation.
- (5) So, rational agency as such requires that we are practically oriented towards the world in some way.

Here, we see that rational activity—and so, the kind of activity that Korsgaard is describing—presupposes that we are practically oriented towards the world in some way. So if we are truly condemned to act, as Korsgaard has argued, so too must we be condemned to have some antecedent practical orientation that makes action possible. And this orientation, whatever it is, would be a fundamental ingredient of rational agency.

But the ARO argument doesn't tell us that we have to have any *specific* practical orientation, let alone the Participant Attitude. So why should we think, with Strawson, that the Participant Attitude in particular, is inescapable?

Well first, it is plausible that our fundamental orientation towards the world is not passive. Insofar as we have any practical point of view, we do not merely find things happening to us and to others. Rather, our fundamental orientation towards the world is

active. But human activity requires the ability for participation, since distinctively human activity is rational. And rational activity is action done for reasons. But the very fact that there are reasons presupposes that there are reasoners. But an agent is a reasoner only insofar as she regards herself as engaged in the game of asking for and giving reasons—a game whose rules are set not by independent, irreducibly normative facts, but by the requirements of human life. In other words, in regarding yourself as a reasoner, you are committed to participation in this feature of human life—that we ask for and offer reasons for action. Of course, you can participate in the game of asking for and giving reasons only if you are practically oriented towards others as owing you reasons and being owed reasons by you, as participants in a shared way of life. So, our fundamental practical orientation is the Participant Attitude.

Taking stock, we've seen three important things.

First, Korsgaard has given us a reason to think that rational activity is practically inescapable for us. Second, the ARO argument offers us a reason to think that rational activity presupposes that we have some (or other) practical orientation. And third, the above reflections suggest that the practical orientation that the ARO argument shows to be presupposed by rational activity is the Participant Attitude. Taken together, this certainly suggests that the Participant Attitude is practically inescapable.

Strawson's ruminations further suggest the inescapability of the Participant Attitude. According to Strawson, "the existence of the [Participant Attitude] itself is something we are *given* with the fact of human society" [Strawson 1974, 23; emphasis added]. Now, here you might worry that there is a gap in Strawson's argument since human society seems to be escapable. If the "inescapability" of the Participant Attitude is given with the fact of human

society, in what sense is it *really* inescapable? After all, I could move out into the desert, eschewing all human involvement until either I die of thirst or starve or manage to eke out a subsistent existence. In any case, haven't I thereby escaped human society?¹⁴⁹

To this challenge I offer two replies. First, I would say that if this is what it takes to escape the Participant Attitude (and thereby escape the rational commitments that emerge from this practical orientation), then if you want it, you can have it. But such a life surely ain't for me. Or anyone I've ever known for that matter. If the "inescapability" of the Participant Attitude really is like this, then while escape from the Participant Attitude is practically possible, it is especially unattractive. And maybe it is so unattractive that it is psychologically impossible that any human person could choose this option. But now Russell's worry reemerges: why should this kind of impossibility or inescapability matter normatively?¹⁵⁰

So alternatively, the Strawsonian might simply stick to her guns: by moving to the desert you only escape from human society in one sense—you are far away from other persons and you have limited social interactions. However, in much more significant sense, you have not escaped human society at all. For starters, your ability to make sense of yourself—to act for reasons—only emerges because you have come from a community of practical reasoners who ask for and offer reasons for action. This isn't simply a genealogical point. That is, it's not just that you learned how ask for and give reasons from someone else

¹⁴⁹ On this view, the inescapability of the Participant Attitude seems more like the inescapability of US law for its citizens. We can escape its jurisdiction, but only by fleeing the country and renouncing our citizenship (if we want to avoid extradition), which, despite the current right-ward turn in American politics, is still an unattractive option for most of us.

¹⁵⁰ By leaving the issue here, I'm not suggesting that there is no plausible reply to Russell's worry as it reemerges in this context. Rather, I simply think the second reply, which I consider below, is adequate.

in human society. Rather, it's that the very standards for what counts as an adequate reason are determined by human society. Such standards are given by the constitutive standards of human society itself.

Moreover, because the isolation that you experience is only geographical, the fact that you have left human society doesn't isolate you from others in further ways. For example, you can't, simply by leaving town, decide to just do whatever you want, consequences for others be damned. And you can, even though you leave town, expect that others won't unreasonably interfere with your decision to do so. So moving to the desert doesn't isolate you from the rational commitments that make human society possible. And if this is correct, then by simply moving away—by isolating yourself geographically—you have not really left human society, at least not in the sense at stake. Nor have you repudiated the Participant Attitude itself (even if you have repudiated many of the features we identify with it) that is given with the fact of human society.

Thus, no matter how far behind we leave human society, we find that we are always oriented towards ourselves and others from the Participant Attitude. Indeed, as long as we regard our actions as not mere bodily movements but as expressions of our agency, we are committed to the Participant Attitude. As long as we “take responsibility” for our actions in the sense that see ourselves as an apt target for the reactive attitudes, we are committed to the Participant Attitude.¹⁵¹ And simply leaving human society in mundane ways—like moving to the desert—does not affect such commitments, commitments that are grounded in who we are as agents. Therefore the Participant Attitude is inescapable for us because it plays a fundamental role in constituting our practical identities. We can't be who we are

¹⁵¹ On this point, see John Martin Fischer and Mark Ravizza [1998].

(practically speaking) without being oriented towards the world in this particular way. But of course, merely showing that the Participant Attitude plays an essential role in structuring our practical identities doesn't yet fully explain its reasons-giving authority. That is, we need to understand the normative significance of practical inescapability. Accordingly, I turn to this question below.

4.2.

If we want to explain the normative significance of inescapability, the question we need to be asking ourselves is this: how do we get from the inescapability of the Participant Attitude to Structuralism's claim that if the Participant Attitude is internal to our agency then motivational states issuing from it have *decisive* reasons-giving authority? For if we can explain how practical inescapability is tied to reasons-giving authority, then we'll be able to make plausible, on broadly Structuralist grounds, an account of moral reasons as instantiating Categoricity and Practical Priority.

Of course, you might be generally skeptical of such a move. Why should the fact that X is inescapable for me matter normatively? This concern is at the heart of Russell's objection to his token-naturalist interpretation of Strawson: why does the fact that I cannot revise my emotional reactions according to the dictates of reason justify those reactions? A similar question could be asked about reasons: why does the fact that I cannot repudiate the Participant Attitude (from the practical point of view) entail that it is a source of reasons for me? These questions are important, since in many contexts, inescapability is normatively irrelevant. The fact that I'm tied to a chair against my will means that the chair is, in some sense, inescapable. But no one thinks that this entails that I have a reason to be in the chair

or that I have a reason to stay in the chair. So why should the fact that the Participant Attitude is a practical orientation from which I cannot escape—at least not while maintaining my status as a rational agent—generate moral reasons?

Again, Korsgaard is helpful.¹⁵² According to Korsgaard’s favored view, metaethical constitutivism (or “constitutivism” for short), because an activity is inescapable for us, the constitutive norms of activity as such—which Korsgaard takes to be the Categorical Imperative—have reasons-giving authority for us.

To illustrate the general move from the constitutive norms of an activity to reasons-giving authority, consider the following mundane example. Insofar as you’re playing chess, the constitutive norms of chess have reasons-giving authority for you. It’s hard to see how anyone could deny this. Thus, if you’re playing chess you have reason to checkmate your opponent’s king, to never move your knight in a straight line, and to keep your bishops on squares of the same color. In such a case, it’s no mystery why the constitutive norms of chess have reasons-giving authority: by choosing to play chess, I *want* these norms to guide my actions.

Of course, you don’t have to play chess—you could choose not to. In which case, you no longer have chess-reasons. So chess-reasons are not categorical. Nor do they have practical priority. But activity as such, Korsgaard stresses, unlike chess, is not optional. As she likes to say, whatever else you’re doing, you’re always acting. And so the norms of acting

¹⁵² Korsgaard [2009].

as such have reasons-giving authority for you, since whatever else you're doing, you're condemned to act.¹⁵³

However, as I suggested in §4.1., we act for reasons only because we are antecedently engaged with ourselves, others, and our environments through the Participant Attitude. That is, we are “condemned to choice” (as Korsgaard puts it) only because we are oriented towards the world as participants. In other words, we are sentenced to action only because we are first condemned to regard ourselves and others not as mere objects, but as participants in a shared way of life—a way of life given by the fact of human society that involves the activities of asking for and giving reasons for action and of holding others responsible for how they act. And if Korsgaard’s constitutivist strategy is successful, perhaps we can point to our inescapable commitment to the Participant Attitude as a basis for moral reasons just as Strawson points to our inescapable commitment to the Participant Attitude as a basis for moral responsibility.

According to *Strawsonian Constitutivism*, it is the Participant Attitude that is practically inescapable.¹⁵⁴ And it is inescapable because this practical orientation plays an essential role in structuring our identities as rational agents. Just as was the case in “Freedom and Resentment,” this is where Strawson’s discussion of the reactive emotions is of particular importance. According to Strawson, we primarily manifest our orientation towards others as fellow participants in human society through our emotional engagement. When we oriented

¹⁵³ David Enoch [2009] uses a chess example to press the above worry—viz., that mere inescapability is normatively irrelevant. It seems to me that Enoch misses the point, since no one has ever claimed that *mere* inescapability is normatively significant, only that *practical* inescapability is normatively relevant to agents like us who have practical perspectives.

¹⁵⁴ Strawsonian Constitutivism is meant to parallel Korsgaard’s constitutivism. But rather than taking the inescapability of action as the basis of moral reasons, Strawsonian Constitutivism takes the inescapability of the Participant Attitude as the basis of moral reasons.

towards someone as an object, our emotional engagement is accordingly muted and infused with objectivity. But when we are oriented towards someone as though she is a participant in our shared way of life, our emotional engagement is must richer. It is not just happiness and sadness, frustration and elation, fear and anticipation that characterize the emotional coloration of the Participant Attitude. It is also gratitude and resentment, esteem and indignation, and pride and guilt. And these emotions, the reactive emotions, take as their intentional object the good or ill will displayed in particular actions. Thus, because these emotions manifest our engagement with ourselves and others through the Participant Attitude, it would appear that part and parcel of what it is to take up the Participant Attitude is to be practically oriented towards our own actions and the actions of others as needing to manifest interpersonally significant forms of respect.¹⁵⁵

If this is correct, then insofar as we are oriented towards the world from the Participant Attitude, we are oriented towards respect as significant. Indeed, this seems fundamental to the Participant Attitude, since this orientation is one of participation, and reciprocal participation—at least reciprocal participation of a particularly meaningful sort—requires mutual respect. And this has important implications for an agent's reasons, since a practical orientation that directs us towards good will and respect as to be prized and directs us towards ill will and disrespect as to be repudiated, would undoubtedly issue in

¹⁵⁵ Now you might be thinking that if what I'm trying to do in this chapter is explain how, on Structuralist grounds moral reasons plausibly instantiate Categoricality and Practical Priority, that I have made a mistake here. After all, Structuralism is a thoroughly reductive analysis of practical reasons and the notion of respect seems to be irreducibly normative.

While this is a legitimate worry, I think that it is too quick to assume that the notions of respect or good will are irreducibly normative. You might think, for example, that what it is to respect someone is nothing more than to treat them in ways that allow for the two of you to interact in reciprocal cooperative relationships. On this conception of respect, it is not irreducibly normative, since what counts as respectful or disrespectful will bottom out in the decidedly non-normative conditions under which reciprocal relationships of cooperation are possible.

motivational states whose objects would be well-promoted by showing others good will and respect. And since good will and respect are at the very *heart* of morality, it would follow that the class of motivational states issuing from the Participant Attitude would explain our moral reasons. Moreover, since these motivational states—e.g., my desire to have friends or Jennifer’s judgment that love is a particularly good thing—issue from a practical orientation that plays an essential role in structuring our practical identities as rational agents, they have reasons-giving authority for all agents, and their reasons-giving authority is decisive. In other words, on Structuralism, it would seem that these motivational states would explain reasons that have the properties of Categoricity and Practical Priority. And if so, the ambitious Humean is vindicated.

4.3.

To conclude, I want to briefly reconsider the General. He is undeniably evil and worthy of blame. All parties agree to those judgments. But the question remains: does he still have too many weighty reasons to countenance evil actions—such weighty reasons that it might not be irrational for him to act in evil ways? According to the modest Humean, while it’s probably the case that the General’s evil is irrational, there is nothing in principle that rules out the possibility of an agent who has such weighty reasons to act immorally that she might be rational for doing so. On this view, it’s unfortunate that it might sometimes be the case that agents can make themselves into creatures who have most reason to act in heinous, immoral ways. Unfortunate, but not impossible.

By contrast the ambitious Humean claims that if this is where we are left, then the Rationalists are correct to worry about the extensional adequacy of Structuralism and other

versions of HTR. So in an attempt to explore the ambitious Humean strategy, I have considered what resources Structuralism might have for grounding the Categoricality and Practical Priority of moral reasons. I ultimately claimed that because the Participant Attitude plays an essential role in structuring our practical identities as rational agents (a claim I lifted from Strawson), any motivational states issuing from this practical orientation will have necessary and decisive reasons-giving authority. That is, these motivational states will explain reasons that instantiate the properties of Categoricality and Practical Priority.

Thus, on this view, even the General would have a decisively weighty reason to refrain from killing the boy. He would not still have too many weighty reasons since, on this ambitious Humean reply, insofar as he has a practical identity at all—that is, insofar as reasons have any significance for him—he is most fundamentally (and essentially!) oriented towards others through the Participant Attitude. And this orientation will ultimately explain why his reasons to kill the boy as punishment or as an example or to instill fear and respect in his serf are necessarily and decisively outweighed. As Strawson suggested, it is an inescapable fact given by human society that good will and respect have the particular significance that they do. And so, it is a given fact for us and for the General that *we* all have moral reasons and that they are decisive.

Selected Bibliography

- Behrends, Jeff and Joshua DiPaulo. [2011]. "Finlay and Schroeder on Promoting a Desire," *Journal of Ethics & Social Philosophy*.
- Buber, Martin. [1970]. *I and Thou*. Touchstone Press.
- Darwall, Stephen. [2006]. *The Second-Person Standpoint: Morality, Respect, and Accountability*. Harvard University Press.
- Davidson, Donald. [2001]. *Essays on Actions and Events*. 2nd ed. Oxford University Press.
- Dostoevsky, Fyodor. [2002]. *The Brothers Karamazov*. Farrar, Straus and Giroux.
- Enoch, David. [2006]. "Agency, Shmagency: Why Normativity Won't Come from What is Constitutive of Agency," *The Philosophical Review* 115: 687-715.
- _____. [2011]. "On Mark Schroeder's Hypotheticalism: A Critical Notice of *Slaves of the Passions*," *The Philosophical Review* 120: 423-446.
- Faulkner, William. [1990]. *Light in August*. Vintage International.
- Ferrero, Luca. [2009]. "Constitutivism and the Inescapability of Agency," in R. Shafer-Landau (ed.), *Oxford Studies in Metaethics IV*: 303-333.
- Finlay, Stephen. [2006]. "The Reasons that Matter," *The Australasian Journal of Philosophy* 84(1): 1-20.
- Fischer, John Martin. [2010]. "The Frankfurt Cases: The Moral of the Stories," *The Philosophical Review* 119: 315-336.
- Fischer, John Martin and Mark Ravizza. [1998]. *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge University Press.
- Frankfurt, Harry. [1969]. "Alternate Possibilities and Moral Responsibility," *The Journal of Philosophy* Vol. 66, No. 23: 829-839.

- _____. [1971]. "Freedom of the Will and the Concept of a Person," *The Journal of Philosophy* Vol. 68, No. 1: 5-20.
- _____. [2004]. *The Reasons of Love*. Princeton University Press.
- Goldman, Alan. [2009]. *Reasons from Within*. Oxford University Press.
- Hieronymi, Pamela. [2005]. "The Wrong Kind of Reason," *The Journal of Philosophy* 102 (9). 437-57.
- Hume, David. [2000]. *A Treatise of Human Nature*. Oxford University Press.
- Jaworska, Agnieszka. [2007]. "Caring and Internality," *Philosophy and Phenomenological Research* 74 (3): 529-568.
- Kant, Immanuel. [1998]. *Groundwork for the Metaphysics of Morals*, trans. Mary Gregor. Cambridge University Press.
- _____. [2003]. *Critique of Pure Reason*, trans. Norman Kemp Smith. Palgrave Macmillan.
- Korsgaard, Christine. [2009]. *Self-Constitution: Agency, Identity, and Integrity*. Oxford University Press.
- Kraut, Richard. [2007]. *What is Good and Why: The Ethics of Well-Being*. Harvard University Press.
- Nietzsche, Friedrich. [1989]. *Beyond Good and Evil: Prelude to a Philosophy of the Future*, trans. Walter Kaufmann. Vintage Books.
- Pereboom, Derk. [2001]. *Living Without Free Will*. Cambridge University Press.
- Quinn, Warren. [1995]. "Putting Rationality in its Place," *Virtues and Reasons: Philippa Foot and Moral Theory*. eds. Rosalind Hursthouse, Gavin Lawrence, and Warren Quinn. Oxford University Press.
- Raz, Joseph. [2002]. *Engaging Reason: On the Theory of Value and Action*. Oxford University Press.

- Robinson, Marilynne. [2012]. *When I was a Child I Read Books*. Farrar, Straus and Giroux.
- Russell, Paul. [1992]. "Strawson's Way of Naturalizing Responsibility," *Ethics* 101: 287-302.
- Scanlon, T. M. [1998]. *What We Owe to Each Other*. Belknap Press of Harvard University Press.
- _____. [2008]. *Moral Dimensions: Permissibility, Meaning, Blame*. Belknap Press of Harvard University Press.
- Schroeder, Mark. [2007a]. *Slaves of the Passions*. Oxford University Press.
- _____. [2007b]. "Weighting for a Plausible Theory of Humean Reasons," *Nous* 41 (1). 138-60.
- _____. [2007c]. "The Negative Reason Existential Fallacy," unpublished manuscript.
- Smith, Angela. [forthcoming]. "Moral Blame and Moral Protest," in D. Justin Coates and Neal A. Tognazzini (eds.), *Blame: Its Nature and Norms*. Oxford University Press.
- Sommers, Tamler. [2007]. "The Objective Attitude," *The Philosophical Quarterly*. Vol. 57 (228): 321-341.
- _____. [2011]. *Relative Justice: Cultural Diversity, Free Will, and Moral Responsibility*. Princeton University Press.
- Strawson, P. F. *The Bounds of Sense*. Methuen.
- _____. [1974]. "Freedom and Resentment," *Freedom and Resentment and Other Essays*. Methuen.
- Tognazzini, Neal A. [2012]. "Blame as a Volitional Activity," unpublished manuscript.
- van Fraassen, Bas. [1980]. *The Scientific Image*. Oxford University Press.
- Vargas, Manuel. [2004]. "Responsibility and the Aims of Theory: Strawson and Revision," *Pacific Philosophical Quarterly* Vol. 85 (2): 218-241.
- _____. [forthcoming]. *Building Better Beings*.

- Velleman, David. [1992]. "What Happens When Someone Acts?" *Mind* 101(403). 461-481.
- Wallace, R. Jay. [1994]. *Responsibility and the Moral Sentiments*. Harvard University Press.
- Watson, Gary. [1975]. "Free Agency," *Journal of Philosophy* Vol. 72, No. 8: 205-220.
- _____. [1987]. "Responsibility and the Limits of Evil: Variations on a Strawsonian Theme," in Ferdinand Schoeman (ed.), *Responsibility, Character, and the Emotions: New Essays in Moral Psychology*. Cambridge University Press: 256-286.
- Williams, Bernard. [1981]. *Moral Luck: Philosophical Papers 1973-1980*. Cambridge University Press.
- Wolf, Susan. [1982]. "Moral Saints," *Journal of Philosophy*. Vol. 79, No. 8: 419-439.