Title
What are practical reasons? Explaining the counting in favor of relation

Author
Johnson, Rachel Ann

Publication Date
2013

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

What are practical reasons?

Explaining the counting in favor of relation

A dissertation submitted in partial satisfaction of the

requirements for the degree Doctor of Philosophy

in Philosophy

by

Rachel Ann Johnson

2013

ABSTRACT OF THE DISSERTATION

What are practical reasons?

Explaining the counting in favor of relation

by

Rachel Ann Johnson

Doctor of Philosophy in Philosophy

University of California, Los Angeles, 2013

Professor Pamela Hieronymi, Chair

It is commonly claimed, in normative ethics, that a reason for action is a consideration that counts in favor of performing an action. There is, however, considerable debate about whether and how this counting in favor of relation is to be explained. That is, there is considerable debate about in virtue of what a consideration counts in favor of performing an action. In this dissertation, I defend the idea that a consideration counts in favor of performing an action in virtue of picking out something about that action that would have value. Chapter 1 examines a prominent strategy for defending this explanation of the counting in favor of relation. This strategy attempts to understand the relation by looking at the role that reasons for action play in distinguishing intentional actions from non-intentional actions. I argue that this distinction is not helpful for understanding the counting in favor of relation, and I suggest that an investigation into this relation should focus instead on a

different distinction—that between actions that agents perform because they regard those actions as "called for" by something about them and actions that agents are merely moved to perform. Chapters 2-4 then examine accounts of the counting in favor of relation that attempt to capture this distinction without appealing to value. They appeal instead to non-normative desires and formal principles of reasoning. I argue that each of these accounts fails to capture the distinction. Furthermore, the ways in which each account fails to capture this distinction make clearer what the distinction is and why, in order to capture it, we need to appeal to value. In the final chapter, I discuss what the arguments presented in Chapters 2-4 show about what taking a consideration to "call for" an action amounts to. I also explain why we need the idea that a reason counts in favor of an action in virtue of picking out something about that action that would have value in order to account for this sense in which, when an action acts for a reason, she takes that reason to call for her action.

The dissertation of Rachel Ann Johnson is approved.

Barbara Herman

Gavin Lawrence

Stephen Finlay

Pamela Hieronymi, Committee Chair

University of California, Los Angeles

2013

Table of Contents

Introduction.....1

I. What is a reason for action?.....1

II. The counting in favor of relation as primitive.....3

III. The Intuitive Thesis.....8

IV. Overview of the dissertation.....11

Chapter One.....17

I. Reasons as intelligible answers to the why question that picks out intentional action.....18

II. Agential control—what a reason explains.....22

III. Agential control and value.....30

IV. The Problem.....34

V. An objection and some clarifications.....44

VI. Moving forward.....50

Chapter Two.....57

I. The shareability of reasons.....58

II. Why think we are required to share reasons?.....61

III. Four commitments of constructivist accounts of reasons for action.....74

IV. Korsgaard's constructivism.....80

V. The Problem.....84

VI. Proposed solution.....89

VII. Why this solution?.....92

VIII. Conclusion.....96

Table of Contents

Chapter Three.....99

I. Street's constructivism: the basics.....100

II. “Counting in favor of”.....105

III. The Problem.....108

IV. The Problem revisited.....115

V. Three options for explaining “counting in favor of”.....122

VI. S-sharing is required part of the unreflective experiential notion of a reason for action.....125

Chapter Four.....134

I. The basic claim: Desires are background conditions on reasons.....136

II. Constitutive explanation.....137

III. How Schroeder fits into my project.....140

IV. Objection 1: Humean reasons are objectively self-regarding.....144

V. Objection 2: Agent-neutral reasons.....154

Conclusion.....173

I. “Calls for”.....173

II. How R accounts for “calls for”.....184

III. Why we need R to account for “calls for”.....187

IV. Final Remarks: How many explanations?.....196

Bibliography.....203

Acknowledgments

I would like to especially thank Pamela Hieronymi for five years of invaluable advice and guidance, both philosophical and professional, and for apparently unshakeable faith that this document would one day come together, without which, it may not have happened. I would also like to thank Barbara Herman, Gavin Lawrence, and Stephen Finlay for reliably helpful advice and challenging criticism. I am grateful to the many members of UCLA's Ethics Writing Workshop from 2009 to 2012 for helpful comments on early versions of most chapters in this document. Finally, I am also grateful to Douglas Lavin for piquing my interest in this topic eight years ago as my undergraduate thesis adviser.

Vita

Rachel Johnson received an A.B. in Philosophy from Harvard University in 2006. She graduated *cum laude* with High Honors in Philosophy and a citation in French. In graduate school at the University of California, Los Angeles, she received a Eugene V. Cota-Robles Fellowship in the 2006-2007 and 2009-2010 academic years and a Dissertation Year Fellowship in the 2012-2013 academic year. She has taken a position as an Assistant Professor at Trinity University in San Antonio, Texas.

Introduction

I. What is a reason for action?

It is now a familiar claim in normative ethics that a reason for an agent[1] to perform some action is a consideration that counts in favor of that agent performing that action, perhaps under a set of circumstances. Such a consideration picks out something about the action, or its consequences, that is to be said for performing it (at least for a particular agent, under a particular set of circumstances). For example, that vitamin A promotes good eyesight counts in favor of agents consuming foods that contain vitamin A (other things equal). That an action will help someone in need counts in favor of performing that action (other things equal). That an action will be pleasant (or something like this) counts in favor of performing it (other things equal).

This notion of a reason for action is used to answer questions about what agents should do. What an agent should do, in a given situation, is determined by the considerations that count in favor of and against various possible actions under that particular set of circumstances. This can be quite complicated. Some reasons have more weight than others. For example, reasons to preserve one's life generally have more weight than reasons to preserve the lives of one's houseplants. Furthermore, some considerations which would, under some circumstances, be reasons for or against performing a particular action may, under different circumstances, be no reasons at all. For example, the fact that someone is dying will, under many circumstances, count in favor of trying to save her.

1 My inclusion of "the agent" in this formulation is not meant to take a position on whether an agent-relative reason relation or an agent-neutral reason relation is primary.

However, if she is terminally ill and has asked that she not be saved, then the fact that she is dying may be no reason at all to save her.[2]

This notion of a reason for action is also used to explain actions that agents actually perform. We can explain why an agent performs some action by citing the reason for which she acts, where the reason for which the agent acts is understood as a consideration that she *takes* to count in favor of acting and which moves her to act.[3] That is, actions often result from some (more or less explicit) process of practical deliberation—some calculation on the agent's part over what she takes to be said for and against acting. When an agent acts, whatever it is about the action or its consequences that moves her to act tends to be something about that action or its consequences that she takes to count in favor of so acting (whether or not she did, or would upon further reflection, determine that, all things considered, she should perform this particular action.)

All of this raises some rather obvious questions: What does it mean for a consideration to count in favor of an action? In virtue of what does a consideration count in favor of an action? What makes a consideration count in favor of an action? Or at least, what must be true for a consideration to count in favor of an action?

Ideally, we should have answers to (at least some of) these questions, if we are going

---

2 I think it is often not entirely clear whether, in such cases, the relevant consideration is no reason at all, or whether it is just a very weak reason. Whether the particular case is best analyzed in the way described is not important. It is merely meant to illustrate the possibility that a consideration that under most circumstances is a reason for a particular kind of action can, under other circumstances, be no reason at all to perform actions of that kind. (Naturally, the difference is explained by the differences between these circumstances.)

3 It is likely that not every account of the role reasons play in explaining actual actions will accept this characterization. Nevertheless, I will be operating with the assumption that there must be some compatibility between an account of reasons for action in normative ethics (one that is used to determine what we should do) and an account of reasons for action in philosophy of action (one that is used to explain actual actions). See, for example, Hieronymi, Pamela. "Reasons for action." *Proceedings of the Aristotelian Society* 111 (2011): 407-427 for a discussion of the possible positions in this area.

to answer questions about what we should do by appealing to considerations that count in favor of (and against) actions. (If we are going to figure out what counts in favor of what, with any confidence, then it would be helpful to know more about what the counting in favor of relation is.) Furthermore, if practical deliberation is characterized by "taking to count in favor of", then an account of the counting in favor of relation must provide a plausible view of practical deliberation. (I.e., answers to the four questions posed above must be such that, when we explain why an agent acts by citing her "reason" for acting, the agent can plausibly be said to take that consideration to count in favor of performing that action.)

II. The counting in favor of relation as primitive

Some philosophers largely avoid these questions by claiming that the counting in favor of relation is primitive. By this they mean that any account we give to explain what it means for a consideration to count in favor of an action or in virtue of what a consideration counts in favor of an action will have to appeal to the notion of counting in favor of itself. There are a couple of different ways that this might happen which yield quite different accounts of reasons for action. The first (and I think what is most commonly meant by the claim that the counting in favor of relation is primitive) claims that there is no general explanation of what makes a consideration count in favor of an action that will apply to all instances of counting in favor of. This is not to say that there are no explanations at all of why considerations count in favor of actions. We can give such explanations in particular cases. But these explanations will appeal to other particular instances of the counting in

favor of relation.

For example, I might explain why the fact that carrots contain vitamin A counts in favor of eating carrots by citing the fact that vitamin A promotes good eyesight and some other consideration that counts in favor of having good eyesight. For example, that having good eyesight will enable me to read subtitles, even when I have to sit in the back of a theater, counts in favor of having good eyesight. This instance of the counting in favor of relation, in turn, can be explained by appeal to other particular considerations that count in favor of being able to read subtitles. And so, the explanation continues, explaining each particular instance of the counting in favor of relation by appealing to other particular instances of the counting in favor of relation.[4]

However, there is another way of taking the counting in favor of relation to be primitive that allows for a general answer to the "in virtue of what...?" question (or, at least, to the "what must be true...?" question) and thus allows for a more general explanation of reasons for action. Suppose I think that a reason for an action counts in favor of performing that action by showing that the action has some particular feature X (e.g., that it promotes the agent's desires or that it promotes utility). If so, then every explanation of why a particular consideration counts in favor of an action will have the same form—it will show that that consideration shows that the action promotes X. However, we might then ask why the fact that some action has a particular feature X counts in favor of performing it (or makes some consideration that reflects this fact count in favor of performing the action). It may be that we cannot answer this latter question without appealing to the idea of counting

4 T.M. Scanlon and Derek Parfit, for example, seem to advocate views of this kind. (See Scanlon's *What We Owe to Each Other* (Cambridge, Mass.: Belknap Press of Harvard University Press, 1998) and Parfit's *On What Matters* (Oxford: Oxford University Press, 2011).)

in favor of. For example, if I claim that the considerations that count in favor of actions are those considerations that show that actions promote utility, then I might claim that they count in favor of actions because they promote utility and utility is to be promoted. Or, I might claim that utility is to be promoted because considerations that show that actions promote utility count in favor of performing those actions. This latter view is one on which the counting in favor of relation remains primitive because we cannot avoid appealing to it in our explanation of what reasons we have. Nevertheless, on the latter view, we have interesting and unified restrictions on what can count as a reason. (It is only those considerations that reveal actions to have feature X that count in favor of actions.)[5]

I am dissatisfied with views on which the counting in favor of relation is primitive in the first sense (there is no general explanation of why considerations count in favor of actions). It strikes me as a last resort when all investigations into the possibility of the general explanation that it denies fail. (This is not to deny that treating the counting in favor of relation as primitive in this sense may well be good enough for a lot of work in normative ethics. We are quite confident that we have certain reasons and about how those reasons relate to each other, and we can do a lot of work figuring out what is right and wrong on that basis. Nevertheless, if there were a unified explanation of the counting in favor of relation, it would surely be useful in the harder cases.)

I am far from alone in being dissatisfied with this view of the counting in favor of relation and there have been a number of recent attempts to raise problems for views of this

5 There is clearly room for intermediate positions—some unification in explanations of why considerations count in favor of actions but not so much so that all reasons pick out a single kind of feature. I will not attempt to draw the line between the amount of unification that is interesting and the amount that is not. At any rate, substantial unification is necessary to get a meaningful alternative to the first kind of reasons-as-primitive view.

kind.[6] For example, we might worry that, if we do not have a unified explanation of why considerations count in favor of actions, we will lack an account of how to weigh reasons against each other and face the possibility of rampant incommensurability of reasons. We might also have epistemic worries about reasons—we might worry that, if reasons are primitive in this way, there will be no plausible explanation of how agents become aware of them. Unfortunately, I think that this kind of reasons-as-primitive view is difficult to refute in this way because (1) it has the ability to answer a variety of particular questions about reasons by appealing to particular instances of the counting in favor of relation, (2) it regards the fact that there are no more general explanations available as an unfortunate fact (it admits that we want more), and (3) it points out, likely correctly, that no one has managed to give a completely unproblematic general explanation of the counting in favor of relation (we might think that this is because it cannot be done). I think that (1) provides ways to address the problems raised for this kind of view. It will not provide the kinds of solutions that we want, but the primitivist about reasons can say that this is just because we will not accept (2). Furthermore, the primitivist can claim that we should accept (2) in light of (3) and in light of how much we explain given (1).[7]

So, although I do not think that the counting in favor of relation can be primitive in the first sense, my dissertation is not intended to show that. Rather, it operates on the assumption that the view that the counting in favor of relation is primitive such that there are no general answers to the questions posed above is unsatisfying and that we should still look

6 See, for example, Ulrike Heuer's "Explaining reasons: where does the buck stop" (*Journal of Ethics & Social Philosophy* 1, No. 3 (2005)) or Pekka Väyrynen's "A Wrong Turn to Reasons?" (*New Waves in Metaethics.* 2010.)

7 See, for example, Scanlon's "Being Realistic about Reasons." (*The Locke Lectures* (2009)) for what I think are pretty persuasive rebuttals to some of the more obvious (and pernicious) problems for this kind of a view.

for an account of reasons that can provide some general explanation of what makes considerations count in favor of actions.

I do not, however, mean to rule out that the counting in favor of relation is primitive in the second sense discussed above. In trying to answer the questions posed above, I am primarily interested in finding unifying restrictions on when a consideration can count in favor of an action. It is not terribly important to me that the account that this yields can be an analysis or reduction of "counting in favor of" such that we need not appeal to this concept itself in explaining what reasons for action are. As a result, I may provide no answer to the first question—the question of *what it means* for a consideration to count in favor of an action. I am at least trying to answer the last question. I.e., I am trying to identify interesting conditions on when considerations can count in favor of actions. Furthermore, I think the second and third questions can be heard in such a way that they are not asking for analyses or reductions. They can be questions that arise for particular considerations—what makes (or in virtue of what does) *this* consideration count in favor of φ-ing? If this is right, I am attempting to answer them too. In thinking that reasons are not primitive in the first sense discussed above, I expect that the answer to these questions will be, in some important sense, the same for all particular considerations that count in favor of actions.[8]

That said, for the sake of brevity, I will, throughout this document, discuss how to "explain the counting in favor of relation" or "give an account of reasons for action". These should be read with the above clarifications in mind.

8 That said, I am not committed to *not* giving an analysis or reduction of "counting in favor of," but whether the view I am advocating can do that is not a question I mean to take up here.

III. The Intuitive Thesis

I will be defending what I take to be an intuitive way of explaining the counting in favor of relation and answering the questions posed above:

> R: A consideration C counts in favor of some agent A performing some action φ in virtue of picking out some feature of A's φ-ing (or its consequences) that would have value.[9]

I think this is a very natural way of thinking about reasons. The question of what one's reasons for performing some action are seems, in ordinary parlance, to be quite close to the question of what is (or what the agent takes to be) *good* about her action. So, I think it natural to suppose that a reason counts in favor of an action—it is something to be said for performing the action—insofar as it show the action to be good in some respect. It is because actions are good in certain respects—because there is something about them that has value—that we have reasons to perform them.

The aim of this dissertation is not to identify what has value. It is to support the idea that we will need to appeal to value(s) to explain the counting in favor of relation. If I am right, then we will need to figure out what has value in order to figure out what reasons we have and in order to fully understand the counting in favor of relation. But that is not my project here. As a result, the notion of value I am working with in the dissertation is quite loose.[10]

---

9 R is meant to give a necessary condition on reasons for action that does some work in explaining the counting in favor of relation, but it need not be a sufficient condition. That is, I do not mean to claim that whenever a possible action would have a valuable feature, there is a reason to perform it. There may be additional constraints on when considerations that pick out features of actions that have value count in favor of actions. (E.g., it may not be the case that every person in the world has a reason to send me $20 to improve my wine collection, though these actions may have some (small) value. I owe this example to Gavin Lawrence.)

10 I take no position, for example, on the issue of whether it is only states of affairs that can be bearers of value or whether other things (e.g., actions, objects) can be valuable as well.

If something has value, on the intuitive view, then it is to be promoted, preserved, or respected.[11] The three disjuncts are not meant to be exclusive. For example, health seems like a prime candidate for something that has value. It also seems likely that health is both to be promoted (e.g., we should bring it about that more people are healthy) and to be preserved (e.g., if an agent is currently healthy, she should take steps to remain that way). We might even think that preserving health is *part of* promoting it. I will not, here, attempt much of a characterization of what it means for something to be "to be respected."[12] I include "to be respected" in my initial characterization of value to try to dispel the concern that R will necessarily yield a consequentialist moral theory. Understanding reasons for action in terms of value need not mean that thinking about what do to is nothing more than thinking about which states of affairs to promote, nor need it require that it always be better to have more instances of value around.[13]

Additionally, if something has value, on the intuitive view, then its being to be promoted, preserved, or respected does not depend on the attitudes that particular agents take towards it. This is, I think, an intuitive idea. I do not get to decide, for example, whether my health is to be promoted. Whether my health is to be promoted does not (at

11 This may not be an exhaustive list of ways of being valuable. The goal here is merely to get a fairly open and intuitive notion of value on the table so as to have something to work with going forward. We will learn a bit more able what value must be if R is to be true in later chapters.

12 To get some idea of what it would be for something to be "to be respected," consider, for example, a Kantian-inspired theory that takes rational nature to be valuable in the "to be respected" sense. This need not mean that we should try to bring a lot of rational natures into the world, but instead could mean that we should (among other things) allow for the free exercise of rational natures. Such a view may call for some promotion of rational nature. E.g., it may be that we should promote the proper exercise of rational nature. But there are limits to what we may do to promote this. E.g., there are at least many cases where it seems wrong to force people to exercise their rational nature properly. Respect for rational nature seems to demand that rational natures be permitted to exercise themselves freely, even when this leads to making mistakes (including perhaps, their failure to respect other rational natures)

13 I am not convinced that we actually need the three disjuncts. It seems to me at least a live possibility that we could account for, e.g., "to be respected" in terms of "to be promoted" if we are precise about what exactly it is that is to be promoted.

least normally) depend on whether I want to be healthy. My health is to be promoted (other things equal) whether or not I want it to be promoted or think that it is to be promoted. If I fail to see this, I make an error.

As stated, this claim is oversimplified. There are ways in which, at least at first glance, it would seem that the fact that something is to be promoted, preserved, or respected can depend on a particular agent's attitudes. For example, what is to be promoted, preserved, or respected for me seems to depend, at least in some part, on my preferences—on what I *like*. In setting ends that are particular instances of values, it seems that agents can make it the case that those particular instances are to be promoted, preserved, or respected in ways that other possible instances of value are not. So, we might think that we need to allow some partial dependence of values on agents' attitudes. Furthermore, sometimes what is to be promoted might itself be an attitude—e.g., that some action will increase a particular agent's sympathy for her fellow human beings might be a respect in which that action has value. In Chapter 4, I will consider the possibility that what is to be promoted just is desire-satisfaction itself (although according to the view in question, that desire satisfaction is to be promoted does not itself depend on any particular agent's attitudes towards it). So, one goal of the dissertation is to clarify what are the ways in which something's being to be promoted, preserved, or respected can depend on a particular agent's attitudes towards it, and what are the ways in which it cannot.

R claims that a consideration counts in favor of an action insofar as it "picks out some feature of that action that would have value." I have said a bit about what it is to have value, but there remain two parts of this claim that also merit some explanation. Reasons, R

claims, cite "features" of actions. "Feature" is meant to be understood broadly. A feature of action is something about the action. It can be something about what the action would itself be or involve, or it can be something about the expected consequences of the action.

"Picks out" is also intended to avoid taking more of a position than necessary on what the content of a reason is prior to my investigation. If a consideration picks out a feature of action, that feature appears in the consideration and is attributed to the action. If a consideration picks out a feature of action that would have value, this need not mean that "value" appears in the consideration and is attributed to the action. Furthermore, if we identify what is valuable—suppose it is utility that is to be promoted—then if a consideration picks out a feature of action that has value, this need not mean that "promotes utility" appears in the consideration and is attributed to the action. All that the consideration must do, in order to pick out a feature of action that has value, is to pick out a feature of action such that the fact that the action has that feature shows that the action promotes utility. That the action has that feature is part of an explanation of why the action promotes utility (and thus has value).[14] (For example, that it is a lie could be a reason not to do something because it, plus a utilitarian argument against lying, explain why the action in question fails to promote utility.)

IV. Overview of the dissertation

In this dissertation, I aim to defend the intuitive thesis (R). That is, I aim to show that we will need to appeal to some notion of the intuitive thesis to explain what makes a

14 This description of the relationship between the content of a reason and the value that grounds reasons is inspired by Mark Schroeder's description of desires as background conditions in *Slaves of the Passions* (Oxford: Oxford University Press, 2007).

consideration count in favor of an action. The basic strategy of the dissertation is to examine accounts of the counting in favor of relation that deny R and show that they fail to capture the relation. Examining the reasons why they fail will clarify what exactly we are trying to explain, when explaining the counting in favor of relation, and reveal some reasons why we will need some version of R in order to explain it.

Chapter 1 examines a strategy for figuring out what reasons for action are that has been commonly used to assess R. Prominent defenders of R use this strategy to support it, and many critics of R seem to think that R must be false because this strategy fails to support it.[15] I will argue that the strategy does not support R, but that the reasons why it does not support R are reasons to think that it is not a promising way to develop an account of reasons for action or of the counting in favor of relation. If so, then the fact that the strategy does not support R is no reason to think that R is false.

The basic argument of the aforementioned strategy is this: Citing the reason for which an action is done reveals the action to be intentional. Intentional actions are those actions over which an agent has a certain kind of control. A reason for action reveals an action to be intentional by revealing it to be under this kind of control. Only certain sorts of considerations can do this, e.g., those that show actions to have value. So, we can discover restrictions on the content of reasons for action—we can learn something about what has to be true for a consideration to count in favor of an action—by investigating the role that reasons for action play in picking out intentional actions.

In Chapter 1, I argue that the notion of control that this strategy takes to characterize

15 This is far from the only reason for denying R. For example, there are skeptical reasons to deny R—facts about value are often taken to be metaethically dubious, whereas facts about, e.g., desires are not.

intentional actions places only minimal restrictions on what considerations can be reasons for action. It does not support R, but it also gives us no reason to think that it will reveal all of the restrictions there are on what considerations can be reasons for action (nor will it reveal all of the interesting restrictions). So, I propose that instead of looking at the distinction between intentional and non-intentional actions to figure out what reasons for action are, we look to the distinction between being moved to act by a consideration one takes to "call for" action and being merely moved by a consideration to act. The former involves some concern for a kind of justification that the latter lacks and that seems to be characteristic of acting for a reason.

Chapters 2, 3, and 4 then examine accounts of reasons for action that attempt (or can be taken to be attempting) to give an account of reasons for action that captures the sense in which acting for a reason involves taking a consideration to "call for" the action, but that also deny R. According to R, a reason for action picks out a feature of the action that has value. (Having value is understood as being to be promoted, preserved, or respected, and being so independent of any particular agent's attitudes.) According to these alternative views, a reason for action picks out a feature of the action that the agent desires, where desires are understood as not responsive to value. (E.g., that I desire something is no indication that it has value or that I take it to have value, nor is it the case that it should be any such indication.) These views then attempt to account for the difference between an agent acting because she takes her action to be called for and acting from a mere desire without appealing to values.[16]

16 I will not explicitly discuss expressivist or other sorts of anti-realist accounts of practical reasoning and reasons for action in the dissertation. I have some reasons for this (footnote continues on next page):

Chapters 2 and 3 concern constructivist accounts of reasons for action. These accounts claim that there are norms that are constitutive of practical reasoning that determine which features of action are eligible to be picked out by reasons for action. An agent makes the fact that an action has one of these features count in favor of performing that action when she desires the feature and chooses to act for its sake in proper deference to these norms.

In Chapter 2, I examine Christine Korsgaard's constructivist account of reasons for action. Through examining her account, I argue that a constructivist account cannot account for a requirement to share reasons (in a particular sense of sharing). This is because the constitutive norms of practical reasoning that constructivist views like hers propose are

In this document, I am trying to think about how we must understand reasons for action if we are to account for the sense in which agents take them to count in favor of, to call for, or to be something normative and positive to be said in favor of performing an action. There is a way of seeing this that is consistent with an anti-realist or expressivist project. As a matter of fact, I think there are normative truths, but to the extent that the account I am looking for is an account of *how we deliberate*, this account of when and why we as deliberators take considerations to call for actions could be accurate without there being any corresponding normative claims that are *true*. Suppose I am right that reasons must be grounded in something that agents regard as to be promoted. Call it P. This could be correct, as a thesis about how we engage in practical deliberation leading to action—the considerations that move us to act because we take them to call for acting could all reveal our actions to promote P—and it yet be the case that claims like "you ought to promote P" are false (we could be systematically mistaken). Alternatively, such claims could have no truth value but rather express a particular sort of approval that is distinctive in part because it is directed at P.

However, I think that expressivist and anti-realist may be ill-disposed to find the sorts of arguments that I am making to be persuasive. Many expressivist and anti-realists take themselves to be, among other things, giving error theories. (See, for example, Simon Blackburn's *Ruling Passions: A Theory of Practical Reasoning* (Oxford: Clarendon Press, 1998) and Allan Gibbard's *Wise Choices, Apt Feelings: A Theory of Normative Judgment* (Cambridge, Mass.: Harvard University Press, 1990).) They grant that we take ourselves to be thinking about normative truths in (at least some cases of) practical deliberation. They have independent reasons for thinking that there are no such normative truths. They then think of their task, in giving an account of practical reasoning, as giving an account that fits as closely as possible with how we see ourselves reasoning without relying on any claims about normative truths. But if this is their strategy, then the sorts of appeals that I make to what will seem like satisfying reasons or justifications or genuine instances of counting in favor of to rational agents will not have the same effect. The account of practical reasoning that they propose is guaranteed to conflict with some of these claims (if it is in fact true that we take practical reasoning to be about normative facts). The expressivist/anti-realist claims that the truth about practical reasoning will be hard for practical reasoners to accept, but that does not change the fact that it is the truth.

So it may be easier for the reader to proceed under the assumption that everything I am doing here is done within a normative realist framework.

norms that determine mere rational permissibility. That is, they determine that a consideration is eligible to be made a reason—that it is permissible to act for the sake of that consideration. On the other hand, modifying Korsgaard's view so that it adopts a version of R (by adopting constitutive norms that specify features of actions as to be promoted, preserved, or respected) will allow the view to include the missing requirement to share reasons.

In Chapter 3, I examine Sharon Street's constructivist account of reasons for action. This chapter has several goals. Street's view is minimalist. She posits few and quite conservative constitutive norms of practical reasoning. As a result, examining problems for her view help make clear (1) what we are trying to explain in explaining the sense in which reasons "call for" or "count in favor of" an action, and (2) why norms that determine mere rational permissibility to be a reason will have difficulty doing this. The chapter concludes with a discussion of the requirement to share reasons that I claim that a constructivist account cannot for in Chapter 2. Using the newly clarified notion of "called for", I argue that an account of reasons for action must include this requirement to share reasons, if it is to capture the sense in which reasons "call for" actions.

At this point, I think we will have strong reasons to think that we cannot account for the sense in which reasons count in favor of actions without appealing to features of actions that are to be promoted, preserved, or respected. However, the possibility remains that what is to be promoted is desire-satisfaction itself. If so, then R might be true, but reasons would still be very much dependent on non-normative desires. Chapter 4 considers an example of this alternative way of understanding reasons in terms of desires—Mark Schroeder's

Humean view on which a reason counts in favor of an action insofar as it partially explains why the action promotes the satisfaction of one of the agent's desires. I argue that his account cannot escape two traditional problems for a Humean view that stem from the dependence that reasons have on non-normative desires according to such views.

I conclude with summaries and discussions of (1) what all of this has shown about the sense in which reasons "call for" actions, (2) how R accounts for this sense in which reasons "call for" actions, and (3) why we seem to need to appeal to value (why we need some version of R) to account for this sense in which reasons "call for" actions. (I also say a bit, in the final section, about an additional commitment that appears in my arguments in previous chapters. This is the idea that the explanation of what makes a consideration count in favor of action must be such that an agent can be aware of it and still both regard her action as counting in favor of acting and comply with all relevant requirements.)

Chapter One

Introduction

Many philosophers try to explain what a reason for action is by looking at the role reasons are thought to play in picking out intentional actions. Put very roughly, the idea is that intentional actions are characterized by a certain kind of agential control. An agent's reasons for action play a role in her exhibiting this kind of control over her actions. And in order to play this role, an agent's reasons must pick out features of her actions of a certain kind. Reasons for action are then considerations that pick out features of this kind. I will call this line of argument the Appeal to Intentionality Strategy.[17] Prominent proponents of R (the claim that a reason counts in favor of some agent A performing some action φ in virtue of picking out some feature of A's φ-ing or its consequences that would have value) think that R is supported by the Appeal to Intentionality Strategy.[18] And, while there are many motivations for denying R, one of them seems to be that R must be false because the Appeal to Intentionality Strategy in fact fails to support it (perhaps because it supports some other account of reasons for action).[19] In this chapter, I will argue that both sides of this debate are

17 There is an extensive body of work devoted to the relationship between intentional actions and reasons for acting, and I do not mean to suggest that all of it exhibits this basic argumentative structure. I do think that the work in this area that seems most relevant to the question of whether reasons for action should be understood in terms of value employs something like this strategy.

18 See, for example, Joseph Raz's "On the Guise of the Good" (*Desire, Practical Reason, and the Good.* Oxford: Oxford University Press, 2010. 111-137) or Sarah Buss's "What practical reasoning must be if we act for our own reasons" (*Australasian Journal of Philosophy* 77, no. 4 (1999): 399-421).

19 See, for example, J. David Velleman's "What Happens when Someone Acts" (*The Possibility of Practical Reason*. Oxford: Clarendon Press, 2000) and *Practical Reflection* (Princeton, N.J.: Princeton University Press, 1989) for a use of the strategy that arrives at a different conclusion about the content of reasons. More generally, R is often taken to be false because it is understood as a thesis about intentional actions —the mark of an intentional action is that, when an agent acts intentionally, she acts for a reason that satisfies R. See, for example, Michael Stocker's "Desiring the Bad: An Essay in Moral Psychology" (*The Journal of Philosophy* 76, no. 12 (1979): 738-753) for a long list of now familiar purported counterexamples to R, so understood.

getting something right and something wrong. Proponents of R are mistaken that the Appeal to Intentionality yields an argument in support of R. But the reasons why this is so also show that deniers of R are wrong to suppose that the lack of such an argument suggests that R is false.

The upshot of all of this is that thinking about what makes actions intentional (at least in the way proscribed by the Appeal to Intentionality Strategy) is not a promising way to figure out what makes something a reason. At the end of this chapter, I will propose an alternative approach which will be the focus of the remainder of the dissertation.

I. Reasons as intelligible answers to the why question that picks out intentional action

The Appeal to Intentionality Strategy for explaining what a reason for action is is made possible by a close connection between reasons for action and intentional action. Since Anscombe, many philosophers seeking to explain what an intentional action is agree that it is an important and illuminating fact about intentional action that, when an agent acts intentionally, she has an answer to the question of why she acts, where this question asks for her reasons for acting. Intelligible positive answers to this question cite the agent's reasons for acting. They cite something about the agent's action or its consequences for the sake of which the agent acts. (For example, "I'm watering the garden" would be an intelligible answer to the question "Why are you pumping water?") Some responses, e.g., "I didn't know I was doing that", are not intelligible as answers to the question. Such an answer may well be an intelligible response to being asked a "Why?" question—it conveys that the question doesn't apply in a familiar way—but it is not intelligible *as an answer* to the

question. It says nothing about the agent's reasons. Sincere responses that are not intelligible as answers to the why question are thought to reveal the action in question to be non-intentional.[20]

Brief reflection reveals that (unsurprisingly) not just anything can be intelligible as a positive answer to the why question. "I love Sophocles" is not (barring some special story) an intelligible answer to the question "why are you making coffee?", where that question asks for one's reasons for acting.[21] This reveals that there are constraints on intelligibility as an answer to the why question that asks for reasons. Because the why question asks for reasons, these constraints on intelligibility as answers to the why question are also then constraints on what counts as a reason for action. They are constraints on the content of reasons for action. We might conclude from the fact that "I love Sophocles" is not (normally) an intelligible answer to the question "why are you making coffee?" that the content of a reason for action must at the very least bear some relation of relevance to the description of the action about which the question is asked. Perhaps the content of a reason must include some further description of the action or its consequences.

So, it looks like we might be able to learn something about what reasons for action are by looking for constraints on intelligibility as an answer to the why question that asks for reasons. That is, we should be able to learn something about what restrictions there are on

---

20 Anscombe, G.E.M. *Intention.* 2nd ed. Cambridge, Mass.: Harvard University Press, 2000. §5. Anscombe also allows the possibility of intelligible negative answers to why questions. "No reason" or "I just felt like it" are such answers. They answer the question in the negative and so grant that it applies to the action. There is disagreement amongst those who adopt the Appeal to Intentionality Strategy about whether there are intentional actions that are done for no reason, and thus whether negative answers to why questions pick out actions as intentional. I will discuss the possibility of intentional action done for no reason later, but for now will only consider positive answers.

21 Raz, Joseph. "When we are ourselves: The active and the passive." *Engaging Reason: On the Theory of Value and Action*. Oxford: Oxford University Press, 1999. 8.

the content of an agent's reasons for action. We might, for example, discover that an intelligible answer to a "why?" question specifies a purpose that the agent believes the action will achieve. Someone who shares my philosophical proclivities might hope to find that an intelligible positive answer to a "why?" question specifies a respect in which the agent takes her action to be good.

Initially, this strategy seems promising. For example, we might note that not every accurate description of an action under which an agent might understand her action is an intelligible answer to the "why" question. "I am moving my arm" is not (usually) an intelligible answer to the question "why are you pumping water". This answer gives a causal explanation: the water is being pulled out of the ground by the pump because my arm is moving in such and such a way. And so, the answer gives a description of my action that I might well believe true of it. Nevertheless, if I gave it as a response to the question "why are you pumping water?", you are likely to take me to not understand what I was asked or to be uncooperative. This reveals something about what the "why?" question asks for and what counts as a reason—the why question isn't asking for a causal explanation of the movements involved in the action (or at least not just any causal explanation). And so, we arrive at the familiar conclusion that not just any cause of an action can be the reason for which it is done.

Eventually, however, we might find ourselves somewhat at a loss as to how to proceed. We might simply make lists of intelligible and unintelligible answers to the why question and then look at what the intelligible answers seem to have in common and how they seem to differ from the unintelligible answers. However, there is huge variety among

intelligible answers to why questions, so this task is daunting. More importantly, even if we did find similarities between intelligible answers to why questions—say we did find that they always pick out something about the action that is good or that the agent believes to be good—we would want an explanation of *why* this is the case, an explanation that shows it not to be coincidence. Thus, we might do better in our investigation if we could find something that we take reasons for action to do or explain, something that we can understand (at least somewhat) independently of intuitions about what counts as a reason in particular cases. This could then guide our investigation of what constrains intelligibility as an answer to the why question. Intelligible answers to why questions would be those answers that can do or explain whatever it is that a reason for action does or explains.

The role of intelligible answers to the why question in picking out intentional actions suggests such a guide. An intelligible answer to the why question asked about a specific action shows that that action is intentional. So, we might then ask what must an answer to the why question must do in order to show that an action is intentional. Presumably there is something about intentional actions that distinguish them from other actions and movements, some feature (or features) that is characteristic of intentional action. So, we might think that, in order for an answer to the why question to reveal an action to be intentional, it must show that the action has some characteristic feature of intentional actions. If we could identify such a feature, something that is distinctive of intentional actions, then we would have a guide to finding further constraints on intelligible answers to why questions. We could then ask: What sorts of answers to why questions reveal an action as having this feature?

II. Agential control—what a reason explains

So what might this characteristic feature of intentional action be? What is distinctive of intentional action? Colloquially, we say that intentional actions are those actions that an agent does "on purpose." They are those actions over which she has a certain kind of control. But what kind of control?

It is commonly thought that intentional actions are those that are initiated and guided by the agent. More specifically, they are often thought to be actions that are initiated by the agent because of something the agent believes about them and that are guided through to completion by the agent via those beliefs. For example, if I pump water intentionally, I might initiate the water pumping because I believe that in doing it, I would be watering the garden. And I might guide that action via this belief. For example, I might adjust my movements so as to pump the water at a rate that I think will water the garden but not flood it.

Now notice that, if I am asked why I am pumping water, my answer to this question, picks out the content of the belief with which I initiate and guide the action. I will answer "why are you pumping water?" with "I am watering the garden". So, we might think that an intelligible answer to a why question reveals an action to be initiated and guided by an agent (and thus to be intentional) by identifying the belief that the agent used to initiate and guide the action. And so, it looks like we may have found the distinctive feature of intentional action that we were looking for to make progress with the Appeal to Intentionality Strategy. It appears to be distinctive of intentional actions that they are initiated and guided by an agent because of something she believes about it. And it looks like answers to why

questions reveal actions to be so guided by identifying the guiding beliefs.

Furthermore, we might think that an agent cannot initiate and guide an action, in the way characteristic of intentional action, with just any belief about what she would be (or is) doing when performing that action. Only beliefs of with certain sorts of contents can play this role. (It at least seems that only beliefs about what an action is or will bring about could make sense as playing this role. But perhaps there are additional restrictions on what sorts of features these beliefs must attribute to these actions or their consequences.) Whatever restrictions there are on which beliefs can play this role, they will also be restrictions on what can be an intelligible answer to a why question and so will be restrictions on what can be an agent's reason for acting.

So, if we could get clearer on how an agent initiates and guides an action with a belief about what she is doing, such that her action is intentional, we might discover whether there are these additional restrictions on the content of this belief (and thus on the content of reasons). At this point, I think it will be helpful to look at a particular use of the Appeal to Intentionality Strategy—that put forward by Joseph Raz.[22] Raz gives a more specific characterization of what an agent must do in order to count as initiating and guiding an action with such a belief that he thinks will yield these additional restrictions. He claims that, in order for an agent to count as initiating and guiding her action with these beliefs, the agent must choose to act on those beliefs rather than to resist the impulse to act.[23]

At first glance, this might seem to be an odd claim. It does not fit the phenomenology of intentional action. When we act intentionally, we do not generally

22 See footnotes 28 and 31 for comparisons to David Velleman's use of this strategy.
23 Raz, "On the Guise of the Good," 116.

consult our urges and then decide whether or not to act on them. We just act. So, I will say a bit about why Raz makes this claim (and in doing so, about what he must mean by it).

Raz thinks the choice of whether to resist the urge to act is necessary to distinguish the agent who acts intentionally from the agent who acts non-intentionally in a particular way, namely the agent who only "felt, and witnessed himself being, propelled toward acting by some psychological condition."[24] We might be tempted at this point to think that Raz is claiming that there is a class of familiar, non-intentional actions that we can only distinguish from intentional actions by pointing to the fact that their "agents" can't explain why they chose to act rather than resist the urge.

But the sorts of familiar non-intentional actions that we might take Raz to have in mind here don't seem to fit this description, nor do they require any appeal to choosing whether to act on urges to distinguish them from intentional actions. Reflexes, for example, are not the result of any belief of the "agent's" about them. We could distinguish an intentional action from a reflexive action by noting that the former, but not the latter is caused by a belief of the agent's about what the action is. Phobic actions might be something closer to Raz's "propulsion by a psychological condition". If I have a phobic fear of sharks, I might non-intentionally scramble to get out of a swimming pool upon seeing a ripple in the water, unable to stop myself and thus not in control of my action. I don't choose whether to act on this urge (as evidenced by the fact that I can't choose not to). In such a case, my acting might even be caused by the right "reason." If there is a shark in the pool, I should be getting out and my getting out *in a panic* would still be caused by the right sort of belief—my belief that there is a shark. (This same belief could cause me to act

24 Ibid., 126.

intentionally, if I do not panic.) Nevertheless, the panicked action seems to be outside of my control in a way that makes it non-intentional.[25]

Fortunately, I think we need not worry about how these cases should be analyzed. This is because we can understand Raz as instead making a conceptual point. That is, I think we should understand him as claiming that we cannot explain an intentional action merely as an action that is caused by beliefs and desires that pick out something else we do by performing it or a purpose to be achieved by it. Our intentional actions are not mechanistic in this way. Something is missing in this explanation—something that accounts for why we describe our intentional actions as things we do rather than as things that happen to us. And this can be true even if most (or perhaps even all) actually existing non-intentional actions can be distinguished from intentional actions in other ways.[26]

The idea is this: If we accept that an intentional action is an action that is initiated and guided by its agent because of her beliefs about what it is, this suggests a simple explanation of intentional action. This belief and my desire for the feature of the action that it picks out ***cause*** me to perform that action. On this picture, for example, what it is for me

25 It might turn out that the content of the belief that causes this panicked action and the content of the belief that causes the corresponding intentional action are not, in fact, the same. For example, perhaps if I exit the pool in a panic, my action is caused by the belief that there is a shark, whereas if I exit the pool intentionally, my action is caused by the belief that exiting the pool would be escaping death by shark, or some other description of the action.

26 That Raz's point is "conceptual" is, I think, clear. He claims that his points are conceptual throughout "On the Guise of the Good." (See, for example, the passage from p. 126 quoted in Section III of this chapter.) Whether it is the point I will attribute to him below is, admittedly, less clear. He presents the claim that an agent's reason for action must reveal her as choosing not to resist the urge to act in an argument against Kieran Setiya's view of reasons for action. Raz claims that Setiya's view cannot distinguish between different kinds of explanations of actions, specifically (1) Jill killed Jim because she was jealous of Jim, (2) Jill killed Jim because she felt a sudden rage; a sudden rush of blood to her head made her do it, and (3) Jill killed Jim in order to inherit Jim's wealth. ("On the Guise of the Good", 126) Only the latter is a reasons-explanation, so if Raz is right, this would be a serious problem for Setiya. Raz thinks that what distinguishes (3) from (1) and (2) is that (a) (1) and (2), but not (3), are consistent with Jill's action being non-intentional and resulting from a loss of control (127) and that (b) only (3) reveals the agent to have chosen not to resist the urge to act (by explaining why she makes this choice). Furthermore, Raz thinks that (b) explains (a).

to initiate my action of pumping water because of my beliefs about what that action would be is for my belief that pumping is watering and my desire to water to cause me to move in such a way that is pumping (and watering).

On such an explanation, the sense in which I control my intentional actions is that they are caused by my beliefs and desires about what I am doing (as opposed to caused by someone else's beliefs and desires or by a psychological mechanism that bypasses propositional attitudes altogether). Such explanations are subject to deviant causal chain worries. That is, there are ways in which beliefs and desires about what I am doing can cause me to act, but not in such a way as to make my action intentional. In Donald Davidson's famous example, an agent wants to escape the danger and burden of holding another climber on a rope. This agent believes that letting go of the rope would free him of said danger and burden. And this agent is so horrified at the realization that he has this belief and desire that he drops the rope.[27] In this case, the agent has beliefs and desires that cite the right sort of thing to be a reason—they cite an outcome of the agent's action that the

---

However, Setiya does seem to have a straightforward way to distinguish (3) from (1) and (2). He can say that only (3) cites a "ground" for Jill's action—some point she saw in it—which is picked out by her reason. (Setiya, Kieran. "Sympathy for the devil." *Desire, Practical Reason, and the Good.* Oxford: Oxford University Press, 2010. 89.) Being intelligible as a "ground" for which an action is done seems to require at least that reasons pick out some feature of the action or its consequences that the agent wants and believes the action to have. And this is enough to distinguish (3) from (1) and (2). So, if Raz's argument is to be any good, he cannot be claiming that Setiya has nothing to say to distinguish (3) from (1) and (2). Instead he must be claiming that what Setiya does have to say is not enough (or, perhaps, that if we think more about what a "ground" must be, in order to make this distinction, we will see that the "ground" of an action reveals the agent to have chosen not to resist the urge to act).

So, I think that, in order to understand why Raz thinks that an agent's reasons for action must reveal her to have chosen not to resist the urge to act, we should take him to be saying something like the following: It is enough to distinguish (3) from (1) and (2) that we understand an agent's reasons as picking out features of actions that she desires. But this is not yet enough to show that an action is intentional. (3) does something more. In the remainder of this section, I attempt to articulate why Raz might think this. That is, I attempt to articulate why an explanation of an action that cites something about the action that the agent wants would leave open the possibility that the agent loses control and acts non-intentionally, and what about (3) rules out this possibility.

27 Davidson, Donald. "Freedom to Act." *Essays on Actions and Events*. Oxford: Clarendon Press, 1980. 79.

agent knows about and desires. And these beliefs and desires cause him to do the action that has the desired outcome. However, the agent does not seem to count as initiating or guiding action via those beliefs and desires and the action seems like an accident. Thus we might conclude that if intentional action is caused by beliefs and desires, then there must be more to the story than this to explain what makes it intentional.

Raz's claim that an agent's reason for an action explains why the agent chooses not to resist the urge to act seems like an attempt to fill in this gap in the explanation of how an agent's beliefs about what she is (or could be) doing lead to intentional actions. When an agent acts intentionally, and actually counts as initiating and guiding her actions, there seems to be a sense in which it is up to her whether her beliefs and desires cause her actions that is not the case for agents in deviant causal chain cases. Even if we have no reason to believe that there is a deviant cause for an action (suppose that, unlike in the climber case, I would not have done otherwise if I had the choice and there is no easily identifiable abnormal cause of my action), Raz might say, the fact that the simple causal explanation above leaves open the possibility of a deviant cause reveals it to be at best incomplete as an explanation of intentional action.

The basic idea is this: If I pump water intentionally, I initiate and guide my pumping water because I want to water the garden and I believe that my pumping water will water the garden. But it is not enough to reveal my action of pumping water to be initiated and guided by me that this belief and desire do in fact cause me to act. If my action is intentional, if I count as initiating it, though I want to water the garden, I do not have to pump water once I realize that this is a way of watering the garden. If I do pump the water intentionally, I do

not merely observe the want cause me to do what I believe will satisfy it. It is up to me whether or not I act on this want. This makes my action a doing rather than something that merely happens to me. This makes it intentional.

For my purposes here, I will accept the basic idea that it is characteristic of an intentional action that it is up to the agent whether or not to resist the urge to act.[28] [29] My objection to the Appeal to Intentionality Strategy will be directed at what will come next: the idea that insofar as not resisting an urge to act is to be up to the agent, the agent must have a particular sort of explanation of why she does not resist the urge to act. She must choose not to resist the urge to act on a particular basis. (An intelligible answer to the why question, then, reveals an action to be intentional by showing that the agent choose not to resist the urge to act on this basis.)

At least at first glance, this seems to be a plausible idea. In ordinary cases where agents seem to decide whether or not to resist urges to act, it does seem that they make this decision on some basis.[30] For example, if I have an urge for chocolate cake, I might choose to resist it because I decide that the object of that urge (chocolate cake, or perhaps its taste)

---

28 Compare to Velleman's view: Velleman thinks that the mark of intentional action is that an agent can adjudicate between her motives for acting. Her (strongest) desires do not merely cause her to do the action that satisfies them. It is up to her whether to act on those desires, and which ones to act on. This is what makes her action intentional and makes her count as an agent. This seems to me to be the same basic idea as Raz's claim that, when an agent acts intentionally, it is up to her whether to resist the urge to act.

29 See Section V of this chapter for further discussion of different ways in which we might understand "up to the agent".

30 I suspect that some motivation for the idea that, when an agent acts intentionally, she has some particular basis for choosing not to resist the urge to act may be coming from a conflation of different ideas of control over one's actions. To illustrate this, I will borrow some terminology from Harry Frankfurt (see *The Importance of What We Care About: Philosophical Essays*. Cambridge: Cambridge University Press, 1988). Suppose that an agent controls her action insofar as she can regard it as hers. Suppose she can regard an action as hers insofar as she can identify with the motives that lead to it. There seem to be different ways in which a person could feel alienated from (fail to identify with) her motives, and subsequently regard her action as not *hers*.

would be bad for me or because I would rather have carrot cake. Or I might choose *not* to resist the urge for chocolate cake because I decide that the taste is worth the negative health effects (at least this once). Davidson's climber (supposing he avoids being startled) might decide to resist the urge to drop the rope and avoid danger to himself because dropping the rope would kill his friend.

On the face of it, the examples suggest a variety of possible bases for the choice of whether to resist an urge to act. An agent might choose whether to resist an urge by thinking about whether the object of that urge is good or bad, about how it effects her ability to satisfy her other desires, or about particular facts about the object or its effects (e.g., dropping would be, or would result in, killing). We might yet think that some of these bases should be analyzed in terms of others (e.g., we might think that what makes an action good or bad should be understood in terms of its effect on her desire-satisfaction). The hope of the Appeal to Intentionality Strategy is that there is a common basis for choosing whether to resist urges to act, or at least that there are some interesting constraints on what could be an

---

Consider for example: (1) I do something that conflicts with my beliefs about what is worthwhile. (2) I do something that undermines ends that are important to me. (3) I do something that is out of character —it is not the sort of thing that I am normally disposed to do. Each of these suggests a different conception of "identification" with my motives and a different sense in which the action is mine. (1) suggests there is a notion of identification on which I can only fully identify with my motives and regard the action they bring about as mine when those motives accord with my beliefs about what is good. (2) permits a weaker interpretation: Perhaps I can identify with my motives and regard the action they bring about as mine so long as they promote (or at least do not undermine) projects that are important to me (where this could be true even if I think that these projects are not worthwhile). (3) suggests that I have certain motivational tendencies and I will have trouble identifying with motives that are unusual for me and regarding the actions they bring about as mine. If I am normally a very calm, deliberative, sensible person, for example, I might find myself feeling uncharacteristically spontaneous and say "I am not myself today."

Perhaps there is a sense in which, when I do something non-intentionally, I do not identify with my motives. If the sense of identification required for an action to be intentional were that associated with (1), (2), or (3), this would suggest a basis for choosing whether to resist the urge to act. E.g., if the sense in which I identify with my intentional actions is (1), then it would make sense that, when I act intentionally, I choose whether to resist the urge to act on the basis of whether it is good. However, I think the arguments I make in Section IV suggest that the minimal notion of identification necessary for intentional action is much weaker than those associated with (1), (2), or (3).

agent's basis for so choosing. The Appeal to Intentionality Strategy claims that an agent's reasons for action pick out the object of urges that she has chosen, on some basis, not to resist. So, if we can figure out what basis an agent uses in making this choice, then we can figure out constraints on what can be intelligibly chosen on this basis, and so, we will have arrived at constraints on what counts as a reason for action.

III. Agential control and value

Raz thinks that, when an agent acts intentionally, she always has the option of resisting the urge to act, but she chooses not to because she takes the features of that action that she desires to be appropriate bases for endorsing that action—she takes them to be features of the action which show the action to be worthwhile (at least in some respect), to be at least partially justified, and to have value in some respect. Thus, intelligible answers to why questions will pick out features of an agent's actions that she takes to have value.[31]

Although it is not the move from the claim that, when an agent acts intentionally, she chooses whether to resist the urge to act on some basis to the claim that that basis must be whether the action would have value that I will be challenging below, I think it will be helpful to say something about how Raz gets to his conclusion. This is in part because I will refer back to considerations that appear in this argument when proposing my alternative strategy later in this chapter. The basic idea behind this move is that, if we describe the

31 Compare to Velleman's view: Velleman claims that the adjudication between motives that is characteristic of intentional action is made possible by a further motive which is constitutive of agency. ("What Happens When Someone Acts," *The Possibility of Practical Reason*, 139.) This further motive is a desire for a certain kind of self-understanding (*Practical Reflection*). So Velleman, like Raz, thinks that, in order to explain intentional action, we need to identify a basis on which agents choose whether to resist the urge to act. But he identifies a different basis for this choice (what promotes this kind of self-understanding rather than what is valuable), and thus proposes different constraints on what counts as a reason for action.

situation an agent finds herself in when choosing whether to resist an urge to act, we will see that the activity the agent is engaged in appears to be one characterized by a concern for justification. In particular, when an agent chooses not to resist an urge to act, she appears to do so because she takes the object of the urge (something about her action) to make the action (at least somewhat) worthwhile.[32]

When an agent decides to not resist urges to perform particular actions, this can have an explanation. The agent can make this choice on some basis. She might take other features of the action into account and weigh them against the feature that underlies the urge. If I am deciding whether or not to resist the urge to water the garden, I might think about whether there is a water shortage, or about whether I have other more important things to do and weigh them against whatever I take to be said for watering the garden. (Perhaps I am growing food.) I might then decide not to resist the urge to act because I think there is more to be said for watering than against it. According to Raz, this weighing of features of the action that count for and against an action just is deliberating about what actions are worthwhile or justified.[33] So, when an agent decides not to resist an urge to act after weighing the object of that urge against considerations to be said against so acting, she plausibly chooses not to resist the urge to act because she regards the object of the urge as having value in some respect that (at least partially) justifies her action.

---

32 Fortunately, for my purposes here, we will not need a clear understanding of what value is. Among other things, I will not attempt to clarify the relationship between an action's being justified, its being worthwhile, and its having value. There is considerable debate about how these concepts are related, whether they all come together, and which ground which. (See for example, Heuer, "Explaining Reasons: Where Does the Buck Stop?" for an argument that justifying reasons must be grounded in independent facts about value.) Raz thinks that justification must be done in terms of value, so for my purposes here, I will suppose that if we can plausibly think of an agent as concerned with justification, then we are thinking about her as concerned with value and with what is worthwhile.

33 Raz, "On the Guise of the Good," 126.

Of course, it seems unlikely that, whenever an agent acts intentionally, she explicitly thinks about features of her action that she might take to weigh against performing it. To see how the agent in such a case could still count as choosing whether or not to resist the urge to act on the basis of whether acting for the sake of the object of the urge is justified, consider the following description of what taking an action be justified (or worthwhile) amounts to. Raz says:

> "To take [the action] to be worthwhile in the required sense, the agent's attitude to the reason for which he acts must be capable of sustaining certain counterfactuals: Had the agent been aware (or had he thought that he was aware) of undesirable features of the action he would have formed a view on whether the features that provide, as he believes, a reason for the action still make it the action to perform in spite of its undesirable aspects. Such a view, and that is another conceptual observation, consists in some judgment on the relative importance of the good and bad features of the action."[34]

So, if we consider an agent who acts intentionally, but who does not explicitly consider whether the objects of her urges are appropriate grounds for action, we might suppose that she became aware of these features. This awareness would likely give her pause. There would a conflict to be resolved. It would not be as if the feature of the action that underlies the urge would immediately drop out of consideration or cease to move her. The agent would (normally) take the object of the urge to oppose the new feature and would thus be forced to engage in the weighing of features. This suggests that the agent regards the object of the urge as having some justificatory force all along. It has weight that could be compared to that of other features of the action in question, were they to exist, and were the agent to become aware of them.

So, to return to our example, suppose I am happily pumping water in order to water

34 Ibid.

the garden without having undergone any explicit deliberation about whether this is worthwhile. Now you come along and tell me there is a water shortage. I might recognize that this fact weighs against my pumping water. But when I do so, it is not as if I suddenly think there is nothing to be said for pumping water. It still matters to me that pumping water will water the garden. This reveals that I regard the object of my urge (watering the garden) as having some justificatory weight to be measured against other considerations like the fact that there is a water shortage.

So it looks like, when an agent chooses not to resist the urge to act, she (implicitly or explicitly) regards the object of that urge as (at least partially) justifying the action. The thought that something is good or justified tends to come with some motivational force, so it is reasonable to suppose that, when the agent chooses to act on an urge that she takes to pick out something about the action that would be good or justify it, she makes this choice, at least in part, *because* she takes the object of the urge to be a respect in which the action is good or that (at least partially) justifies it.[35] Her reason for acting identifies the object of this urge. So, her reason for acting identifies something about her action that she takes to be good.[36]

As a matter of fact, Raz's claim about the relationship between reasons and value is

35 This step from the conclusion that, when an agent acts intentionally she regards the object of the urge to act as at least partially justifying her action to the claim that, when she chooses not to resist the urge she makes this choice *because* she takes its object to at least partially justify the action may be problematic. We might think that even if the agent took the object of the urge to justify the action, she could still choose to act on the urge on some other basis. Since I will be raising problems for an earlier stage in Raz's argument, I will not discuss this further here.

36 This is not to say that she has a belief that something about her action is good. Raz thinks that an agent might have a mistaken idea of what it is to be good (likely an overly stringent or overly moralized idea of goodness). This could lead her to take something about her action to be good (insofar as it has some justificatory force in her practical deliberation) without realizing that to do so involves seeing one's action as good in some respect. ("On the Guise of the Good," 114.)

more complicated than I have described it here. He takes the claim that a reason for action identifies something about an action that is good to be a statement of a norm, rather than a statement about individual reasons.[37] So, rather than claiming that an agent's reasons always identify something about her action that she takes to be good, Raz is claiming at least that a reason for action is the sort of thing that *should* identify something about an action that is good. I suspect Raz also means to claim something stronger: when acting for a reason, the agent takes the consideration that moves her to act to be the sort of thing that should pick out something good about her action.

This shift to understanding the result of the Appeal to Intentionality Strategy as a claim that expresses a norm could be done with any alternative basis on which an agent chooses not to resist the urge to act when she acts intentionally. Doing so will make it more difficult to raise counterexamples for the resulting view (more on this in the next section). However, I think that the problems I raise for the Appeal to Intentionality Strategy apply regardless of which way we understand its conclusion. In the next section, I will present each of these problems first as it arises for the simpler version of the conclusion and then briefly discuss why understanding the conclusion as a norm will not solve the problem.

IV. The Problem

The obvious way to object to Raz's view is by producing counterexamples—cases where an agent seems to act intentionally and yet regard nothing about her actions as good.[38] (The same strategy is available for other uses of the Appeal to Intentionality Strategy that

37 See, for example, Raz's discussion of anomic actions in "Agency, Reason, and the Good," *Engaging Reason,* 31-4.

38 See, for example, Velleman, "The Guise of the Good," 117-22.

come to different conclusions about the basis of the choice of whether to resist the urge to act. Whatever this basis is, we can attempt to construct cases where an agent seems to act intentionally but regard nothing about her action as endorse-able on that basis.) I have several reasons for not going this way. The first is that I do not think the cases are very clear. Establishing whether the features of an action that move an agent, in a given case, are good likely requires a worked out theory of value, or of what a person could conceivably regard as valuable. Furthermore, even agreement on what sorts of considerations would plausibly move an agent in a particular case seems hard to come by. Take a classic purported counterexample: I do something because I think it is bad (smash crockery, defy God, waste my life, etc). Some will say that I am moved to act by the very badness of the act. Others will say that to the extent that we can see this as intentional action, I must be moved to do what I take to be bad because of something further—perhaps I see my action as an exercise of ultimate freedom or as a statement against unjust conditions that make the pursuit of the good too difficult. These further motivations pick out something about the action which is (or can plausibly be taken by an agent to be) good. So, I do not think that cases like this are likely to give us a satisfying refutation of the Appeal to Intentionality Strategy. If anything what they do is illustrate which side of several debates a person is already on.

The second reason to avoid looking for counterexamples is that, if we understand the conclusion of the Appeal to Intentionality Strategy as expressing a norm, the view is much less susceptible to counterexamples. For example, Raz takes his claim that an agent's reason for φ-ing indicates a respect in which she takes φ-ing to be good to state something like a

constitutive norm of acting for a reason. On his view, part of what it is for a consideration to be a reason for action is to be subject to a norm that directs that the consideration pick out something about the action that is good. Norms can be violated without thereby ceasing to be norms. So, cases where an agent's reason fails to pick out something about her action that is good do not show that the norm does not hold (nor do they show that the norm is not constitutive of acting for a reason). Furthermore, Raz claims that some purported counterexamples actually presuppose his thesis (understood as a norm).[39] Cases like those just discussed of acting for reasons that I know to be bad, because they are bad, involve deliberate defiance of Raz's thesis. The agent makes a feature of an action that she knows cannot constitute a reason to act (and in fact constitutes a reason against so acting) into her reason for acting, not only on purpose, but also precisely because she thinks it is not a reason. She is determined to act for reasons that are not reasons at all. A value inversion case like this seems to derive its meaning or significance from the validity of the norm. If it were not constitutive of a reason to φ that it should pick out some good feature of φ-ing, value inversion cases would lose at least some of their defiant character, and one might think, thereby lose their point.[40]

More importantly, I am not going to object to Raz's view by raising these kind of counterexamples because I think his view is subject to a more general problem, a problem that applies to the Appeal to Intentionality Strategy in general and not only to his (or any other) particular use of it. As such, it is a problem that can tell us something about how we

39 Raz, "Agency, Reason and the Good," 31-4.

40 Depending on how we understand the norm, apparent counterexamples may still present themselves. E.g., if Raz means that, when an agent acts intentionally she *takes* herself to be subject to a norm that says that her reasons should identify features of her action that are good, then we might look for actions where an agent acts intentionally, but fails to take herself to be subject to this norm.

should go about looking for an account of reasons for action.

My criticism is aimed at the claim that an agent must have some basis for not resisting the urge to act if not resisting the urge is up to her. That is, I aim to show that it is not a necessary condition on resisting an urge being up to an agent that she do so on any particular basis (that it is good, whimsical, bad, whatever), or on any basis at all. If this is the case, then we will have lost grounds for thinking that we can get interesting restrictions on the content of reasons for action from the idea that reasons for action reveal actions to be intentional by showing that it was up to the agent whether to resist the urge to act.

I think there are (at least) two reasons to think that an agent's not resisting an urge to act need not be done on any basis in order for it to count as up to her such that the resulting action is intentional.

The first is that there are ordinary cases where a kind of baseless choice occurs. And in these cases, we seem comfortable saying that the outcome is up to the agent who makes the baseless choice. Consider a case where an agent picks between options that she judges to be equally good (e.g., two cans of soup that are exactly alike). First, note that this is not a straightforward counterexample to Raz's conclusion of the sort discussed above. The agent in this case may have an explanation of why she takes *some* can of soup. After all, soup is good in a variety of respects. So, it seems that the agent can have the sort of explanation Raz wants for why she does not resist the urge for soup that reveals her action of taking a can of soup to be intentional. We need not take Raz to be claiming that, in order to count as acting intentionally, an agent must have an explanation of why she chooses not to resist an urge that shows (or even purports to show) that the action she performs is better

than the alternatives. It is enough that there is *something* to be said for the action on the basis of which she chooses not to resist the urge.

However, in the explanation of intentional action that the Appeal to Intentionality Strategy suggests, a basis for not resisting the urge to act is supposed to be needed to account for the sense in which whether or not to resist the urge is up to the agent. In the case of picking between cans of soup, it is not only whether to resist the urge for soup that is up to the agent. It also seems to be up to the agent whether to pick can A or can B. The agent who picks between can A and can B may not have an explanation of why she picked one over the other. That is, she may not decide to go for one over the other on any basis. But this does not lead us to believe that it was not up to her whether to go for can A or can B. We can say that it is up to her just because she had two options and she could have gone for the other one.

If this choice can count as up to the agent even if the agent makes it on no basis, then it would seem that whether to resist an urge to act could also count as up to the agent even if the agent has no basis for going one way or the other. It may well be that most of our choices not to resist urges are made on some basis, but this does not appear to be necessary for the outcome to count as up to the agent. Furthermore, taking the conclusion of the Appeal to Intentionality Strategy to be a norm will not solve this problem. The reasons for thinking that the agent can have no basis for which way she picks in these kinds of cases are reasons to think that there is no correct choice. (Can A and can C are exactly alike.) Yet it seems to be up to the agent whether to go for can A or can B even though there is (and she takes there to be) no correct choice, and thus there is no norm that tells the agent which way

to go.[41]

The second reason to think that an agent's not resisting an urge to act need not be done on any basis in order for it to count as up to her is a regress argument. To make this argument, we must first note that the fact that the agent has a belief that her action is good in some respect that is picked out by her reason cannot itself reveal the action to be intentional. This belief is just another belief. If Raz is right, then it is a necessary step in the causal chain leading to action along with the belief that the action has some feature F and a desire for F. However, simply adding this belief to the causal chain does not itself make the resulting explanation of action any less mechanistic. It merely makes this chain a causal mechanism that includes a belief that the action is good in some respect. Whatever worries we had about the original causal chain being mechanistic still apply.[42] We would still need an explanation of how the agent gets from her belief that an action would be good to setting

41 There are reasons to doubt that cases of picking contain baseless choices. We might think, for example, that picking can A over can B can count as up to the agent only insofar as the agent takes the two cans to be equally good when going for A instead of B. This would allow us to say that she is still concerned with or responsive to the right basis for choice when going for A instead of B. It just so happens that, in this case, the basis provides no answers. I suspect that this concern for justification is present in many instances of picking, but I hope that the second problem for the Appeal to Intentionality Strategy will undermine the suggestion that it is necessary for the picking to count as up to the agent.

Alternatively, we might think that if the agent goes for one or the other, she must have some urge for the one she goes for. Perhaps this urge itself has no basis. But we might nevertheless think that, once the agent has the urge (given that soup is good and that each can of soup is equally good), that she has an urge for one is itself something to be said in favor of picking that one. If so, the agent can choose not to resist the urge on the grounds that under these sorts of circumstances (there are two equally good options), the presence of an urge for one of the options is something to be said in favor of that option. (See, for example, Ruth Chang's "Can Desires Provide Reasons for Action?" and Jonathon Dancy's "Enticing Reasons" in *Reason and Value: Themes from the Moral Philosophy of Joseph Raz* (Oxford: Oxford University Press, 2004. 56-118) for discussion of this view.) I am suspicious of this approach because I doubt that there is any error in going for the other option, but I will not take up the task of defending that position here.

42 Setiya makes a similar point: "If there can be a wrong sort of connection between the belief that one is doing φ because p and one's doing it, there can be the wrong sort of connection between doing φ and a belief about it's justification. Problems about the right connection between attitude and action, as in cases of causal deviance, cannot be solved by adding more beliefs, whose relationship to what one is doing will be equally problematic." ("Sympathy for the Devil," 91.)

herself to do it. When an agent acts intentionally, her belief that her action is good or justified is not a special sort of belief that, when she has it, automatically causes her to perform the action (or even to intend to). And if it were, we would seem to have no more reason to call such actions intentional than we do in deviant causal chain cases. So, if we continue the line of thought that led us to think that intentional actions are those actions in which it is up to the agent whether to resist the urge to act, we will arrive at the conclusion that if the action is intentional, then it is up to the agent whether to do what she takes to be good, just as it is up to her whether to do what she has an urge to do.

So, the belief that an action is good in some respect does not itself explain the sense in which it is up to the agent whether or not to act on an urge when she acts intentionally. Such a belief cannot be a sufficient condition for intentional action. At best, then, this belief is a necessary condition on an agent's not resisting the urge to act counting as up to her. Raz must claim that if the agent's acting (and thus not resisting the urge to act) is not done on the basis of beliefs about whether and in what respects the action she wants to perform is good, then whether or not to resist the urge it is not really up to her.

However, the sorts of considerations that lead Raz to think that a choice of whether or not to resist the urge must be characteristic of acting intentionally may force us to think that this choice, or some further choice that underlies it *must* be made without a further basis. Recall that the belief that not resisting an urge to act would be good is a belief like any other belief and that adding it to a causal mechanism through which beliefs result in action does not itself make that process less mechanistic (and thus the right sort of process to bring about an intentional action). It is still up to the agent to decide whether to do what she

takes to be good. And if we try to account for the agent's choice of whether to do what she takes to be good by appeal to some further basis (B) for that choice, a regress threatens. If the choice to act on the belief that something is good is made because the agent takes acting on the belief that something is good to be B, then, in order to avoid merely housing a causal mechanism in which the beliefs that the action is good and that choosing to act on the belief that the action is good is B cause the action, one must choose whether to make choices that are B. And so on.

The way to stop the regress seems to be to admit that at some point the agent just chooses to choose on one basis or another and makes this choice on no further basis. And thus we might think that, if we pursue Raz's idea that acting intentionally requires choosing whether or not to resist the urge to act, we will be forced to conclude that the sort of choice that does the job—that distinguishes intentional action from deviant causal chain cases and the workings of a merely causal mechanism—is a choice that is made on no further basis.

But once we admit this, it is not clear why we must think that the first choice—the choice of whether or not to resist the urge to act—*must* be made on one basis or another. Perhaps it often is. But it is not clear why it would have to be. If we accept that other choices can be up to the agent without the agent having some basis or explanation for going one way or the other, then we would seem to have no reason to think that this particular choice must be made on a basis, let alone a specific basis (e.g., the agent's beliefs about what is good) in order to count as up to the agent such that the resulting action is intentional.[43]

---

43 Simon Blackburn points out a similar regress in an argument for the Humean conclusion that facts and principles of reasoning will never be able to, by themselves, move the will and result in action. ("Practical Tortoise Raising." *Practical Tortoise Raising and Other Philosophical Essays*. Oxford: Oxford University Press, 2010. 7-25.)

Taking the conclusion of the Appeal to Intentionality Strategy to express a norm might allow us to avoid the regress, but I think that it does so in a way that does not vindicate the strategy. To see this, first recall that there are a couple of ways that we might understand the relevant norm. We might take the conclusion of the Appeal to Intentionality Strategy to express a norm that the agent takes herself to be subject to. On this understanding, Raz's conclusion would be that, when an agent acts intentionally, it is up to her whether to resist the urge to act, and she chooses not to resist the urge in the recognition that she should do so on the basis of what is valuable (even if she in fact chooses on some other basis. Alternatively we might take the conclusion of the Appeal to Intentionality Strategy merely to express a norm to which the agent is subject (whether or not she take this to be the case whenever it is up to her whether to resist her urges to act). On this understanding, Raz's conclusion would be that, when an agent acts intentionally, it is up to her whether to resist the urge to act and she should make this choice on the basis of whether the object of that urge is valuable.

Either of these understandings seems to provide a way to avoid the regress. The regress arises for the original view because the belief that the object of an urge is good is not enough to show that the agent who acts on that urge is not a mere causal mechanism. If she is to avoid being a mere causal mechanism, there must be something else in the process of practical deliberation that is up to her. The obvious contender is whether to resist urges on the basis of what is good. But if this further thing is up to her, and "up to her" requires a basis for choice, then we need another basis for choice, and the regress gets going.

On the other hand, if we understand the conclusion of the Appeal to Intentionality

Strategy as a norm, we might lose the threat of the agent's being a causal mechanism. Neither the belief that she should decide whether to resist urges on the basis of what is good, nor the fact that she should do so, threaten to make her a causal mechanism. It is still up to her whether to decide as the norm says she should. If something's being up to her requires that there be a norm specifying the correct basis for choice, then there needs to be a norm governing her decision of whether to decide as the norm says she should. If the decision of whether to follow the norm is a separate decision from the decision of whether to resist the urge to act, then the regress gets going. (There will need to be a norm governing this second decision, a decision of whether to follow that second norm, a norm governing this third decision, and so on.) On the other hand, if the decision of whether to resist the urge to act and the decision of whether to follow the norm that governs the first decision are in fact the same decision, then we already have the norm that we need to satisfy the necessary condition on something's being up to the agent—the agent should decide whether to resist urges on the basis of what is good. There is no regress.

However, although there is no regress, it also does not seem that the norm is doing anything to explain the sense in which the action is intentional. That work is done by the fact that the agent has two options—resisting or not resisting the urge—and it is up to her which she goes for. We have gotten no further in understanding what this means by positing a norm that specifies the correct way to make this choice. So, while there may be a correct way to make this choice, I do not think we can get this from the concept of "up to the agent" itself. That is, we cannot discover it by thinking about what makes an action intentional (or at least not in the way that the Appeal to Intentionality Strategy suggests).

V. An objection and some clarifications

In my explanation and criticism of the Appeal to Intentionality Strategy, it may seem that the strategy is committed to an implausible conception of the sense in which intentional actions are "up to" agents. That is, it may seem that the strategy denies that there can be a causal explanation of intentional action in terms of the agent's mental states and commits itself to the claim that, in order to explain intentional action, we will need to appeal to some notion of the agent as having a kind of radical free will. Somewhere in the process of choosing whether to resist an urge she must perform a pure act of will that has no cause, but that has effects in the world (it leads to action). If this is the notion of "up to the agent" that is at work in the Appeal to Intentionality Strategy, then my criticism of the strategy seems to say that part of the reason that we cannot get interesting restrictions on the content of reasons out of the strategy is that, in order to explain intentional action, we need to appeal to some primitive notion of this pure act of free will. And so, my criticism would seem to commit me to understanding intentional action in terms pure acts of will and deny the possibility of a causal explanation of intentional action. As a result, my criticism might seem to be as implausible as the view that it criticizes.

I think there are two ways of understanding "up to the agent" in the Appeal to Intentionality Strategy which avoid this concern. Nonetheless, each of these understandings comes with a corresponding understanding of my criticisms that preserves their force.

The first way of understanding "up to the agent" is more or less as I just described it, but it is taken to be the way "up to the agent" is understood from the agent's perspective when she is deliberating about what to do. This, at least on the face of it, need not conflict

with the claim that there is a causal explanation of intentional action in terms of the agent's mental states. This causal explanation is just not the explanation that an agent has access to from the deliberative perspective (though she may well have access to it from some other perspective, e.g., the perspective she takes when studying psychology).

On this way of understanding "up to the agent", the Appeal to Intentionality Strategy seeks to give an explanation of intentional action from the agent's perspective. My criticism of the strategy then boils down to something like the following: The Appeal to Intentionality Strategy adopts a notion of "up to the agent" such that something's being up to the agent cannot be fully explained by citing some mental states that cause it. But the Appeal to Intentionality Strategy then proposes to explain the sense in which an action is "up to the agent" by citing some mental states that cause the action—it merely adds another belief to this causal story. The Appeal to Intentionality Strategy proposes to understand intentional action as action that is caused by a desire for X, a belief that the action will bring about X, and a belief that X is B, where B is whatever basis of endorsement is constitutive of intentional action. If the proposed explanation works, as an explanation of intentional action from the agent's perspective, it is only because the proposed explanation is not in fact entirely causal. The notion of "up to the agent" must appear somewhere in this story about the agent's mental states.

The second way of understanding "up to the agent" is perhaps more interesting. We might take the Appeal to Intentionality Strategy to accept that the process of assessing the objects of urges, and resisting or not resisting those urges on the basis of that assessment is itself a causal process involving the agents mental states. We might think that the sense in

which an action is up to the agent just is that this particular causal process is possible. This is a way of saying that the sense in which an intentional action is up to the agent is just that the agent can step back and reflect on her motives for acting, and whether or not she performs that action can depend on the results of that reflection.[44] She can decide not to act on her urges and desires (or which ones to act on) in light of her other beliefs and desires. But the assessment of these urges in light of other beliefs and desires is still understood as a causal process.

I think this is the most plausible way of understanding the Appeal to Intentionality Strategy. What is distinctive of intentional action, on this account, is this possibility of assessment and the process of assessment is a causal one involving the agent's mental states. However, I think that what does the work of "up to the agent" in this story is the possibility of reflecting on one's motives and assessing them on the basis of other motives and beliefs. We can identify this possibility without specifying any particular basis for that assessment. It might turn out that human agents do or should assess their motives on some particular basis (e.g., what is valuable). (The claim that we should assess our motives on some particular basis seems more plausible than the claim that we do.) But this connection between intentional action and the correct basis for assessment does not seem to me to be a *conceptual* one. The ability to step back and assess one's motives on some basis does not seem to require that there be a particular or correct basis for that assessment. There may be a correct basis for this assessment, but whatever the correct basis for assessment is, that it is the correct basis for assessment does not seem to me to follow from this sense in which

44 This assessment need not, I think, require that the agent reflect on the fact that she has certain urges. It could merely involve reflecting on the objects of these urges. (The agent need not represent her mental states to herself in order to count as engaged in this assessment.)

intentional actions are "up to" their agents.

Consider for example, two hypothetical agents: The first agent always decides whether or not to resist her urges on some basis, but that basis varies. Sometimes it is what is good, sometimes it is what is beautiful, sometimes it is what is whimsical, sometimes it is what is yellow, and so on. Suppose further that there is no rationale for the varying of the basis. (She has not, for example, decided that alternating between these bases for urge-resisting is good.) The second agent, on the other hand, never has a basis for deciding whether not to resist her urges. An urge presents itself to her. And she always says "sure". Both agents are deeply strange (perhaps the first more so than the second). However, both strike me as conceivable and as acting intentionally. There may be much that is unintelligible about their actions and presumably many of their actions are irrational. But because both agents' actions pass through some process of assessment where the agent can decide whether or not to resist the urge to act in light of her other beliefs and desires, they seem intentional. (In the second case, the agent decides to forgo this assessment.)[45]

This may seem too quick. If an intentional action is "up to the agent" insofar as it is possible for the agent to assess her motives for acting and decide whether or not to do what they move her to do, then there may be an argument from the claim that, when an agent acts intentionally, it is up to her whether to resist the urge to act, to the claim that there is *some* standard that governs this decision, some correct basis on which to make the choice of whether to resist the urge to act. There are a couple of lines of though that I think might lead

45 I do not mean to suggest that all intentional actions involve the explicit assessment or refusal to assess that appears in these examples. The thought is that if even the over-intellectualized versions of these cases fail to support the claim that, when an agent acts intentionally, she chooses not to resist the urge on a particular basis, then the claim is in serious trouble.

us to this conclusion. But I think that both will ultimately direct us to look outside the idea of "up to the agent" to figure out what this standard is (and what reasons for action it would yield).

The first starts with the idea that our capacities tend to have functions. They (or we) can do better or worse at performing those functions. A big part of our capacity for practical deliberation is the ability to assess or adjudicate between our motives. So, it would seem that this capacity must have a function at which it (or the agent) can do better or worse. If we can do better or worse at assessing our motives, then there must be standards of correctness that capture this. So, we might think that there must be standards of correctness for the assessment of or adjudication between motives.

However, this does not yet, I think, give us any reason to think that we can figure out which standards these are just by thinking about what the possibility of the assessment of or adjudication between motives is. We might think that, in order to figure out what the standards are, we would need to look at least at the whole human of which these capacities are a part. If there are better and worse ways for a human life to go, it would make sense that these set the standards for assessing one's motives.

The second line of thought (that might lead us to think that we can get from the claim that, when an agent acts intentionally, it is up to her whether to resist the urge to act to the claim that there is *some* standard that governs this decision) stems from the idea that motives can conflict. In fact, we might think that the easiest way to see that agents can assess their motives (perhaps the most basic case of assessment) is by looking at cases where the agent has a desire to do something that will frustrate other of her desires. The agent must decide

what to do, and can decide not to follow one desire because it will frustrate the other. She can assess the object of one desire on the basis of her other desires. We might think that we cannot conceive of this conflict without taking there to be a correct way to resolve it, or at least better and worse ways of resolving it. There must be some standards that tell the agent how to decide which desire to act on.

Again, some argument along these lines might be enough to get us from the possibility of assessing or adjudicating between one's motives to the idea that there must be some standard that governs this assessment/adjudication. However, I think it unlikely that this line of thought will help us figure out which standard this is. This is not to say that I think there is nothing to be learned about this standard from looking at how agents assess or adjudicate between motives. (In fact, this is more or less what I suggest we do in the next section.) I just think that all that the Appeal to Intentionality Strategy needs to account for the agential control distinctive of intentional action is the possibility of assessment. And this possibility seems consistent with just about any basis or standard for assessment.

If the Appeal to Intentionality Strategy can get us to the idea that there must be some standard (or standards) that governs the choice of whether or not to resist an urge to act, but has no way to determine which standard this is, then it will not yield interesting restrictions on reasons for action. The restrictions on the content of reasons for action in a view like Raz's arise from restrictions on the appropriate basis for the choice of whether to resist the urge to act. E.g., Raz thinks that the appropriate basis for this choice is what is good, so an agent's reasons for action reveal her to have made this choice (and thus reveal her actions to be intentional) by picking out something about the agent's action that is (or that the agent

takes to be) good. If thinking about what makes an action intentional leads us to think that there must be *some* correct basis for making the choice of whether or not to resist an urge to act, but fails to determine which basis this is (and, I think, also fails to narrow down the possibilities), then the Appeal to Intentionality Strategy loses its ability to reveal interesting restrictions on the content of reasons for action.

VI. Moving forward

So, the Appeal to Intentionality Strategy for discovering restrictions on the content of reasons for action by thinking about the role that reasons play in the way agents control their intentional actions will get us only very limited restrictions on the content of reasons for action. If, in order to act intentionally, an agent must choose whether or not to resist the urge to act, then that urge must be directed at some description the agent takes to be true of the action. That is, if whether to resist the urge to act is up to the agent, the agent must have some conception of what the action is that can distinguish it from her other option—doing nothing—such that there is some other option that it was up to her whether to go for. The agent's reason picks out this description under which she sees her action. This suggests a very weak restriction on the content of reasons for action: an agent's reason for action must identify some description that she takes to be true of her action. But this is unsurprising. In fact, it is more or less the idea we started with in Section I.

We might conclude from the difficulties with the Appeal to Intentionality Strategy that the only restriction on the content of reasons for action is the very minimal one just discussed—that a reason for an action must pick out some description that the agent believes

applicable to her action. I think that drawing this conclusion now would be much too quick. For example, nothing I have said here shows that it is not true that reasons for action pick out respects in which actions are good (or taken to be good by their agents.) What I have tried to show is that, *if* this is true, it is not because acting for the sake of the good is part of, or a necessary condition on, exhibiting the kind of control over one's actions that an agent exhibits when, and insofar as, she acts intentionally. This notion of control seems to be useless for discovering restrictions on the content of reasons. Nonetheless, we might think that reasons for action (and intelligible positive answers to why questions) do more than just reveal actions to be intentional.

Furthermore, it seems that intelligibility as an answer to the why question that asks for reasons for action (and thus intelligibility as a reason for action) is constrained by more than the minimal constraints on intelligibility that the Appeal to Intentionality Strategy supports. Suppose I want a saucer of mud.[46] I walk out into the yard and scoop some mud into a saucer. I respond to the question "why did you get yourself a saucer of mud?" with "in order to have a saucer of mud". Further questioning reveals no further story about why I want to have a saucer of mud beyond perhaps "it's mud!" If we take the constraints on the content of reasons for action revealed by the Appeal to Intentionality Strategy to be all of the constraints that there are, then it would seem that I act for a reason in this case. My reason is something like "it will get me a saucer of mud". My reason is (or alludes to) some description of my action that I take to be true of it, and it moves me to act. Yet, I think it is at least highly dubious whether "it will get me a saucer of mud" is an intelligible reason for

46 The example is famously Anscombe's (*Intention*, §37), but I make no claim to be using it in the same way as she does (or, for that matter, to be using it differently).

acting, in the absence of some further story about why I want a saucer of mud. As a result, it would seem that there are additional constraints on intelligibility as an answer to the why question that asks for reasons (and thus additional constraints on the content of reasons for action). The Appeal to Intentionality Strategy is not promising as a way to identify these constraints, so we will have to look elsewhere to find them.

The Appeal to Intentionality Strategy tried to identify constraints on intelligible answers to the why question by thinking about how to distinguish intentional actions from non-intentional actions. I propose that we look to a different distinction to guide us in looking for constraints on intelligible answers to why questions (and thus constraints on the content of reasons). We might think that there are, at least in principle, two varieties of intentional actions: those that an agent performs because she takes there to be something about her action that *calls for* performing them and those that an agent performs just because she feels like it (without thinking that there is anything to be said for them or at least without being motivated by anything one takes to be said for them). The former are done for reasons. And in these cases, an agent's reason for action picks out what it is about her action that she takes to call for so acting.

Raz came to the conclusion that agents choose not to resist urges to act because they take the objects of those urges to be good in some respect by noticing that these choices are often characterized by a concern for justification. When an agent chooses not to resist the urge to act, it is often because she takes the object of that urge (something about the action) to be something like a piece of justification for the action. It is not necessarily a complete or sufficient justification, but the agent takes it to be something to be said for performing the

action. It is something that she takes to have weight that opposes things to be said against performing the action (or that she would take to have weight against things to be said against performing the action were she to become aware of them). In addition, this weight is normative. E.g., the agent takes it to be appropriate that something to be said for performing the action opposes those things to be said against it. This opposition is not merely a conflict in motivation. It is an appropriate conflict in motivation. It would be a kind of error to fail to recognize it and to be moved entirely for or against the action. When the agent sees the object of an urge in this way and then acts and pursues the object of that urge, she is motivated in part because she thinks the object of the urge "calls for" pursuit.

We might think that it is in those cases where an agent does not resist urges to act and exhibits this concern for justification that she acts for reasons. In those cases where she acts (and it is up to her whether to resist the urge), but where she has no such basis for not resisting the urge, she acts intentionally, but for no reason.[47] If this is plausible, it opens up a new strategy for looking for restrictions on the content of reasons. We might investigate this "concern for justification" that agents seem to exhibit when deciding what to do. We might ask whether there are constraints on what sorts of features of actions can intelligibly be regarded as something to be said for actions, as pieces of justification for actions, as "calling for" actions. These constraints will be constraints on what counts as a reason for

47 Given my criticisms of the Appeal to Intentionality Strategy, I think we do not have grounds for thinking that these are the only possibilities for intentional action. I suggested in Sections IV and V of this chapter that the idea that, when an agent acts intentionally, it is up to her whether to resist the urge to act is consistent with the agent having a wide variety of bases for resisting or not resisting the urge. On the other hand, I think that the sense in which agents take their reasons to "call for" for actions requires more than that they have *some* basis for this choice. Not every basis will yield "reasons" that can intelligibly be regarded as pieces of justification for the action, or at least not as the sort of justification that we seem to take reason to give. So, there can, at least in principle, be a category of intentional actions where the agent has a basis for not resisting the urge to act, but cannot intelligibly take the object of that urge to "call for" so acting.

action.

I do not think that this notion of "called for" can be fully specified independently of this investigation into when agents can intelligibly regard considerations as calling for actions. As a result, the notion of justification associated with "called for" is quite loose at this point in my project. Part of the project of explaining what a reason for action is, is making precise what this notion of justification is.

At least at this point, I would like to leave open the possibility there are multiple senses or kinds of justification for actions. This is to allow that reasons can intelligibly be taken to call for actions (and so at least partially justify them) even when they provide no justification (in some other stronger sense of justification) for those actions. I have in mind here cases like those Setiya presents as counterexamples to the claim that, when agents act for reasons (or at least have intelligible positive answers to the why question), they take those reasons to be normative—to justify those actions rather than merely explain them. Setiya describes a case in which an agent flees a burning building, abandoning his family inside. He claims that this agent acts for a reason—because the house is on fire. However, under these circumstances, while this consideration gives ground for, and makes sense of, the agent's action, it does not justify the action at all.[48] When one's family is in a burning house, the fact that the house is on fire just is not anything to be said for fleeing. Or at least Setiya thinks so, and I think the claim is at least somewhat plausible. Under certain (often dire) circumstances, an agent's reasons change. Considerations that would normally call for actions should not be considered at all.

There are many ways in which we might explain an action, and thus make sense of

48 Setiya, "Sympathy for the Devil," 90.

an action (e.g., a causal explanation that appeal to the agents mental states, a causal explanation that appeal to facts about the way the agent was brought up, a causal explanation that appeals to facts about subatomic particles). But a reasons-explanation, even on Setiya's view, even in the fire case, is not just *any* explanation of the action. A reason for an action gives a ground for the action or makes sense of it from the agent's perspective by identifying some point to be achieved by the action. It seems likely that not just anything can be the ground for an action. We might question whether the fact that the house is on fire does anything to justify fleeing the house, when doing so would be abandoning one's family, if we take justification to be an explanation of what the agent ought (morally, all things considered) to do. Nevertheless, that the house is on fire does seem to give ground for or make sense of the action in a way that not just any consideration could. (By contrast, consider, for example, an agent who flees the house "because it is Tuesday".) Furthemore, there seems to be something normative about considerations that give this "ground." It may be true that the fact that the house is on fire is not something to be said in favor of fleeing with respect to the question of what the agent ought, all things considered, to do. Nevertheless, there seems to be some sense of appropriateness such that the fact that the house is on fire makes fleeing appropriate, in a way that the fact that it is Tuesday does not. Or, at least there seems to be some sense of appropriateness such that the fact that the house is on fire is the right sort of consideration to appear in an explanation of why fleeing is appropriate, whereas the fact that it is Tuesday is not.

So, I think it fair to say that, in attempting to explain what makes a reason call for an action, we are attempting to give an account of what makes a consideration give ground for

or make sense of action. There is some notion of appropriateness (and thus, I claim, some notion of justification) at work here, which the idea of "calling for" is supposed to capture. It is, at least initially, an open question whether this notion of appropriateness/justification is the same as that which appears in explanations of what we, all things considered, morally ought to do.

In my discussion of Raz's view, in this chapter, I treated reasons-justification and value as if they come together. I think that they do, but this is a substantive position. In the rest of this dissertation, I will look at accounts of reasons for action that take themselves to capture this distinctive feature of reasons—that a reason for action "calls for" that action, that it is something normative to be said in favor of that action—without appealing to value. That is, I will look at views that attempt to account for the sense in which reasons call for actions without identifying anything that is to be promoted, preserved, or protected (independent of any particular agent's attitudes towards it). I will try to show that these accounts of reasons for action make it difficult to see how an agent could intelligibly regard her reasons as calling for her actions.

Chapter Two

Introduction

In Chapters 2 and 3, I will consider constructivist accounts of reasons for action and the explanation that they provide of what makes a consideration count in favor of performing an action. On a constructivist view, an agent creates (or constructs) a reason for action—i.e., she makes it the case that a particular consideration counts in favor of performing a particular action—when she chooses *in the right way* to perform an action because of some feature of that action that appeals to her.

In this chapter, I attempt to undermine constructivist accounts of reasons for action that take on four characteristic commitments by showing that three of these commitments together prevent constructivist views from adequately accounting for a key feature of reasons for action. This key feature is that reasons are shareable (in a particular sense), i.e., that agents can and should act for the sake of the same features of action. I argue that constructivists cannot account for (this sense of) the shareability of reasons by illustrating how this problem arises for Christine Korsgaard's constructivist account of practical reasons. Because the problem for her view stems from its constructivist commitments, it should also arise for other constructivist views.

Section I of this chapter describes two senses in which Korsgaard thinks we are required to share reasons. She says little to defend one of these requirements. Because this chapter aims to show that her view cannot accommodate this requirement, Section II contains three somewhat speculative attempts to support the requirement within her view. (I

take these arguments to provide some support for the requirement independently of her view as well.) Section III describes the four commitments that are characteristic of a kind of constructivist account of reasons for action—one that denies R. Section IV describes Korsgaard's view and identifies its constructivist commitments. Section V raises a problem for her view's ability to allow for one of the requirements to share reasons that arises because of its acceptance of three of these commitments. Section VI proposes a substantial modification to her view that gives up the first and third commitments and that can solve the problem. Finally, Section VII argues that this modification is the most plausible of the available ways to solve the problem—i.e., that we should solve the problem by giving up the first and third commitments, and not by giving up some other combination of the commitments.

I. The shareability of reasons

Constructivists think that, when an agent (permissibly)[49] acts for the sake of a particular feature of her action, she acts for a reason. Her reason, then, cites or picks out this feature. For example, I might go to a coffee shop for the sake of buying coffee (and my reason for going to the coffee shop would be that I can buy coffee there). Or, I might climb Kilimanjaro for the sake of the beautiful view from the top (and my reason for climbing would be that there is a beautiful view from the top).[50] Following Korsgaard, I will refer to the particular feature of an action for the sake of which an agent (permissibly) acts as her

49 An agent may also count as acting for a reason in at least some cases when she acts *impermissibly* for the sake of a particular feature of that action. I take no position in this chapter on whether, when, and why this happens.

"purpose" in performing that action.[51]

Reasons for action, on Korsgaard's view, are "shareable". For two agents to share a reason is for them to (permissibly) act for the sake of the same feature of actions. For example, agents A and B would share a reason if each performs an action because he or she thinks it will promote B's successfully climbing Kilimanjaro. The actions they perform for that reason might be of the same kind or of different kinds. For example, A and B might go on a hike together to help B train, or B might hike and A might give B a map of Kilimanjaro. In either case, A and B each perform the action that they perform because it will promote B's successfully climbing Kilimanjaro. The claim that reasons are "shareable," for Korsgaard, means more than that it is possible for agents to share reasons. For Korsgaard, the claim that reasons are shareable means that reasons are to be shared—that agents are *required* to share reasons. The basic idea here is that a reason is the sort of thing that makes claims on all rational agents (under the right circumstances). For example, that an action will save a life is a reason for every agent to perform that action, other things equal (under those circumstances where it is true that that action will save a life).

Korsgaard thinks that we are required to share reasons in two senses, a weaker and a stronger sense. These senses differ with respect to the explanation of *why* agents make the same feature of actions to be their reason for acting. For an agent A to share a reason in the

---

50 In the examples I discuss in this chapter, I will only loosely characterize the features of actions that appeal to agents and that are cited by their reasons for acting. It will not matter, for the arguments I make in this chapter, which in a variety of related features of action appears in the correct formulation of a reason. E.g., it will not matter whether the correct formulation of an agent's reason for acting cites the fact that her action will promote one of her ends, or whether it cites something about the action that makes it the case that the action will promote one of her ends.

51 Korsgaard, C.M. "Acting for a Reason." *The Constitution of Agency: Essays on Practical Reason and Moral Psychology*. Oxford: Oxford University Press, 2008. 218. She also uses "end" to refer to the same thing. "Purpose" strikes me as less likely to seem theory-laden.

weaker sense (henceforth, to "w-share" a reason) with another agent B is for A to make the fact that an action would promote B's pursuit of some purpose P her reason for performing some actions that have this feature *just because A recognizes that P is a permissible purpose of a fellow rational agent*.[52] That is, A might act for the sake of promoting B's pursuit of P just because A recognizes that P is a permissible purpose of a fellow rational agent. (Similarly, A might make the fact that an action would undermine B's pursuit of her purpose P her reason to refrain from performing actions that have this feature just because A recognizes that P is a purpose set in the right way by a fellow rational agent.) Korsgaard thinks that the permissible purposes of other rational agents give each of us these reasons to promote, or at least to refrain from interfering with, those agents' pursuit of those (permissible) purposes just because the agents' choices to pursue those purposes are proper expressions of rational agency.[53] Thus, she thinks we are required to w-share reasons.

On the other hand, when A shares a reason in the stronger sense (henceforth, to "s-share" a reason) with B, A acts for the sake of the same feature of actions (F) as does B and A finds F interesting and appealing for its own sake and not *merely* because another rational agent has permissibly made it her purpose.[54] This means that A finds F interesting and appealing on her own because of something about F itself and not merely because of the interest that others take in it. (For, example, if F is that an action will involve a beautiful experience, A can s-share reasons that cite F if she finds beautiful experiences interesting or

52 Korsgaard, C.M. "The reasons we can share: An attack on the distinction between agent-relative and agent-neutral values." *Creating the Kingdom of Ends*. Cambridge: Cambridge University Press, 1996. 308. (Endnote 31).

53 See, for example, Korsgaard, C.M. "Aristotle and Kant on the source of value," *Creating the Kingdom of Ends*, 241. I will not discuss Korsgaard's argument for this requirement here. I do think that some requirement of this kind is plausible whether or not Korsgaard's argument for the requirement is persuasive.

54 Korsgaard, "The reasons we can share," 308. (Endnote 31).

appealing because of what they are and not because someone else has chosen to pursue them.) As a result, when A pursues F, whether by promoting B's pursuit of F or by pursuing F on her own, she does so (at least in part) because of what she finds interesting about F and not just because B permissibly acts for the sake of F.

Finding F appealing because of something about F itself includes finding F appealing because of other more general features of F (that are not themselves appealing only because others find them to be so). For example, If F is a specific purpose, e.g., B's climbing Kilimanjaro, A might find F interesting because she finds something about F interesting—e.g., the beautiful vistas and physical challenge that it will involve. If she does so, she will count as s-sharing reasons to promote B's climbing Kilimanjaro even if the particular project of climbing Kilimanjaro (as opposed to another way of experiencing beautiful vistas and physical challenge) leaves her cold.

II. Why think we are required to share reasons?

Korsgaard claims that s-sharing reasons (as well as w-sharing reasons) is required as part of valuing the rational nature of others.[55] That is, she thinks that we should act for the sake of the same features of actions as do others, and we should find those features interesting or appealing because of what they are and not merely because other agents permissibly make them their purposes. A requirement to s-share reasons has some intuitive plausibility. We tend to think of many reasons as to be shared in this sense. For example, we tend to think that someone who never takes the fact that an action will lead to a beautiful experience (be it seeing a beautiful vista or hearing a symphony) to be a reason to perform

55 Ibid., 290.

that action (or who does so, but only because other agents pursue beautiful experiences) has made an error by not being moved by something that is worth pursuing (because of what it is not merely because others pursue it). This is a reason that we all have. And we have it because of something about what beautiful experiences are. To miss this, and to think that beautiful experiences are beautiful only because some people (perhaps ourselves) permissibly pursue them, we might think, is to not "get" what is good about a beautiful experience.

However, as will become clearer in Section III, this intuitive idea seems to be in tension with the constructivist idea that an agent creates a reason when she chooses in the right way to act for the sake of a particular feature of her action. Korsgaard does not say much about why she nevertheless takes on a requirement to s-share reasons. In this section, I will sketch three ways that she might support this claim. I hope these arguments provide the reader with some reasons to think that we should not abandon the requirement to s-share reasons in the face of the problem I eventually argue that it poses. It is also, I think, of interest to see why Korsgaard would take this to be a necessary feature of reasons for action given the difficulties that it poses for her account. I will revisit these arguments in Chapter 3 to further support a requirement to s-share reasons.

The first way that Korsgaard suggests of supporting a requirement to s-share reasons comes from considering the attitudes we take towards those with whom we w-share, but do not s-share, reasons. She presents the requirement to s-share reasons as a response to the idea that to merely w-share another agent's reasons involves taking a particular objectionable attitude towards that agent. Korsgaard acknowledges in a footnote that taking ourselves to

have reasons to promote the purposes of others, merely because they are the purposes of others, can have a patronizing ring: "You like it, so I guess we'll have to count it as good".[56] This, she thinks, expresses an attitude that we should not have towards each other. To avoid this attitude, we should strive to find the features of actions that others make their purposes interesting for their own sakes (and thus to "count them as good" for the same reasons these other agents do).

Why would merely w-sharing reasons involve objectionable patronizing attitudes? To begin with, it is not obvious that a "you like it, so we'll count it as good" thought is itself problematic. (The "I guess" seems to be doing a lot of rhetorical work in the original formulation.) The fact that an agent takes an interest in something does explain part of why it gives us reasons. An agent's purpose of climbing Kilimanjaro, for example, does not give us reasons merely because it has features that interest us for their own sakes (e.g., beautiful vistas and perseverance). It also gives us reasons because the agent made it her purpose. She could have chosen to pursue some other purpose with these features, but given that she choose this one, it gives us reasons that those other possible purposes do not. So, Korsgaard's thought must not be that the "you like it, so we'll count it as good" thought is objectionable in itself, but rather that it should not express the *only* reason for my counting another agent's purposes as good. It should not be my only basis for taking them to be reason-giving.

What would be wrong if my only reason for taking another's purpose to be reason-giving were my recognition that it is the permissible purpose of another rational agent? I do not think that the answer to this question is clear. It seems true that, when we find ourselves

56 Ibid., 307. (Endnote 31).

in this position with respect to another agent, it tends to involve thinking less of their activity, and thus, perhaps, by extension, thinking less of them. Consider the person who spends significant time, energy, and resources collecting a particular sort of trivial-seeming object (TO). Suppose we have talked to her extensively about her hobby, and we have determined that there really is nothing about it that interests us. (E.g., TO's are not beautiful, or historically significant, or worth a lot of money.) I think our resulting attitude towards this person is likely to make both us and her uncomfortable. We might worry that our attitude towards her project might be patronizing. We might catch ourselves thinking that her collecting "isn't really worthwhile", though perhaps it is "harmless". This could seem like a way of not taking her seriously.

To regard the purpose of another as reason-giving merely because it is her purpose inevitably involves some disavowal of that purpose. It is to withhold a further endorsement of the purpose—there is nothing about the purpose itself that seems to call for action. Furthermore, this withheld endorsement is one that we do tend to give our own purposes. If I am climbing Kilimanjaro, this purpose does not interest me merely in virtue of being the purpose of rational agent. It interests me because of something about it—beautiful vistas, perseverance, etc. In light of this difference between the way I regard my own purposes and the way I regard the purposes of others, we might take this withholding of this further endorsement to say something like "that may be okay for you to do, but I have more important concerns". My more important concerns are purposes that interest me because of something about what they are, not because some rational agent made them her purposes. They are more important because I can take there to be two explanations of why they are

reason-giving—one that appeals to the fact that they are the permissible purposes of a rational agent and one that appeals to whatever it is about these purposes themselves that attracts me.

That we feel this way, when our reasons are not s-shared, suggests that we want our reasons to be s-shared. Moving from this to the conclusion that we fail to properly respect someone if we do not s-share her reasons strikes me as difficult. It seems pretty clear that we want other people to share our interests. It is less clear that other people wrong us if they fail to do so. It is not clear that our attitudes towards the TO collector involve us taking our own purposes to be more important than hers in any objectionable way. We just find our purposes more interesting. As will become clearer later in this chapter, Korsgaard thinks that, when purposes are valuable, this is because rational agents have made them their reasons for acting (in the right way). So, if we think that the purposes of another rational agent are of less value than our own purposes, we are failing to acknowledge her equal status as a value creator. However, it might be that taking the purposes of other agents to be reason-giving just in virtue of being the permissible purposes of other agents is all that is required to recognize their equal status as value creators. If so, then as long as we respect the TO collector's right to pursue her own purposes (e.g., refrain from interfering with her pursuit of those purposes without good reason), our failure to find her purposes interesting will not be objectionable.

Nevertheless, while I think that it is unclear why exactly we would have to think that "you like it so we'll count is as good" is an objectionably patronizing attitude to have towards another rational agent, I do think that there may be something behind Korsgaard's

worry that does support a requirement to s-share reasons. In the TO collector case, although it is not clear that our attitudes toward this agent are patronizing, it may still be natural to take the "you like it, so we'll count it as good" thought to express a criticism. It suggests that the agent has made a mistake about what is worth doing, while allowing that her pursuit of this purpose nonetheless has some value merely in virtue of being her purpose (at least as long as that the purpose is not grossly objectionable). (This is not yet to say that the TO collector does something wrong. It may be permissible to occasionally do worthless things. Nor is it to say that we can permissibly interfere with the TO's activity. Proper respect for a rational agent may require that he be allowed to pursue worthless purposes, if he so chooses, other things equal.)

If the TO collector objects to our having this attitude, this need not be because he takes it to be patronizing. (Believing that another agent has made a mistake cannot automatically be *objectionably* patronizing. After all, this belief might be true. And it need not, I think, be patronizing even if it is false.) It seems more natural to think that in this kind of case, the TO collector objects because he thinks there is in fact more to be said about why TO collecting is good. It is not merely because he has made it his purpose. It is because of the appealing features of TOs that we have failed to recognize. In short, cases where the "you like it, so we'll count it as good" thought is likely to arise seem like cases that it is natural to interpret as disagreements about what is worth doing.

It looks like this disagreement presupposes that reasons should be s-shared. We think that the TO collector's collecting is not really worth doing because we do not s-share her reasons—because we do not take the features of TOs that move her to collect them to be

anything to be said for collecting TOs apart from the fact that they (permissibly) move her to do so. If so, then we might think that reflection on situations where the "you like it, so we'll count it as good" thought is likely to arise reveals that we take there to be a requirement to s-share reasons in our everyday use of reasons. This might seem to leave us with two options: include such a requirement in our account of reasons or accept that we are systematically in error about what reasons are in our everyday use of reasons. The former option seems more reasonable.

Of course, this only suggests that we take there to be a requirement to s-share reasons if we cannot explain this kind of disagreement with the TO collector without appealing to s-sharing. It is, admittedly, somewhat difficult to find cases where agents do not s-share reasons, where this seems to make some criticism or disagreement appropriate, and where it is clear that the problem is not disagreement about whether the purpose in question, or the lengths taken to pursue it are in fact permissible. We might think, for example, that if we have a problem with the TO collector's activity, it is not because the pursuit itself is silly because there is nothing about TO collecting that is interesting for its own sake. Rather, it is because, given all of the duties the TO collector has, she should be doing something else (e.g., promoting the common good, spending time with her family, etc.) I think it consistent with the letter and spirit of Korsgaard's view (apart from the requirement to s-share reasons) to think that there are (or in a better, but not yet perfect, world would be) permissible pursuits that fulfill no duties. An agent's moral duties, in principle, leave her with a large space in which to pursue her own projects—purposes that interest her independently of her duties (more on this in Section IV). This leaves room for agents to permissibly pursue

projects that other agents see nothing interesting about.

In the same footnote, Korsgaard suggests a second line of thought supporting a requirement that we s-share reasons. She worries that if you view my purposes as giving you reasons only because they are the (permissible) purposes of a rational being, then in accepting help from you (or permitting you to curtail your pursuit of your own purposes to accommodate my pursuit of mine), I might be using you as a mere means and thus wronging you. This would be a problem because Korsgaard thinks (plausibly) that agents are required to (at least sometimes) help each other pursue their purposes. But if you are required to help me, then it must be permissible for me to accept the help that you are obligated to offer, at least under some circumstances.[57]

To see what Korsgaard might have in mind here, I think it is helpful to think not about what might be wrong with my accepting help in pursuing a particular purpose, but rather about whether I may pursue that purpose in the first place knowing that doing so will generate this obligation for others. If I am considering whether to pursue a purpose that I know will be of no independent interest to another agent, I know that, in acting for the sake of that purpose, I will be making a demand on that other agent to take my purpose to be reason-giving. (It will give her reasons to help, or at least to refrain from hindering, my pursuit of that purpose). And I cannot help but make the demand because it is created whenever a rational agent chooses to (permissibly) pursue a purpose.

At first glance, this thought seems less than promising. As a human being, there is a sense in which I make demands on the other human beings all of the time that do not seem problematic. That is, the other human beings are morally required to treat me (and refrain

57 Ibid.

from treating me) in certain ways. If one of these requirements is to help other people pursue their purposes, then we need not think that the demand that I make on others by having purposes is any more objectionable than any of these other demands (e.g., the demand not to be murdered or the demand not to be lied to.)

Nevertheless, we do tend to think that a genuine demand should be justifiable to the agent who is subject to it. There are some differences between familiar moral requirements and the demand to help an agent pursue a purpose, when the relevant reasons are w-shared but not s-shared, that might suggest that this latter demand is not justifiable to the agent who is subject to it. I seem to have a kind of control over the demand to help me pursue my purposes that I do not have over the demand, e.g., not to murder me. Although it is not up to me whether agents face a general requirement to help me pursue my purposes, it is up to me which purposes to pursue, and so it is up to me which purposes other agents must, as a result, take to be reason-giving. If I climb Kilimanjaro, now other agents must (sometimes, other things equal) take the fact that an action would promote my climbing Kilimanjaro to be a reason to perform that action. If I decide to collect trivial-seeming objects, then other agents must (sometimes, other things equal) take the fact that an action would promote my collecting trivial-seeming objects to be a reason to perform that action. On the other hand, I have no say in the content of other agents' reasons not to murder me. The fact that an action would be a murdering of me is a reason not to do it regardless of which purposes I choose to pursue (at least supposing they are not tremendously evil).

If other agents face demands to help me pursue purposes whether or not they find those purposes independently interesting, this means that, when my purpose is of no

independent interest to this other agent, there is little I can say to her beyond "it is permissible and I want it" to justify that demand. But this justification seems inadequate for making the demand. In acting for the sake of such a purpose, it looks like I will be, in a way, forcing other people to help me to get what I want, just because I want it.

In placing this demand, I am making other agents an extension of my own agency. (Similarly, a tool can be an extension of my own agency.) These other agents must now work (to some extent) to pursue my purposes. And so, it begins to look like I might be using them as a mere means. I may not intend to do this. I need not decide to pursue purposes with the intention of making other people into tools to promote my own ends. (I may not think about the demands that my purposes place on others at all.) Nevertheless, I can (and probably should) know that my purposes make demands on other rational agents and that which purposes I decide to pursue, and thus which purposes I direct others' agency to pursue, is up to me. Given this, we might think that these other agents should have some say in which purposes their agency will now be directed towards. I should not make this decision without some consent or cooperation on their part.[58]

Cooperation and consent seem too demanding. An agent can hardly be expected to actually consult with all the other agents when deciding which purposes to pursue. But we might nonetheless think that I should worry about the sorts of considerations just discussed when I am considering pursuing purposes that will not interest others. Perhaps I should

58 Admittedly, this may all seem less pressing given that the demand to help others pursue their purposes may not be very demanding. On Korsgaard's view, our reasons to help others pursue their purposes may be easily outweighed. Perhaps if it costs us nothing, we have reasons to help. And perhaps the level of cost to our own purposes that justifies interfering in another's pursuit of his purposes is quite low. However, whether or not the demand to help others pursue their purposes treats the agents subject to it as mere means (or something like this) does not seem to depend on how taxing the demand is. And so, I think, the fact that the demands are weak will not get us out of the problem.

consider what the interests of others happen to be and pursue purposes that will interest those other agents independently of their being the permissible purposes of a rational being. That way, when these other agents then face a demand to take my purposes to be reason-giving, it is at least not a demand to promote a purpose that they regard as having no claim on them independent of my wanting it and deciding to pursue it. I could cite those common interests in justifying the demand. If these other agents then pursue this purpose for me, they would also be pursuing a purpose of their own when pursuing my purposes. This would make them, I think, seem less like mere extensions of my agency when pursuing my purposes.

But this seems like the wrong solution. It constrains the interests an agent can permissibly pursue to those that the other agents around her also happen to have. This is an unsettling thought as it has the potential to require massive conformity and squelch diversity of interests.[59] Furthermore, it seems to merely relocate the problem. An agent would now face demands to refrain from pursuing interests that do not also happen to interest those around her. However, the justification that these others can offer this agent for this demand will appeal to something like "I (permissibly) do not want that." If "it is permissible and I want it" was an insufficient justification for a demand, then this new justification should also seem insufficient.

On the other hand, if there is a requirement to s-share reasons, the problem seems not to arise. Whenever I decide to pursue a purpose, I create a demand that others take my purpose as reason-giving. Because we are also required to s-share reasons, we are also

59 This at least goes against the spirit of Korsgaard's account. For example, in "The reasons we can share" she speaks with admiration of "our distinctively human capacity to take an interest, and to find something interesting, in whatever we find around us" (290).

required to find the same features of action interesting for their own sakes (and not merely because they are the permissible purposes of rational agents). Thus I can always cite these common interests in justifying the demand. That is, I can justify the demand by citing interesting features of the purpose itself and not merely by citing that fact that I want it and have permissibly chosen to pursue it.[60]

The third line of thought supporting a requirement to s-share reasons is suggested by Korsgaard's argument against private reasons in *Sources of Normativity* (*SN*).[61] She claims that, in ordinary exchanges in which we appeal to reasons and the purposes they cite, we take the reasons and purposes of others, to make claims on us. When we deliberate together about what to do (e.g., scheduling a meeting) and when we demand that others treat us in certain ways (e.g., not harm us), we take our own reasons and purposes to make claims on them, and we also, under normal circumstances, take their reasons and purposes to make claims on us. I might, for example, take the fact that you have a doctor's appointment at a particular time to be a reason not to schedule a meeting with you at that time. (And this need not be because I think you won't show up. It can be because I take your doctor's appointment to be a reason that bears on when we should schedule a meeting, and thus bears on what I should do, just because it determines what meeting times are consistent with your pursuit of your purposes.) That this is how we use the concept of a reason for action suggests that part of what it is to be a reason is to be the sort of thing that makes claims on

60 We might worry that this merely pushes the problem back a step. Now choosing to pursue a purpose creates a demand that other agents find that purpose interesting for its own sake. And the justification that the agent can produce for *this* demand is little more than "it is permissible and I want it". While this is not the problem I raise for Korsgaard's view in this chapter, I do think the second way of supporting the requirement to s-share reasons may eventually undermine her way of accounting for it. (See Chapter 3, Section VI for further discussion of this second argument in support of a requirement to s-share reasons.)

61 Korsgaard, C.M. "Lecture 4." *The Sources of Normativity.* Cambridge: Cambridge University Press, 1996. 132-145.

multiple agents and that can thus be shared.

The argument in *SN* does not make a distinction between the weaker and stronger sense of shared reasons. As a result, it does not clearly establish whether the claims that the reasons of others make on us are claims to take others' pursuit of their own purposes as reason-giving merely in virtue of the fact that they are the permissible purposes of rational agents or whether they are claims that we make the relevant features of action our own purposes that we pursue and that interest us for their own sake. However, I think that the argument can be extended to support the latter claim. When we ask for a person's reasons for an action, we seem to be looking for, and to only be fully satisfied with, an answer that reveals the action as having some feature that we ourselves have made a purpose, or at least find interesting, because of something about that feature itself. In normal cases, we either receive this answer and find ourselves satisfied, or if we do not, we tend to assume that we just do not have all the facts in order to avoid finding another's action un-intelligible. (For example, "Really?" seems to be a natural response to the TO collector who, when extolling the virtues of TOs, cites nothing about them that interests us independently interesting him.) That we do this suggests that the sort of intelligibility that we are striving for, when asking agents for their reasons in trying to make sense of their actions, is intelligibility in light of common interests that we pursue for their own sake.

The sorts of Wittgensteinian considerations that Korsgaard appeals to in the argument against private reasons in *The Sources of Normativity*—that our use of the concept "reason" gives us some information about what reasons are—suggest a route from here to the claim that we *should* s-share reasons. The aim of our reasons-discourse, of our asking

for agent's reasons and giving our own, appears to be intelligibility of our actions in terms of some purpose that we all pursue for its own sake. If our everyday reasons-discourse is a good guide to what a reason for action is—i.e., if a reason for action is the sort of thing that can play the role in explaining and justifying actions that we take reasons for action to play in our reasons-discourse—then it looks like a reason is the sort of thing that is to be s-shared. If this is true, then, if we use reasons, we should s-share them simply because this is part of what reasons are.

In short, it looks like a basic strategy discussed in Chapter 1 supports a requirement to share reasons. We take reasons to be answers to a particular why question. There are constraints on what can intelligibly be answers to this question. If we look at our use of reasons in response to the why question, we will discover that intelligible answers to why questions cite purposes that we find interesting because of something about what they are and not merely because they are the permissible purposes of rational agents. Thus, intelligible answers to why questions cite reasons that are s-shared. And thus, a reason seems to be the sort of thing that is s-shared.

III. Four commitments of constructivist accounts of reasons for action

This chapter concerns constructivist accounts of practical reasons that have four key features.[62] The first (**C1**) has two parts.

62 I will argue in the next section that Korsgaard takes on C1-C4. Another prominent constructivist who endorses these commitments is Sharon Street. (See for example "Constructivism about Reasons." *Oxford Studies in Metaethics* 3 (2008): 207-245.) I will discuss her view and what I think it can tell us about Korsgaard's view and constructivist accounts of reasons for action in general in Chapter 3.

C1a: A reason for an action picks out a feature of that action that the agent desires.

C1b: Desires are non-normative and are thus not subject to rational revision.

C1a states that the considerations that count in favor of actions cite features of actions (or their effects) that agents find appealing—that they desire. Desires, on this picture, seem to be conceived of as something like bare urges.[63] According to **C1b**, it is not correct or incorrect to desire a particular feature of action. Desires determine inputs for the process of practical reasoning. The process of practical reasoning is governed by standards of correctness that determine when it is correct or incorrect to act for the sake of a desired feature of action. But desires themselves are not subject to revision according to these (or any other) standards. In short, agents cannot reason their way into or out of having particular desires.[64]

However, agents do not have reasons for actions merely in virtue of desiring particular features of those actions. The second key commitment of constructivist accounts of practical reasons (**C2**) explains what more is needed.

C2: An agent makes a feature of an action count in favor of performing that action when she chooses, in proper deference to norms that are constitutive of the activity of practical reasoning, to act for the sake of that feature.

The constitutive norms of practical reasoning are those norms to which an agent must exhibit a certain responsiveness in order to count as engaged in practical reasoning at all.[65] To see roughly what this amounts to, consider an example from theoretical reasoning: An

63 I am skeptical that this stripped-down notion of desire accurately captures a useful or familiar notion of desire, I will not address that issue here.

64 This is an oversimplification. Reasoning can transfer the motivating force of a desire even though it cannot create it. For example, if I desire ice cream and know that shop X has ice cream, I might desire to go to shop X as a result. And if I were to learn that shop X were closed, my desire to go there could disappear as a result of the recognition that going there is no longer a way to satisfy my desire for ice cream. However the desire for ice cream itself (or some further desire that yields it) does not originate as a result of rational deliberation.

agent who is manipulating and changing her beliefs, but who is completely indifferent to rules of inference like *modus ponens*, does not seem to count as engaging in theoretical *reasoning*. So, we might think that such inference rules are constitutive norms of theoretical reasoning. Constructivists about practical reasons think there are additional norms like this for practical reasoning. Different constructivist views make different claims about what these norms are. (Candidate constitutive norms are, for example, norms directing an agent's reasons-takings to be consistent, instrumental principles, prudential principles, and moral principles.)

Constructivist accounts like Korsgaard's take the constitutive norms of practical reasoning to determine rational permissibility to act for the sake of particular feature or eligibility to be a reason for a particular action. This leads to the third key commitment of such accounts (**C3**).

C3: The constitutive norms of practical reasoning determine only eligibility to be made a reason. In many cases, they determine that more than one feature of actions is eligible to be made a reason for action.

For example, I might find myself in a set of circumstances today such that either going to see a movie for the sake of entertainment or grading papers for the sake of my students' learning would accord with the constitutive norms of practical reasoning. (Either option is morally permissible, either option is a reasonably effective way to promote my particular set of ends, etc.). This means that there is nothing to be said, in light of the constitutive norms, against performing one of these actions for one of these reasons. But there is also nothing to

65 Different views propose different descriptions of what it means to be indifferent or responsive to constitutive norms of reasoning. They attempt to find some middle ground between explicitly representing the rule to oneself in all instances of reasoning and accidentally complying with the rule, while leaving room for an agent to make mistakes, perhaps knowingly, and still count as engaged in the relevant form of reasoning.

be said, in light of the constitutive norms, *for* performing one of these actions for one of these reasons.[66] Doing so is merely permissible. (By contrast, suppose, for example, there were a constitutive norm that specified a value, e.g., utility, that a reason to φ would indicate that φ-ing would promote. On this view, in every case, there would be something to be said in light of the constitutive norms for performing a particular action for a particular reason—doing so would promote that value.)[67]

And finally, the fourth key commitment of constructivist accounts of practical reasons (**C4**) is a constraint on what could be a genuine requirement of practical reasoning.

C4: If X is a genuine requirement of practical reasoning, then it must be possible for a rational agent to comply with X using reasoning alone (given sufficient time and accurate information).

C4 is motivated, I think, by the idea that all standards of correctness applicable to practical reasons are standards for how to engage in practical reasoning (not for some other activity) and an "ought implies can" intuition.

Constructivist accounts of practical reasons are appealing for several reasons. They are thought to have certain metaethical advantages. This is because the only normative facts that they posit are standards of correctness that are supposed to be necessary in order to

66 In some cases, it may turn out that only one feature of possible actions is eligible to be made a reason in that particular case. I might, for example, find myself in a situation where all but one action for one reason would significantly undermine the ends that I have set for myself. Assuming some instrumental principle is a constitutive norm of practical reasoning, there will be something to be said in light of those norms for performing this one action for this one reason—namely that so acting is the *only* eligible thing to do.

67 Constructivists like Korsgaard tend to describe the third commitment differently. They claim that the constitutive norms are "formal". This means that these norms do not direct reason-takings to have any particular content. (This also contrasts with a norm that specifies a value, e.g., utility, that a reason to φ would indicate that φ-ing would promote.) The "formal" constitutive norms that Korsgaard and other constructivists propose do determine mere eligibility to be a reason, but it is not clear that formal norms *must* do so. A principle directing an agent to act on her strongest desire, for example, might count as formal and yet identify one feature of actions in each case that should be the agent's reason. Because it is the fact that the constitutive norms determine mere eligibility to be a reason that creates the problem for Korsgaard's view, I will not discuss the formality of the norms in this paper.

describe the activity of practical reasoning in which we all engage. If such facts are necessary to describe this familiar activity, then we have strong reasons to think that there are such facts. Constructivist accounts also provide a simple and attractive explanation of why reasons for action motivate an agent to act: Reasons motivate an agent to act because they are based in her desires.

Constructivist accounts that take on **C1**-**C4** seem set up to take reasons to be fundamentally subjective.[68] According to **C3**, the constitutive norms rule a wide variety of features as eligible to be made reasons for action. Which features are made reasons for action, then, seems to be up to particular agents in two ways. An agent's reasons depend on (1) what she happens to desire (**C1a**) and (2) on which of these desired and eligible features of actions she chooses to act for the sake of (**C2**). We might take this subjectivity to be a serious problem with constructivist accounts because it conflicts with an intuitive picture of practical reasons. As suggested in the previous section, we tend to think that there are some general features of actions that count in favor of actions for all of us regardless of what we happen to desire and the features of action that we choose to act on. Features of action pertaining to beauty and health come to mind, as well as morally significant features of action (e.g., that an action will help or harm someone else).

Korsgaard attempts to account for this intuitive objectivity of reasons while retaining the four commitments by claiming that the constitutive norms of practical reasoning yield a requirement to s-share reasons. In this paper, I argue that her attempt to do this faces a problem, created by **C1**, **C3**, and **C4**. Because the problem is created by these general

68 Sharon Street, for example, advocates a constructivism that makes this claim in "Constructivism about Reasons."

commitments and not by her particular interpretation of these commitments, I take this to show that any account that takes on these commitments will have difficulty accounting for this sense in which reasons are shared on the intuitive picture.

The Basic Problem is as follows: In order to s-share reasons, agents must make the same features of actions their reasons for action by choosing to act for the sake of them. By **C1a** reasons pick out desired features of actions. So, in order to s-share reasons, agents would seem to have to desire the same features of actions. A requirement to s-share reasons, then, would seem to yield a requirement to have particular desires. By **C4**, a genuine requirement must be such that an agent can, in principle, comply with it using reasoning alone (given enough time and information). However, by **C1b**, an agent cannot reason her way into having particular desires.

Korsgaard thinks that she can solve this problem by allowing agents to act for the sake of particular features of actions just because they recognize that doing so is required (more on this in the next section). In Section V, I will argue that her commitment to **C3** prevents her from solving the problem in this way. I then argue in Section VI that the problem can be solved by modifying the view in a way that abandons **C3** (and that once we abandon **C3**, **C1** is no longer plausible). I conclude by briefly discussing some reasons for thinking that, while giving up **C4** could also solve the problem, this is a much less plausible solution. First, however, Section IV will briefly explain Korsgaard's account of reasons for action and identify its constructivist commitments.

IV. Korsgaard's constructivism

On Korsgaard's view, a reason for an act specifies the "purpose"[69] for which an agent performs that act.[70] Put roughly, an act is what the agent does and a purpose is why she does it—it is the point she sees in what she does, that for the sake of which she performs the act. Purposes often cite further ends for the sake of which an agent performs an act. For example, I might walk to a coffee shop to get coffee, to meet a friend, or to work. However, Korsgaard's conception of a purpose is not limited to that of a *further* end to which the act is a means.[71] A purpose can be any feature of the act that makes it attractive to the agent.[72] It is the object of a (non-normative) desire that is not subject to rational revision.[73] And so, Korsgaard accepts **C1**.

Korsgaard also takes on **C2**. On her view, that an act A promotes a purpose P counts in favor of A-ing and becomes a reason for A-ing when the agent chooses to A for the sake of P in proper deference to norms that are constitutive of practical reasoning. Korsgaard

69 Some terminological clarification is in order here. In "Acting for a Reason," Korsgaard claims that the term "reason" properly applies to principles that cite both the agent's act and her purpose in acting (e.g., "I will perform act A for the sake of purpose P" rather than to the specification of purposes (220-2, 225). However, she also claims that, when an agent chooses in the right way to act for the sake of a purpose, she (correctly) takes that purpose to count in favor of that act (227-8). Because of this, I think the more standard terminology I use here—where the "reason" is the thing that counts in favor of an act—is a fair translation. (Unfortunately, a thorough defense of this choice would take this chapter far afield.) Also, Korsgaard takes the term "act" to refer to the thing that an agent does described independently of the purpose for which she acts. She takes the term "action" to refer to the act-done-for-a-purpose complex. ("Acting for a Reason," 219). I use the two terms interchangeably here to refer to the thing that an agent does described independently of the purpose for which she acts.

70 Korsgaard, "Acting for a reason," 218.

71 Korsgaard grants that there are interesting distinctions between the ways in which purposes can be related to acts, but she remains committed to a very broad description of this relation: Any time an action is guided by the application of a concept to a particular, or any time it is guided by a conception of what the agent is doing, it will count as done for a purpose. This will allow for purposes that some might not clearly count as ends (e.g., "for the fun of it," "for its own sake"). (Korsgaard, C.M. "The Normativity of Instrumental Reason," *The Constitution of Agency*, 28.)

72 Korsgaard, "Acting for a reason," 214.

73 For example, in "The Myth of Egoism," Korsgaard describes practical reasoning as operating on "givens" proposed by desire (*The Constitution of Agency*, 87).

takes Kant's categorical imperative to capture the constitutive norms of practical reasoning.[74] This means that the agent can make the fact that A promotes P a reason for A-ing if and only if she can will that everyone act on the principle "I will do A for the sake of P" without willing a contradiction.[75]

In taking the categorical imperative to be the constitutive norm of practical reasoning, Korsgaard takes the categorical imperative to tell us something about what it is to be engaged in practical reasoning. She thinks that thought patterns that exhibit no responsiveness to the categorical imperative will not count as instances of practical reasoning.[76] This claim may seem implausible because most practical reasoners don't know what the categorical imperative is. However, there is a way of understanding the claim that alleviates this worry: We might think that, when an agent exhibits no concern for whether her acts are justified, she does not count as engaging in practical reasoning. We might then understand the categorical imperative as a formalization of this idea of justification that agents operate with everyday.[77]

Korsgaard also understands the categorical imperative to be a norm that determines rational im/permissibility. Her particular claims about what it does and does not permit make clear that, in many cases, when an agent acts for a reason, her reason is merely one

74 Korsgaard thinks that the categorical imperative is the only principle that is constitutive of practical rationality. However, she also thinks that a variety of other principles that we might have thought were also partly constitutive of practical rationality (e.g., some instrumental principle directing the agent to take the means to her ends and some principle(s) directing the agent to give some consideration to her other ends when deciding what to do) are *aspects* of the categorical imperative. ("The Normativity of Instrumental Reason," 68)

75 For example, I cannot will without contradiction that everyone act on the principle "I will tell you that I will pay you back without intending to do so in order to get you to lend me money." Put crudely, this is because if everyone acted on such a maxim, "I will pay you back" would not lead anyone to believe it and thus would no longer be a possible way to get you to lend me money.

76 See, for example, "Realism and Constructivism in Twentieth-Century Moral Philosophy," *The Constitution of Agency,* 321 or the discussion of the reflexive structure of rational motivation in "Acting for a reason" (213-4).

among many features of action that the categorical imperative rules as eligible to be made reasons. The categorical imperative rules out actions that are morally wrong, and perhaps actions that grossly undermine an agent's purposes. This leaves the agent with many options when deciding what purposes to pursue. Korsgaard takes it to be a virtue of her account that it grants agents a wide variety of eligible options when deciding which purposes to pursue.[78] Thus she takes on **C3**. I think it is relatively uncontroversial that Korsgaard thinks that it must be possible, in principle, for an agent to reason her way into compliance with a requirement, and thus takes on **C4**, so I will not take the time to defend this attribution here.

There is much debate about how exactly the categorical imperative should be understood and about what it can and cannot tell us about what we should do. My argument in this paper does not require a precise understanding of the categorical imperative that would take a position on these issues. It requires only that the categorical imperative be understood to be the constitutive norm of practical reasoning, that it be understood to determine rational im/permissibility, and that it be understood such that any requirement to make this or that feature one's reason for acting follows from the categorical imperative (given a set of circumstances). There may be independent reasons to think that each of these claims is false. However, my aim here is to show that Korsgaard's view faces a problem *even if* she is right that we can understand the categorical imperative so that all three claims

77 We might wonder whether the categorical imperative is the best way to formalize this notion. I will not attempt a direct discussion of this issue here. However, we might take the arguments in favor of a requirement to share reasons to suggest that such a requirement is necessary to capture this everyday notion of justification. If so, then insofar as I can show that Korsgaard's view fails to account for the s-shareability of reasons, I suggest that the categorical imperative, as Korsgaard understands it, is not the best formalization of this notion of justification.

78 This is suggested, e.g., by the role that she takes these choices to play in constituting an individual's identity in *The Sources of Normativity* and in her discussion of ambitions in "The reasons we can share" (284-91).

are plausible.

As stated above, unlike other constructivists who take on **C1**-**C4**, Korsgaard attempts to account for the intuitive objectivity of reasons by including a requirement to s-share reasons in her view. In order to see how she does this, it is first necessary to understand a further commitment of her view. Although she endorses **C1**, and thus thinks that many reasons originate in non-normative desire, she also thinks that agents can be moved to act for a particular reason just by the recognition that the categorical imperative requires acting for that reason. In those cases where the categorical imperative rules out all but one feature of action as ineligible to be made a reason, the agent can recognize this and be moved to act for the sake of this feature just because she sees that so doing is required. No desire for the feature is necessary for this motivation. [79] Furthermore, the categorical imperative requires that agents adopt a few necessary ends, e.g., helping others and developing one's talents. On Korsgaard's view, agents are able to recognize that the categorical imperative requires the adoption of these ends (and taking them to be reason-giving), and do so on that basis. No separate desire to help others is required. Korsgaard's view, then, seems to allow that there are two different sources of the motivation behind reasons—non-normative desire and some sort of purely rational motivation.

79 Korsgaard has to allow this second source of motivation because she takes the categorical imperative to yield a variety of ordinary moral requirements. We are often morally required to do things that we may not have any desire to do (e.g., stop what we are doing to help someone who is in need). By **C4**, agents must be able to reason their way into compliance with requirements, so agents must be able to reason their way into acting for the sake of helping another. Agents cannot reason their way into desires, so agents must be able to act for the sake of this purpose without desiring it, but instead merely because they recognize that the categorical imperative requires so acting.

V. The Problem

Recall, from Section III, the Basic Problem for accounting for a requirement to s-share reasons while holding **C1**, **C3**, and **C4**: According to **C4**, a genuine requirement must be such that it is possible for an agent to reason her way into compliance with it (given sufficient time and information). However, it will not always be possible for an agent to reason her way into s-sharing reasons with other agents. This is because s-sharing reasons sometimes requires having particular desires, but according to **C1**, agents cannot reason their way into desires.

S-sharing reasons will require having particular desires when the reason to be s-shared is of a particular kind. On Korsgaard's view, a reason is created when an agent (permissibly) chooses to pursue a particular purpose. Some of these purposes follow from other purposes (e.g., a purpose of buying boots might follow from a purpose of climbing Kilimanjaro). Some purposes are necessary—they are required by the categorical imperative (e.g., helping those in need). But there is a third way that reasons are created: Because the categorical imperative rules a variety of features of actions as eligible to be reasons (**C3**), the purposes that remain started out as merely attractive (desired) features of actions that the categorical imperative ruled eligible to be made reasons. Reasons that cite these purposes are created when an agent chooses to pursue them. In order for an agent to adopt a purpose of this latter kind (and thus be able to s-share reasons that cite it), it would seem that the agent must have a desire for the relevant feature of action. A requirement to s-share a non-necessary reason, then, appears to yield a requirement to have a particular desire. Thus, the Basic Problem arises.

I think that, in the end, Korsgaard's view is subject to the Basic Problem. However, she has a way to try to avoid it. Recall that she allows two sources of motivation for reasons. Many reasons originate in desires. But an agent can act for the sake of a particular feature of an action even if she has no desire for it when (and just because) she recognizes that the categorical imperative requires so acting. If she does this, she will count as reasoning her way into pursuing that purpose, bypassing desire. This suggests a solution to the Basic Problem. If there is a requirement to s-share reasons, then an agent should be able to pursue a particular purpose F for its own sake when and just because she recognizes that she should s-share reasons that cite that purpose.

Ultimately, however, this does not provide a way to reason one's way into s-sharing reasons. The thought that s-sharing each other's reasons is required might, for example, get an agent to make beauty, testing one's physical abilities, and perseverance into purposes that she *pursues* independently of furthering the projects of the other agents who established those reasons in the first place. (For example, A could recognize that she should pursue beautiful experiences independently of B's pursuing them and, as a result, go to an art museum without B.) However, while the pursuit of a purpose set in this way may be independent of the agents who first pursued this purpose, the agent's *interest* in the purpose is not. A will remain interested in beautiful experiences only because she recognizes that B has permissibly made them his purpose. As a result, she will fall sort of s-sharing reasons that cite this purpose.

Put roughly, the problem is this: if an agent were to set purposes for herself in the way just suggested, something like "they're interesting because you find them interesting"

will still appear in the explanation of why she is interested in those purposes. If so, then she is only w-sharing the reasons that cite these purposes after all. In order for an agent to s-share a reason (and to avoid the objectionable attitudes discussed in Section II), the sort of interest she must have in the purpose picked out by that reason must be something like "it's interesting for its own sake." It is what the first agents who had these interests were on to, whether or not anyone else shared the interests.

What exactly prevents an agent (A) from finding the purpose of another (B) interesting for its own sake because A recognizes that the categorical imperative yields a requirement to s-share reasons? Suppose A takes productivity in a particular profession to be a central interest and purpose in her life. A thinks that beautiful vistas distract from the pursuit of this purpose and that perseverance is only worth pursuing (in fact, is only interesting) when it yields this productivity. Therefore, A sees nothing interesting about climbing Kilimanjaro itself. However, A sees that B has chosen, in the right way, to climb Kilimanjaro and A can thus accept that B's purpose of climbing Kilimanjaro gives her reasons, not because it is productive, but simply because B chose it in the right way.

In order for A to get from here to making the experience of beautiful vistas and the exhibition of perseverance (the features of climbing Kilimanjaro that appeal to B) purposes that she has and pursues for their own sakes so that she will comply with the requirement to s-share reasons, A needs to find a way to remove a reference to B's choice of purposes from *some*[80] explanation of why she has and pursues these features. This would allow A to be

80 Of course, in many cases, A's interest in a feature will have many sources—different aspects of the feature itself that interest her and perhaps the fact that others pursue it. S-sharing reasons that cite that feature does not require that *none* of her interest in the feature depends on its being the permissible purpose of others. It merely requires that there is *some* explanation of some of her interest in the feature that does not depend on its being the permissible purpose of another but instead cites something about the feature itself.

interested in and pursue a feature of action that B has made a purpose for its own sake and not merely because it is B's purpose. The way to do this, on Korsgaard's picture, is for A to recognize that she should take an interest in the features of actions that B makes his purposes for their own sake.

However, this recognition itself involves a recognition that B's purposes give A reasons, and should interest A for their own sake, *because B made them her purposes*. Since beauty and perseverance are non-necessary purposes, A only faces a requirement to s-share reasons that cite them because B happened to make those features of action his purposes. In order to recognize that she is under this requirement, A must recognize the role of B's choice in creating it. (In order to determine which features of action the requirement to s-share reasons demands that A make her purposes, A must know which features of action B has made his purposes and thus which reasons she is required to s-share.) So, the recognition that s-sharing the reasons that cite these features is required necessarily involves the very reference to another's making that feature a purpose that we are trying to get rid of. If this recognition that s-sharing is required explains A's interest in a particular feature of action, then that explanation will reveal that A's interest depends on B's having made this feature his purpose.

So, if A makes beautiful vistas and perseverance, and not merely B's pursuit of them, her purposes (if A pursues them in her own life whether or not they continue to interest B) because A recognizes that she should do so, in order to s-share reasons they cite, then it might look like A is interested in those features for their own sake. (After all, she can pursue them herself whether or not B does.) But the explanation of why these features interest A is

not in fact free of references to B making them his purposes. A made beautiful vistas and perseverance her purposes, not because A found them interesting because of something about those features themselves, but because A took herself to be obligated to make them hers *because they were B's purposes.* So, a reference to B's making those features his purposes remains in A's explanation of her interest in them.

So, the route via rational motivation to adopting purposes so as to s-share reasons that cite them is in this way problematically infused with references to other agents' purposes. This means that the only route left to adopting a non-necessary purpose so as to s-share reasons that cite them seems to be to first desire that feature of actions for its own sake. This in turn means that, in order to s-share reasons, agents will be occasionally required to desire particular features of actions. Since desires are not subject to practical reasoning, practical reasoning would not provide a way to comply with such a requirement. And because of **C4**, the view cannot then claim it as a requirement. In other words, the Basic Problem persists after all.

This should not be entirely surprising. I suggested above that in striving to take an interest in a particular feature of actions for its own sake, we are striving to have the same sort of interest in these features as did the first agents who had these interests. But in cases of non-necessary purposes, it seems that the first agents' interests could only be desires for these features of action (for their own sake). Because these purposes are non-necessary, they are among a wide variety of features of actions that the categorical imperative rules as eligible to be a reason (**C3**). It is not their eligibility to be a reason that interests agents in them. (After all, many options may be eligible, but not interesting to an agent). So agents'

interest in these features must be that of having a desire for them. There does not seem to be anything else in Korsgaard's story of acting for a reason that it could be.

VI. Proposed solution

The problem for Korsgaard's view arises because the categorical imperative rules a variety of features of actions as eligible to be made reasons (**C3**), leaving room for an agent's desires to play a key role in determining which reasons she has. This suggests a solution to the problem: Suppose that rather than ruling a variety of features eligible to be made reasons for action, the categorical imperative specified general, foundational features of actions that all agents should make their purposes which can then ground more particular purposes (regardless of what agents happen to desire). If we could understand the categorical imperative in this way, then we would understand the reasons that we should s-share with other agents as reasons that we each have just in virtue of being practical reasoners. Korsgaard's claim that the categorical imperative requires us to value our capacity for practical reasoning and its proper exercise in ourselves and others suggests a way in which we might do this:[81] We might understand the purposes that we should have so as to s-share reasons that cite them as necessary purposes (like that of helping others) that we are required to have because of their relationship to our rational nature—our capacity to set ends and to act in proper deference to the constitutive norms of practical rationality. What these purposes have in common and what makes them purposes that we should have is that they promote, preserve, respect, or are part of the well-functioning of this capacity.

The ability to persevere and to overcome obstacles, for example, is presumably a

81 See, for example, "Aristotle and Kant on the source of value," *Creating the Kingdom of Ends*, 241.

virtue for the rational agent because it makes her more effective at exercising her rational nature—at achieving the purposes she pursues whatever they may be. It also opens up a wider variety of purposes as reasonable for her to suppose she could achieve, and therefore as reasonable for her to pursue. It makes her more likely to be able to go through with odious moral obligations. For these reasons, we might think, the ability to persevere and overcome obstacles is a purpose that we should all have in virtue of being practical reasoners. And because of this, we have reasons to engage in activities that express, maintain, and develop the willingness and ability to persevere and overcome obstacles because these actions promote and preserve that value. (This leaves open the possibility that one agent may choose to climb mountains, another to do mathematical proofs, and another to master the violin all for the sake of this same value.)

This modified view provides a clear explanation of how we can reason our way to s-sharing reasons. On Korsgaard's view, if agents A and B do not already s-share some particular reason, there is no way for A to reason her way into sharing that reason with B such that a reference to B's having that reason would not appear in the explanation of why A has the reason (or vice versa). This problem does not arise for the modified view because, on the modified view, A and B can each correctly choose to pursue that purpose because they recognize its relationship to his or her own (or anyone else's) rational nature. In this explanation of why each of us must have a given purpose, there is no necessary reference to any individual whose interest in, or non-normative desire for, that feature of action, and whose subsequent choice to act for the sake of that feature, make that feature an interest or a purpose we should all have. So, there is no danger of such a reference necessarily appearing

in the explanation of why an agent has that purpose.

On this picture, the basic purposes that ground agents' more particular purposes as well as their practical reasons do not originate in desire. This solves the Basic Problem—s-sharing reasons no longer requires having desires for the relevant features of action. If this is the case, however, it requires giving up, or at least significantly modifying, **C1**. Reasons would no longer originate in non-normative desires. Instead they would follow from the value of rational nature. On the modified view, an agent comes to have the purposes that ground reasons not because she happens to want them, and chooses to pursue them in the recognition that doing so accords with the value of rational nature. Instead, she comes to pursue them because she sees that they are *called for* by the value of rational nature.

Desires or preferences might retain a limited role in this picture. Perhaps an agent may choose whether to pursue the value of perseverance by climbing mountains, by doing mathematical proofs, or by mastering the violin on the basis of which activity is an object of her (non-normative) desires. And in choosing to make one of these projects her purpose rather than another, she might create reasons for herself—reasons to pursue this specific project—that she did not have before.[82] But these reasons, unlike Korsgaard's, would be grounded in the necessary purpose of perseverance (and ultimately in the value of rational nature) as well as in the agent's choice based on desire.

82 Or at least she can change the relative weight of her existing reasons. For example, perhaps, for any beautiful experience, an agent has *a* reason to pursue it because beauty is a necessary purpose. However, adopting a project of pursuing a particular beautiful experience (e.g., the beautiful vistas of Kilimanjaro) makes her reason to pursue this beautiful experience weightier than those other reasons. Alternatively, perhaps an agent has a reason to pursue some beautiful experiences because beauty is a necessary purpose, but she has no reasons to pursue any given particular beautiful experience, until she makes it her purpose.

VII. Why this solution?

The problem for Korsgaard's view is caused by **C1**, **C3**, and **C4**. Because Korsgaard takes the categorical imperative to rule a variety of features of action eligible to be made reasons for action (**C3**), which features an agent makes her reasons depends on her desires (**C1a**). (This appeal to desires is necessary to explain why reasons motivate agents to act. The mere recognition that acting for the sake of some feature of action is eligible does not itself move agents to act for the sake of that feature because an agent can often just as well comply with the categorical imperative, in such a case, by doing nothing.) But once many reasons depend on desires in this way, a requirement to s-share reasons cannot avoid requiring that agents desire particular features of action. Agents cannot reason their way into desires (**C1b**). But by **C4**, a genuine requirement must be such that it is possible for an agent to reason her way into compliance with it, at least given adequate time and information. So, **C1**, **C3**, and **C4** rule out a requirement to s-share reasons.[83]

In Section VI, I suggested that the problem could be solved if we re-interpret the categorical imperative so as to give up **C3** and that, if we do this, we should give up or substantially revise **C1** as well. Much more work would need to be done in order to show that this modified view yields a plausible account of practical reasons. We would need to show that this re-interpretation of the categorical imperative is plausible as a constitutive norm of practical reasoning and that it would yield the right reasons. (It is not clear, for

83 C2 does not play a direct role in creating the problem. (It may play an indirect role insofar as it motivates C4). C2 adds a level of subjectivity to reasons insofar as an agent chooses whether or not to act for the sake of an eligible desired feature of action. But the problem for Korsgaard's view would arise even if we eliminated this choice. Suppose all that determined whether an eligible feature of action were a reason for that action for a particular agent was whether or not the agent desired it. There would still be non-necessary reasons such that, in order to have one of these reasons, an agent must desire a particular feature of actions.

example, how reasons that cite beauty would follow from the value of rational nature.) My aim here is not to defend the particular account of practical reasons sketched in Section VI but instead to suggest that, whatever norms a constructivist proposes as the constitutive norms of practical reasoning, these norms must be understood in such a way that avoids **C3**, if the view is to account for the shareability (in the stronger sense) of reasons. Section VI aims to show that giving up **C3** allows the view to account for this shareability. In this section, I will briefly explain why this seems to be the most plausible way to account for this shareability while retaining some constructivist commitments.

The seemingly simplest way to modify Korsgaard's view so as to solve the problem would be to revise **C1** so that agents could reason their way into desires after all. The idea here would be that an agent could have a desire for a particular feature of actions upon recognizing that she should have this desire. This revision, however, will not solve the problem of Section V. In order to s-share some non-necessary reasons, an agent must find the relevant features of actions interesting for their own sakes (and not merely because they are reasons that were set in the right way by other rational agents). To do this, the agent must desire those features for their own sakes. The view that revises **C1** suggests that agents do this by desiring a feature F of actions for its own sake because F is the permissible purpose of another rational agent. But this looks impossible. If the agent's reason for desiring F is that F is the permissible purpose of another agent, then she does not desire F for its own sake.

The remaining alternative for trying to solve the problem is to give up **C4**. This would allow there to be requirements compliance with which requires more than practical

reasoning (given sufficient time and information) and would thus permit a requirement to s-share reasons that yielded a requirement to have particular desires.[84] I think that giving up **C4** is much less plausible than giving up **C3**. I will not attempt to provide a thorough defense of this claim here, but the basic idea is this: Some modification or other to Korsgaard's view is called for if her view is to include a requirement to s-share reasons. However, I think that a view that gave up **C4** would have difficulty justifying this requirement to the agents subject to it. And I take it to be a condition on being a plausible requirement that it be justifiable to those who are subject to it.

On a modified view that gives up **C4**, the requirement to s-share reasons yields a requirement to have particular desires. As a result, when one agent chooses to act for the sake of a particular desired feature of action, she places a demand on other agents that they desire that feature of action (or some more general feature of action in light of which this feature is desired) for its own sake. The justification that an agent has to offer others for this demand is something like "it's permissible and I want it". However, this seems to be the wrong kind of consideration to justify a demand to desire something for its own sake. A demand to desire a particular feature of action for its own sake seems to require a justification that cites something to be said *about that feature itself* that reveals it to be interesting or appealing on its own, and not just in virtue of being the object of another agent's desires that she has permissibly pursued.

Furthermore, an "it's permissible and I want it" justification seems particularly insufficient as justification for a demand to desire particular features of action when we

84 This would require a change to **C1b**: Desires would not be subject to rational revision but would be normative.

consider what this demand requires agents to do. Because agents cannot reason their way into having desires, in order to acquire a desire that she does not already have, an agent is presumably going to have to put herself through some sort of psychological conditioning. A requirement to do this, I think, intuitively comes with a fairly high justificatory burden.

While we might be inclined to accept such a requirement in a few cases of morally significant desires (e.g., a desire for the happiness of one's children), a modified view that gives up **C4** would include such a requirement for every desire that some agent happens to have and chooses to act on in proper deference to the constitutive norms of practical reasoning. Thus, agents would be required to take non-rational measures to shape their own psychology to ensure that they desire all of the same features of actions as do others. The level of intrusion into and non-rational manipulation of our psychology that this would demand is at least disturbing. It also seems likely to be difficult to justify to the agents of whom it is supposedly required given that the reason for it is not that there is something about those features themselves that is interesting or appealing or worth desiring, but rather that some other agents happen to be motivated by them.[85]

As a result, giving up **C3** and modifying **C1** seems to be the most plausible way to account for the shareability of reasons while retaining some constructivist commitments.

85 A view that denies **C4** could give an explanation of why we should desire particular features of actions that appeals to something about those features themselves. But this sort of explanation suggests that it is something about the features and not the agents' permissible interest in them (and subsequent choice to pursue them) that that makes them reasons. Thus, in order to make this sort of explanation available, I suspect we would need to deny **C1** and **C3** (and perhaps **C2** as well). This sort of view might well be worth considering in its own right, but it would bear little resemblance to constructivist views. And importantly, it takes on the commitment that I will want to draw from all of this—that an explanation of why the fact that an action has some particular feature counts in favor of an action will have to appeal to something about that feature itself and not merely to agents' attitudes towards it.

VIII. Conclusion

In summary, we have seen that three characteristic commitments of constructivist views prevent these views from regarding reasons as to be s-shared. Constructivists think that, in many cases, the constitutive norms of practical reasoning determine that there are several features of actions for the sake of which it is permissible to act under the circumstances (**C3**). Which features of action are picked out by an agent's reasons depend on which she desires (**C1a**). In order to act for the sake of the same feature (F) of action as does another agent, then, it would seem that both agents must desire F. Agents cannot reason their way into desires (**C1b**), but agents must be able to reason their way into requirements (**C4**). So, it would seem that constructivists cannot allow a requirement to share reasons.

Korsgaard's view attempts to provide a way for agents to reason their way into sharing reasons and thus allow a requirement to share reasons. She claims that in those cases where an agent is required to act for the sake of particular feature (F) of action, she can be moved to act for the sake of F just by the recognition that she is rationally required to act for the sake of F. I argued in Section V that, while agents could, in this way, reason their way into acting for the sake of the same features of actions, agents could not, in this way, reason their way into finding those features of action interesting for their own sake. Because of the way in which agents create reasons on a constructivist view (**C1-C3**), recognizing that acting for the sake of F is required must (often) involve recognizing that so acting is required because another agent chose, in the right way, to act for the sake of F.

Yet, as Section II suggests, we have reasons to think we are under a requirement to s-

share reasons—to act for the sake of the same features of actions and to regard those features as calling for actions because of what they are, and not merely because some rational agents made them their reasons. In Section VI, I propose a modification to Korsgaard's view that can allow for this requirement. This modification bears an important similarity to Korsgaard's failed solution; agents are able to act for the sake of the same feature (F) of action just because they recognize that they are required to act for the sake of F. They are thus able to reason their way into compliance with a requirement to share reasons that cite F even if they do not desire F. However, the explanation of why an agent must act for the sake of F, on the modified view, is not that some other agent chose to act for the sake of F and sharing reasons is required. If F is reason-giving, F is a way in which actions promote, preserve, or respect the value of rational nature. It is because of this fact about F that F is reason-giving and not because of the attitudes that others happen to have towards F.

The modified view allows a requirement to s-share reasons by giving up **C1a** and **C3**. It was the fact that, in many cases, the constitutive norms of practical reasoning rule several features of actions as eligible to be made reasons (**C3**) that left room for a particular agent's desires to determine which reasons she has (**C1a**). This in turn made the agent's choice to act on those desires part of the explanation of why the features of actions for the sake of which an agent acts are reason-giving for others. In order to remove this problematic appeal to agents' desires (giving up **C1a**), we need to make the constitutive norms of practical reasoning, on their own, sufficient to determine what reasons agents have. This is possible if the constitutive norms specify a value (or values) that reasons for action

indicate that actions would promote, preserve, or respect instead of merely determining which features of actions are eligible to be made reasons and letting an agent's desires determine the content of her reasons within that constraint.

Chapter Three

Introduction

In Chapter 2, I argued that a constructivist view that takes the constitutive norms of practical reasoning to be norms of permissibility cannot allow for a requirement to s-share reasons. In this chapter, I will argue that this is a serious problem for constructivist views (and indeed for any view which denies R and instead takes the content of reasons to be given by particular agents' non-normative desires and constrained by "formal" principles of practical reasoning).

This argument is somewhat roundabout. Recall that at the end of Chapter 1, I suggested that, rather than attempt to find restrictions on the content of reasons for action by looking thinking about the difference between intentional and non-intentional action, we should attempt to find restrictions on the content of reasons for action by looking at the difference between actions that an agent does because she takes some consideration to "call for" those actions, and those that she does, e.g., just because she feels like it. In this chapter, I will examine a constructivist account of reasons for action—that proposed by Sharon Street—that explicitly takes itself to be attempting to make this distinction. I will raise problems for Street's view that, I think, help clarify what the target is for an account that attempts to capture what it is for a reason to "call for" (or to count in favor of) an action. And finally, I will return to the arguments for a requirement to s-share reasons discussed in Chapter 2. I will argue that these arguments show that reasons for action must be to be s-shared, if they are to call for actions.

I. Street's constructivism: the basics

Street presents her constructivist view as an account of correctness in normative judgment. For Street, normative judgments are, one way or another, judgments about reasons.[86] They can be judgments about what particular reasons an agent has in a particular situation—i.e., judgments that some consideration C is a reason for some agent A to perform some action φ. They can also be judgments that express principles of practical reasoning. These principles can be quite general, e.g., some principle of instrumental rationality. Or they can be quite specific, e.g., some principle telling me how I ought to weigh the effects of some action on my well-being and the effects of that action on my mother's well-being, in a particular situation. Normative judgment can be judgments about value, e.g., that health is valuable. However, Street takes judgments about value to be reducible to judgments about reasons. The claim that health is valuable simply means that there are reasons to promote and preserve health. Normative judgments can be moral judgments—e.g., that it will benefit me and I can get away with it is no reason (or at least not a sufficient reason) to steal from the poor. They can also be more mundane reasons judgments—e.g., the fact that ice cream is tasty is a reason to eat it.[87] When an agent's normative judgment that C is a reason to φ is correct, C is in fact a reason (for that agent) to φ. So, in giving us a way of determining

86 Street, Sharon. "Constructivism about Reasons."

87 To describe all of these examples as "judgments" is a bit misleading. Each of them *could* be an explicit judgment—I could judge that the fact that ice cream is tasty is a reason to eat it or that the effects of an action on my well-being and the effects of that action on my mother's well-being should be weighed in a particular way. However, we might take the category of "normative judgment" to include more than explicit judgments. We might think that agents frequently act for reasons, and in some sense takes the relevant considerations to be reasons, without making any explicit judgment "X is a reason to φ". Being moved to act by some consideration will count as taking that consideration to be a reason to φ when the agent's motivation by that consideration is responsive to other normative judgments. (The notion of responsiveness to normative judgments and the sense in which agents take considerations to be reasons when acting will be discussed further in Section II of this chapter.) This agent is then committed to the normative claim "X is a reason to φ" even if she makes no explicit judgment to this effect.

which normative judgments are correct, Street is giving us a way of determining what reasons we actually have (and, as we shall see, an account of what makes something a reason for action).

Street describes her view as follows:

> [T]he fact that X is a reason to Y for agent A is constituted by the fact that the judgment that X is a reason to Y (for A) withstands scrutiny from the standpoint of A's other judgments about reasons.[88]

The basic idea of Street's view is that an agent's normative judgment is correct when (1) she makes it, (2) in making it she is responsive to some norms demanding that it be consistent with these other judgments, and (3) it is consistent with her other normative judgments. (1) and (2) are related. An agent will not count as making a normative judgment if it is not the case that, in making it, she is responsive to other normative judgments. That is, for an agent to count as making a normative judgment N, it must be the case that, were she aware that N conflicted with other of her normative judgments, she would be moved to give up N (or some other normative judgment) in order to resolve this conflict. These other normative judgments also, according to (3), determine whether any given normative judgment is correct. (N is correct if N is consistent with the agent's other normative judgments.)

Street's view is motivated by metaethical skepticism. She thinks that there are no standards outside the practical point of view of a particular agent—i.e., the point of view of the agent thinking about what reasons she has and what she should do—with which to assess her normative judgments.[89] However, the only standards available *within* the practical point

88 Ibid., 223.

89 See Street, Sharon. "A Darwinian Dilemma For Realist Theories Of Value." *Philosophical Studies* 127, no. 1 (2006): 109-166.

of view of a particular agent are her normative judgments themselves.[90] Fortunately, there is some necessary common ground between agents for this assessment. It is provided by norms that are constitutive of practical reasoning. These are normative judgments that an agent must make, if she is to count as making normative judgments and engaging in practical reasoning at all.[91]

To see why there are such necessary normative judgments, first note that the kinds of assessments that Street takes to make up an assessment of whether a normative judgment is correct are of an ordinary and familiar kind. For example, if an agent thinks that the fact that someone is drowning is a reason to save him (other things equal), and thinks that she only has reasons to perform actions that are in her self interest, we think that she should give up (or revise) at least one of these judgments. It at least appears that we do not need to appeal to anything but her own normative judgments to explain why this is. The fact that she makes the first judgment commits her to taking herself to have a reason to help each drowning person that she encounters. The fact that she makes the second judgment commits her to taking herself to have a reason to φ only when φ-ing is in her self-interest. In cases where she encounters a drowning person and believes that saving the person is not in her self interest, she is committed both to having a reason to save the drowning person and to having no reason to save the drowning person. Because this pair of commitments is inconsistent, we take this agent to have made a mistake. At least one of her normative judgments must be incorrect.

This assessment suggests that the agent is subject to a norm that states (roughly) that

90 Street, "Constructivism about Reasons," 220.
91 Ibid., 228.

she should not judge both that she has a reason to φ and that she has no reason to φ. It is because the agent is committed to norms like this that it is possible to assess her reasons in light of each other. Furthermore, is is because the agent is committed to norms like this that she can regard a normative judgment as *normative*. Normative judgments set standards of correctness. This would be meaningless in the absence of some norm demanding consistency between normative judgments. We could hardly regard X as a standard, as *normative*, if a person could hold X and hold its negation without making an error. An agent making a normative judgment must see that judgment as normative (more on this in Section II). As a result, an agent is committed to some consistency norm merely in virtue of making a normative judgment at all. We might think that this basic consistency norm is contained in or entailed by any other normative judgment just insofar as that judgment is normative.

Street puts the point in a slightly different way. She claims that those normative judgments that we must make, in virtue of making any normative judgments at all, are those that we cannot, in full awareness, violate and still count as making a normative judgment.[92] I think this claim captures the same idea. To see something as normative is to see it as setting a standard. We cannot see something as setting a standard without being in some way responsive to the norms which permit it to function as a standard.

Street thinks that there are two kinds of norms that fit this description. So, she thinks that there are two kinds of constitutive norms of normative judgment. They are instrumental standards and, what I will call, standards of consistency. Street claims that an agent who judges that she has conclusive reason to adopt end E and, at the same time, in full awareness, judges that she has no reason to take what she recognizes to be the necessary

92 Ibid.

means to E, fails to make genuine normative judgments.[93] Thus, some principle of instrumental rationality must be a constitutive norm of normative judgment. As discussed above, Street claims that we also cannot judge that X is a reason to Y and, at the same time, in full awareness, judge that X is not a reason to Y. Nor can we judge that only facts of some kind K are reasons to Y, judge that Z is not a fact of kind K, and at the same time, in full awareness, judge that Z is a reason to Y.[94] So, there must be a constitutive norm of normative judgment that reflects each of these facts. (These would be among what I have called "standards of consistency.")

Because these are the only two kinds of norms that Street takes to be constitutive of practical reasoning, on her view, which normative judgments are correct for a given agent (and thus which reasons an agent has) is going to depend, in large part, on what particular normative judgments he makes in the first place. For example, an agent who frequently takes himself to have reasons to promote the well-being of others is going to wind up with a very different set of correct normative judgments than will an agent who takes himself to only have reasons to promote his own well-being. Street thinks that there will not in fact be significant discrepancies between agents' most basic reasons and values, but this is a contingent matter—a result of our common evolutionary history.

Once we have these constitutive norms we can assess the normative judgments that an agent actually makes against each other and it becomes possible for an agent to be in error about what reasons he has. When an agent becomes aware that some combination of his normative judgments violates one of the constitutive norms, he should abandon some of

93 Ibid., 227.
94 Ibid., 229.

those judgments to dispel the conflict.[95] When there are no longer any conflicts of this kind within his set of normative judgments, those normative judgments are correct and the agent in fact has the reasons that he judges himself to have. Before the process is complete (which is, in all likelihood, most of the time), we can refer to the reasons that an agent would take himself to have, once the process were complete, as the reasons he actually has. Those considerations that he now takes to be reasons, but would not take to be reasons after completing the assessment process, he mistakenly takes to be reasons. Thus, although the correctness of an agent's normative judgments is determined by his normative judgments, he can nonetheless make a false normative judgment. When he does so, the false normative judgment is the result of a failure to sufficiently attend to the relations that hold between his relevant normative judgments.[96]

II. "Counting in favor of"

Before turning to my criticism of Street's view, I will say a bit more about how exactly she identifies the constitutive norms of practical reasoning and what exactly she takes them to explain (i.e., what she takes a reason for action to be).

Street starts with an unreflective experiential notion of taking something to be a reason to φ. She claims that we have a particular kind of experience of things in the world as "counting in favor of", "calling for", or "demanding" certain responses (actions) on our parts.[97] It is an experience of being moved to perform some action by our recognition that

95 Street thinks that, in cases of conflict, the agent should give up those judgments that are less deeply his. This is "a function of how strongly [he] holds the normative judgments in question and how close the center of her total web of normative judgments they lie" (234-5).
96 Ibid., 229-30.
97 Ibid., 239-40.

some state of the world demands so acting. For example, I might be moved to get out of the way by my recognition of the fact that a car is speeding towards me.[98] This "normative experience" is supposed to be something we can all recognize from our own lives and distinguish from other kinds of mental states. E.g., while the normative experience involves representations of the state of the world—e.g., that the car is speeding towards me—it is unlike mere belief in that it is motivating.

In particular, we are supposed to be able to recognize the normative experience as distinct from other kinds of motivating states like mere desire. For example, Street claims, a Civil War soldier facing a leg amputation might have two motivational experiences of distinctly different kinds: First, he has no desire that his leg be cut off. In fact, he probably desires that it not be cut off and is subsequently moved to resist the operation. Second, he recognizes that amputation is "called for" by the threat of gangrene and is motivated to permit the operation by this recognition. The difference between these two states, Street claims, is that the latter state (the reason-taking state) comes with other commitments that are reflected in the agent's motivational tendencies. For example, the reasons-feeling with respect to life (taking myself to have reason to live) seems to force the reasons-feeling with respect to amputation (the means to life) under the right circumstances (recognition of amputation as the necessary means to life). A desire to live, on the other hand, does not seem to force any desire for amputation, whatever the circumstances.[99]

98 Ibid, 240.

99 Ibid., 230-1. Street attributes the example to Korsgaard ("The Normativity of Instrumental Reason," 238). I think this example may be somewhat misleading. The agent's motivation against having his leg cut off need not be entirely unresponsive to normative concerns. (At least under normal circumstances, there are lots of reasons not to want one's leg cut off.) And given that the agent *is* moved to have his leg cut off (he must be if he takes himself to have a reason to do so), it is not so clear that we would ordinarily say that he lacks a desire to have his leg cut off. So, as in Chapter 1, I suggest we take the case to make a conceptual point. Whether or not we frequently (or ever) experience motivation that is in no

The constitutive norms of practical reasoning, on Street's view, articulate the commitments that distinguish normative judgments from other motivational states. She takes a constitutive norm of normative judgment to be a norm such that if, on a particular occasion, an agent, in full awareness, fails to be moved to comply with the norm, then she fails to make a normative judgment at all. Street takes this to be an analytic or conceptual point about reasons. She goes so far as to draw an analogy with the concept "parent". Just as if a person does not have a child, they are not a parent, so too if a person, in full awareness, fails to meet a constitutive norm of practical reasoning, with respect to a particular purported normative judgment, then they have not made a genuine normative judgment.[100] If this is true, we can identify the constitutive norms by using our unreflective experiential notion of taking something to be a reason to identify such failures.

So, according to the picture of practical reasons that this yields, an agent takes a consideration C to be a reason for her to φ when she is moved by C to φ and that motivation is responsive to the constitutive norms of practical reasoning (and subsequently to her other normative judgments). This is supposed to explain the sense in which when an agent acts for a reason she is moved by a consideration that she takes to count in favor of her action.

Similarly, on Street's view, a consideration C *is* a reason for an agent A to φ when either (1) A takes C to be a reason for her to φ and the judgment C is a reason to φ accords

---

way responsive to our other normative commitments, motivation by a reason does seem to be characterized by this responsiveness. As an agent's motivational states exhibit less of this responsiveness, I think we become less and less inclined to think they are reasons-takings (or at least non-deviant reasons-takings).

100 Ibid., 227-8. This is over-simplified. There might be cases where an agent makes a normative judgment that, when combined with another of her normative judgments, violates a consistency requirement. She might be fully aware of this conflict and yet fail to alleviate it because she does not know which judgment to give up. We might still count her as responsive to the constitutive norms (and thus count her as making normative judgments, albeit incorrect ones) as long as she regards this state of conflict to be a problem to be resolved.

with the constitutive norms of practical reasoning (and her other normative judgments) or (2) "C is a reason (for A) to φ" follows from A's other normative judgment and the constitutive norms.[101] This is supposed to explain the sense in which considerations do in fact count in favor of actions (when they do).[102]

III. The Problem

In what follows, I will attempt to show that Street's account of normative judgment fails to fully account for the very thing that she sets out to explain—the unreflective experiential notion of taking a consideration to count in favor of an action. I will do this by showing that the constitutive norms she proposes do not account for the sense in which we take reasons to "count in favor of" actions. If so, then an agent's taking some consideration C to "count in favor of" φ-ing must be accounted for in terms of C's motivating the agent to φ. However, if we attempt to do this, we either (1) lose a meaningful distinction between

101 Technically, this will only be true once an agent has resolved all conflicts within her normative judgments.

102 At this point, it should be clear that Street holds the second, third, and fourth commitments of a constructivist account of reasons for action discussed in the previous chapter. (I think that it is also clear that she must take on the first commitment. Why this is so will become clear in my criticism of her view in Section III.) On Street's view, an agent's judgment that some consideration is a reason to perform some action φ is correct when this judgment accords with the constitutive norms of practical reasoning and the agent is responsive to these norms in making it. So, Street is committed to **C2**. In claiming that all normative assessments must be made within the practical point of view of the agent to whom they apply, Street seems committed to **C4**.

C3 states that the constitutive norms of practical reasoning determine mere eligibility to be a reason and in many cases rule multiple features of actions eligible to be made reasons. On Street's view, the constitutive norms of practical reasoning are only some consistency requirements and an instrumental principle. As such, they permit a wide variety of sets of normative judgments. Once an agent has made some other normative judgments, these other judgments further constrain the sets of normative judgments that are eligible for her. Some particular normative judgments, when combined with the constitutive norms, could even rule only one set of normative judgments as eligible. Suppose, for example, that an agent makes the normative judgment that there is most reason to do whatever action will maximize utility. This, combined with the constitutive norms will determine the agent's reasons in every case. However, the reduction in eligible sets of normative judgments, in a case like this, comes from the normative judgments that are up to the particular agent, not from the constitutive norms of practical reasoning. So Street is committed to **C3**.

being moved by a reason and being moved by a desire, or (2) we have to appeal to the idea of "counting in favor of" itself in order to preserve this distinction, and thus fail to give an explanation of this idea.

As discussed in the previous section, Street's account of normative judgment aims to distinguish between being moved by a consideration one takes to count in favor of, call for, or demand an action from being merely moved to act. Her account of normative judgment has two parts. For an agent to judge that a consideration C is a reason to φ, the agent must (1) be motivated to φ by C and (2) the corresponding judgment "C is a reason to φ" must be responsive to the constitutive norms of practical reasoning (a principle of instrumental rationality and some consistency norms). Since this responsiveness to constitutive norms is supposed to distinguish normative judgments from other motivating states, it is natural to look to the constitutive norms first for an understanding of the sense in which reasons "count in favor of," "call for," or "demand" actions.[103]

At first glance, the constitutive norms may seem to do the job. For many agents, in many cases, that some consideration C is a reason to φ follows from some subset of his normative judgments and the constitutive norms. When an agent recognizes this, he may feel that φ-ing is called for by C and can cite those other normative judgments as an

103 We might have thought that the mere presence of a motivating consideration does the job of distinguishing motivation by a reason from mere desire. When an agent acts for a reason there is some consideration—some proposition about the action or its effects—that moves her to act. When she acts from a mere desire there is no such consideration. She is simply and directly moved to act. I think that this distinction is not tenable. When an agent acts from mere desire, she still represents her action to herself in some way, and she is moved towards the action under that description. So, this representation seems to pick out a consideration that moves her to act. Perhaps, in many cases of acting for reasons, the agent represents her action one way (e.g., going to the grocery store) and the consideration that moves the agent is something further (e.g., that I can get milk at the grocery store). But this does not seem to be necessary for acting for a reason. Doing something for its own sake can still be acting for a reason. I might watch a sunset because doing so is watching a sunset or help someone in need because doing so helps someone in need.

explanation of this judgment. For example, if I take the fact that there are carrots at the grocery store to be a reason to go to the grocery store, this reason-taking might follow from a number of my other reasons-takings. I might take myself to have reasons to preserve my eyesight (which I know that eating carrots will help me do). My eyesight is useful for the pursuit of a variety of my ends which, in turn, may be supported by reasons. E.g., eyesight is useful for reading. I might take the fact that reading increases knowledge to be a reason to read. And so on. If I make all of these normative judgments, then it will follow from them, plus the instrumental principle, that I should take the fact that there are carrots at the grocery store to be a reason to go there.

So, we might think that I take a reason to φ to count in favor of φ-ing insofar as I take the corresponding judgment (C is a reason to φ) to be required by my other normative judgments, including the constitutive norms (and find myself at least somewhat motivated to φ because of this). However, while Street thinks that, when an agent judges that he has a reason to φ, this judgment should be in compliance with the constitutive norms and consistent with his other normative judgments, she does not claim that this judgment must *follow* from other of the agent's normative judgments.[104] That is, Street permits agents to make substantive additions to her set of normative judgments. Thinking about what reasons she has need not merely be a matter of thinking about what follows from the normative judgments she has already made.[105] But, if an agent can correctly judge that a consideration counts in favor of φ-ing in cases where that judgment is consistent with, but does not follow

104 I will consider what a view that did take on this commitment would look like in Section IV of this chapter.

105 We might think this allowance is necessary if practical reasoning is ever to get off the ground. Otherwise, we need an explanation of how an agent comes to have a sufficiently complete set of normative judgments, in the first place, so as to be capable of the complex practical reasoning that people actually engage in.

from his other normative judgments, then we cannot understand "counting in favor of" as following from some of the agent's other normative judgments. Reasons will count in favor of actions when this is not the case.

Nor does it seem that we can understand "counting in favor of" as merely being (or being taken to be) consistent with the agent's other normative judgments. The consistency required by Street's constitutive norms is quite minimal. If I judge that C is a reason to φ, I am not committed to thinking that I should φ or that there are not many good reasons against φ-ing. I violate the consistency norms only if I also make judgments like "C is not a reason to φ", "No facts of kind K (of which C is a member) are reasons to φ", or "I have no reason to φ". Similarly, with respect to the principle of instrumental rationality, if I judge that I have conclusive reason to adopt end E, I am not committed to taking the necessary means to E at a given time. I might have good reasons not to, at least at that time. In fact, it may turn out that, at every particular time, I have good reasons not to take the means to E. All I am committed to, in adopting E, is having *a* reason to take the necessary means to E.[106]

So, the kind of consistency required by the constitutive norms is not necessarily difficult to achieve. An agent who makes a lot of normative judgments, especially if many of those normative judgments concern the relations between other normative judgments, (e.g., judgments regarding how to weigh conflicting values), as well as the agent who makes judgments that all normative judgments should be grounded in some particular thing (e.g., utility), may find the consistency norms quite demanding. On the other hand, an agent who

106 It might be that if I take myself to have conclusive reason to adopt E at time t and end E can only be accomplished at time t, then I must take myself to have conclusive reason to take the necessary means to E at time t. However, in the case of ends the pursuit of which extends over some length of time and must be integrated with the pursuit of other ends, no similarly strong claim can be made.

makes few normative judgments, or at least few normative judgments about the relations between his other normative judgments, and who makes no judgment that all normative judgments should have some particular ground, is not likely to find it difficult to comply with Street's constitutive norms. This latter agent's normative judgments will not, together, rule much out.

That compliance with the constitutive norms need not be difficult is not itself significant. However, thinking about the agent for whom compliance will be easy provides a clear illustration of why the constitutive norms cannot themselves account for the sense in which reasons count in favor of or demand actions. Suppose, for example, that I make two sorts of normative judgments: I take myself to have reasons to stay alive and I take myself to have various moral reasons (not to lie, not to harm others, to help others, etc.) I do not judge that I should only do what I have these reasons to do. But, at least so far, these are the only activities that have ever struck me as called for. My set of normative judgments is thus consistent with my taking myself to have a wide range of reasons for a wide range of actions. Such an agent can recognize that taking any of several considerations to be reasons to φ would be consistent with her other normative judgments. She can recognize that several considerations are eligible to be normative judgments. Yet, this provides us with no reason to think that she will take any of those considerations to call for φ-ing. The mere fact that taking C to be a reason to φ would be consistent with her other normative judgments will not move an agent one way or the other with respect to the question of whether to φ.

Nor should it. The fact that taking C to be a reason to φ would be consistent with an agent's normative judgments might be a necessary condition on C's counting in favor of φ-

ing, but it leaves open the question of whether C counts in favor of or against φ-ing (or if it does not count in either direction with respect to φ-ing). And if it does not address this question—if it doesn't itself add any additional *positive* force to C—then it cannot itself show that C calls *for* or demands φ-ing. So, counting in favor of cannot merely be a matter of being consistent with an agent's other normative judgments. An agent cannot take a consideration to count in favor an action insofar as she takes doing so to accord with the constitutive norms. If this is the case, then Street's constitutive norms cannot, at least on their own, explain the sense in which a consideration counts in favor of φ-ing when it is a reason to φ.

If Street's constitutive norms cannot explain the sense in which taking a consideration to be a reason to φ is taking it to count in favor of φ-ing, we must look to the motivating aspect of normative judgments to understand "counting in favor of". If reasons do not count in favor of actions insofar as they comply with the constitutive norms, then they must count in favor of actions insofar as they motivate the agent. After all, when a number of possible normative judgments are consistent with an agent's normative judgments, then an agent that is responsive to the constitutive norms can regard them all as eligible. The constitutive norms will provide no way for the agent to decide which normative judgments to make, and which reasons he has. We need some further explanation of how some of these eligible considerations go from being merely eligible to be reasons to counting in favor of or demanding actions. We need something else to explain how a consideration becomes something to be said *for* an action. On Street's view, if the constitutive norms cannot explain this, there is only one other option. The eligible

considerations that count in favor of or demand actions are the ones that actually motivate the agent.[107] And because the constitutive norms could not themselves explain the sense in which these considerations demand actions, it looks like Street must claim that if I take C to count in favor of φ-ing, then the counting in favor of appears to amount to my being moved by C to φ.

At this point, Street's view seems to be faced with a dilemma. Either the motivation present in acting for reason (apart from responsiveness to the constitutive norms) is mere[108] motivation—motivation of the same kind as that present in a desire—or it is some other kind of motivation. If being moved by C to φ, (apart from responsiveness to the constitutive norms) is mere motivation, then we seem to have already ruled it out as explaining the sense in which agents take reasons to count in favor of actions. Street presented this sort of mere motivation as the contrast class to motivation by a reason. She describes the task of explaining motivation by a reason as the task of distinguishing taking a consideration to count in favor of an action from merely being moved by it to act. So, if we must now appeal to the same sort of motivation present in a desire to do the work of explaining "counting in favor of," we will have failed in this task. On the other hand, we might think that being moved by C to φ, when acting for a reason, apart from responsiveness to the constitutive

107 This reveals Street's account to be committed to the first commitment of constructivist accounts of reasons for action discussed in Chapter 2. Because the constitutive norms determine mere eligibility to be a normative judgment, we must appeal to something else to explain why an agent makes some eligible judgments rather than others. The only available explanation seems to be that she is simply moved towards some and not others. Because Street takes the constitutive norms to account for the normativity of reasons, this motivation must be non-normative.

108 I say "mere" motivation here to suggest that this motivation contains no recognition of anything normative. This is not to deny that there might be other interesting differences between ways of being motivated and that many ways of being motivated or kinds of motivation may be quite complex. Street, however, is concerned with the difference between normative motivation and other motivation, so we need not worry about these additional differences here.

norms, is a distinct kind of normative motivation. This would succeed in distinguishing motivation by a reason from mere motivation, but if we go this way, then Street's view is not giving us an explanation of normative motivation. She would merely be saying that, when an agent acts for a reason, she takes a consideration to count in favor of acting and is responsive to the constitutive norms. The question of what taking a consideration to count in favor of acting amounts would remain.

I am not going to say much about the second horn of this dilemma. It seems to leave open two possibilities for Street's view. We might take it to suggest that the notion of counting in favor of is primitive and can be given no further explanation. Street's view would then serve simply to identify some principles of reasoning, responsiveness to which is necessarily connected to (but does not explain) having the attitude of taking a consideration to count in favor of an action. On the other hand, we might think that the need to appeal to the idea of counting in favor of, on Street's view, merely shows that her view is incomplete. She may not have given us a complete list of the constitutive norms. If the list were complete, the constitutive norms would themselves specify the sense in which reasons count in favor or, call for, or demand actions. I will discuss this possibility in greater detail in Section V. First, however, I will return to the first horn of the dilemma

IV. The Problem revisited

Facing the first horn of the dilemma, I think it worth considering the possibility that the problem only arises because I have been making "counting in favor of" out to be more than it is. Despite everything I have said, Street's view does succeed in making *a* distinction

between acting from mere desire and acting for a reason—an agent acts for a reason in just those cases where she is moved by to act by a consideration and is responsive to the constitutive norms. Furthermore, this distinction takes into account a basic idea that an appeal to "counting in favor of" in describing practical reasons is supposed to capture. Motivation by a reason is normative—it involves some concern for what is appropriate or correct. The constitutive norms represent a standard of correctness. (And this standard of correctness does seem to apply to practical reasoning, whether or not we think there is more to correctness in practical reasoning than it delineates.) So, responsiveness to the constitutive norms does capture *some* concern for what is appropriate or correct.

In my criticism of Street's view, I suggested that (at least part of) what is normative about a reason must be that it is in *favor* of an action, that it is something *positive* to be said for performing an action. Or to put the point another way: The constitutive norms may determine that C is eligible to be taken to be a reason to φ, but leave open the question of whether C counts in favor of or against φ-ing (or whether C is neutral with respect to the question of whether to φ.) Whatever determines the answer to this question—whatever makes it the case that a consideration counts in favor of rather than against an action (or says nothing either way)—must be (and seem to the agent to be) normative. I took this to be an important part of what the claim that a reason counts in favor of, calls for, or demands an action is meant to capture. However, we might take Street's view to suggest that this is too strong. If all we need to do, in order to pick out the line between normative motivation and mere motivation, is to show how normative motivation exhibits some concern for some consistency norms (that is plausibly present when an agent acts for reasons), then perhaps

the notion of counting in favor of that I am appealing to in my criticisms of Street's view is simply overblown.

In this section, I will argue that this minimal conception of normative motivation does not seem to capture the unreflective experiential notion of taking a consideration to count in favor of acting that Street is after. We can construct cases in which I think it is clear that an agent will not have the familiar unreflective experience of taking a consideration C to count in favor of φ-ing even though she is moved to φ by a consideration C and the judgment that C is a reason to φ is consistent with her other normative judgments (and so should count as a normative judgment on Street's account). Suppose I happen to note that the grocery store across the street has ten aisles and find myself motivated by this fact to go to the grocery store. Suppose I have no further story to tell (in light of my other normative judgments, desires, or anything else) about why the fact that the grocery store has ten aisles counts in favor of my going there. (E.g., it is not the case that I have noted the grocery stores with more than six aisles tend to sell my favorite chocolate.) Suppose that the judgment "the fact that the grocery store has ten aisles is a reason for me to go to the grocery store," is both consistent with my other normative judgments (including the constitutive norms) and responsive to the constitutive norms (if I recognized a conflict, I would be moved to dispel it). I do not, for example, also judge that I should only go to the grocery store when I have completely run out of food. Street would have to say of this case both that I have taken the fact that the grocery store has ten aisles to count in favor going to the grocery store and also that this normative judgment is correct (in this case). The fact that the grocery store has ten aisles is a reason for me to go to the grocery store.

Looking at the case independently of Street's view, it is, I think, fairly obvious that the fact the grocery store has ten aisles does not count in favor of going to the grocery store, and that, although this fact may (conceivably[109]) move me, I will not regard it as counting in favor of, calling for, or demanding that I go to the grocery store, even if I am responsive to the constitutive norms. Or, at least, this is obvious if we insist that there is no further story to be told about the case. In the absence of such a story, it seems to make no difference whether I go to the grocery story in light of the fact that it has ten aisles. That is, there seems to be an answer to the question of whether this consideration counts in favor of, against, or is neutral with respect to going to the grocery store: it is neutral (and, therefore, it is not a reason for going). If I know that this is the answer, then, while I may be moved by the fact that the grocery store has ten aisles, I do not think I can intelligibly (in full awareness) take this fact to count in favor of, call for, or demand acting. I might go along with the motivation, perhaps in the knowledge that doing so creates no conflicts with my normative judgments, but it does not seem that I would take it to be a reason to act.

If there is anything I take to be said for this action, it seems to be that I am moved to do it and not that the store has ten aisles. We might think that "I feel like it" (or something like this) is reason to act, at least in some cases. (Perhaps this is how we should understand action done on a whim.) But this does not yet explain why some *other* consideration, i.e., that the grocery store has ten aisles, counts in favor of going to the grocery store. On the

109 I grant that there is some reason to doubt that such a thought could really move me to act (absent some further story), but if it turns out that there are constraints on what sorts of considerations can intelligibly (as opposed to merely familiarly) be taken to move an agent to act that are not captured by Street's constitutive norms, then this could actually support the idea that she has not adequately captured the idea of "counting in favor of" that is present in acting for a reason. It seems likely that the reason that we doubt that such a thought could really move me is that we think that I am generally moved by what counts in favor of actions (or by what I take to count in favor of actions).

other hand, we might think that the fact that I am moved towards some action by some consideration makes it the case that that consideration counts in favor of that action. Both of these possibilities involve normative judgments that a particular agent could make and thus be committed to following: e.g., "that I feel like it is a reason to do what I feel like doing" or "whenever I am moved to act, I have a reason to act." However, on Street's view, neither of these are normative judgments that an agent *must* make, so we cannot assume that the agent in the ten aisles case makes them. As a result, it remains difficult to see how this agent could regard the fact that the grocer store has ten aisles as calling for or demanding going to the grocery store.

Of course, the ten aisles case is strange and not likely to actually arise. It is at least unusual for an agent's reasons to stand apart from each other in the way that this example does. Usually, when I take a consideration to count in favor of an action, I have some story that appeals to my other reasons that would seem to fill in what is missing in this example. Suppose I am in search of my favorite chocolate. I have a variety of reasons for doing this. It is tasty. It has some health benefits. If I use it as a reward, I might be able to motivate myself to finish a chapter of my dissertation. In light of these reasons, and the fact that grocery stores with more than six aisles are more likely to carry it, the fact that the grocery store across the street has ten aisles may begin to seem to count in favor of my going there.

So, we might think that there is a relatively easy fix to Street's view that can rule out obviously implausible cases like the first grocery store case above. We might add to her list of constitutive norms a norm that requires a normative judgment to be not merely consistent with the agent's other normative judgments, but also to be supported by those other

normative judgments. It will then be in the justification for any given reason-taking that appeals to the agent's other reasons-takings that we will be able to see how the agent can regard the relevant consideration as counting in favor of or demanding the action.[110]

While an agent's normative judgments do tend to stand in these sorts of justificatory relations to each other, I think it unlikely that adding a principle demanding this to the constitutive norms of practical reasoning will solve the problem for Street. To see this, let us think about how one of these justifications in terms of the agent's other normative judgments would end. Either it must bottom out in a normative judgment that itself has no further justification, or the justification must circle back on itself. In the former case, this foundational normative judgment will be subject to the problem I raise for Street's view. The constitutive norms will not always rule one way or the other on whether the consideration it picks out counts in favor of an action. There will be no story that appeals to the agent's other normative judgments to reveal the consideration as calling for the action. It just moves the agent. Thus, this case is just like my taking the fact that there are ten aisles in the grocery store to be a reason to go there.

Admittedly, the normative judgments in which such justifications are likely to bottom out will not seem strange in the way that the original ten aisles examples does. Suppose, for example, that I take the fact that something is tasty to be a reason to eat it (or perhaps more broadly, the fact that an experience is pleasurable to be a reason to pursue it), and I have no further reasons-takings that support this one. This is a familiar normative judgment. We probably do not doubt that this consideration counts in favor of actions.

110 Street herself may have something like this in mind. She talks of agents having webs of normative judgments and determining which normative judgment to give up in cases of conflict by thinking about how central they are to this justificatory web. ("Constructivism about Reasons," 235.)

However, on Street's view, it is in fact no different from the original ten aisles case. If the original ten aisles case contains no reason-taking, but this more familiar case does, there must be something else to be said about why this consideration counts in favor of acting that can distinguish the two cases. If Street's account was sufficient, the original ten aisles case should be unproblematic as well.

On the other hand, suppose that an agent's set of normative judgments contains no foundational normative judgments. Instead, his set of normative judgments is such that every particular normative judgment gets some support from some other normative judgments. His set of normative judgments forms a justificatory web. For simplicity, I will say that such a set of normative judgments is "coherent." (By contrast, a set of normative judgments that is merely "consistent" is one in which there are no normative judgments that conflict with other normative judgments, but where there are normative judgments that receive no support from other normative judgments.)

Nevertheless, I think it plausible to assume that there are a wide variety of possible coherent sets of normative judgments. Demanding that an agent's set of normative judgments be coherent adds a restriction on the kinds of relations that these normative judgments can have to each other, but it tells us nothing about what the content of those normative judgments should be. (At least in the absence of a rather impressive argument to the contrary, we have no reason to think that there is only one possible coherent set of normative judgments.) Thus, the constitutive norms of practical reasoning that Street proposes rule a variety of *sets* of normative judgments as eligible for an agent to make. The constitutive norms then fail to make it the case that any particular set of considerations

counts in favor of actions. Which set of reasons an agent has will then depend on which of these sets of considerations happens to move him.

I think this suggests that the problem for Street's view can still arise globally for the agent whose set of normative judgment is coherent. If the constitutive norms provide no correct way to choose between two coherent sets of reasons, then I must appeal to the fact that one happens to move me in order to explain why I have those reasons. This makes it seem to me that the motivating force behind particular reasons is just this mere motivation. The fact that I have a coherent set of normative judgment reveals what merely motivates me to be complex. However, the fact that any given reason in this set receives some support from other reasons that are part of this set that merely motivates me does not seem to get me out of the realm of mere motivation and into seeing these considerations as counting in favor of actions. In fact, if I get caught up thinking about the support my reasons offer for each other, and begin to think that they call for or demand actions in a way that considerations that merely move me do not, I need only remind myself that there is nothing to be said in favor of the whole complex of motivation, as opposed to other possible coherent complexes, and I will have trouble regarding my reasons as demanding actions (where this means something more than that they move me to act).

V. Three options for explaining "counting in favor of"

If Street's view is, in this way, unable to account for the sense in which a reason counts in favor of, calls for, or demands an action, we seem to be left with (at least) three options to explain this.

The first is to adopt a version of R—the claim that a consideration counts in favor of an action insofar as it picks out something about the action that is valuable. That is, we could identify values—features of actions or their effects that are to be promoted, preserved, or respected—that reasons pick out. An agent would take a consideration to count in favor acting insofar as it picks out something about the action that is to be promoted, preserved, or respected. That the action has this feature that is valuable provides something to be said in favor of the action. This account would provide a way to answer the question of whether a consideration counts in favor of, against, or is neutral with respect to φ-ing. That the action has a valuable feature counts in favor of, calls for, or demands action because the feature is to be promoted, preserved, or respected. As discussed in Chapter 2, this claim could be consistent with some aspect of a constructivist strategy for explaining practical reasons. We might try to specify a constitutive norm (or norms) that identify what is to be promoted, preserved, or respected and direct an agent's reasons-takings to reflect this.

However, at this point, there would seem to be a much simpler strategy available to solve Street's problem that requires no investigation into what is valuable. In objecting to Street's view, I have suggested that if an agent notes the fact that there is nothing to be said in favor of one set of normative judgments rather than another, and subsequently notes the fact that which normative judgments he makes depends on which considerations happen to move him, then we cannot plausibly understand him as taking these considerations to count in favor of actions (as opposed to merely moving him). Faced with this claim, the reader may have thought that the fact that the agent actually has one of many eligible sets of normative judgments makes a normative difference. It may be true that I merely happen to

be moved by the considerations that move me, but we can account for the sense in which they count in favor of actions by appeal to the idea that these are the considerations that do actually move me, that are *mine*.

There are a couple of ways to do this. The first is to claim that the fact that a consideration moves an agent to φ is itself be something to be said in favor of φ-ing (or, if not itself something to be said in favor of φ-ing, it grounds the fact that the consideration is something to be said in favor of φ-ing). On this kind of view. it may be true that there is nothing to be said in favor of being moved by one particular consideration rather than another, but once you are moved, to put the point crudely, you ought to do what you are moved to do, and the fact that you are moved explains why you ought to do it (other things equal). (This is crude, in part, because "ought" is likely too strong. You may well make no error in not doing what you are moved to do.) Despite explaining the sense in which reasons count in favor of actions in terms of "mere motivation," this view allows us to preserve a distinction between a considerations merely moving an agent to φ and an agent's taking a consideration to count in favor of φ-ing: An agent can be moved to φ by a consideration. Or, an agent can be moved to φ by a consideration in some recognition that φ-ing is appropriate, given that she is so moved. I will consider a view of roughly this kind (Mark Schroeder's view) in Chapter 4.

The other option is to grant that motivation by a consideration itself (perhaps given compliance with and responsiveness to the constitutive norms) is not enough to make that consideration count in favor of acting, but to claim that the agent can bridge the gap between a consideration's permissibly motivating an agent to φ and its counting in favor of φ-ing

through an act of will that is responsive to the constitutive norms. We could claim that an agent *makes* a consideration count in favor of acting when it moves her to act and she chooses to act on it in proper deference to the constitutive norms. Korsgaard takes this position. I think that arguments presented in Chapter 2 can be used to show that this last option will not work.

In the remainder of this chapter, I will explain why. The basic idea is that, in order to get a notion of "counts in favor of", "calls for" or "demands" action that fits with the unreflective experience of acting for a reason, we will need an account of reasons that allows for a requirement to share reasons in the stronger sense. Chapter 2 purported to show that a constructivist view can only do this if it includes norms that specify values in the constitutive norms of practical reasoning.

VI. S-sharing is required part of the unreflective experiential notion of a reason for action

I think that the arguments in favor of a requirement to share reasons in the stronger sense (Chapter 2, Section II) suggest that an account of reasons for action must include a requirement to share reasons in the stronger sense if it is to provide a sense in which a reason "counts in favor of", "calls for", or "demands" an action that fits with the unreflective experience of acting for a reason. That is, in order to capture the unreflective experience of acting for a reason, we need to regard a reason as the sort of thing that is to be shared in the strong sense. We need to regard a reason for action as picking out a feature of that action that makes claims on agents because of what it is and not merely because some rational agent has permissibly made it her purpose and acted for its sake.

In this section, I will explain why I think that a requirement to share reasons in the stronger sense is necessary to capture the unreflective experience of acting for a reason. If I am right about this, it supports there being such a requirement. In Chapter 2, I aimed to show that a constructivist view can only allow for a requirement to share reasons in the stronger sense if it includes norms that specify values in the set of constitutive norms of practical reasoning. So, if this requirement is necessary to capture the unreflective experience of acting for a reason, this suggests that the problems I raise for Korsgaard's constructivism are quite serious and general. In failing to include a requirement to s-share reasons, an account of reasons fails to account for a very basic feature of the unreflective experience of acting for a reason, and thus, we might think, fails to account for a very basic feature of reasons. So, if I am right about the problem for Korsgaard's account, there is strong reason to think that a plausible constructivist account of reasons must include norms that specify values in the constitutive norms of practical reasoning.[111] Such a view will account for the sense in which a reason counts in favor of, calls for, or demands an action in terms of these features of actions that are to be promoted, preserved, or respected. A reason will count in favor of acting because it shows that action to have some feature that is to be promoted, preserved, or respected. This eliminates the need to appeal to an act of will to *make* a consideration count in favor of an action (except in the limited way suggested in Section VI of Chapter 2), and thus eliminates the need to appeal to an act of will to account

111 If I am right, then I think that talking about a "requirement to share reasons" is a bit misleading. It suggests that agents need to pay attention to what reasons other agents have and be sure to have them too. But this is an artifact of the kind of constructivist view I am rejecting. If agents create reasons by choosing (in the right way) to act for the sake of desired features of action, then if reasons are to be shared, agents will have to pay attention to what reasons others have in order to know which they must share. On the kind of view I am proposing, if two agents fail to share basic reasons, one of them is likely in error. But this will be because this agent has failed to recognize which features of actions call for acting and not because her reasons do not match the other agent's.

for the sense in which reasons count in favor of, call for, or demand actions.

The three arguments that I present in Section II of Chapter 2 in favor of a requirement to share reasons in the stronger sense all point to a difference that arises on a constructivist account between how an agent sees her own purposes as reason-giving and how she sometimes sees the purposes of others as reason-giving. This difference seems to be an implausible consequence of a constructivist view. A requirement to s-share reasons would eliminate the difference. (Unfortunately, as I argued in Chapter 2, it is not an option.)

On a constructivist view, a consideration comes to count in favor of an action when acting for the sake of that consideration is rationally permissible (it accords with the constitutive norms of practical reasoning) and an agent is moved by that consideration to act, in proper deference to those constitutive norms. When an agent acts for a reason, the consideration that moves her to act picks out some feature of her action or its consequences that attracts her to the action. The agent's interest in this feature tends to be because of something about it. (An agent who buys boots for the sake of climbing Kilimanjaro might be interested in climbing Kilimanjaro for the sake of something about it—e.g., its beautiful vistas.) Because of this, from the agent's perspective, these desired features of her actions seem to call for those actions because of something about what they are. The agent will have a story to tell about why these features call for actions and it is a story about the features themselves. (The fact that buying boots will help me climb Kilimanjaro counts in favor of buying boots because of something about what climbing Kilimanjaro would be—e.g., that it would involve the experience of beautiful vistas.)

On the other hand, on a constructivist view like Korsgaard's, an agent's purposes are

reason-giving for other agents as well as for herself. An agent A's purposes may not give another agent B reasons to perform the same actions as they give A, but they will give B some reasons—at least reasons not to interfere with A's pursuit of his purposes, and probably reasons (other things equal, which may rarely happen) to help A pursue those purposes as well. When B acts for a reason stemming from A's purposes, B is moved by a consideration that picks out A's purpose. B's interest in A's purpose may not be because of anything about the purpose itself. E.g., B may find nothing about A's climbing Kilimanjaro appealing (beauty leaves him cold, he has no particular interest in the satisfaction A will get out of the endeavor). B's interest in A's climbing Kilimanjaro stems merely from B's recognition that A's climbing Kilimanjaro is a purpose of a rational agent. So, from B's perspective, the fact that an action will promote A's climbing Kilimanjaro calls for performing that action because B happens to have been interested in climbing Kilimanjaro and (permissibly) made it her purpose.

Recall that the unreflective experience of acting for a reason reveals that a reason for an action is supposed to be normative for an agent—it is supposed to make a kind of demand on an agent. This demand is weak.[112] Many reasons do not make it the case that the agent should perform the actions that they call for. There may be other reasons that outweigh them, and at least for most reasons, at least in most particular circumstances, it may well be permissible not to act on them even in the absence of any reasons for not acting. At any given time, I have a variety of reasons to perform a variety of actions. (At any given time there are probably a lot of worthwhile things that I might do.) If I do nothing, even in the

112 I am describing this demand as quite weak for the sake of argument. If it turns out to be true that, for example, we *should* always do what the balance of reasons calls for doing, this will not be a problem for my arguments.

absence of reasons to do nothing, it is not clear that I have made any error. Nevertheless, a reason remains something normative and positive to be said for actions. I should take it into account in deciding what to do, especially when I am moved by reasons that oppose them. (And perhaps if I never act for the sake of a particular reason-giving feature of action, e.g., beauty, and do not have a long list of reasons why it was never a good idea, I have made an error.)

Each of the arguments in favor of a requirement to s-share reasons suggests that the explanation available to an agent B who does not s-share reasons with agent A for why the fact that an action will promote A's purposes calls for performing that action seems like the wrong sort of explanation for this demand. That is, they raise doubts that B will regard purported reasons of this kind as having a genuine claim on her—as counting in favor of, calling for, or demanding actions. (And thus, they raise doubts that the explanation of these reasons captures the unreflective experience of acting for a reason.)

This is easiest to see in the second argument. The basic idea here was that, when an agent acts to pursue a permissible purpose, he creates a demand that other agents take his purpose to be reason-giving. They must now take themselves to have reasons to help him pursue this purpose (or at least to curtail their own pursuit of their own purposes so as to not interfere with his pursuit of this purpose). If the agent can cite nothing about his purpose (aside from the fact that he has permissibly made it his purpose) that other agents have made their purposes, then it looks like he would not be able to provide a satisfying justification for this demand to the agents that are supposed to be subject to it. In the absence of this justification, it does not seem that there can be a genuine demand. In Chapter 2, this was

presented as support for a requirement to s-share reasons. S-shared reasons would provide the missing justification and allow for the demand.)

Alternatively, we might deny that there are any such reasons provided by the purposes of others for those agents who fail to find those purposes interesting independent of the other agents' interest in them.[113] (Street does this.) However, aside from being a rather large bullet to bite, I do not think that this will solve the problem. On a constructivist view, where the constitutive norms of practical reasoning are norms of permissibility, *all* reasons in fact count in favor of actions insofar as they pick out a rational agent's (permissible) purposes. The normativity of a reason for action stems from the fact that it is acted upon (in the right way) by a rational agent. That is, to become a reason to φ, a consideration that picks out a desired feature of φ-ing, and for which it is permissible to φ, comes to count in favor of φ-ing when an agent makes it her purpose in acting (in proper deference to the constitutive norms of practical reasoning that determine on which considerations it is permissible to act). This means that the explanation of why a particular feature of action calls for an action from the point of view of an agent (the one that appeals to what that feature is) is illusory.[114] The desired features of an agent's actions that her reasons pick out do not count in favor of her actions because of something about what they

113 We might still happen to be interested in certain particular agent's successfully pursuing their purposes, and get some reasons to help those particular agent's this way. For example, I might be interested in my sister's pursuit of her purposes without finding anything about the purposes themselves interesting, and yet have that interest in her purposes be based in something other than the mere fact that another rational agent has made them her purposes. My interest in her purposes is based in the fact that someone I happen to care about has made them her purposes.

114 There might be partial explanations of why a particular feature of action counts in favor of performing it that do appeal to what that feature is, but the explanation will ultimately bottom out in an appeal to the fact that some rational agent made this feature of action her purpose. E.g., the fact that buying boots will promote my climbing Kilimanjaro might count in favor of buying the boots because of Kilimanjaro's beautiful vistas, but this is only the case because some rational agent (permissibly) made experiencing beautiful vistas their purpose.

are. They count in favor of her actions because they are the permissible purposes of a rational agent (her).

But if this explanation of why a consideration counts in favor of an action does not seem sufficient to explain the demand involved in a reason, when the reason in question stems from another agent's purposes that do not interest me independently of being the (permissible) purposes of a rational agent, then it should also be insufficient to explain the demand involved in a reason when this reason stems from my own purposes (which do interest me independent of their being the (permissible) purposes of a rational agent). If I cannot regard the former sort of reasons as counting in favor of actions, knowing the constructivist explanation of what this amounts to, then it does not seem that I could regard the latter sort of reason as counting in favor of actions either. And thus, the constructivist view seems to fail to capture the unreflective experience of acting for a reason. And it seems to do so because it makes certain features of actions count in favor of performing those actions because some rational agent has happened to make them purposes. In the unreflective experience of acting for a reason, agents regard features of action as counting in favor of performing those actions because of something about the features themselves.

The first and third arguments in favor of a requirement to s-share reasons get at this point a bit differently. Unlike the second argument, they do not ask us to consider the demands that the purposes of others make on us when we have no interest in those purposes aside from their being the purpose of a rational agent. However, they do ask us to contemplate the two contrasting explanations discussed above of why a consideration counts in favor of action.

The first argument discussed cases where an agent is thought to be acting for the sake of a purpose such that, from our perspective, there is nothing to explain why that purpose counts in favor of acting except that it is the (permissible) purpose of a rational agent. I suggest that it is natural to think in this kind of situation that someone has made a mistake about what is worthwhile. That is, he need not have made a mistake about what is permissible, but he has made a mistake in thinking there is something to be said for this action. We think the agent has made a mistake, insofar as he takes himself to have reasons to act. On the other hand, he likely thinks that we are making a mistake in failing to see what it is about his purpose that calls for action. (I suggested this as a liberal interpretation of the claim that an agent would take offense at our judgment.) But we all seem to agree that if it were the case that the only thing to be said for acting for the sake of that purpose were that it is the permissible purpose of a rational agent, then the purpose would not call for the action and the agent would not have a reason for acting.

The third argument takes a reason to be an intelligible answer to a particular why question and suggests that an answer that (as far as the recipient can tell) merely shows that an action promotes the permissible purpose of a rational agent does not seem to be an intelligible answer to this question. (Or at least, it tends to prompt further why questions that ask the agent to explain why she made it her purpose). This is supposed to show that, in our ordinary use of reasons, we take them to count in favor of actions because of something about the features of actions that they pick out and not merely because a rational agent made those features their purposes. It seems reasonable to suppose that our ordinary use of reasons tracks the unreflective experience of acting for a reason. If there are constraints on

intelligible answers to why questions, then we might think that they are constraints on what sorts of considerations can intelligibly play the role of a reason in that unreflective experience. So, the third argument suggests that a reason counts in favor of an action because of something about the feature of actions that it picks out and not merely because a rational agent made that features her purpose.

So, I think, the three arguments in favor of a requirement to s-share reasons presented in Chapter 2 suggest that the unreflective experience of acting for a reason—the unreflective experience of taking a consideration to count in favor of, call, for, or demand acting—is an experience of taking a feature of an action to count in favor of that action because of something about what the feature is and not merely because some rational agent (permissibly) made that feature of actions her purpose. However, constructivist accounts of reasons for action that take the constitutive norms of practical reasoning to be norms of permissibility make all reasons ultimately count in favor of actions in virtue of picking out the (permissible) purposes of a rational agent. Korsgaard, tries to make a requirement to s-share reasons a consequence of the constitutive norms, and thus, I think, tries to capture the unreflective experience of acting for a reason. I argued in Chapter 2 that the account of reason-creation which explains how considerations come to count in favor of actions, on a constructivist view, rules out such a requirement.

Chapter Four

Introduction

This chapter will concern Humean accounts of reasons for action. A Humean account of reasons for action explains an agent's reasons for action in terms of her desires. On a simple Humean view, if an agent desires X (where X is some state of affairs or perhaps an action itself), she has a reason to do what will bring about X (perhaps assuming that her desire does not depend on any false beliefs). If an agent lacks a desire for X, then she may have no reason to do what will bring about X, at least given that there is no Y such that the agent has a desire for Y and bringing about X will help bring about Y. Different Humean views can identify different contents for reasons for action. For example, we might think that a reason for action is the desire for X itself, that it is the fact that the agent has the desire for X, that it is the fact that φ-ing would satisfy the desire for X, or that it is the fact that φ-ing would promote X (with no mention of the fact that the agent desires X in the reason itself).

I will not, in this chapter, attempt a refutation of all Humean accounts of reasons for action. I will not consider views that take the content of a reason for action to be the desire itself, the fact that the agent has the desire, or the fact that the action will satisfy the desire. (I agree with the many people who point out that our ordinary reasons-discourse suggests that reasons pick out facts about the world and not (usually) agents' attitudes or facts about those attitudes.) Furthermore, the arguments I have made in previous chapters, I think, rule out some other possible Humean views. Street's view, for example, is a Humean account of

reasons and my arguments against it, I think, can have rather broad application. In Chapter 3, I suggested that, if we want to account for the sense in which, when an agent acts for a reason, she is moved in part because she takes her reason to "count in favor of," "call for," or "demand" acting (as opposed to merely moving her to act), and if we want to account for this in terms of desires that are not themselves responsive to what is to be promoted, preserved, or respected, then our account of reasons for action will have to appeal to something normative other than norms of rational permissibility. As a result, it seems that, if we want to explain reasons for action in terms of desires, we will need to appeal to a claim that desire-satisfaction is to be promoted (or something like this) in order to capture the sense in which reasons for action call for actions.

So, in this chapter, I will consider a prominent Humean account of reasons for action that, I think, meets this criteria. I will examine Mark Schroeder's view that desires are background conditions on reasons, i.e., that a consideration counts in favor of an action in virtue of showing that the action will promote the satisfaction of one of the agent's desires.[115]

Humean theories of reasons are subject to a variety of traditional objections stemming from the role desires play in explaining reasons. These objections suggest that Humean accounts of reasons for action conflict with deeply held beliefs about reasons. Some Humeans simply bite the relevant bullets. On the other hand, Schroeder thinks that if we understand desires as background conditions on reasons, we can address these objections and defend a plausible Humean theory of reasons. The theory he proposes is, thus, a Humean theory of reasons without the sting, if you will. In what follows, I will discuss two of these traditional objections, how Schroeder proposes to solve them, and argue that each of

115 Schroeder, Mark. *Slaves of the Passions.*

these solutions either fails to remove the "sting" that the objection points out or is untenable in its own right.

This falls short of what we might want, in light of previous chapters. It falls short of showing that an account of reasons for action that appeals to the idea that desire-satisfaction is to be promoted cannot adequately explain the sense in which an agent regards a reason as calling for or demanding her action. This is is part, I think, because the view adopts a version of R and proposes desire satisfaction as the value that grounds reasons (see Section III of this chapter). It thus avoids the problems I have raised for constructivist views. Nevertheless, I think that insofar as my arguments are successful in showing that the view cannot escape some traditional concerns for Humean theories of reasons, it should leave us deeply unsatisfied. This gives us some reason to think that desire-satisfaction is the wrong value to ground reasons. A view that uses it to explain reasons gives up too much.

I. The basic claim: Desires are background conditions on reasons.

On Schroeder's view, a reason for a particular agent to perform a particular action is a consideration that counts in favor of that agent performing that action. This consideration is usually an ordinary fact about the world like "there will be dancing at the party" rather than a fact about the agent's desires.[116] (Thus, the sorts of considerations that are reasons are the sorts of considerations that we normally think of agents as thinking about when deciding what to do.) The agent's desires are not themselves her reasons, nor are they part of her reasons. Instead they are background conditions on those reasons being reasons. That is, the agent's desires are a necessary and fundamental part of the explanation of why a given

116 Ibid., e.g., 27.

consideration is a reason for the agent for whom it is a reason to perform a particular action. For example, that there will be dancing at the party is a reason for Ronnie to go because Ronnie desires to dance. The same consideration is not a reason for Bradley to go to the party because Bradley lacks the desire to dance.[117]

A desire explains why a consideration is a reason for a particular agent to perform a particular action when that consideration reveals a way in which that action will promote the satisfaction of that desire. More specifically, a consideration is a reason for an agent to perform a particular action when the truth of that consideration is part of an explanation of why the agent's performing that action would promote the satisfaction of one of her desires.[118] (Promotion, for Schroeder, is a very weak relation. An action promotes the satisfaction of a desire when it increases the likelihood that the desire will be satisfied relative to a baseline of doing nothing. If the desire is more likely to be satisfied if the agent performs an action than it is to be satisfied if the agent does nothing, then that action promotes the satisfaction of that desire.[119]) So, a consideration is a reason for an action when and because it explains how that action promotes the satisfaction of one the agent's desires. Thus, for any given reason, an explanation of why it is a reason—an explanation of why it counts in favor of acting—will have to appeal to the agent's desires.

II. Constitutive explanation

On Schroeder's view, an agent's desires explain why she has the reasons she has because all it is to be a reason is to be a consideration that reveals a way in which an action

117 Ibid., 1.
118 Ibid., 59.
119 Ibid., 113.

will promote one of the agent's desires. Schroeder's account of the normative is reductive in that he attempts to give constitutive explanations of normative concepts in non-normative terms. He does this by first endorsing the popular view that all normative properties and relations can and should be understood in terms of reasons (a normative relation).[120] If this is right, then the only normative concept we need to account for in non-normative terms, in order to give a reductive account of the normative, is that of a reason. Schroeder's account of reasons proposes a constitutive explanation of reasons in terms of desires. He takes desires to be psychological states that can be understood in non-normative terms. So, on his view, the basic normative concept—a reason—is explained by reference to non-normative psychological states.[121]

So, on Schroeder's view, there is one fundamental normative thing. It is a relation between an agent X, an action A, and a consideration C (i.e., C *is a reason for* X to do A).[122] The relation holds when there is some p such that X has a desire for p, and "the truth of C is part of what explains why X's doing A promotes p".[123] The agent X, the action A, the consideration C, and the object of desire p can all be understood in non-normative terms. For example, the consideration merely cites a fact about the world like that there will be dancing at the party. Similarly, the desire cites some activity or state of affairs like dancing that can be non-normatively described, as can the agent's motivation with respect to it. Neither C or p need include any reference to goodness, aptness, appropriateness, and the like. That is, neither C or p themselves contain anything normative. When C has normative

120 Ibid., 81.
121 Ibid., §4.4.
122 Ibid., 17.
123 Ibid., 59.

force—when it is a reason and counts in favor of A-ing—this is because it stands in a particular relation to X and A in light of p.

On Schroeder's view, there is no further normative thing that explains this reason relation. Its relata can be understood without reference to anything normative, but the relation itself is normative. This sort of relation between these non-normative relata, just is normative, and it is the fundamental normative thing in terms of which all other normative facts, relations, etc. are to be explained. There is nothing more to being normative than being, or being explicable in terms of, this relation.

This constitutive explanation allows Schroeder to avoid one traditional kind of objection to Humean accounts of reasons.[124] Such accounts are often taken to explain why I have reasons to do what will promote my particular desires by appealing to a reason to pursue the satisfaction of one's desires that underlies these particular reasons (or to some other normative fact, e.g., that the satisfaction of one's desires is good). But then, in order to avoid circular explanation, such views must grant that there is at least one reason (or other normative fact) that is not grounded in desires.[125] Once you accept one such reason (or other normative fact), you need some principled reason for claiming that it is the only such reason (or other normative fact). Such principled explanations are hard to come by. However, if all it is to be a reason is to be a consideration that reveals a way in which an action would promote one or more of the agent's desires, there is no need to posit a further reason (or

124 Ibid., §3.2.

125 Note that this looks rather like points I made in Chapter 3. I do think that Schroeder's solution to this problem is illuminating. A similar problem may arise for other specifications of the value that explains reasons. Although I think that Schroeder is wrong to make desire-satisfaction the value that explains reasons, the form of his solution to this problem may well solve these problems for other values as well. For whatever value V explains reasons, if all it is to be a reason is to be a consideration that reveals a way in which an action would promote V, then there is no need to posit a further reason (or other normative fact) to promote V to explain this.

other normative fact) to promote desires to explain this reason, and thus the problem can be avoided.[126]

III. How Schroeder fits into my project

All of this makes Schroeder's view significantly different from the constructivist views discussed in Chapters 2 and 3. Although these views have in common that they attempt to explain reasons for action in terms of non-normative desires, non-normative desires play a very different role in Schroeder's theory than they do in a constructivist view. Furthermore, I think, they are not "non-normative" in the same sense. On a constructivist view, reasons for action pick out desired features of actions. Reasons also count in favor of actions. That is, they are something normative to be said in favor of performing certain actions. However, desires are non-normative in at least two senses. First, we can give a non-normative description of the psychological state of desire. The contents of desires, as well as the attitude that a desiring agent takes towards that content, can be described in non-normative terms. Desires do not "aim at" the good or represent their objects as good. Second, that an action has some desired feature is not itself something normative to be said for performing that action. The constructivist attempts to get normativity into reasons by appealing to norms of permissibility. I have argued in earlier chapters that this strategy cannot succeed.

Schroeder, on the other hand, accepts the first sense in which desires are non-normative, but denies the second. Put a bit crudely, on Schroeder's view, that an action has some desired feature just is something normative to be said in favor of performing that

126 Ibid., 59-60.

action. In fact, this is all there is to being something normative to be said in favor of performing an action. (That an action has some desired feature reveals the action to promote the satisfaction of some desire—the desire for that feature. For a consideration to reveal this just is for it to stand in the reasons-relation with an agent and an action.) So, Schroeder escapes the burden of having to provide some *further* explanation of how considerations that pick out desired features of actions count in favor of those actions. This was the explanation that got the constructivist into trouble. Thus, it would appear that Schroeder has provided an explanation of reasons which bases them in desires, but avoids the problems I raised for constructivism.

It is not clear that Schroeder's view denies R (the thesis that a reason for an action counts in favor of performing that action insofar as it picks out something about the action that has value). I think there is a way of understanding Schroeder in which he adopts a version of this claim (at least given the quite broad conception of "value" I have been working with so far). (This will also help make clear how Schroeder avoids the problems I raised for constructivist views.) In Chapters 2 and 3, I suggested that, in order to account for the sense in which reasons for action call for, count in favor of, or are something normative to be said in favor of performing actions, we will need to identify something that is to be promoted, preserved, or respected that would ground these reasons. Reasons for action would count in favor of actions insofar as they revealed actions to promote, preserve, or respect what is to be promoted, preserved, or respected. Schroeder's view seems to fit this description. What is to be promoted is whatever an agent desires. Considerations count in favor of actions insofar as they explain how those actions make it more likely that an agent

will satisfy those desires.

So, if to be valuable is to be promoted, and on Schroeder's view something is to be promoted by some agent insofar as it is the object of that agent's desire, then it looks like on Schroeder's view objects of desire have value (at least, they have the kind of value we are looking for to ground reasons). And they have that value in virtue of being objects of desire. So, it looks like, on Schroeder's view, to be the object of a desire is all it is to have value (in light of what it is to be a reason).

The reader might wonder whether a view like Schroeder's has already been ruled out in Chapter 3. In Section VI of Chapter 3, I argued for the conclusion that, in order to regard the fact that an action has a particular feature as counting in favor of, calling for, or demanding an action, an agent must regard that feature of that action as one that makes claims on agents because of what it is and not merely because some rational agent has desired it, permissibly made it her purpose, and acted for its sake. The very basic point of the arguments for this claim was that "I (permissibly) want it" does not seem like the right sort of explanation of why some action is called for. This seems to be the basic explanation that Schroeder is offering, so it would seem that, if the arguments in Chapter 3 are right, we already have grounds for thinking that desire-satisfaction itself cannot be the value that grounds reasons for action.

I think the arguments I make in Chapter 3 (and their predecessors in Chapter 2) do put pressure on Schroeder's view. However, as stated in those chapters, they rely on the constructivist idea that an act of will is a key part of the explanation of how a consideration becomes a reason. (That an action has some feature F does not become a reason to perform

that action whenever an agent desires F and performing the action for the sake of F is permissible. It becomes a reason to perform that action when these conditions hold *and* the agent decides to perform the action for the sake of F.) For example, the argument that failing to s-share reasons with another agent may result in treating him as a mere means seems to rely on the idea that it is up to the agent whether to act for the sake of any given desired feature of action and thus create a demand that others help her pursue this purpose. If reasons are not up to agents in this way, it will be less clear that such demands are insufficiently justified by the relevant desires because these demands may no longer be clearly different from more familiar moral claims.[127]

Nevertheless, I think that the two traditional objections to Humean accounts of reasons that I discuss below are motivated by the same sort of concern that the arguments I give at the end of Chapter 3 express—namely, that particular agents' desires, are not the right sort of thing to explain why considerations count in favor of, call for, or demand actions (whether or not agent's can decide when they do so). (Moreover, because Schroeder presents himself as responding to these objections, he seems to grant that they have some merit.) In arguing the Schroeder's responses do not succeed, I mean to suggest that a Humean account of reasons will have difficulty avoiding this concern.[128] However, I will not be providing additional reasons to think that it is, in fact, a legitimate concern (though I

127 I think that the first and third arguments in favor of a requirement to s-share reasons are less dependent on the role that an agent's choices of whether to act for the sake of a particular desired feature play in a constructivist account of reasons for action. Thus, they are more likely to be able to apply directly to Schroeder's view. However, they are also, I think, more susceptible to a different kind of criticism. One might think that they do little more than state some rather obvious truths about how we ordinarily think of our reasons. As such, few are likely to deny them, but they are not likely to move someone who thinks that our ordinary thoughts about reasons are not a particularly good guide to the actual explanation of what reasons for action are to change her view.

128 I say "suggest" because I have not canvassed all Humean views. It may well be that there are solutions to these problem that I have neither thought of nor run across.

suppose one might be inclined to take it more seriously if it cannot easily be avoided).

IV. Objection 1: Humean reasons are objectively self-regarding

Many claim that a Humean theory of reasons requires that, when deciding what to do, an agent must consult her own desires, thus making some "I want…" claim a necessary premise in every bit of practical reasoning. Some think this makes correct practical reasoning objectionably self-regarding.[129] The basic idea of this objection (henceforth, rather awkwardly referred to as "the self-regarding objection") is that, because the Humean theory of reasons makes an agent's reasons depend on her desires, it requires agents to have and be moved by thoughts (in particular thoughts about their own desires and what will promote them) that it seems objectionably self-centered to have and be moved by when thinking about what to do. If I am considering whether to help a drowning man, for example, we might think that I should be thinking about him and what he needs ("he needs help"), rather than about myself and what I want ("I want to help him").

Or at least it seems that I should not be *required* to have and be moved by the latter thought in order to act in order to help the drowning man. It should be possible for me to be moved to act just by thinking about the needs of others and not by thinking about my own desires. Nagel expresses such a worry in *The Possibility of Altruism*. He thinks it important that a theory of reasons for action account for the possibility of altruistic motivation—i.e., the possibility that an agent is moved to act by "the direct influence of one person's interest on the actions of another, simply because in itself the interest of the former provides the

129 Ibid., 24-27.

latter with a reason to act."[130] This means that the self-regarding objection need not claim that having thoughts about one's desires when deliberating about what to do is itself always objectionably self-regarding, but rather it need only claim that a theory of reasons and practical reasoning that *requires* that one do this is an objectionably self-regarding view of practical reasoning.

For my purposes here, I will initially take the self-regarding objection to be somewhere in between these two positions. I take it to claim that the Humean theory of reasons rules out the possibility of altruistic motivation. I also take it to claim that, in some cases (helping those in need being a prime example), an ideal or well-functioning practical reasoner will not have and be moved by thoughts about her own desires.[131] This is a weak gloss on the idea that having such thoughts would be "objectionable" that, I hope, remains neutral on the question of how exactly the agent who does have and is moved by thoughts about her own desires is criticizable (e.g., whether she acts wrongly, whether she exhibits some character flaw, or even if we just think that altruistic motivation is to be praised).

Schroeder takes his view to be able to avoid the self-regarding objection because he thinks that, when an agent decides what to do, at least in ordinary cases, only her reasons appear as premises in her practical deliberation.[132] Desires, because they are background conditions, explain why that agent's reasons are her reasons, but they do not themselves appear in that deliberation. The agent does not deliberate about her desires when deciding

130 Nagel, Thomas. *The possibility of altruism*. Oxford: Clarendon Press, 1970. 80.

131 More accurately: in some cases, an agent's action may be over-determined, but one of her motivations should be non-self-regarding. This would allow an agent to, for example, help the drowning man both because she wants to help him and because he needs help. As long as her non-self-regarding recognition that he needs help is enough to move her to act, she need not fall short of ideal agency just because she also happens to have a self-regarding motive for helping.

132 Schroeder, *Slaves of the Passions*, 37.

what to do. Instead her desires determine what reasons she actually has and thus also determine whether that deliberation is correct. But when she deliberates about what to do, she need only consider her reasons. When Ronnie thinks about whether to go to the party, he can have the thoughts "there will be dancing at the party" and "so I should go". He does not have to entertain the thought "and I want to dance" to arrive at this conclusion. Similarly, when I think about whether to save the drowning man, I can have the thoughts "he needs help" and "so I will help him" without having the thought "and I want to help him".

Now, does this resolve the self-regarding objection? At first glance, it may seem to do so. Schroeder's view provides a plausible story of what agents think about when deliberating about what to do, at least in many ordinary cases, that does not include thoughts about desires. If practical deliberation, on Schroeder's view, does not involve thinking about one's desires, then his view would seem to neatly avoid the self-regarding objection. It might be objectionably self-centered to think about one's desires (in some cases) when thinking about what to do, but Schroeder's agent does not. Problem solved.

Unfortunately, I do not think the problem can be so easily solved. Explaining why will take some time: Suppose that an agent *can* reliably figure out what she should do and do it without thinking about the desires that underlie the reasons that determine what she should do. On Schroeder's view, where desires are background conditions, it is still because an agent has a particular desire that she these reasons to do what will promote it, and that she should do whatever it is that she should do. Moreover, an agent can understand this. Suppose, for example, that I tell Ronnie that he should go to the party. He asks why. If desires are background conditions on reasons, I can answer this question by citing a reason

and not by citing a desire—I can say that he should go to the party because there will be dancing at the party. So, it might look like there is an available explanation of why Ronnie should go to the party that appeals to, e.g., facts about the world and not to his desires. So far, so good. However, Ronnie can ask why the fact that there will be dancing at the party makes it the case that he should go (or that he has a reason to go). (Or more familiarly, when I say he should go because there is dancing, he can say "So what?") If desires are background conditions on reasons, I can correctly answer this question by saying that it is because Ronnie desires to dance. The apparently desire-free explanation of what Ronnie should do turns out to be grounded in facts about Ronnie's desires.

Of course, Schroeder would not deny any of this. So far, I have just stated his view. Nonetheless, I think that this leaves his view in an uncomfortable position, if we accept that there would be something objectionable or non-ideal about being moved to act by thoughts about one's own desires. (If Schroeder does not accept this, then he's not addressing the self-regarding objection, he is denying that it is something to worry about in the first place.[133]) The possibility of altruistic motivation seems to depend on a kind of ignorance or lack of attention. It seems plausible to suppose that a good practical reasoner ought to be able to know why she has the reasons she has. (We might even think that she should have this knowledge because it will make her better at discovering the reasons she has, and subsequently acting correctly.) So, it is plausible to suppose that a good practical reasoner can know that what reasons she has are determined by her desires. (And presumably, in particular cases, she can know which desires ground which reasons and which facts about

133 I will eventually suggest that this is what he should do. I think that the heart of the self-regarding objection is a concern that Schroeder's view cannot accommodate, but also one that does not move him.

what she should do.) Yet, in order to act out of direct concern for the interests of others, she must not think about her desires when deciding what to do. So, in order to act out of direct concern for the interests of others and avoid objectionable or non-ideal thoughts, she must avoid thinking about why she has the reasons she has. She must ignore the full story of why she has the reasons she has.

This strikes me as implausible. It is hard to see how it could be objectionable, or in some way less than ideal, to consider the full story of why one has the reasons one has (especially because this just is the full story of why one should do what one should do) when thinking about what to do. If this full story were very long, such that an agent would not have time to both consider it and act, then it might be correct for her to consider only some part of it. But the full story on Schroeder's view is fairly simple and easy to grasp, so this cannot be his answer to the question. Similarly, it may well be more efficient for agents to omit consideration of the full story in many cases. This might make it objectionable for an agent to have a practice of considering the full story in every case and suggest that agents should adopt at least some shortcuts for practical reasoning. However, it is a long way from here to the claim that it is objectionable for an agent to *ever* consider the full story when engaged in practical deliberation (in the relevant sorts of cases, e.g., helping someone in need).

Furthermore, there seem to be good reasons to think that an agent *should* consider this full story when deliberating about what to do. Later in this chapter, I will discuss in some detail how particular reasons are determined. One upshot of this discussion will be, unsurprisingly, that in many cases, this story is quite complex. If so, then we might think

that, in complicated and difficult cases (e.g., cases where there are many relevant and conflicting desires), an agent will have a better chance at getting it right about what reasons she has, and subsequently about what she should do, if she considers the full story of why she has the reasons she has. By reflecting on what desires she has and how they might be promoted or thwarted by various possible actions, she will do a better job of discovering what reasons she has. In short, we might think that in hard cases, it would behoove the agent to consider the full story (or at least a bigger part of it than usual) of what she should do in order to work out what to do. At the very least, it seems that doing so would hardly be non-ideal.

In addition, suppose we take seriously the idea that, when an agent acts for reasons, she takes her action to be called for by her reasons and is (at least partially) motivated by this. That is, she understands her reason as normative. She is not merely moved to act, she is moved to act, at least in part, by some recognition that her reason makes her action appropriate or (at least partially) justified. However, a consideration does not, by itself, make an action appropriate or (at least partially) justified. A consideration can do this, on Schroeder's view only insofar as it is part of an explanation of how the action promotes one of the agent's desires. It may well be possible, or even normal, to act without explicitly considering this explanation. But, if acting for a reason involves taking a consideration to call for acting, then it hardly seems that it can be *objectionable* to contemplate (or even just *better to avoid* contemplating) the thing that makes the consideration call for acting when acting for the sake of that consideration.

To be fair, I think that Schroeder would respond by appealing to a distinction

between practical deliberation and something like theoretical reflection on what one ought to do. The thought here is that only the consideration of reasons (and not what grounds or explains them) is properly considered part of practical deliberation. Thinking about why we have these reasons, or what makes them reasons (and thus thinking about a more complete story of what we ought to do), is a kind of theoretical reasoning. It is like thinking about why World War I happened or what are the fundamental building blocks of the universe. Deliberating about what to do and thinking about why we have the reasons we have are two separate activities.[134] An agent may well engage in both when thinking about what to do (theoretical reasoning about practical reasons may often accompany practical deliberation), but we might be able to address the worry I just raised for Schroeder's view if it is only practical deliberation that needs to be free of thoughts about one's desires. Such thoughts would be objectionable if they appeared in practical deliberation, but this is not to say that it is better if do not have them at all when thinking about what to do. As long as they appear only in corresponding theoretical reasoning about what we ought to do and why, they are not objectionable. This would allow the agent to think about the full story of what she ought to do and yet allow what motivates her to act to remain free of thoughts about her own desires.

I am not going to take a position here on whether this is an important or interesting distinction. Even if it is an important and interesting distinction, I nevertheless do not see why one would think that it would be objectionable to have thoughts about one's own desires appear in one's practical deliberation, but not in one's theoretical reflection on the complete justification of one's actions that that practical deliberation reflects. With respect

134 Schroeder describes the distinction as that between thinking about normative reasons and thinking about the explanatory reasons that make them normative reasons. (Ibid., 37.)

to this issue, at least, the distinction seems arbitrary.

For example, the distinction might purport to separate the parts of an agent's deliberation that motivate actions and the parts that do not. (Practical deliberation operates on considerations that move the agent to act. Theoretical reasoning does not.) If this is the case, we might think that it is objectionable to be *moved* to help others by thoughts about your own desires. However, since it is only reasons that actually move agents, if agents have other thoughts about their desires, these other thoughts will not be objectionable.

However, if this is the idea behind the distinction, then it seems rather tenuous when we reflect again on the fact that there can and, at least in some cases probably should, be interaction between these two ways of thinking about what to do. My theoretical reasoning about what reasons I have and why ought to be able to inform my practical deliberation. I might come to discover new reasons that I have (or that I do not have reasons I have previously taken myself to have) through this theoretical reasoning. My practical deliberation should reflect these conclusions. It may well be possible for this kind of theoretical reasoning about what reasons one has to have no influence on one's practical reasoning. (I might form beliefs about what reasons I have, but be utterly unmoved by these conclusions.) Yet, it also seems clearly possible that thinking about what grounds reasons, so as to figure out what reasons I have, can influence my practical reasoning. When it does, it seems artificial and arbitrary to say that only the *conclusions* of theoretical reasoning (e.g., "that he needs help is a reason to help him") or their practical analogs (e.g., "he needs help" when it moves me to help him) can count as motivating an agent's action.

Similarly, if we take seriously the idea that part of what it is for an agent to act for a

reason is for the agent to be moved to act in part by the recognition that the consideration calls for acting, then this will cast doubt on the possibility of neatly relegating thoughts about desires out of practical deliberation. If practical deliberation is just deliberation with reasons, but acting for a reason involves motivation by the recognition that a consideration calls for acting, then if an agent knows why a consideration calls for acting, it seems arbitrary to claim that the agent is not moved by this knowledge as well as by the consideration itself. It might be possible to be moved by some recognition that one's reason calls for action without knowing the full story of what that amounts to, but it seems strange to say that if you *do* know what a consideration's calling for an action amounts to, you are moved by the recognition that a consideration calls for acting, but not also by the recognition of what that amounts to.

As a result of all of this, I think that Schroeder should simply deny that there would be something objectionable or non-ideal about considering one's desires and being moved by thoughts about one's desires when deliberating about what to do. I think it doubtful that making desires background conditions on reasons will firmly separate thoughts about an agent's desires from what motivates the agent to act. But more importantly, even if it does, I think that it will not satisfy anyone who was bothered by the self-regarding objection to begin with.

Why would we have found it objectionable or non-ideal for an agent to be moved to, e.g., help others, by thoughts about her own desires in the first place? I suspect it is because we think that the agent who does this fails to recognize why the interests and needs of others matter, why they make claims on her. If I am moved to pursue the interests and needs of

others by thoughts about my own desires, then I make a mistake (and a self-centered mistake at that) in thinking that the interests and needs of others matter because I desire them. In fact, they matter because of something about those other agents. The self-regarding objection, then, is not really a concern about what agents think about or are moved by when, e.g., helping others. It is a concern about what makes the interests and needs of others matter. But if this is the concern behind the self-regarding objection, then making desires background conditions on reasons cannot alleviate it. On Schroeder's view, whether or not an agent considers them or is moved by them, desires play a fundamental role in determining what reasons she has and thus what she should do. It is because the interests of others promote the satisfaction of *my* desires that considerations like "he needs help" make claims on *me*. Finding a way in which they can do this, but an agent need not be aware of it (or moved by it), is not going to satisfy someone who thinks that an agent's desires are not the right sort of thing to make the interests and needs of other agent's matter.

Furthermore, if Schroeder denies that there would be something objectionable or non-ideal about considering or being moved by thoughts about one's own desires when, e.g., helping others, this resolves a number of the worries I raised above. I suggested above that, if the complete story of what one ought to do appeals to facts about what promotes one's desires, it is rather mysterious what would be objectionable or non-ideal about considering this full story when deliberating about what to do. But the need for this to be objectionable only arises, on Schroeder's view, because Schroeder wants to grant that there is something to the self-regarding objection. If he denied that there would be something objectionable or non-ideal about considering or being moved by thoughts about one's own desires in the first

place, then we need no explanation of why considering this full story would be objectionable or non-ideal and we need not solve what I suspect is an unsolvable mystery.[135]

V. Objection 2: Agent-neutral reasons

Agent-neutral reasons are reasons that are reasons for all agents and that seem to be necessarily so. They are not reasons for everyone merely because we happen to have many of the same desires. They are reasons that have what Schroeder calls "strong modal status". This means that these reasons do not depend on any particular desires.[136] An agent would have these reasons even if she had a completely different set of desires. Many moral reasons are commonly thought to be agent-neutral. For example, we tend to think that we each have reasons to help those people that we encounter that are in need of help. Additionally, we tend to think that we each have these reasons even if we do not care about the well-being of the particular people in need that we happen to encounter (indeed even if we desire that they suffer).

The strong modal status of agent-neutral reasons is commonly thought to be a problem for Humean accounts of reasons. If reasons depend on an agent's desires, then it is not clear how there could be reasons that an agent would have even if she had completely different desires. Schroeder is optimistic that his view can explain the strong modal status of agent-neutral reasons (so much so that he takes his later account of the weight of reasons

135 This is not to say that Schroeder should give up the distinction between reasons and background conditions (though I think some of my points put pressure on the distinction if it is meant to mark a sharp distinction between motivating considerations and considerations that explain why those motivating considerations are appropriate). My main purpose here is just to say that that distinction is not likely to alleviate the concern that a Humean theory of reasons makes practical reasoning objectionably self-regarding.

136 Ibid., 106.

to depend on agent-neutral reasons[137]). An agent-neutral reason, he claims, is massively over-determined. It is a reason that can be explained by any possible desire.[138] Thus, there is no particular desire that an agent must have in order to have an agent-neutral reason. She must only have some desire or other.

There is a sense, then, in which agent-neutral reasons follow from what it is to be an agent. To be an agent, one must have some desires. (Without the motivation they provide, one could not act at all.) Yet, as soon as you have a desire (any desire whatsoever) you have agent-neutral reasons just in virtue of having some desire. This contrasts with a view like Velleman's which claims that it is constitutive of being an agent that you have a particular desire (a desire for a sort of self-understanding), and whatever agent-neutral reasons there are follow from this necessary desire.[139] Schroeder remains neutral on the question of whether the agent-neutral reasons are reasons that all agents have or whether they are merely reasons that all human agents have.[140] That is, he remains neutral on the question of whether they follow from any possible desire of any possible agent or only from any possible desire of any human agent.[141]

137 On Schroeder's view, the weight of a reason is determined by reasons to place more or less weight on other reasons. (This allows him to avoid taking on the often thought to be problematic position that on a Humean account of reasons, the strength of a reason is proportional to the strength of the desire that grounds it.) However, Schroeder thinks that *all* reasons to place more or less weight on other reasons are agent-neutral. (Ibid., 142.) This means that his account of the weight of reasons depends on his account of agent-neutral reasons. In this section, my arguments will focus on ordinary agent-neutral reasons for action (as opposed to reasons to place weight on reasons). I believe that the problems that I raise for these reasons are general and will arise for weighing reasons as well. However, showing this would require an extensive exposition of Schroeder's account of weighing reasons (and the problems are not interestingly different), so I will not take the time to do this in this chapter.

138 Ibid., 109.

139 Ibid., 107. See also Velleman, *Practical Reflection*

140 Schroeder, *Slaves of the Passions*, §6.4.

141 For example, it might promote any possible desire to be in good physical health. But what promotes good physical health for a human agent depends on what the human body is like, not the mere fact that a human agent has desires. So, the fact that certain substances contain vitamins or protein might be a reason for humans to eat them, because they will promote health. However, this same fact may be no reason for another sort of desire-haver, e.g., an alien or a robot, if they are differently constructed.

Schroeder gives an example of a theoretical reason that can be explained this way. We can see that each person has a reason to believe any proposition only if it is true by appealing to her desires, but without having to appeal to the content of any particular desire. For any given desire an agent may have, there will be a variety of things that it would be useful for her to know, or at least not have false beliefs about, in order to satisfy that desire. In turn, there will be a variety of facts that are relevant to that first set of facts, so it would be useful for her to know or at least to not have false beliefs about those facts as well. (And so on.) As a result, we might think that, for any given proposition, no matter what an agent desires, she will be in a better position to satisfy her desires if she believes that proposition only if it is true. So, whatever an agent desires, believing a proposition only if it is true, promotes the satisfaction of that desire. And so, whatever an agent desires, she has a reason to believe a proposition only if it is true.[142]

Of course, the agent-neutral reasons that the anti-Humean tends to be most concerned with are moral (and thus practical) reasons. These are, for example, reasons to help those in need or reasons not to kill, steal, lie. To show that agents have agent-neutral reasons to do (or not to do) these things, we need to show that, just as believing a proposition only if it is true tends to promote the satisfaction of desires, whatever they happen to be for, so do helping others, not killing, not stealing, and not lying tend to promote the satisfaction of desires, whatever they happen to be. If I will be more likely to fulfill my desires, whatever they are, if I help those in need, refrain from killing, stealing, lying, etc., then I have reasons to help those in need to refrain from killing, stealing, lying etc just in virtue of having some desire or other (but not in virtue of having any particular

142 Ibid., 113-5.

desire that those actions promote).

Before attempting to show how such actions promote my desires, whatever they may be, it will be useful to examine some (relatively) easy cases of practical reasons that would seem to admit of this kind of explanation. We might think that each of us has an agent-neutral reason to promote our own health. A healthy person is more likely to be able to promote his or her desires, *whatever they may be* (if for no other reason because she will be alive longer to do so.) So, it would seem that, as long as an agent has some desire or other, she has an agent-neutral reason to promote her own health. Similarly, we might think that we have agent-neutral reasons to develop our talents. It is plausible to suppose that (at least many) talents are complex sets of capacities. As such, they are versatile—they can be used in the service of a wide variety of desires. Even the process of developing talents itself, is likely to cultivate useful abilities—e.g., the ability to work at a complex, difficult, and sustained long-term project. So, it is plausible to suppose that having talents, and having the additional capacities likely to come out of a project of developing one's talents, is going to make an agent more likely to satisfy whatever desires she happens to have, than she would be if she did not have those talents and capacities. Therefore, it does seem that we have a reason to develop our talents, and a reason that does not depend on our having any particular desire, but rather that is explicable by any possible desire.

So, how might we explain agent-neutral reasons to help others? Why might we think that helping others promotes any desire whatsoever? At first glance, this seems implausible. Presumably, ignoring the needs of others, when they do not fulfill any of one's particular desires, frees up more time, energy, and resources to spend pursuing the satisfaction of one's

particular desires, whatever they may be. We might try to get around this problem by pointing out that we are all better at satisfying our desires when we live in some system of social cooperation. We might then argue that requirements to help others and prohibitions against killing, stealing, lying, etc. are principles that govern either the best system of social cooperation or every system that is in fact beneficial. Schroeder thinks that views like Scanlon's and Rawls' can be understood in this way, so we might be able to import a view like theirs into Schroeder's picture to determine which agent-neutral reasons we actually have.[143]

However, this explanation is subject to a fairly obvious problem. We might worry that an agent could have a set of particular desires that are inconsistent with social cooperation. If so, then following principles that promote social cooperation will not promote this set of desires, and thus reasons to do so do not follow from *any* possible desire. At least at first glance, this seems like a viable possibility. I might at least take myself to want to forgo the benefits of social cooperation in exchange for not having to take on its burdens. Such a set of desires may be rare, or even non-existent, in actuality, but it need only be possible to upset the argument that we have agent-neutral moral reasons.

This argument against Schroeder's view's ability to account for moral agent-neutral reasons appeals to the content of a single particular possible desire (i.e., the desire to forgo the benefits of social cooperation), and claims that actions that promote social cooperation or that are required by rules necessary to sustain social cooperation will not promote the desires of a person whose set of desires is organized around this desire to live apart. If so, then these actions do not promote every possible desire (after all, they do not promote this

143 Ibid., 116-7.

one). If that is the case, then there can be no agent-neutral reasons for these actions.

This kind of counterexample relies on a natural understanding of "promote any possible desire". On this understanding, if an action would promote any possible desire, then for any particular desire we can think of, it will be true that performing the action would make an agent more likely to satisfy it than would not performing the action. If this is how "promote any possible desire" is to be understood, then I think that it is highly doubtful that there are any such reasons. Perhaps actions that preserve one's life might be plausible candidates to promote every possible desire—after all, remaining alive seems to be a precondition on satisfying a desire, whatever it is for. But then of course, an agent might desire to die. I think it likely that for any action that one proposes as a candidate for making an agent more likely to satisfy her desires, whatever they happen to be, we will be able to find some particular desire, or set of particular desires, that this action will not promote.

We might conclude from this that there just are not any agent-neutral practical reasons (or that there are very few, and they are not the ones we might have wanted). Schroeder could go this way. However, it leaves his view no better off that other Humean views with respect to the concern that a Humean theory of reasons cannot account for the agent-neutrality of certain moral reasons. On the other hand, I think there is another way of understanding "promote any possible desire" that has a better chance at yielding agent-neutral reasons (and may better capture Schroeder's actual examples). In Schroeder's examples of explanations of agent-neutral reasons, some notion of agential efficacy functions as something like a fundamental value in grounding agent-neutral reasons. Those actions which we have agent-neutral reasons to do are those actions which make us better

desire-satisfiers in general. They are actions which cultivate, improve, or increase an agent's capacities and the resources available to her for satisfying desires, whatever they may be. Using this understanding of "promote any possible desire," we might say something like the following to resurrect moral reasons: Forgoing the benefits of social cooperation deprives an agent of a variety of useful resources for desire-satisfaction (most notably, the help of others). Social cooperation provides an agent with these resources. Social cooperation makes it possible for an agent to promote a wider range of desires. So, social cooperation makes an agent a better desire-satisfier in general, even if it does impede the promotion of a few possible particular desires.

I think this agential-efficacy-as-a-value understanding of "promotes any desire" is more promising that the previous understanding as a resource for explaining agent-neutral reasons. However, I think it nonetheless creates at least three serious problems for Schroeder's view.

The first of these problems is that there seem to be circumstances in which, on Schroeder's view, agents have exactly the opposite agent-neutral reasons from what we would want or expect. Suppose we can produce an argument showing why social cooperation tends to promote agential efficacy. Suppose further that we can show how certain kinds of (moral) behavior promote social cooperation such that we all have reasons not to lie or steal, reasons to help those in need, etc. It nevertheless seems that this sort of explanation of agent-neutral moral reasons might run into trouble in particular cases. It seems subject to something like an objection to rule consequentialism: it is plausible to suppose that there are situations where my desires are more likely to be satisfied if I refuse

to help those in need (or kill, steal, lie, etc). Furthermore, we do not need to appeal to particular desires to construct such cases.

For example, suppose I know that my friend has a habit of keeping large amounts of cash in her wallet and that she does not keep careful track of where she spends it. I realize that I can occasionally take $20 from her wallet and she will never notice. Whatever my desires, an additional $20 will likely improve my chances of satisfying them, and so, it would seem that I have an agent-neutral reason to take the money. It may be true that, most of the time, stealing and lying hinder the satisfaction of my desires, whatever they may be. While this may give me reason to have a general policy of not stealing or lying, it seems to have no bearing on what reasons I have in a case such as this where I know that *this* theft and *this* lie won't hinder my ability to satisfy my desires and instead will actually improve it.

This leaves us with a very different set of agent-neutral reasons than we might have wanted. We would seem to have agent-neutral reasons to steal and lie in just those cases where we can (1) improve our agent efficacy by doing so, in part because (2) we can be assured that we will not get caught. But these are precisely the cases (or at least they are among them) where agent-neutral reasons are most important. When we think that a theory of reasons should preserve agent-neutral reasons, especially agent-neutral moral reasons, the reasons we have in mind are reasons not to steal and lie when we could benefit from doing so and will not get caught. The reasons we have in mind are the reasons that remain when we do not have self-interested reasons to treat others as we ought.

Before moving on to the second and third problems for Schroeder's account of agent-neutral reasons, I will pause and note that a similar problem to that just raised for agent-

neutral moral reasons would seem to arise for theoretical reasons as well. In the case of reasons for beliefs, I think this problem can be solved. However, this solution is not available in the practical case. This gives us additional reasons to think that the first problem I have raised for Schroeder's account agent-neutral reasons is in fact a serious problem.[144]

Even the example of a theoretical reason that Schroeder discusses seems subject to the kind of counterexamples discussed above. There may well be cases where having a few false beliefs will make me more likely to satisfy my desires. For example, they say that pessimists and cynics tend to have more accurate beliefs about themselves and the world than do more optimistic people, but the motivational tendencies that result from those accurate beliefs can lessen the pessimist's ability to satisfy desires. (Furthermore, it seems reasonable to suppose that this can happen regardless of what the content of the pessimist's desires happens to be.) So, we might think that a view like Schroeder's would have to accept that it is not the case that, for all propositions, there is a reason to believe it only if it is true. Instead, whether or not there is a reason for a particular agent to believe a particular proposition only if it is true would depend on how believing that proposition only if it were true would affect the likelihood of that agent's satisfying his desires whatever they may be.

Fortunately, there seem to be at least two features of belief formation that count in favor of taking the general usefulness of true belief to ground a standing reason to believe any given proposition only if it is true. The first is that beliefs persist. If I form a belief that

144 This is in part because I suspect that Schroeder has drawn his account of agent-neutral reasons from thinking primarily about reasons for belief. (His examples in Chapters 6 and 7 certainly seem to suggest this.) However, I think that there are important disanalogies between reasons for belief and reasons for action that undermine this strategy, or at least Schroeder's use of it.

some proposition p is true and act on that belief, I do not normally then immediately give the belief up. I continue to believe p for some time, and that belief may well influence other actions of mine. I may forget about p or I may give up my belief that p upon encountering new evidence, but, for most beliefs, once I form them, they remain my beliefs for some time. It seems likely that this is part of the point of having beliefs—to store information so that it can be accessed again at a later time.

If beliefs persist in this way, and if I form a false belief because I deem it to be useful in a particular situation, then that false belief may create difficulties for my desire satisfaction later on down the road. I might try to avoid this by attempting to give myself the useful false belief only for the duration of the necessary practical deliberation and then give it up immediately afterward so as to avoid its potentially harmful consequences. However, significant obstacles to doing this are created by the second feature of belief formation that supports the claim that believing a proposition only if it is true promotes the agent's desires.

This second feature of belief formation has to do with the sort of control we have over our beliefs. In order to identify a situation where the general tendency for true beliefs to be useful does not hold, we would have to first determine the usefulness of the relevant proposition for desire satisfaction and then decide whether to believe it or to not believe it on that basis. But we can't seem to form our beliefs about what is true based on what is useful for desire satisfaction, or at least we cannot do this at will.[145] This suggests that we cannot operate with anything other than the general reason to believe a proposition only if it

145 See, for example, Hieronymi, Pamela. "Two Kinds of Agency." *Mental Actions*. Oxford: Oxford University Press, 2009.

is true when we actually engage in theoretical reasoning and belief formation on the basis of theoretical reasons. And so, while this is not Schroeder's argument for the existence of a reason to belief a proposition only if it is true, it does seem to suggest that this conclusion is correct and captures something fundamental about what belief is.[146]

It also suggests that it would be quite difficult to get around the persistence of beliefs by planning to give them up when they are no longer useful. One could not expect to give up a belief upon recognizing that it is no longer useful. So, to make this happen, an agent would have to plan to make sure that she was presented with overwhelming evidence of the falsity of the belief at just the time when it becomes no longer useful. She also must contrive to forget the plan before its execution (or it would provide reasons that undercut that evidence). There might, nonetheless, be a few cases where this is possible, but I think it safe to assume that in almost all cases, it is not likely to be an effective method of desire-

146 At this point, one might wonder if we have simply undermined Schroeder's explanation of the reasons we have to believe propositions only if they are true altogether. That explanation appeals to the usefulness of true belief to explain why we should believe propositions that are true, but I have just said that we cannot form beliefs on the basis of considerations about their usefulness. In fact, however, Schroeder avoids this problem (at least at first glance) because on his view, facts about usefulness will be background conditions on reasons for believing propositions only if they are true—agent's do not deliberate about them, but they explain why the considerations that the agent deliberates about are reasons. Schroeder's treatment of the example suggests that theoretical reasons are explained in terms of practical reasons. Presumably most reasons for belief are reasons because they, to put the point roughly, suggest that a given proposition is likely to be true. One might have thought that Schroeder would present this sort of idea as an analysis of what it is to be a reason for believing a proposition. Instead he seems to think that there is a further reason that explains why such considerations are reasons for belief—a practical reason for forming beliefs that are true. My guess is that this is motivated by a concern for unified explanation —practical reasons and theoretical reasons are all reasons and so they should be explained in the same way if possible. However, the resulting theory provides interestingly dis-analogous pictures of theoretical and practical justification. Theoretical justification seems to have two steps in a way that practical justification does not. Put crudely, an agent's reasons for believing p justify that belief insofar as they suggest that p is true. And the reason for this is because believing p only if p is true promotes the agent's desires. On the other hand, an agent's reasons for A-ing justify A-ing insofar as they suggests that A-ing will promote the agent's desires. There is nothing analogous to truth playing an intermediary role in the justification of A-ing. However, it seems at least as good, from the standpoint of unified explanation and theoretical economy, to claim that theoretical reasons and practical reasons are grounded in separate things (truth and desire promotion, respectively) than it is to claim that they are grounded in one thing, but that grounding takes two steps for one and one step for the other.

promotion. So, the idea that we have an agent-neutral reason to belief a proposition only if its true seems pretty plausible, despite the fact that there are particular cases where particular false beliefs might promote any possible desire.

However, no analogous moves seem to be available in the practical case. Regarding the first point: While it may be true that there are many cases where lying, for example, may have unforeseen consequences, there seems to be no principled reason why there could not be situations in which I have very good reason to believe that this particular lie will have no harmful consequences. My actions do not persist in the same way as do my beliefs. In telling this lie, I am not committing myself to telling others or to having reasons to tell others. Or if I am, I am committing myself only to telling other lies or having reasons to tell other lies in identical situations—and those, being identical, will have no bad consequences. In telling a lie under this particular set of circumstances, I take no stand on whether lying in general promotes my desires. It may be true that our psychology is set up so that, if we are comfortable lying in this sort of situation, we may be more likely to lie in others, and so, more likely to lie when that will get us into trouble. But this psychological tendency does not seem to be part of what taking oneself to have a reason or acting on a reason is, in the way that the persistence of belief seems to be part of what a belief is. If I were master of my own psychology such that I could lie only when it is genuinely useful, I would not cease to count as acting for reasons. On the other hand, if I were master of my psychology such that I could regard propositions as true at will, having initially known them to be false, only as long as regarding those propositions as true was useful, I think it is doubtful whether my attitude towards those propositions would really count as belief.

The second feature of belief formation that may help Schroeder defend the idea that we have reasons to believe propositions only if they are true against the charge that this will be true only in some cases was the idea that we cannot form beliefs at will because we take them to be useful. The reasons for which we form beliefs at will must provide some indication that those beliefs are true. No such restriction arises in the practical case. In order to identify a situation where a particular lie escapes the general tendency for lying to hinder the satisfaction of one's desires—a situation in which there may be no reason against lying, we have to determine the usefulness of the relevant lie for desire satisfaction in the given situation (taking into account foreseeable later consequences) and decide whether to act or not act on that basis. In the theoretical case, this way of deliberating requires that the agent form beliefs for the wrong sort of reasons. On the other hand, it is a perfectly normal bit of practical deliberation—usefulness for desire satisfaction is precisely the sort of consideration that explains why we have practical reasons on Schroeder's view.[147] So, the wrong kind of reasons problem that arose in the belief case doesn't apply to the practical case.

There is a second kind of problem that arises for the agential efficacy as a value understanding of Schroeder's view of agent-neutral reasons. Suppose we can identify determinate actions that would make an agent a better desire-satisfier in general—actions that promote the value of agential efficacy. It seems plausible to suppose that performing these actions and promoting agential efficacy may well inhibit one's ability to promote one's

147 Schroeder does think that practical deliberation does not normally explicitly consider the agent's desires (and so does not consider propositions like "lying will increase the likelihood of desire satisfaction). However, it is plausible to suppose that we can describe my deliberation in the lying case above in such a way that I do not explicitly consider my own desires and the likelihood of their satisfaction. I just think "I can take the money and not get caught" and the desires that the money will help satisfy are the background conditions that explain why this is a reason.

particular desires. To return to an example from above (though in this case, perhaps, somewhat cliched): Having money tends to promote agential efficacy. For many desires, you are more likely to be able to satisfy them if you have more money rather than less. This might seem to ground an agent-neutral reason (or at least a very widely shared agent-relational reason) to have lots of money. However, the pursuit of money can get in the way of promoting the satisfaction of any of the particular desires that we wanted the money to be able to promote. We might take this example either as an instance of the problem (money seems like a plausible candidate for an all-purpose desire-satisfaction-promoting resource) or as an analogy (we might take money to be a stand-in for any all-purpose desire-satisfaction-promoting capacity or resource). Either way, the example suggests that, by pursuing the value of agential efficacy, we might make ourselves excellent desire-satisfiers, in principle, but quite poor desire-satisfiers in practice. We might cultivate a variety of abilities conducive to desire-satisfaction, but rarely be able to use them to satisfy our particular desires because we will be too busy cultivating the capacities.

This makes the view look self-defeating. The value of agential efficacy is grounded in the fact that particular desires are to be promoted. So, if the pursuit of agential efficacy actually tends to inhibit the promotion of particular desires, then it would seem to be, at best, groundless and, at worst, in irreconcilable conflict with the normative facts which ground it.

The problem is compounded because on Schroeder's view, an agent-neutral reason cannot be outweighed by an agent-relational reason.[148] If this is so, then it is not clear how one could set aside the promotion of the value of agential efficacy in order to promote the satisfaction of particular desires. Perhaps there might be reasons to sometimes place weight

148 Ibid., 143.

on reasons to satisfy one's particular desires (whatever they may be) even in the face of competing agent-neutral reasons. These reasons to place more weight on one's particular desires might still count as agent-neutral because they do not depend on the content of any of these particular desires—they would hold regardless of what particular desires an agent has, as long as the agent has some particular desires or other. However, if we go this way, then Schroeder may lose his explanation of why agent-neutral reasons cannot be outweighed by the wrong sort of personal considerations, e.g., my reason to help you could be outweighed by reasons arising from the fact that I do not like you. We would at least need a principled way to distinguish the cases where it seems objectionable to allow reasons stemming from particular desires to outweigh agent-neutral reasons (e.g., permitting my dislike of someone to outweigh my moral reasons with respect to him) from those where it is not clearly objectionable (setting aside the pursuit of the resources and capacities that make up agential efficacy in order to satisfy some particular desires). It is not clear what the relevant differences could be that would explain this distinction, on Schroeder's view, because all agent-neutral reasons are explained in the same way—by appeal to what promotes agential efficacy.[149]

To make matters worse, there is a third sort of objection to this explanation of agent-neutral moral reasons. We might think that explaining why I have agent-neutral moral reasons in terms of what will increase my ability to promote my desires, whatever they happen to be, is simply the wrong sort of explanation for moral reasons. Even if it gets us the right answers about which agent-neutral reasons we have, it gives us the wrong sort of

149 A thorough discussion of this problem would require extensive exposition of Schroeder's account of the weight of reasons that I will not undertake here.

justification for those answers. Although moral reasons, on this view, will not depend on that agent's happening to have any particular desires, a version of the self-regarding worry still seems to apply. Schroeder's view, as I have described it, makes moral reasons all about me in the end—every agent-neutral reason for an action X has the same explanation—*I* will be more likely to satisfy *my* desires, whatever they are, if I do X. The fact that I must now treat others in certain ways that are beneficial for them looks like a happy accident. (It is a happy accident that what promotes my desires turns out to be what is beneficial to them.) We might think that morality can't be merely a happy accident.

Recall from Section IV that Schroeder takes himself to avoid objections of this general sort—objections that claim that his view makes reasons objectively self-regarding—by making the explanation of reasons that appeals to the agent's desires part of the background conditions on reasons and not part of the reasons themselves. A moral agent, then, does not actually think "I will do X because I will be more likely to satisfy my desires, whatever they are, if I do X" when deciding to act on the right reason. Thus, she need not think about herself or her own desires when, for example, deciding to help someone. She need only think about the fact that the person needs help. Nevertheless, I suggested, we might think that the problem that the self-regarding objection raises is not really one about whether or not agents think about their own desires, but rather about what makes the needs of others matter. Reasons determine what we ought to do, so if reasons are explained in terms of what promotes the satisfaction of the agent's desires, then what an agent should do is explained by an agent's desires. The reason something matters for what an agent should do is because it bears on the satisfaction of her desires. However, we might think, the reason

that your needs matter for what I should do is not merely that my taking your needs into account promotes my desire satisfaction or agential efficacy. It is (at least in part) that they are the needs of another human being.

Schroeder's account of agent-neutral reasons might seem to avoid this problem because his explanation of why an agent has any given agent-neutral reason appeals to "any possible desire" or "agential efficacy" and not to the agent's particular desires. So, the explanation of what an agent ought to do that appeals to agent-neutral reasons will contain claims like "it promotes the satisfaction of any possible desire" or "it promotes agential efficacy," but need not contain any "I want it" claims. Unfortunately, this solution is illusory. As I said above, my agent-neutral reasons to, e.g., help you are explained in terms of what will promote the satisfaction of any possible desire *for me* or what will promote *my* agential efficacy. We might think that my agent-neutral reasons to help you should have something to do with your needs, or at least it seems that they should have something to do with what will promote any possible desire *for you* or with *your* agential efficacy. If we do think this, then Schroeder's agent-neutral reasons seem to be no less subject to the self-regarding worry than do his agent-relational reasons.[150]

So, even if I am wrong about the first and second problems with Schroeder's account of agent-neutral reasons, the third problem suggests that his view will only get us the right answers about what reasons we have. (It will also grant those answers greater modal

150 In fact, we might even think that an agent-neutral reason to help others is more subject to self-regarding worries than is an agent-relational reason to help others. (Schroeder uses "agent-relational" instead of the more familiar "agent-relative.") An agent-relational reason to help others could be grounded in a desire for the well-being of others. Although the reason would still be based in an agent's desire, at least this agent exhibits some concern for the well-being of others in desiring the well-being of others. On Schroeder's view, an agent-neutral reason to help others is based in no such concern for others because the explanation for this reason appeals to what promotes any possible desire (of the agent) or the agent's agential efficacy.

security than other Humean theories—agents will not be able to avoid having moral reasons by failing to have the relevant particular desires.) However, worries that Humean theories of reasons fail to properly account for moral reasons need not be understood as merely extensional worries. They are not eased by getting the right answers about what reasons we have. Nor, I think, are worries that Humean theories of reasons fail to properly account for moral reasons primarily modal worries. They need not be eased by removing dependence on particular, contingent desires. We might think of them as worries about why it is that we have claims on each other to be treated in certain ways. On Schroeder's picture, it seems to me unavoidable that your claims on me, whatever they are, can only arise because you are useful to me in my project of satisfying my desires. (And these desires are not subject to any further standards except those stemming from my other desires and perhaps some value of (my!) agential efficacy. So, no recognition of your independent value is sneaking in via the content of my desires.) The idea that persons are bearers of value that make claims on others and not merely on themselves is a fundamental and pervasive idea in much of our moral thinking and not one that I think we should set aside unless we absolutely cannot come up with a plausible account of practical reasoning that supports it.

So, if my arguments in this chapter are right, Schroeder's view fails to address two traditional worries for Humean accounts of reasons for action. His view bases the (in particular moral) claims that others seem to make on us in facts about what promotes the satisfaction of our own desires. Thus, his view conflicts with what I think is a fairly strongly held belief that the (moral) claims that others make on us stem from something about these

other people—their value or their needs—that is independent of their usefulness for our desire-satisfaction. And it appears that Schroeder's view may fail to allow for agent-neutral reasons. (It appears that his account either gets the wrong answers or is self-defeating.) None of this forces Schroeder to give up his basic view that desires are background conditions on reasons for action. However, I think, it does give us reasons to look for alternative views.

Conclusion

I. "Calls for"

In Chapter 1, I suggest that, in giving an account of the counting in favor of relation, and thus in giving an account of reasons for action, we are (likely among other things) attempting to capture the sense in which, when an agent acts for a reason, she takes some feature of her action (or that the action has that feature) to call for or demand performing that action and is motivated to act in part by that recognition, as well as by the consideration itself.[151] This contrasts with an agent's being merely moved to act by some feature of her action without concern for whether that feature calls for action. (I suspect that there are actions that meet this latter description (actions that agents perform for no reason), but for the purposes of my argument here, this contrast class may be purely conceptual—there may be no actual instances of it.

My arguments in Chapters 2, 3, and 4 against views that attempt to capture this feature of acting for a reason, while denying R, help clarify what this taking a consideration to "call for" an action amounts to. They help us see when, and under what conditions, we can intelligibly regard an agent (perhaps ourselves) as taking her actions to be called for. We begin with the idea that, when an agent takes a consideration to "call for" an action, this means that the agent takes a reason to be normative.[152] This means at least that, in acting for

151 That the agent is moved to φ by the consideration and the recognition that it calls for an φ-ing does not require that she does, or even intends to, φ. She might have stronger motivations to something else. Or she might simply choose not to φ.

152 This claim is, I think, pretty uncontroversial and held by each of the views discussed above (Raz, Velleman, Setiya, Korsgaard, Street, and (probably) Schroeder), though their conceptions of exactly what it means vary. (E.g., Raz thinks we need R to account for it whereas Street things that we can account for it with a notion of responsiveness to minimal consistency norms.

a reason, an agent takes her action to be subject to standards of appropriate or correctness. In taking a consideration to call for an action, an agent takes it to be part of an explanation of how her action is appropriate or correct according to those standards—she takes it to be part of a justification of her action according to those standards. So, in order to explain the sense in which an agent takes her reasons to call for her actions, and thus, in order to explain the counting in favor of relation, we need to specify this standard.

Admittedly, there are significant complications for figuring out what this standard is by thinking about the reasons for which agents actually act. These complications might lead us to doubt that any such standard exists.[153] To begin with, this taking a consideration to be part of an explanation of how one's actions are appropriate or correct according to these standards is likely (at least usually) not explicit. An agent probably does not represent the standards to herself and come to believe that her reason for acting explains how her action meets those standards, whenever she acts for a reason. However, I think that there is a sense in which the agent nonetheless takes her reason to be part of such an explanation. Caching this out will likely require appealing to the agent's dispositions to deliberate with the considerations she takes to call for acting (see Chapter 1, Section III) or some notion of "responsiveness" to the standards that determine when actions are appropriate or correct (see Chapter 2, Section III or Chapter 3, Sections I and II). E.g., the agent will take the considerations she takes to call for her action to have weight against considerations that suggest that her action is not appropriate or correct, if she becomes aware of them. Or, when taking a consideration to call for acting, an agent takes it to be part of an explanation of how

153 That is, we might doubt that there is any standard that agents recognize when acting for reasons. There could, in principle, still be standards of correctness that apply to reasoning and to action.

that action meets some standards insofar as she would be moved to resolve conflicts with those standards were she to become aware of them. (These need not be different phenomena.)

Furthermore, in some cases, an agent can intelligibly take a consideration to call for an action when, as a matter of fact, that consideration fails to be part of an explanation of how the action is appropriate or correct because the action in question is in fact inappropriate or incorrect. For example, the consideration that moves the agent to act might be a reason that is in fact outweighed by other reasons. In this case, this consideration is still part of an explanation of *which* action is appropriate or correct (it just so happens that it is not the one that the agent performs), so the consideration is still part of an explanation of the appropriateness or correctness (or lack thereof) of the action that the agent does perform. So the consideration is still of the right sort to be a reason for action.

Alternatively, it may turn out that the agent's action is not appropriate or correct according to the relevant standards and the consideration that she takes to count in favor of her action is not part of the explanation of why this is the case because, under the circumstances, that consideration should not be considered. (E.g., maybe the fact that the house is on fire is not a reason to flee if one's family is still inside.) In such a case, the agent makes a mistake in taking this consideration to call for her action, but it is an intelligible mistake. There are a couple of ways in which this might happen. The easier case is one in which the agent is unaware of (or inattentive to) the relevant facts. (E.g., she forgets that her family is in the house.) In this case, if what she knew (or attended to) and only what she knew (or attended to) were true (e.g., if her family were not in the house), then her reason

would be part of an explanation of how her action is appropriate or correct. So, I think it plausible to suppose that this agent can intelligibly take her reason to be part of an explanation of why her action is appropriate or correct.

The harder case is one in which the agent is aware (and even attentive to) the facts that make it the case that her consideration is not part of this explanation. (Suppose she's aware that her family is still inside the house when she flees.) I think such cases are likely to be borderline cases—e.g., it will not be entirely clear whether the agent who flees the burning house takes the fact that the house is on fire to call for fleeing or whether it merely moves her. Inosfar as she does take it to call for fleeing, in order to explain this, we would likely have to appeal to some notion of what would, under normal or not too distant circumstances, be part of an explanation of why fleeing would be appropriate or correct. (The basic idea here is that it is intelligible to take this consideration to be part of an explanation of why fleeing is appropriate or correct, even in the knowledge that it is not, because, but for one fact, it would be.)[154]

So, when an agent takes a consideration to call for acting, she regards that consideration as normative. That is, she regards it as part of an explanation of how her action is appropriate or correct, according to some standard. However, my arguments in later chapters (especially Chapter 3) suggest that there is more to taking a consideration to call for an action than this. When an agent takes a consideration to call for an action, she

154 Such a case would have to involve some inconsistency on the part of the agent—she seems to both take a consideration to call for acting and believe that it does not—but such inconsistency seems possible. We need not think that an agent be completely and wholeheartedly responsive to the standards that govern reasons, in order to count as taking a consideration to call for acting, even given full awareness of the relevant facts. (However, cases involving this inconsistency may be deviant reason-takings and I do not think that they will be particularly helpful in illuminating what it is for a consideration to call for or count in favor of an action.)

takes the consideration to be something normative and *positive*, something normative to be said *for* the action. This claim gets at two ideas—one about the role that the consideration plays in the explanation of how an action is appropriate or correct, and one about what kind of appropriateness or correctness is involved in the idea that reasons call for actions.

The first sense in which, when an agent takes a consideration to call for an action, she takes it to be something normative to be said *for* the action is that she takes it to be not merely part of an explanation of how her action is appropriate or correct, but rather to be a positive part of this justification. This should be fairly obvious after the discussion of difficult cases above. An explanation of why some action (φ) is appropriate or correct may contain considerations that suggest that φ-ing is *not* appropriate or correct and considerations that undercut or outweigh considerations that suggest that φ-ing is appropriate or correct. The considerations that suggest that φ-ing is not appropriate or correct do not seem to call for acting (though they are nonetheless part of an explanation of why φ-ing is correct). In fact, they do quite the opposite—they call for not φ-ing.[155]

The second sense in which, when an agent takes a consideration to call for an action, she takes it to be something normative to be said *for* the action is more interesting. When an agent takes a consideration to call for acting, she takes the consideration to be something

155 Throughout this document, the reader may have wondered how a view that takes an account of reasons to be an account of what "calling for" or "counting in favor of" amount to would say about negative reasons. For example, what would it say about a reason against acting—e.g., the fact that φ-ing would harm someone is a reason *not* to φ. I think there are a couple of options here. I do not take a position on which is correct. We might think that reasons against actions are simply considerations that call for not performing them. A reason against φ-ing is something normative and positive to be said for not φ-ing. This way, negative reasons can be understood in terms of "calling for." On the other hand, we might think that figuring out what a reasons "calling for" an action amounts to will yield a complementary understanding of "calling against" an action. E.g., if a reason calls for an action insofar as it shows that performing that action will promote, preserve, or respect X, then we might think that a reason "calls against" an action insofar as it shows that performing that action will destroy X, disrespect X, prevent X from coming about, etc.

normative that directs her towards acting, not merely something that reveals the action to be permissible or shows that she would be making no error in performing the action. In other words, the standards of correctness according to which a reason for action (partially) explains why that action is appropriate or correct have to do more than just determine whether the action is rationally permissible. (For example, they must determine more than that the action does not undermine the agent's ends or, perhaps, that it is morally permissible.)

This is because that an action is permissible is not, normally, something we take to call for performing that action (More accurately, it is not something we take to show that there is some consideration that calls for performing that action). (See Chapter 3) There may be some cases where a particular action is the only permissible action available. In such cases, that the action is the only permissible action available may call for performing it. (Again, more accurately, a consideration that is part of an explanation of why that action is the only permissible action available may call for performing that action.) However, standards of permissibility tend to rule, in most cases, that a variety of possible actions are permissible (and that not performing each of those actions is also permissible).[156] There may even be cases where a consideration that is part of an explanation of why an action is permissible is also part of an explanation of how not performing that same action is permissible.[157] In light of these facts, that a consideration shows an action to be permissible does not address the question of whether the consideration counts for, against, or remains

156 A view that purports to give standards of permissibility, but, in every case, rules that all but one action is impermissible strikes me as one that adopts a version on R despite its protestations. It likely involves some version of the claim that one should do what is best, and in identifying what is best, effectively identifies something as to be promoted, preserved, or protected.

neutral with respect to performing the action. A consideration that shows an action to be permissible does explain how the consideration makes an (albeit weak) claim or demand on the agent to perform that action. (See Chapter 3 Section VI). Such a consideration seems to be a reason why it is okay to perform the action, but it is not yet a consideration that calls *for* performing the action. If it is to call for action, to put the point crudely but intuitively, it must suggest that performing the action is a good idea.

Finally, at the end of Chapter 3 (and to some extent in Chapter 4), I argued for the claim that, for an agent to intelligibly take a consideration to call for an action, she must take that consideration to be something normative to be said for an action because of something about what the feature of action it picks out is and not merely because of the motivating attitudes (the desires or reason-takings) of herself or other particular rational agents. (Recall that to s-share a reason with another agent is to take the same feature of action to count in favor of acting and for one's interest in that feature to be because of something about what the feature itself is and not merely because some rational agent has taken an interest in it and (permissibly) acted for its sake. I presented three arguments in Chapter 2 for a requirement to s-share reasons. In Chapter 3, I argued that we can take these arguments to show that, when an agent takes a consideration to call for an action, she takes it to be something normative to be said for an action because of something about what the feature of action it picks out is and not merely because of the motivating attitudes of herself or other agents.)

This claim too requires clarification. First of all, the appeal to what a feature of

157 This is possible on Street's view (see Chapter 3, Sections III and IV). However, it may not be possible on every view in which the standards of correctness or appropriateness relevant to reasons-explanation determine mere permissibility. E.g., such a view might require (for rational permissibility) that at reason identify some end of the agent's that the action will promote. If φ-ing will promote end E, then it is pretty unlikely that not φ-ing will promote E as well (though it may promote some other end).

action "is" (in the claim that a reason calls for an action in virtue of something about what the feature of action it picks out *is*) may be somewhat misleading. It is not meant to suggest that reasons for action must count in favor of actions because of something about their essences, or because of some explanation of what fundamentally constitutes that feature. It is merely meant to provide a fairly loosely understood contrast to a feature's calling for an action in virtue of some agent's desiring it or taking it to be a reason. That is, it is meant to suggest that the explanation of why an action is appropriate or correct, of which a consideration is a part when it calls for action, must be about what the contents of a reason are about (the action and its consequences) and not about the motivating attitudes that particular agents have towards those things. E.g., an explanation of why the fact that an action would promote my climbing Kilimanjaro calls for that action should involve facts about what climbing Kilimanjaro will involve—beautiful vistas, feats of strength, and the like—rather than my, or anyone else's, desires for any of these things.

Furthermore, I want to leave open the possibility that the value that grounds reasons is something that seems to have little to do with what the particular features of actions that appear in everyday reasons are. E.g., in Chapter 2, I consider the possibility that this value is some notion of rational nature itself. This would seem to have little to do with what, e.g., climbing Kilimanjaro is. The idea here is that the fact that an action would promote my climbing Kilimanjaro calls for an action because it is part of an explanation of how my action promotes, preserves, or respects rational nature. That explanation does not rely on facts about particular agent's attitudes towards climbing Kilimanjaro, towards its various features, or towards rational nature. Instead, it relies on something about what climbing

Kilimanjaro would involve. (E.g., that climbing Kilimanjaro would be an exercise in perseverance explains why it promotes, preserves, or respects rational nature.)

The claim that this explanation does not rely on facts about particular agents' motivating attitudes is also over-simplified and caching it out precisely is tricky. There are ways in which the motivating attitudes and reason-taking activity of particular agents does seem to create reasons (such that those motivating attitudes will appear in an explanation of why the relevant considerations count in favor of acting). In the remainder of this section, I will discuss some of these cases and explain how they can be understood so that they do not undermine the basic idea that agents take considerations to call for an action because of something about what the features of action they pick out are and not merely because of particular agents' motivating attitudes. (This should also help make clearer what this claim about "calls for" means.)

For example, as mentioned in earlier chapters, in setting one end rather than another, an agent may be able to make it the case that she has reasons to pursue that end rather than others. An agent has reasons to pursue the project that she has actually undertaken that she does not have to pursue the possible projects that she has not undertaken (other things equal). Or at least it seems that, in choosing to pursue one project rather than another, an agent makes it the case that her reasons to pursue the project she has chosen are much stronger than her reasons to pursue the project that she does not choose (again, other things equal). So, it looks like agents setting ends have something like the reason creating power that Korsgaard attributes to them. If so, then it seems that an agent's motivating attitudes (in this case end setting—something like reason-takings) must play a role in explaining these

reasons (or the differences in their weights).

Related to this, it is plausible to suppose that an agent can set one end rather than another, at least in some cases, on the basis of what she prefers. We might think that we can explain this by reference to something about the preferred end—e.g., that it will bring the agent pleasure—and thus avoid appealing to the agent's desires. However, it also seems at least initially possible that we will have to understand pleasure in terms of preference or desire—e.g., something is pleasurable because the agent wants or prefers it. If so, then an agent's motivating attitudes (her preferences) play a role in explaining the difference in the reasons she has to do what she prefers and those she has to do (or not to do) that which she does not prefer.

In either case, the explanation of the agent's reasons can achieve the necessary independence from particular agents' motivating attitudes if the appeal to these attitudes is playing a secondary sort of role in this explanation. That is, in order to regard the fact that an action will promote E as calling for performing that action, an agent must take there to be *some* explanation of why E is called for that does not depend on particular agents' motivating attitudes. There must be some explanation of why the agent's end is called for that does not depend on her, or any other agent's, happening to want it (or wanting something else that it will bring about). (Perhaps E promotes another of the agent's ends (G). In this case, if the fact that an action promotes E is to call for action because of E's relation to G, then there must be some explanation of why G (or some yet further end that G promotes) is called for that does not depend on any agents' motivating attitudes towards G. (See Chapter 2, Section VI)

Similarly, if the agent prefers some purpose P over some other purpose O, and the fact that an action promotes P is to call for performing that action, there must be some explanation of why that an action promotes P calls for that action that does not appeal to her preference for P. (This is not to say that this explanation must reveal pursuing P to be *better* than pursuing O, merely that there is something normative and positive to be said for pursuing P other than the agent's preferring it (or something else that P will bring about).

Furthermore, even in cases where an agent acts for the sake of a particular feature of action, and there is nothing to be said for the action except that she desires it and it is permissible (and perhaps, that the agent is properly responsive to the norms that determine that the action is permissible), her so acting may create a variety of reasons. For example, we might think that it gives other agents reasons not to interfere with her action and perhaps even some reasons to help her. (And while I have never been particularly moved by such claims, we might think that it gives her reasons to continue to pursue this purpose—to finish what she started.)[158]

Again, I think that there is a way of understanding these cases such that reasons achieve the sort of independence from particular agents' motivating attitudes that I am suggesting is necessary for them to intelligibly be regarded as calling for actions. The idea here is that, when an agent acts for the sake of a particular feature (F) of action, and there is nothing to be said for the action except that she desires it and it is permissible, *that the action has F* does not call for that action (nor can the agent be intelligibly regarded as taking

158 Similarly, I think, some reasons might pick out particular agent's motivating attitudes as the point of the actions that they call for. For example, that an action would cultivate sympathy in some agent—that it would make some agent more likely to desire the well-being of her fellow human beings—might be a reason to perform that action.

it to count in favor of action, at least if she is aware that this is all that can be said for the action). She acts for no reason. The explanation of *our reasons* with respect to her action are grounded in something other than the feature for the sake of which she acts. For example, we might think that our reasons not to interfere in others' actions stem from facts about the value of free activity. I.e., we might think that the fact that an action would undermine this agent's pursuit of F calls for us to not perform that action because the agent's pursuit of F is free activity and not because the agent desires (or even pursues) F.

It may turn out that a particular agent's desires (or choices to pursue particular purposes) may appear somewhere in an explanation of why particular reasons not to interfere in her activity call for not interfering because these desires partially determine what this activity is (and what its point is) and thus what, specifically, would interfere with it. However, in the explanation I suggest above, a desire is not a fundamental part of the explanation of why these considerations call for (or against) acting. It is not the basic fact that explains why we have reasons not to interfere with (or reasons to help) the agent in the activity that she does for no reason.[159]

II. How R accounts for "calls for"

R provides a straightforward explanation of how a reason calls for an action that

159 Alternatively, we might think that these cases shows that desire-satisfaction is one sort of thing that can give rise to reasons. That is, we might think that *some* considerations count in favor of actions merely because particular agents desire the features of action that they pick out. I take my arguments in Chapter 3 to suggest that this at least cannot be the basic explanation of all instances of taking a consideration to call for acting. (In other words, if we are thinking in terms of value, as I suggest we should, then we might grant that desire-satisfaction is a value. Nevertheless, I think we have reason to deny that it is *the* value that grounds all reasons. (Alternatively, we might think that the value of desire-satisfaction is grounded in some further, unifying value. However, this is not required in order to deny that it is the fundamental value that grounds all reasons.)

captures the above claims about what is required for an agent to intelligibly regard a consideration as calling for an action. If R is true, then there is some general feature of action that is to be promoted, preserved, or respected. (Nothing I have said here shows that there could not be several such features, but for simplicity, I will discuss the case where there is only one. Everything I say in this section will apply to the possibility of multiple basic general values as well.)

The first characterization of the sense in which, when an agent acts for a reason, she takes her reason to call for acting was that, when an agent acts for a reason, she takes her action to be subject to standards of appropriate or correctness and takes her reason to be part of an explanation of how her action is appropriate or correct according to those standards. (I.e., she takes it to be part of a justification of her action according to those standards.) If a consideration counts in favor of an action in virtue of picking out something about the action that has value, then this can be easily explained. The standards that determine whether an action is appropriate or correct specify a value—they specify something that is to be promoted, preserved, or respected. So, a reason for action is part of an explanation of how an action is appropriate or correct (to the extent that it is) according to this standard because it picks out something about the action that has this value, and thus shows that, in performing the action, the agent would be promoting, preserving, or respecting that which is to be promoted, preserved, or respected (other things equal).

The second characterization of the sense in which, when an agent acts for a reason, she takes her reason to call for acting was that, when an agent acts for a reason, she takes the consideration to be something normative and *positive*, i.e., she takes it to be something

normative to be said *for* the action. This had two meanings. First, it meant that, when an agent takes a consideration to call for performing an action, she takes it not merely to be part of an explanation of how her action is appropriate or correct, but rather to be a positive part of this explanation. She takes to be a part of this explanation that suggests that the action *is* appropriate or correct (rather than a part that suggests that it is not, but which might be outweighed). If R is true, this is easily explained. If R is true, then the standards that determine whether an action is appropriate or correct specify something that is to be promoted, preserved, or respected. Correctness or appropriateness in action then is (at least in part) doing what promotes, preserves, or respects what is to be promoted, preserved, or respected. So, a reason for action is part of an explanation of why an action is appropriate or correct because it picks out something about the action (or its consequences) that is to be promoted, preserved, or respected. That an action has some feature that is to be promoted, preserved, or respected is a part of this explanation that suggests that the action *is* correct (rather than that the action is incorrect.)

The second meaning of the claim that, when an agent takes a consideration to call for performing an action, she takes that consideration to be something normative to be said for the action was that, when an agent takes a consideration to call for acting, she takes the consideration to be something normative that directs her towards acting, not merely something that reveals the action to be permissible. Again, R provides a straightforward explanation of this. A reason for action picks out something about the action that is to be promoted, preserved, or respected. To say that something is to be promoted, preserved, or respected is to say more than that it is permissible. It is something normative and positive to

be said for the action. So, if a reason reveals an action to have something about it that is to be promoted, preserved, or respected, it reveals that there is something normative and positive to be said for the action (whether or not the action is permissible).

Finally, the third characterization of the sense in which, when an agent acts for a reason, she takes her reason to call for acting was that, when she acts for a reason, she must take the relevant consideration to be something normative to be said for an action because of something about what the feature of action it picks out is and not merely because of the motivating attitudes (desires or reason-takings) of herself or other rational agents. Again, I think that R can provide an explanation of this. However, in order for it to do so, there must be restrictions on what can be the values that reasons pick out. In short, if the values that reasons pick out are themselves explained in terms of particular agents' motivating attitudes, then the explanation that R provides of what it is to take a reason to call for action will (problematically) depend on these attitudes. A reason is something normative and positive to be said for an action because it reveals something about the action to be to be promoted, preserved, or respected. So, as long as we need not appeal to particular agents' motivating attitudes in the explanation of why (or what it is for) that feature of action to be to promoted, preserved, or respected, then we will have achieved the motivating attitude-free explanation that we are looking for.

III. Why we need R to account for "calls for"

So, R can account for the sense in which, when an agent acts for a reason, she takes her reason to call for her action. But why do we *need* R in order to account for this?

If we accept R, then we have an explanation of why the features of actions that reasons pick out are something normative and positive to be said for actions. They reveal actions to promote, preserve, or respect what is to be promoted, preserved or respected. We also have an explanation of which of the many possible features of action can be cited by reasons—we have an explanation of where the content of reasons comes from. For example, if beauty has value, then a variety of other features of actions (of varying levels of specificity) have value—e.g., the experience of beautiful vistas, climbing, Kilimanjaro, and turning right at the big rock, or the creation of beautiful objections, painting a picture, and adding three more brush strokes. These features of actions are cited by reasons under the right circumstances. That an action has one of these features counts in favor of performing it because, under the circumstances, it explains how the action promotes what is to be promoted (in the example, beauty).

If we deny R (if we want to avoid grounding reasons in values) then we need an alternative explanation of where the content of reasons comes from. We might take a reasons-as-primitive view and think that certain considerations just present themselves to us as calling for actions without further explanation. If we want to avoid this sort of reasons-as-primitive view, then it seems that the only available option to fill in the content of reasons is some notion of desire. That is, it would seem that the only thing left that can determine which of the many possible features of action can be in the content of a reason for action are agents' desires. A reason for an agent to perform some action would then pick out something about that action that the agent desires.

However, if we deny R, then a desire cannot be normative in an important sense. A

desire cannot be a perception or registration of value. To desire something about an action or its consequences cannot involve representing that thing as to be promoted, preserved, or respected. Nor can desire aim to track what is to be promoted, preserved, or respected in the way we might say that belief aims to track truth. (A claim that beliefs aim to track the truth need not claim that beliefs represent anything as true though it may require that to have a belief is to "take" something to be true.) If desires were in some such way responsive to value, then the claim that a reason for action picks out something that the agent desires will commit us to R. (This is not to say that to desire X does not involves regarding X as having some value. It is just to say that, if it does, then a view of reasons that bases reasons in desires is not a view that denies R. If desires involve regarding their objects as having value, then a reason may well count in favor of an action because it picks out something about that action that the agent desires, but this will likely be true because desires pick out things that agents' regard as having value. Reasons would pick out desired features of action because both reasons and desires track (or aim to track) values), but it would be the value and not the desires that yield something normative and positive to be said for actions.)

So, if we are to deny R, we seem to need to appeal to desires that are non-normative in the sense just discussed to give us the content of reasons. On this view, a desire for X is merely a state of being moved to bring X about that is subject to no standards of correctness or intelligibility.[160] (Nothing is an inappropriate object of desire or unintelligible as an

160 This is the conception of desire that Korsgaard, Street, and Schroeder have (that Schroeder has this conception of desire may be somewhat contentious given his remarks about desires as salience-takings of reasons in Chapter 8 of *Slaves of the Passions*). Nevertheless, it might be implausibly strong. We might think there are *some* constraints on appropriateness as an object of desire. For example, perhaps desiring some state of affairs involves taking that state of affairs to be possible. (This might distinguish desires from wishes.) We might even think that desires are subject to some weak consistency constraints. (I think this claim is rather dubious. It seems intelligible to desire many conflicting things in a way that it may not be intelligible to take oneself to have many conflicting reasons.) However, I think we need not

object of desire, though it may well be an inappropriate thing to *pursue* or an unusual or unfamiliar thing to want.)

But if desires are non-normative in this way, then we need something further to explain how reasons are something normative and positive to be said for performing an action (and thus how they call for actions). How to do this is suggested by the characterization just given of the sense in which desires are non-normative. If nothing is an inappropriate object of desire, then, in order to show how reasons that pick out desired features of actions suggest that actions are appropriate, it seems that we must look to standards that determine which desires it is appropriate to act on (perhaps under a set of circumstances). Desires propose features of actions to agents as potential reasons, and we have some further standards for determining which of these proposed reasons actually count in favor of actions.

Those who deny R tend to take these standards to be norms that determine when acting for the sake of some desired feature of actions is permissible. I have suggested (at this point, repeatedly and at length) that such standards will not get us reasons that are something normative and *positive* to be said for actions. That an action is permissible does not itself call for performing the action. That an action is permissible is not itself something normative and positive to be said for the action. So, if a desire does not represent an action to an agent as having something normative and positive to be said for it, then adding the fact that it is rationally permissible to act for the sake of the desired feature of action is not going

worry much here about whether there are such standards for desire. If we deny R, then desires cannot be subject to standards demanding that they track value. But we might take any of the standards that determine when it is appropriate to act on desires that a view that denies R must later appeal to, in order to explain the sense in which reasons call for acting, and make it a standard on desires themselves. My points should apply regardless of whether we take these standards to apply to the desires themselves or to what we do with the desires.

to get us anything normative and positive to be said for so acting.

(This is not to say that standards of rational permissibility have no place in determining which considerations are something normative and positive to be said for actions. For example, suppose we had a view on which desires do pick out features of action that are to be promoted and in some way represent or register these features as such. We might then have a view in which the fact that these features of actions are to be promoted in general is not itself enough to show that they call for actions. We might think that they only call for action—they are only something normative and positive to be said for an action—if it is also rationally permissible to perform that action for the sake of those features. For example, we might have a view on which beauty is a value, but also on which that an action would yield a beautiful experience only calls for action when acting for the sake of the beautiful experience would be permissible under the circumstances. So, that pushing someone off a cliff would get be a better view of a beautiful vista may show that my action promotes a value, but, on this sort of view, it will not call for the action because killing for the sake of a beautiful view is not permissible. On such a view, rational permissibility serves as a necessary condition on "calling for" but it still does not explain how, when a consideration does call for an action, it is something normative and *positive* to be said for the action. We still appeal to the fact that certain features of action are to be promoted (when permissible) to explain this.)

So, standards that determine which desired features of action are rationally permissible to act for the sake of are not enough to get us reasons out of non-normative desires. That a feature of action moves an agent to bring it about and that it is rationally

permissible to act for the sake of that feature do not show that the fact that the action has that feature calls for the agent's performing it. So, it would seem that we need to add something to the standards that determine which desired features of action for the sake of which it is rationally permissible to act in order to make compliance with those standards into something that is something positive to be said for performing an action. (Or more accurately, we need to add a standard such that compliance with those standards will make the fact that an action has some particular feature something positive to be said for performing the action.)

I suspect that any standard that can do this will amount to a claim that something is to be promoted, preserved or respected. It is not clear what else this standard could be. It is not clear what else could be or yield anything positive to be said for actions except for claims identifying features of actions as to be promoted, preserved, or respected. And if so, then it seems that we cannot avoid taking on some version of R to explain the counting in favor of relation. There is, of course, the possibility that this merely demonstrates a lack of imagination on my part. However, it at least seems that it is quite difficult to answer, without appealing to value, the question of, out of all the normative things to be said with respect to an action, which are the positive ones. Obviously, the positive normative things to be said are not that the action is wrong, or that it is unwise, or that it undermines one's own (permissible) ends. Less obviously, they are not that this action is permissible or that it violates no requirements, moral, prudential, or otherwise. (Unless of course the requirement in question is a requirement to do only what promotes some value, in which case if an action does not violate the requirement, this is because it promotes that value.) The positive

normative things to be said with respect to an action seem to be things like that the action will bring about something that is to be brought about, that is to be kept about, that is to be admired, that is to be appreciated. But these just seem to be ways of saying that the action or something about it is to be promoted, preserved, or respected. (Or maybe some of these should be added to the list of "to be"s that describes the notion of value. Either way, what is positive to be said with respect to an action seems to have something to do with value.)

However, this will not yet get us the sort of R that I suggested in the beginning of this document that I want to defend. It will not yet get independence from particular agent's motivating attitudes for reasons. What is to be promoted could just be desire-satisfaction (Schroeder). Or, what is to be promoted could be whatever an agent permissibly decides to promote (Korsgaard). In either case, the resulting account of reasons does not just appeal to non-normative desires and standards of rational permissibility. Each includes some further standard or normative claim: promote desire-satisfaction or promote what agents permissibly decide to promote. E.g., we might think that an agent A's desire for a particular feature of action F plus the fact that A's φ-ing for the sake of F is rationally permissible is not enough to show that the fact that φ-ing has F calls for A's φ-ing. Yet, we might think that A's desire for F plus the fact that A's φ-ing for the sake of F is rationally permissible *plus the fact that desire-satisfaction is to be promoted* is enough to show that the fact that φ-ing has F calls for A's φ-ing.

In the preceding chapters, I have discussed a variety of thoughts that support the claim that for an agent to intelligibly take a consideration to call for an action, she must take that consideration to be something normative to be said for an action because of something

about what the feature of action it picks out is and not merely because of the motivating attitudes (the desires or reason-takings) of herself or other particular rational agents. If this is true, then we will need a version of R in which the value that grounds reasons is something other than desire-satisfaction or whatever agents permissibly decide to promote.

In Chapter 2, I discussed three arguments in favor of a requirement to s-share reasons which I then argued in Chapter 3 support a motivating attitude-free understanding of "calls for" (and which thus support a motivating attitude-free notion of value and of R). The first of these arguments claimed that the intelligibility of certain sorts of disagreements—disagreements about what is worthwhile—seems to presuppose that reasons do not fundamentally depend on particular agents' motivating attitudes towards the features of action that they pick out, but instead depend on motivating attitude-independent values.

The second of these arguments claimed that the fact that some feature of action is permissibly desired by some particular agent ("because I want it and it is permissible") does not seem to be an adequate justification for the claim that reason citing that feature makes on agents. This is clearest when the agent receiving the demand is not the agent who has the desire, but we might think that this suggests that "because I want it and it is permissible" does not justify a demand on myself anymore than it does on anyone else. It does not itself make my action called for. Finally, the third of these arguments claims that chains of the why question that asks for reasons (posed to cooperative recipients) tend to bottom out in claims about value, admissions that the actions in question were done for no reason, or unintelligible answers.

In Chapter 4, I discuss two common challenges to a Humean account of reasons.

(Most of the chapter is spent explaining Schroeder's solutions to these challenges and why they do not work, but the challenges themselves provide additional support for a motivating attitude-free understanding of "calls for.") The first challenge is that, at least for some reasons, a particular agent's desires seems to be the wrong sort of thing to think about when deliberating what to do and thus the wrong sort of thing to explain why the agent has the reason. This problem is most pressing with respect to reasons to help others—some of my reasons to help you seem not to depend on *my* desires and the likelihood that helping you would promote their satisfaction. Related to this, the second challenge is that there are at least some reasons that agents seem to have regardless of their (or anyone else's) particular motivating attitudes—the agent-neutral reasons. So, it would seem that at least some reasons cannot be grounded in the value of the promotion of particular agent's particular desires. (Even in Schroeder's view, I suggested, agent-neutral reasons end up being grounded in a general value of agential efficacy and not in particular desires.)

Each of these arguments, I think, points out that in our ordinary use of reasons we do not regard our reasons as calling for actions because of any particular agents' motivating attitudes towards them. When we deliberate about what to do, we do not think about our attitudes towards features of actions and their consequences. We just think about those features of actions and their consequences. (We think about the beauty and physical challenge of climbing Kilimanjaro and not our desires for these things.) When we offer our reasons to each other as justifications for our actions, or make claims about reasons in asking that we be treated in certain ways, we talk about features of actions and the motivating attitude-free values that they promote rather than about our (or other particular

agents') attitudes towards those features. (Of course, we do also sometimes talk about our desires. However, we might think that either the desire-talk or the value-talk is a shorthand or sloppy way of talking about the other. We might then take part of the point of the arguments for a motivating attitude-free notion of "calls for" to be to show that the value-talk is the more accurate of the two ways of speaking about our reasons.)

In short, in our actual use of reasons for action, we seem to regard them as calling for actions in virtue of something about what those actions are or will bring about and not as calling for actions because of any particular agents' motivating attitudes towards them. If this is mistaken, then we have to give up much of the everyday conception of reasons and reasons-justification. I find this to be a persuasive reason to think that a reason calls for an action in virtue of something about what those actions are or will bring about and not as calling for actions because of any particular agents' motivating attitudes towards them. (And, for the reasons discussed above, I think that this pushes us to adopt a version of R where value is understood to be not fundamentally dependent on particular agents' motivating attitudes).

IV. Final Remarks: How many explanations?

However, others like Korsgaard and Schroeder seem to think that they can (1) preserve the ordinary picture of practical reasoning on which "calls for" does not seem to depend on particular agent's motivating attitudes, (2) explain the counting in favor of relation in terms of particular agents' motivating attitudes, and (3) do this without claiming that agents are ordinarily in any error about their reasons. The basic idea behind their

approach is that there is a distinction between the explanations of which reasons for action are a part (explanations of what an agent should do) and the explanation of why reasons count in favor of actions. (Schroeder makes explicit that he takes there to be this distinction, but I think that Korsgaard can be understood as doing something like this as well. It is because I find this to be an intriguing possibility that I have spent so much of this dissertation discussing their views.)

The explanations of which reasons are a part are explanations of what the agent should do in a particular situation. Why the thing to do in any given situation is the thing to do in that situation is explained in terms of the considerations that count for and against so acting and how they weigh against each other, and how they support or undermine each other. These considerations cite features of actions and not (usually) particular agents' motivating attitudes. This makes a reasons-explanation of what we should do free of references to particular agents' motivating attitudes. This seems to capture the basic idea behind the claim that our everyday conception of reasons and practical reasoning does not support reasons grounded in particular agents' motivating attitudes—*what agents should do* is not (perhaps with some exceptions) determined by what motivating attitudes they or any other particular agents happen to have. Furthermore, this view does not claim that the ordinary conception of reasons and practical reasoning is in error. On this view, if φ-ing is the thing to do, then φ-ing really is the thing to do because that it has features F, E, and G counts in favor of φ-ing (and there are no considerations that count against acting that are sufficient to outweigh or undermine F, E, and G).

On the other hand, Korsgaard and Schroeder think that particular agent's motivating

attitudes appear in a different and further explanation—an explanation of why the reasons that determine what agents should do count in favor of (or against) actions. The Korsgaard/Schroeder thought seems to be that particular agents' motivating attitudes can play a fundamental role in this explanation without infecting the motivating attitude-free reasons-explanation of what agents should do. Thus, this explanation of the counting in favor of relation need not conflict with the everyday conception of reasons and practical reasoning on which what agents should do is not determined by these attitudes. (This is reflected in Korsgaard's thought that an agent can take an interest in a feature of action for its own sake because she recognizes that she is required to do so when another rational agent (permissibly) makes it his purpose in acting. It is also reflected in Schroeder's thought that a Humean account of reasons can avoid the charge that it makes practical reasoning objectively self-regarding by making desires background conditions on reasons.)

I suspect that this distinction is untenable. (I also think that this issue deserves further consideration and it may be that what I have to say on the subject amounts to little more than a declaration that I do not see the distinction.) Reasons explain why we should do what we should do by counting in favor of (or against) actions. When I should φ, this is because there are considerations that count in favor of φ-ing (that outweigh considerations not to φ, etc). So, whatever explains what makes a consideration count in favor an action seems to be part of an explanation of what we should do. It is part of what makes it the case that we should do what we should do, what grounds the claims that reasons makes on agents. If I should φ because some consideration C counts in favor of φ-ing, and a reason counts in favor of an φ-ing because it shows φ-ing to be X, then it seems true that I should φ

because φ-ing would be X. (At least it seems natural to say this. If we explain the counting in favor of relation in terms of desires, it does not seem infelicitous to say that I should φ "because I want F".)

I may not need to think about X in order to figure out or explain what I should do, but X seems to be already at work in this explanation of what I should do whether or not I think about it. It is already in the notion of counting in favor of because what it is for a consideration to count in favor of φ-ing is just for it to show φ-ing to be X. It is probably correct that the explanation of why the consideration counts in favor of acting does not appear in the consideration that count in favor of acting. (I am not claiming that the reason for φ-ing is "φ-ing would have value.") Nevertheless, I think it hard to see how X is not part of the explanation of what I should do. So, if some X seems like the wrong sort of consideration to explain what we should do, then I find it hard to see why it would matter whether that X appears in the content of a reason or in the explanation of why that reason counts in favor of an action. Either way it is explaining what we should do.

More importantly, even if we think that there might be an important difference between a reasons-explanation of what we should do and an explanation of what makes considerations count in favor of actions, it seems incoherent to have certain kinds of conflict between the two explanations.[161] I suspect that such conflicts are inevitable if we try to

161 By contrast, anytime I have a reason to φ, there are a variety of necessary conditions on my having this reason that may be part of some explanation of why I have the reason, that I need not consider whenever thinking about what to do, and that do not seem to be in any conflict with the idea that my reason counts in favor of acting because of something about what the feature of action it picks out is. E.g., that carrots promote good eyesight counts in favor of my eating carrots. In order for this to be true, I must be alive. In order for this to be true, a variety of biological facts must be true. My heart must be pumping adequately oxygenated blood. My neurons must be firing. Etc. But these facts do not seem to be in any conflict with an explanation of why I should eat carrots that appeals to what good eyesight is and can do for me.

explain the counting in favor of relation in terms of particular agents' motivating attitudes. So, appealing to the distinction to suggest that the apparent conflict between these two explanations is not worrisome will not work. This is what happens on Korsgaard's view.

On Korsgaard's view, when that an action has some desired feature of action is a reason for an agent to perform that action, the agent takes the feature (or something about it) to make a claim on her for its own sake (because of something about what it is). Then, Korsgaard claims that this feature makes a claim on the agent for its own sake because that feature is permissibly pursued by some rational agent. But, I think, this simply cannot work. Either the feature makes the claim for its own sake, or it makes the claim because it is permissibly pursued by some rational agent. It cannot do the former because it does the latter. We cannot explain "for its own sake" by appealing to anything other than what "it" is. (Similarly on Schroeder's view, we want agents to see at least some considerations as reasons for action simply because of the features of action that they pick out. E.g., I should see the fact that you need help as making some claim on me to help you without thinking about whether I want to help you. He proposes to explain the counting in favor of relation by claiming that a consideration counts in favor of an action just because it shows that the action will promote the satisfaction of one of the agent's desires. So, it would seem that she is trying to say that the needs of others make claims on us for their own sake because they promote the satisfaction of our own desires.)

This possibility for problematic conflict between the two explanations is reflected in a constraint on explanations of the counting in favor of relation that I appeal to, somewhat implicitly, throughout the dissertation. This constraint is that the explanation of why or in

virtue of what a consideration counts in favor of an action must be such that an agent can be aware of it and continue to regard a consideration as calling for an action, if it in fact does so. (This constraint appears in the argument that Korsgaard cannot account for a requirement to s-share reasons in Chapter 2, Section V, and the argument that Schroeder cannot address the objectively self-regarding objection to a Human account of reasons in Chapter 4, Section IV. It also appears in my argument against Street's view (Chapter 3, Sections III and IV) in the idea that an agent would not be able to continue to regard a consideration as calling for an action if she knows that her reasons-justifications for actions bottom out in considerations that merely move her to act.)

Reasons for action show why an action is appropriate or correct (to the extent that it is). The explanation of what makes considerations count in favor of actions, whatever that is, is part of the complete story of why a given action is appropriate or correct (to the extent that it is). Actual agents do not always (or perhaps often, or even ever) consider this complete story of why their actions are appropriate or correct in light of their reasons. However, motivation by a reason is distinct in that it involves some recognition that there is such a story and that one's reason is a part of it. Because this reasons-motivation contains this recognition, an agent should be able to know the complete story and still retain this motivation (still take considerations to count in favor of actions when they in fact do).

The constraint is explained by the fact that there is no meaningful distinction between the two kinds of explanation discussed above, or at least that these explanations cannot conflict. However, we may come to discover the facts about reasons-explanation that ground the constraint by first observing that practical reasoners seem to be under the

constraint and thinking about why this would be. That is, for the most part, what I have done with the constraint in previous chapters. I have observed that there seems to be no sharp divide in practical deliberation between the two kinds of thinking—thinking about what to do and thinking about why considerations count in favor of actions. The two can bleed into each other, in particular because the latter can influence the former. Agents can ask themselves questions about why their reasons make claims on them and the answers to these questions can change which actions these agents take to be called for and what they take to call for acting. That agents can do this creates problems for views that attempt to preserve the everyday conception of reasons as calling for actions because of something about the features of actions themselves that those reasons pick out and yet explain the counting in favor of relation in terms of particular agent's motivating attitudes. Once an agent sees both explanations, she will see that they conflict and will have to give up one or the other.

Bibliography

Anscombe, G. E. M. *Intention*. 2nd ed. Cambridge, Mass.: Harvard University Press, 2000.

Blackburn, Simon. "Practical Tortoise Raising." *Practical Tortoise Raising and Other Philosophical Essays*. Oxford: Oxford University Press, 2010. 7-25.

Blackburn, Simon. *Ruling Passions: A Theory of Practical Reasoning*. Oxford: Clarendon Press, 1998.

Boyle, Matt and Lavin, Douglas. "Goodness and Desire." *Desire, Practical Reason, and the Good.* Oxford: Oxford University Press, 2010. 161-201.

Bratman, Michael E. *Intention, Plans, and Practical Reason.* CSLI Publications, 1999.

Broome, John. "Is rationality normative?" *Disputatio* 2, no. 23 (2007): 1-18.

Broome, John. "Reasons." *Reason and value: Themes from the moral philosophy of Joseph Raz.* Oxford: Oxford University Press, 2004. 28-55.

Brunero, John. "Reasons as Explanations." *Philosophical Studies* (2012): 1-20.

Burge, Tyler. "Primitive Agency and Natural Norms." *Philosophy and Phenomenological Research* 79, no. 2 (2009): 251-278.

Buss, Sarah. "Weakness of Will." *Pacific Philosophical Quarterly* 78, no. 1 (1997): 13-44.

Buss, Sarah. "What practical reasoning must be if we act for our own reasons." *Australasian Journal of Philosophy* 77, no. 4 (1999): 399-421.

Chang, Ruth. "Can Desires Provide Reasons for Action?" *Reason and Value: Themes from the Moral Philosophy of Joseph Raz.* Oxford: Oxford University Press, 2004. 56-90.

Clark, Philip. "Aspects, Guises, Species, and Knowing Something to be Good." *Desire, Practical Reason, and the Good.* Oxford: Oxford University Press, 2010. 234-244.

Crisp, Roger. "Value, reasons and the structure of justification: how to avoid passing the buck." *Analysis* 65, no. 285 (2005): 80-85.

Dancy, Jonathan. "Enticing Reasons." *Reason and Value: Themes from the Moral Philosophy of Joseph Raz.* Oxford: Oxford University Press, 2004. 91-118.

Dancy, Jonathan. "Should we pass the buck?." *Royal Institute of Philosophy Supplement* 47 (2000): 159-173.

Davidson, Donald. *Essays on Actions and Events*. Oxford: Clarendon Press, 1980.

De Sousa, Ronald B. "The Good and the True." *Mind* LXXXIII, no. 332 (1974): 534-551.

Finlay, Stephen. "Explaining Reasons," *Deutsches Jahrbuch für Philosophie*, Vol. 4. Hamburg: Meiner, 2012: 112-26.

Finlay, Stephen. "The reasons that matter." *Australasian Journal of Philosophy* 84, no. 1 (2006): 1-20.

Foot, Philippa. "Goodness and Choice." *Virtues and Vices and Other Essays in Moral Philosophy*, University of California Press, 1978. 132-147.

Foot, Philippa. *Natural Goodness.* Oxford: Oxford University Press, 2001.

Foot, Philippa. "Reasons for Action and Desires." *Virtues and Vices and Other Essays in Moral Philosophy*, University of California Press, 1978. 148-156.

Frankfurt, Harry G. "Freedom of the Will and the Concept of a Person." *The Journal of Philosophy* 68, no. 1 (1971): 5-20.

Frankfurt, Harry G. *The Importance of What We Care About: Philosophical Essays*. Cambridge: Cambridge University Press, 1988.

Geach, Peter T. "Good and Evil." *Analysis* 17, no. 2 (1956): 33-42.

Gibbard, Allan. *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge, Mass.: Harvard University Press, 1990.

Herman, Barbara. "Bootstrapping." *Contours of Agency.* Cambridge, MA: The MIT Press, 2002.

Heuer, Ulrike. "Explaining Reasons: Where Does the Buck Stop?" *Journal of Ethics & Social Philosophy* 1, No. 3 (2005).

Heuer, Ulrike. "Raz on Values and Reasons." *Reason and Value: Themes from the Moral Philosophy of Joseph Raz.* Oxford: Oxford University Press, 2004. 129-152.

Hieronymi, Pamela. "Reasons for Action." *Proceedings of the Aristotelian Society* 111 (2011): 407-427.

Hieronymi, Pamela. "Two Kinds of Agency." *Mental Actions*. Oxford: Oxford University Press, 2009. 407-427.

Hursthouse, Rosalind. “Arational Actions.” *The Journal of Philosophy* 88, no. 2 (1991): 57-68.

Kolodny, Niko. “Why be Rational?” *Mind* 114, no. 455 (2005): 509-563.

Korsgaard, Christine M. *The Constitution of Agency: Essays on Practical Reason and Moral Psychology*. Oxford: Oxford University Press, 2008.

Korsgaard, Christine M. *Creating the Kingdom of Ends*. Cambridge: Cambridge University Press, 1996.

Korsgaard, Christine M. *Self-Constitution: Agency, Identity, and Integrity*. Oxford: Oxford University Press, 2009.

Korsgaard, Christine M. *The Sources of Normativity*. Cambridge: Cambridge University Press, 1996.

Lavin, Douglas. “Practical Reason and the Possibility of Error.” *Ethics* 114, no. 3 (2004): 424-457.

Lawrence, Gavin. “Reason, Intention, and Choice: An Essay in Practical Philosophy.” *Royal Institute of Philosophy Supplement* 54, no. 1 (2004): 265-300.

Nagel, Thomas. *The Possibility of Altruism*. Oxford: Clarendon Press, 1970.

Nagel, Thomas. *The View from Nowhere*. New York: Oxford University Press, 1986.

Parfit, Derek. *On What Matters*. Oxford: Oxford University Press, 2011.

Parfit, Derek. *Reasons and Persons.* Oxford: Oxford University Press, 1984.

Quinn, Warren. “Putting rationality in its place.” *Morality and Action*. Cambridge: Cambridge University Press, 1993. 228-55.

Quinn, Warren. “Rationality and the human good.” *Social Philosophy and Policy* 9, no. 02 (1992): 81-95.

Railton, Peter. “How to Engage in Reason: The Problem of Regress.” *Reason and Value: Themes from the Moral Philosophy of Joseph Raz.* Oxford: Oxford University Press, 2004. 176-201.

Raz, Joseph. *Engaging Reason: On the Theory of Value and Action*. Oxford: Oxford University Press, 1999.

Raz, Joseph. “The Myth of Instrumental Rationality.” *Journal of Ethics & Social Philosophy* 1, no. 1 (2005): 1-28.

Raz, Joseph. “On the Guise of the Good.” *Desire, Practical Reason, and the Good.* Oxford: Oxford University Press, 2010. 111-137.

Raz, Joseph. *Value, Respect, and Attachment.* Cambridge: Cambridge University Press, 2001.

Scanlon, T.M. “Being Realistic about Reasons.” *The Locke Lectures* (2009).

Scanlon, T.M. “Reasons, A Puzzling Duality?” *Reason and Value: Themes from the Moral Philosophy of Joseph Raz.* Oxford: Oxford University Press, 2004. 231-246.

Scanlon, T.M. *What We Owe to Each Other.* Cambridge, Mass.: Belknap Press of Harvard University Press, 1998.

Schapiro, Tamar. “The Nature Of Inclination.” *Ethics* 119, no. 2 (2009): 229-256.

Schroeder, Mark. “Reasons and Agent-Neutrality.” *Philosophical Studies* 135, no. 2 (2007): 279-306.

Schroeder, Mark Andrew. *Slaves of the Passions*. Oxford: Oxford University Press, 2007.

Setiya, Kieran. “Explaining Action.” *Philosophical Review* 112, no. 3 (2003): 339-393.

Setiya, Kieran. *Reasons Without Rationalism*. Princeton, N.J.: Princeton University Press, 2007.

Setiya, Kieran. “Sympathy for the Devil.” *Desire, Practical Reason, and the Good.* Oxford: Oxford University Press, 2010. 82-110.

Stampe, Dennis W. “The Authority of Desire.” *The Philosophical Review* 96, no. 3 (1987): 335-381.

Stocker, Michael. “Desiring the Bad: An Essay in Moral Psychology.” *The Journal of Philosophy* 76, no. 12 (1979): 738-753.

Stocker, Michael. "Raz on the Intelligibility of Bad Acts." *Reason and Value: Themes from the Moral Philosophy of Joseph Raz.* Oxford: Oxford University Press, 2004. 303-332.

Street, Sharon. “A Darwinian Dilemma For Realist Theories Of Value.” *Philosophical Studies* 127, no. 1 (2006): 109-166.

Street, Sharon. “Constructivism about Reasons.” *Oxford Studies in Metaethics* 3 (2008): 207-45.

Street, Sharon. “What is Constructivism in Ethics and Metaethics?” *Philosophy Compass* 5, no. 5 (2010): 363-384.

Tappolet, Christine. “Through Thick and Thin: Good and its Determinates.” *Dialectica* 58, no. 2 (2004): 207-221.

Thompson, Michael. *Life and Action: Elementary Structures of Practice and Practical Thought*. Cambridge, Mass.: Harvard University Press, 2008.

Väyrynen, Pekka. “A Wrong Turn to Reasons?” *New Waves in Metaethics*. 2010.

Velleman, James David. *Practical Reflection*. Princeton, N.J.: Princeton University Press, 1989.

Velleman, James David. *The Possibility of Practical Reason*. Oxford: Clarendon Press, 2000.

Watson, Gary. “Free Agency.” *The Journal of Philosophy* 72, no. 8 (1975): 205-220.

Williams, Bernard. “Internal and external reasons.” *Moral Luck*. Cambridge University Press, 1981, 101-113.