

CONTRASTIVE REASONS

by

Justin Snedegar

A Dissertation Presented to the
FACULTY OF THE USC GRADUATE SCHOOL
UNIVERSITY OF SOUTHERN CALIFORNIA
In Partial Fulfillment of the
Requirements for the Degree
DOCTOR OF PHILOSOPHY
(PHILOSOPHY)

May 2013

Copyright 2013

Justin Snedegar

*to Emmy, for love and support, and for always sharing my intuitions
and to my parents, for always being proud of me*

Acknowledgements

This dissertation, in its current form, was conceived about two years ago. But work towards it really began after my first year of graduate school at USC, after a seminar on deontic modality with my eventual advisor, Mark Schroeder. I spent much of my second year working with Mark, developing a contrastive view of deontic modals. I also spent a lot of time talking to Steve Finlay about this work, and we eventually co-authored a paper on the topic. So Mark and Steve have been with me from the earliest stages of this dissertation.

Thanks to Steve for always being very encouraging, and for, at the same time, offering sharp criticisms. A common experience was getting a set of comments, which began with something like ‘This is great! I think you’re basically right about everything’, followed by pages of very serious criticisms and suggestions. I find myself agreeing with Steve more than probably any other philosopher, and it was great to have someone who thought about things from much the same angle as I do around. Jake Ross has also been involved from early on. Back when I thought I would write a much longer dissertation, including lots of stuff on deontic modals, he encouraged me to focus on reasons. His work also influenced me a great deal.

I also benefited from talking with many of the other faculty members at USC. I will just mention a few in particular. Thanks to Scott Soames, for playing such a big role in my development early on in my career at USC, and for discussions about the material here in particular. Ralph Wedgwood joined the faculty at USC after I was well into this project, but showed a lot of enthusiasm for my work, and provided useful feedback. Thanks also to Barry Schein, both for teaching me most of what I know about formal semantics and for serving as an outside member on my committee.

Thanks to all of the graduate students for contributing to the great community at USC; thanks especially to Indrek Reiland, Julia Staffel, Josh Crabill, Alida Liberman, Lewis Powell, Johannes Schmitt, and Sam Shpall have all contributed

significantly. Ben Lennertz has read most of my work, and talked with me about philosophy, fantasy baseball, and life (in that order) frequently over the last five years. I have learned a lot over the last few years from talking with Shyam Nair. And Justin Dallmann helped me with several parts of this dissertation, especially with the probabilities in Chapter 4. Thanks to him, and to Amanda Dallmann, for feeding Emmy and me many Friday nights and for being great friends.

Outside of USC, two people especially have been extremely encouraging and helpful. Fabrizio Cariani has given me great feedback on much of my work, including big chunks of material from this dissertation. Walter Sinnott-Armstrong has been more consistently enthusiastic about my work than anyone, including me. He has given me very insightful comments on virtually every chapter of this dissertation. Besides Mark and Steve, he has probably had the biggest influence on the following pages. Thanks are also due to the philosophy department at West Virginia University, between 2005 and 2008, when I was an undergraduate there—especially to Sharon Ryan and Andy Cullison, who taught me how to do philosophy and encouraged me to pursue it.

My greatest professional debt is to my advisor, Mark Schroeder. I’ve been working very closely with him since the beginning of my second year of graduate school. He bought into the project from the start, and his dedication to it was crucial in helping me see it through. Mark has influenced the dissertation immensely; this will be obvious from the beginning. The ideas here are mine, but I would not have had most of them if Mark hadn’t dragged them out of me. His own work has served, at different times, as part of my arsenal, as a target, and as an inspiration. I can’t imagine a better or more invested advisor; thanks, Mark.

Now some personal debts of gratitude. Thanks to my parents for supporting me, encouraging me, and being proud of me. Thanks to my dad, my work has been read by more bikers from West Virginia than is probably the norm for work in metaethics. And my debt to my mom is simply immeasurable. Finally, thanks most of all to my best friend, sounding board, gourmet chef, and wife, Emmy. She has been the best possible partner. Throughout the last five years in Los Angeles, she has been unwaveringly supportive and loving. Without her, this dissertation wouldn’t exist; lots of more important things would also be missing.

Table of Contents

| | |
|---|-----|
| Acknowledgements | iii |
| Abstract | vii |
| Chapter 1 CONTRASTIVISM AND REASONS | 1 |
| 1.1 Reasons and Contrastivism | 2 |
| 1.2 Sinnott-Armstrong’s Argument | 7 |
| 1.3 The Plan | 19 |
| Chapter 2 REASON CLAIMS | 21 |
| 2.1 A Simple Argument | 21 |
| 2.2 A Stronger Argument | 24 |
| 2.3 Contrastivism | 28 |
| 2.4 Negative Reason Existentials | 32 |
| 2.5 Looking Forward | 37 |
| Chapter 3 FAVORING | 38 |
| 3.1 Why Resist Contrastivism? | 38 |
| 3.2 Shallow Contrastivism | 39 |
| 3.3 Favoring | 44 |
| 3.4 Contrastive Reasons and Favoring | 48 |
| 3.5 Looking Forward | 51 |
| Chapter 4 PROMOTION | 53 |
| 4.1 A Problem for the Contrastivist | 54 |
| 4.2 Promotion | 56 |
| 4.3 Contrastive Reasons and Promotion | 64 |
| 4.4 Providing the Constraints | 67 |
| 4.5 Promotion as Probability-Raising | 72 |
| 4.6 Non-Promotional Reasons | 74 |
| 4.7 Conclusion | 76 |

| | | |
|-------------------|--|-----|
| Chapter 5 | INTRANSITIVITY | 78 |
| 5.1 | Transitivity and Reasons | 79 |
| 5.2 | Intransitivity | 82 |
| 5.3 | Two Kinds of Intransitivity | 89 |
| 5.4 | Contrast-Sensitive Importance | 90 |
| 5.5 | Remaining Questions | 95 |
| 5.6 | Conclusion | 104 |
| | | |
| Chapter 6 | WITHHOLDING | 106 |
| 6.1 | Withholding Belief and Contrastive Reasons | 106 |
| 6.2 | Withholding Intention | 115 |
| 6.3 | A Non-Contrastive Alternative | 122 |
| 6.4 | Wrap Up | 127 |
| | | |
| References | | 129 |

Abstract

This dissertation develops and defends a contrastive view of normative reasons, according to which they are fundamentally reasons for one thing rather than another, or for one thing out of some range of alternatives. Reasons for something relative to one set of alternatives may not be reasons for that same thing relative to some other alternatives. This is in opposition to nearly all existing work on reasons, which assumes without argument—even without mention—that reasons are reasons for actions *simpliciter*, independently of what we’re comparing the action to.

I argue that this contrastive view is supported by (i) our *talk* about reasons, (ii) the idea that reasons *favor* the things they are reasons for, and (iii) the idea that reasons involve the *promotion* of things like desires, goals, or values. These arguments build, step by step, to a detailed version of contrastivism that can both claim advantages over the traditional non-contrastive picture and avoid problems that face less well-developed versions of contrastivism.

One thread running through the dissertation is an exploration of the *depth* of contrastivism. Chapter 2 argues that contrastivism goes at least as deep as our use of the word ‘reason’—this word (as used to ascribe normative reasons) should be given a contrastive treatment. Chapter 3 argues that contrastivism actually goes deeper, to the normative reason, or favoring, relation that normative philosophers are concerned with. Chapter 4, though, argues for a somewhat moderate contrastive theory. The normative reason relation is indeed contrastive, but we need not—in fact, we’d better not—adopt the deepest possible form of contrastivism.

Given the central place of reasons in normative philosophy, we should expect that contrastivism will have upshots throughout areas like ethics, practical rationality, and epistemology. In the final two chapters of this dissertation, I start to explore this idea. This is just a start, but it suggests that exploring further applications for contrastivism is likely to be a fruitful area for future research.

Chapter 1

CONTRASTIVISM AND REASONS

When I moved to Los Angeles I lived in Thai Town, so I was surrounded by good Thai restaurants. One of my favorites is Red Corner Asia; so the fact that I'm in the mood for Thai food seems like a reason to go to Red Corner Asia. But Spicy BBQ also serves delicious Thai food. So is the fact that I'm in the mood for Thai food actually a reason to go to Red Corner Asia? It's not clear—after all, I could just as easily get Thai food by going to Spicy BBQ. But now suppose that lots of restaurants are closed tonight, and that my only two options are Red Corner Asia and Paru's Indian restaurant. Then it seems obvious that the fact that I'm in the mood for Thai is a reason to go to Red Corner Asia. Here's a very natural way to describe what's going on: the fact that I'm in the mood for Thai food is a reason to go to Red Corner Asia *rather than* go to Paru's, but it's *not* a reason to go to Red Corner Asia rather than go to Spicy BBQ. Of course, you might think that this doesn't answer the original question: is it a reason to go to Red Corner Asia, or not? But again, that question can seem hard to answer in the situation described, in which I could also get Thai food by going to Spicy BBQ. The *contrastive* question of whether it's a reason to go to Red Corner Asia rather than go to Spicy BBQ, though, is easy to answer. This is a special case of a very general observation: often some consideration seems to be a reason to perform an action rather than some actions, but not a reason to perform it rather than some other actions.¹ A theory of reasons needs to explain this.

In this dissertation, I defend a theory of reasons that takes this observation seriously. On this theory, called *contrastivism* about reasons, reasons are always reasons for one thing rather than another, or more generally reasons for one thing

¹See Sinnott-Armstrong (2006), ch. 5 and Ross (2006), ch. 9.

out of a certain set of alternatives, instead of reasons for things *simpliciter*, as has been traditionally assumed. So the reason why the non-contrastive question in the previous paragraph is hard to answer is that it is *underspecified*: to know whether r is a reason to ϕ , we have to know ‘Rather than what?’. Obviously nothing I’ve said yet establishes that this theory is true. Establishing that is the goal of this dissertation. To do so, I’ll develop a detailed version of contrastivism, motivating each step along the way by arguing that contrastivism gives us the best way to make sense of important and puzzling phenomena involving reasons. Since reasons are so important in philosophy—talk about reasons shows up in any field in which *normativity* plays an important role—this dissertation will include discussion from several fields, including ethics, practical reasoning, metaethics, and epistemology.

1.1 Reasons and Contrastivism

This dissertation advocates contrastivism about reasons. To get clear on just what this view is, we need to get clear on two things: reasons and contrastivism.

1.1.1 Reasons

The theory I will develop here concerns *normative* reasons. These are considerations, usually taken to be facts or true propositions, that count in favor of or against doing various things: performing actions, believing propositions, and having a variety of other attitudes.² For example, the fact that I’m in the mood for Thai food is a reason for me to go to Red Corner Asia, and the fact that you are sleeping in the next room is a reason against turning the volume on my stereo up high. These reasons are what are often called *pro tanto*: they count in favor of or against performing various actions (or having various attitudes) with a certain strength or weight. There may be some reasons for me to go to Red Corner Asia—I am in the mood for Thai food, it is nearby, and so on—and some reasons against doing so—it is crowded at this time of day, it is expensive, and so on. These *pro tanto* reasons for and against the action contribute, in some way, to what I ought to do.³ Reasons for

²This talk of ‘counting in favor’ is ubiquitous, but see especially Scanlon (1998). Hieronymi (2005) thinks that thinking of reasons primarily in these terms leads to problems, though even she agrees that reasons do count in favor of the things they are reasons for.

³Broome (2004), for example, uses the term ‘*pro tanto*’; Dancy (2004), on the other hand, uses ‘contributory’ to talk about the same kinds of reasons.

an action may outweigh the reasons against it, or vice versa; similarly, the reasons for one action may outweigh the reasons for another action. Thus, contrastivism is not a thesis, at least not primarily, about what there is “reason to do, overall”, or what there is “decisive reason to do”. It is rather a thesis about these *pro tanto* reasons, that contribute to what there is reason to do, overall.

Normative reasons are distinct from two other sorts of reasons. First, they are distinct from *motivating* reasons—the reasons for which agents do the things they do. These are cited in explanations of the actions of agents. Of course, it is plausible that the facts which are an agent’s motivating reasons could also be normative reasons for that agent. But the motivating and normative reason relations are distinct. Second, normative reasons are distinct from *explanatory* reasons, which we cite in explaining phenomena more generally (plausibly, motivating reasons are a species of explanatory reasons). For example, the reason why the sink is leaking is that it needs a new washer. Motivating reasons and explanatory reasons more generally do not count in favor of or against doing anything; they are rather explanations why things happen. Besides a short discussion of explanation near the end of this chapter, I will not have much more to say about these kinds of reasons; thus, when I talk about reasons, I will usually not include the qualifier, ‘normative’.

1.1.2 Contrastivism

I will be almost exclusively concerned with contrastivism about reasons. But writers have proposed contrastive theories of several important philosophical concepts. Contrastivism about explanation, for example, goes back—at least in explicit form—to van Fraassen (1980).⁴ Staying in the philosophy of science, philosophers have more recently offered contrastive theories of confirmation.⁵ The basic idea is that some fact *e* explains some other fact *p* rather than another fact *q*, or confirms some hypothesis *h* rather than some other hypothesis *j*, instead of explaining *p* or confirming *h simpliciter*. More recently contrastive theories of knowledge—the view that knowledge is always knowledge of *p* rather than *q*—have become popular.⁶ Blaauw (2012) has argued that knowledge is contrastive because belief is

⁴See also Lipton (1990); Hitchcock (1996).

⁵Chandler (2007, 2013). See Fitelson (2012) for critical discussion.

⁶See Morton and Karjalainen (2003); Schaffer (2004, 2005b, 2007, 2008).

contrastive. Sinnott-Armstrong (2004, 2006, 2008) has argued instead that justification is contrastive. In metaphysics, Schaffer (2005a, 2012) has argued that causation is contrastive. And, perhaps most importantly for my purposes, several writers have proposed contrastive theories of ‘ought’.⁷ In the most general terms, contrastivism about some property or relation R is the view that R includes an argument place that must be filled by a set of alternatives. So R only holds or fails to hold of some number of things relative to sets of alternatives.

Contrastivism about some concept is often motivated by showing how judgments involving the concept appear to be *question-relative*, so that the truth value of a claim involving the concept can vary depending on a relevant question. For example, Schaffer (2007) has argued that to know that p is to know that p as the answer to a particular question, salient in the context. Thus, knowing that p when the salient question is Q does not guarantee that you know that p when the salient question is Q' . And Cariani (fc) describes his contrastive view of ‘ought’ as the view that ‘ought’ is relative to a deliberative question: the question of what to do. The move from question-relativity to contrastivism is a natural one because a standard view of questions in formal semantics treats them as partitions on (some subspace of) logical space, or as a set of alternatives with each alternative corresponding to one cell of the partition.⁸ So on this view, question-relativity just is alternative-relativity.

There’s a question about how to implement this idea: how is the set of alternatives, or question, provided? Sometimes it’s provided explicitly. Most obviously, the alternatives can be provided explicitly using a ‘rather than’ clause, as in ‘There’s a reason to go to RCA rather than Paru’s’. Here the alternatives are clearly going to RCA and going to Paru’s. Since these kinds of ascriptions make the alternative very obvious, I’ll appeal to them frequently. Similarly, some writers have argued that knowledge-*wh* ascriptions, like ‘I know when to hold ’em’ and ‘I know what you did last summer’, relate an agent and a question, where the *wh*-clause provides the question explicitly (‘When to hold ’em?’, ‘What did you do last summer?’).⁹ Of course, even if we are given the question, we need to know the possible answers:

⁷See Sloman (1970); Jackson (1985); Finlay (2009, ms); Finlay and Snedegar (fc); Cariani (2009, fc); Snedegar (2012); Kierland (ms).

⁸See Hamblin (1958); Groenendijk and Stokhof (1997); Higginbotham (1993, 1996) for this sort of view of questions.

⁹See Schaffer (2007); Stanley (2011). Ascriptions that explicitly include a question are very natural for knowledge and explanation ascriptions, but not for reason, ‘ought’, or causal ascriptions.

which ones count as relevant, and how are we to individuate answers? Moreover, many ascriptions do not do anything to explicitly introduce—or even let on that there are—alternatives, as in ‘You know that I love you’ and ‘The fact that I love you is a reason for you to love me back’. What are the sets of alternatives relative to which we should interpret these sorts of claims?

The most natural view is a version of *contextualism*, according to which the argument place is filled by some contextually relevant set of alternatives. There are various ways to make a set of alternatives relevant. For example, using intonational stress, as in ‘There’s a reason to *drive* to the store’ (letting italics mark stress), makes relevant a set of alternatives in which the alternatives differ in (the object corresponding to the) stressed item, e.g. {drive to the store, bike to the store, walk to the store}. If I say instead, ‘There’s a reason to drive to the *store*’, the relevant alternatives might be {drive to the store, drive to church, drive to campus}.¹⁰ Other features of the context can make an alternative salient, as well. For example, Schaffer (2007); Cariani (fc); Finlay (2009); Finlay and Snedegar (fc) all appeal to questions under discussion. If I’m deliberating between going to Red Corner Asia, going to Spicy BBQ, and going to Paru’s, the natural set of alternatives relative to which an ascription like ‘The fact that my guest loves Indian food is a reason to go to Paru’s’ is to be evaluated is {go to Red Corner Asia, go to Spicy BBQ, go to Paru’s}. Thus, on this simple contextualist view, the speaker’s intentions will play a large part in determining the set of alternatives relative to which we should evaluate reason ascriptions.¹¹ I will have more to say about this in the next chapter when I develop a contrastive account of reason ascriptions.

There are two interesting features of these sets of alternatives. First, they are not necessarily exhaustive of logical space, or of all of the things that it’s possible for the agent to do.¹² Some alternatives are left out, for one reason or another

¹⁰See Dretske (1970); Rooth (1992); Schaffer (2005b, 2007, 2008).

¹¹I don’t think this speaker-contextualist version is necessary. For example, we might have a view on which the set of alternatives is provided by the context of assessment, or perhaps even by the relevant agent’s (not necessarily the speaker, and not necessarily in the speaker’s context) circumstances, if we develop things in the right way. But I’ll develop the contextualist version of the view.

¹²It’s not quite right to talk about questions or alternatives partitioning (part of) logical space if you treat the elements of the set as actions, as I’ve been doing, rather than as propositions. So if we take this idea of question-relativity seriously, there are two options. First, we can say that, strictly speaking, reasons are reasons for propositions (i.e. a reason for the proposition that you ϕ) rather than for actions (i.e. a reason for you to ϕ). (Compare Finlay (2006, ms).) Second, we can say that reasons really are reasons for actions, but in our *model*, hold the agent fixed so that

(e.g., they are irrelevant).¹³ Second, the options are divided up at a more or less fine-grained level of detail. For example, one particularly coarse-grained set would be something like {go to Red Corner Asia, don't go to Red Corner Asia}, while a more fine-grained one would be {go to Red Corner Asia, go to Spicy BBQ, go to Paru's}. This is what Yalcin (2011) and, following him, Cariani (fc) call *resolution-sensitivity*, the idea being that we partition logical space at higher or lower resolutions. Either or both of these features can be exploited to solve various puzzles.¹⁴ I'll end up making use of both of these features in my version of contrastivism about reasons. Given the standard treatment of questions I mentioned earlier, originally due to Hamblin (1958), it's easy to see how both of these features could come along with question-relativity. Non-exhaustivity comes along because questions need not partition the entirety of logical space—sometimes they partition only a subspace. Resolution-sensitivity comes along simply because the question has to partition this subspace at some resolution or other. Which possibilities are left out as irrelevant and the resolution at which the relevant possibilities are partitioned will both depend, among other things, on context (on the contextualist version of contrastivism). They may also be determined explicitly by *wh*-complements (in the case of 'knows', for example) or by 'rather than' clauses.

So contrastivism about reasons is the thesis that reasons to ϕ are always reasons to ϕ relative to, or out of, some particular set of alternatives. The set of alternatives need not be exhaustive of all of the things the agent could possibly do, and the alternatives in the set must be divided up at some particular resolution. This captures the following ideas. First, when we deliberate about what to do, we don't always do so over every possible alternative; that's non-exhaustivity. Second, we can only consider the relevant alternatives at some level of detail or other; that's resolution-sensitivity. Which set is relevant can vary depending on features of the

we can treat actions as propositions (i.e. the set of worlds in which the relevant agent performs the relevant action). See Cariani (fc) for discussion of this last strategy.

¹³Compare the relevant alternatives theory of knowledge in Lewis (1996); Dretske (1970).

¹⁴Skorupski (2010) explicitly endorses a view of reasons for action on which they are relativized to what he calls 'choice sets': "When we think of think of sufficient reasons for doing something [...] we assume some exhaustive partitioning of actions, within which one and only one action must be done. Call an exhaustive partitioning of actions a choice set; 'inaction' and likewise the action of further deliberation are members of this set" (p. 38). Skorupski's view does not count as a contrastivist view as I've defined it above, because he does not accept non-exhaustivity. But he does accept resolution-sensitivity, so as far as that goes, I'm sympathetic. But since I think we need to adopt non-exhaustivity to gain lots of the advantages of contrastivism, I reject Skorupski's view.

context. This set must be supplied before we can say whether r is a reason to ϕ or not, since it can only be or fail to be a reason relative to a set of alternatives. Arguments for contrastivism claim that as the set of alternatives shifts, the reasons for and against a given action can vary—what was a reason to ϕ relative to one set may not be a reason to ϕ relative to a different set.

1.2 Sinnott-Armstrong’s Argument

I have just been discussing contrastivism as a general philosophical program which has been applied to several important concepts like knowledge, justification, ‘ought’, and explanation. One thing to notice is that all (or at least almost all) of the concepts for which contrastivism seems plausible and fruitful are intimately related to reasons. So a tempting thought is that contrastivism about reasons is really central, and underwrites contrastivism about all of these other concepts. In fact, Walter Sinnott-Armstrong has exactly this thought:

In my view, epistemologists need contrasts in their analyses of knowledge and justified belief because knowledge and justified belief require reasons [...], and reasons are always reasons for one thing as opposed to another. [...] The same rationale will then apply to any other kind of reasons, and many philosophical issues concern reasons. Epistemologists study reasons *for belief*. Moral philosophers investigate moral reasons *for action*. Aestheticians explore reasons *to like* or *to value* certain art, music, and so on. Philosophers of science analyze explanations, which give reasons *why* events happen. And so on. Because reasons are central to all these areas of philosophy, and all reasons are relative to contrasts, all these areas of philosophy can benefit from introducing a new place for contrasts into the relations used in analyses.¹⁵

This suggests a kind of argument to the best explanation: the best explanation for why these other concepts—justification, obligation, explanation, and so on—are contrastive is that they inherit it from the contrastivity of reasons.

While I agree both that reasons are contrastive and that this is a tempting thought, I do not think this is a very compelling argument for contrastivism about

¹⁵Sinnott-Armstrong (2008), pp. 257-258.



reasons. To show why, I will call into question both the necessity and the sufficiency of contrastivism about normative reasons for contrastivism about the normative notions of justification and ‘ought’. Moreover, the arguments will involve the introduction of plausible views that accept contrastivism about reasons without accepting contrastivism about the other normative concept (‘ought’ or justification), or vice versa. All of this should make us suspicious of the inference Sinnott-Armstrong seems to draw.

1.2.1 Reasons and justification

Contrastivism about justification is the view that you can only be justified in believing that p out of, or relative to, a certain set of alternatives Q . Contrastivism about justification leads naturally to contrastivism about knowledge, since knowledge requires justification, and contrastivism about knowledge has become a major player in the epistemological literature.¹⁶ Nevertheless, I’m going to focus on contrastivism about justification here since justification is more directly tied to reasons (the relationship between knowledge and reasons is plausibly mediated by justification). One motivation for adopting contrastivism about justification is that it seems to offer a solution to the infamous closure paradox. It seems that Moore is justified in believing that he has hands. And he’s justified in believing that if he has hands, then he’s not a handless brain-in-a-vat. But it also seems that he’s not justified in believing that he’s not a handless brain-in-a-vat. The contrastivist can say that Moore is justified in believing that he has hands rather than, say, flippers. But he’s not justified in believing that he’s not a handless brain-in-a-vat rather than that he has hands. In other words, he’s justified in believing that he has hands out of some sets of alternatives, but not out of others. In this section, I’m just going to assume for the sake of discussion that justification is contrastive. So now what we need to do to establish contrastivism about reasons, along the lines suggested by Sinnott-Armstrong, is (i) show how justification is related to reasons, and (ii) show how this relationship requires reasons to be contrastive if justification is.



So what is the relationship between justification and reasons? A very intuitive first answer is that you can only be justified in believing a proposition if you have sufficient reason to believe it. But (i) this is stated non-contrastively, and (ii)

¹⁶See Morton and Karjalainen (2003); Schaffer (2004, 2005b, 2007).

we need some idea of what counts as *sufficient* reason. Sinnott-Armstrong offers the following principle, inspired by the relevant alternatives theories of knowledge offered by Dretske (1970) and Lewis (1996):

Ruling Out: *s* is justified in believing that *p* out of *Q* iff *s* has sufficient reason to rule out all the other alternatives in *Q*.¹⁷

It might be a bit more natural to state things in terms of having sufficient *evidence* to rule out the alternatives. But since I'm talking about reasons, and since any piece of evidence is very plausibly a reason for belief (though it might not be true that all reasons for belief are evidence), I'm going to state things in terms of reasons. One important thing to keep in mind is that having sufficient reason to rule out some alternative *q* does not require having conclusive reason to believe $\neg q$; it may be enough if your reasons for belief make *q* very unlikely, for example. A second important caveat is that it isn't enough, according to this principle, to have more reason to believe *p* than any other alternative in *Q*, since that still might not be enough to rule out the other alternatives.

Note that, though this principle does make justification, and so whether you have sufficient reason to believe that *p* relative to sets of alternatives—you might have sufficient reason to believe that *p* relative to *Q* but not relative to *Q'*—it does not yet make reasons contrastive. So why think that contrastivism about justification requires contrastivism about reasons? Here's a case (borrowed from Dretske (1970)) that I think illustrates the motivating idea well. Suppose you're at the zoo. You see a black and white striped horse-like animal in front of you. And suppose that this visual evidence is the only relevant reason for belief. Well, relative to the set of alternatives {believe it's a zebra, believe it's an elephant, believe it's an ostrich}, it seems that you have sufficient reason to believe it's a zebra—your reasons for belief are sufficient to rule out that it's an ostrich or an elephant. But relative to the set of alternatives {believe it's a zebra, believe it's a cleverly disguised mule}, it seems that you don't have sufficient reason to believe that it's a zebra. In fact, it looks like the fact that was your reason for belief in the previous case, that you were having such-and-such visual experience, is just not even a reason in

¹⁷Sinnott-Armstrong is also sympathetic to the idea that being justified in believing that *p* relative to *Q* requires having some positive reason to believe *p*, as well as having sufficient reason to rule out the alternatives. I'll ignore that addendum here, because I don't think it affects either his or my argument. It is also relevant, I think, that since the alternatives in *Q* are supposed to be mutually exclusive, evidence for *p* will likely be evidence *against* the other alternatives.



this case—the fact that you see a black and white striped horse-like animal is not a reason to believe it’s a zebra rather than a cleverly disguised mule. That would certainly explain why you aren’t justified relative to this new set, since justification requires having sufficient reason for belief. So here’s a principle relating contrastive justification and contrastive reasons:

Contrastive Justification, Contrastive Reasons: *s* is justified in believing *p* out of *Q* iff *s*’s reasons for belief relative to *Q* are sufficient to rule out every member of *Q* except for *p*.

We can explain why justification is contrastive by (i) pointing out that justification requires having sufficient reason to believe, and (ii) reasons for belief are only reasons relative to sets of alternatives.

But is this the only, or even obviously the best, way to explain why justification is contrastive? I don’t think so. Suppose that reasons aren’t contrastive. Here’s a story about how justification could still be contrastive. Non-contrastive reasons favor believing certain propositions independently of any set of alternatives. And, importantly, some non-contrastive reasons favor ruling out some propositions—that is, believing their negation—independently of any set of alternatives. Then we can give the following principle:

Contrastive Justification, Non-Contrastive Reasons: *s* is justified in believing that *p* out of *Q* iff *s* has sufficient reason to rule out all the alternatives in *Q* except for *p*.¹⁸

This thesis appeals to non-contrastive reasons, but still allows justification to be contrastive. To know whether *s* is justified in believing *p*, we still have to know what sets of alternatives we’re talking about, so we know which alternatives have to be ruled out by *s*’s reasons. This principle, of course, is just the original idea captured by **Ruling Out**. Reasons weigh in favor of believing or disbelieving certain propositions, but only certain propositions are relevant in a given context.

Above I appealed to the zebra/cleverly disguised mule case to motivate thinking that reasons are contrastive. The intuition was that the fact that you’re having such-and-such visual experience is a reason to believe that it’s a zebra relative to {believe it’s a zebra, believe it’s an elephant, believe it’s an ostrich}, but not

¹⁸We could add that *s* needs some positive reason to believe that *p*.

relative to {believe it's a zebra, believe it's a cleverly disguised mule}. I think there's something to this claim: this kind of case does at least give us *prima facie* evidence that reasons are contrastive. But note that this evidence is independent of contrastivism about justification. Further, I think there is something the non-contrastivist about reasons can say to explain this kind of case: it's true to say that r is a reason to believe that p when the relevant set of alternatives is Q only if r is a stronger reason to believe p than to believe any of the other alternatives in Q . Since the fact that you're having such-and-such visual experience is not a stronger reason to believe that it's a zebra than to believe that it's a cleverly disguised mule, according to this view, it's not true to say, 'It's a reason to believe it's a zebra', when believing it's a cleverly disguised mule is a relevant alternative. On the other hand, this claim would be true if the relevant alternative were, say, believing that it's an ostrich. (I'll say more about this idea in Chapter 3 when I discuss the view I call *shallow contrastivism*.)

I've argued that contrastivism about justification does not require contrastivism about reasons. Now I'll argue that contrastivism about reasons does not require contrastivism about justification.

Suppose again that reasons are contrastive. Further, suppose I'm a skeptic—I don't think that you can ever be justified in believing anything. I can still accept that reasons are contrastive: you might have good reason to believe that the animal is a zebra rather than an ostrich. But I just don't think this matters for justification, because to be justified in believing that it's a zebra, you have to rule out all of the possible alternative hypotheses. So the contrastive reasons that matter in determining whether you're justified in believing that it's a zebra will be {it's a zebra, it's an ostrich, it's a cleverly disguised mule, it's a hallucination, ...}. And since you don't—and couldn't—have sufficient reason to believe that it's a zebra out of this set, you aren't justified. Or perhaps I'm more liberal in applying justification: I think that the set that matters in determining whether you're justified is always just what Sinnott-Armstrong (2004, 2006, 2008) calls an "everyday" contrast class, like {it's a zebra, it's an ostrich, it's an elephant}. That is, I think the skeptic makes unreasonable demands on the alternatives you have to rule out to be justified. Since you do have sufficient reason to believe that it's a zebra relative to this everyday set, your belief is justified (full stop). According to either of these views,



then, reasons are contrastive, but justification is not—there is a particular, fixed set of alternatives that is relevant for determining whether or not you are justified.

1.2.2 Reasons and ‘ought’

Contrastivism about ‘ought’ is the view that ‘ought’ claims are always to be interpreted as relative to some particular set of alternatives. Whenever an agent ought to perform some action, there is some set out of which she ought to perform it. Several writers have argued that ‘ought’, or what agents ought to do (to put it non-linguistically), is contrastive, for a variety of reasons. For the purposes of this section, I’ll assume that contrastivism about ‘ought’ is well motivated. So now what I need to do is (i) sketch the relationship between ‘ought’ and reasons and (ii) evaluate whether this relationship demands that one of the concepts is contrastive if the other is.

The most natural thesis about the relationship between ‘ought’ and reasons says that you ought to ϕ iff you have most reason to ϕ . This thesis fits nicely with an attractive picture of deliberation. To decide what you ought to do, you just weigh up all your reasons to determine which option is best supported by them; that’s the one that you ought to do.¹⁹ The contrastivist about ‘ought’, though, cannot accept it as stated since she thinks that we need to relativize claims about what agents ought to do to sets of alternatives. And since these alternatives need not be exhaustive, we can’t simply talk about what you have most reason to do *simpliciter*, either, in trying to state the relationship between what you ought to do (relative to a certain set of alternatives) and reasons. It seems that we need to look at your reasons to do things, somehow restricted or relativized to the set of alternatives to which the ‘ought’ claim is relativized.

One way around this problem is to relativize reasons, as well as ‘ought’, to sets of alternatives—that is, to be contrastivists about reasons. This lets us retain an attractive picture of deliberation, analogous to the standard one I mentioned above. We begin with some relevant set of alternatives, Q , and then simply deliberate

¹⁹There are complicated questions about how to determine what you have most reason to do—how reasons interact with one another. See Dancy (2004); Broome (2004); Schroeder (2007), chapter 7, and Horta (2007, 2012) for discussion about how your reasons determine what you ought to do, or how your *pro tanto* reasons determine what you have *most* reason to do.

over the alternatives in Q . So it's easy to see why one might think that adopting contrastivism about 'ought' requires adopting contrastivism about reasons.

But this is too fast. It's possible to give a very plausible picture on which 'ought' is contrastive but reasons are not. Begin with a traditional non-contrastive picture of reasons, on which they are considerations that favor certain actions with a certain weight, some stronger or weightier than others. Collect all of the agent's reasons and weigh them up (in some way). This will give us an overall ranking of all the possible options open to the agent. We can even divide these options up as finely as you want. Maybe the resolution at which we divide them is the resolution at which the agent's intentional actions can make a difference: I can intentionally bring it about that I put between 10.55 and 10.56 gallons of gas in my car, but I can't intentionally bring it about that I put between 10.555 and 10.556 gallons, for example.²⁰ If that's right, then we can have one option in which I put between 10.55 gallons and 10.56 gallons in, and one in which I put between 10.56 and 10.57 gallons in. But we don't have more fine-grained options than this. This doesn't mean that I'll always deliberate between options divided up at even this high of a resolution; I'll say more about that shortly.

So now what we have is an overall ranking of the possible options, according to how my reasons weigh in favor of them. Now in any particular context of deliberation, only some of these options will be relevant, and I'll only deliberate between those relevant options at some particular resolution. This will generate our set, Q . Then we can give the following principle:

Contrastive Ought, Non-Contrastive Reasons: s ought to ϕ out of Q iff s has more reason to ϕ than to perform any other alternative in Q

We essentially pull out the alternatives in Q from our overall ranking, retaining the order of these alternatives. Non-exhaustivity is secured, since Q need not include all the options that are ranked in the original ranking. Still, the options that do make it into Q will occupy a particular place on the original ranking, and the order between them will be retained in Q . We just appeal to this ranking, provided by my non-contrastive reasons, to determine what I ought to do.

That's non-exhaustivity. Here's how to capture resolution-sensitivity. If the alternatives in Q are divided up at a lower level of resolution than the options in the

²⁰Compare the discussion in Cariani (fc); Kierland (ms).

original ranking are—so that my reasons rank the options at a higher resolution than the resolution at which I deliberate between them—we group them in the obvious way. For example, if I’m deliberating between putting more than ten gallons of gas in the car and putting ten gallons or less in, the alternatives will be grouped in that way: {put in 10 gallons or less, put in more than 10 gallons}. We then have a few options in determining how these coarse-grained options are ranked. We could rank them based on (i) the average position of the more fine-grained alternatives they subsume, (ii) the rank of the highest or lowest fine-grained option they subsume, or (iii) the fine-grained option I would actually perform, were I to decide on the coarse-grained option. I will not try to settle which of these is best; the point is just that there are plausible ways for a contrastivist about ‘ought’ to retain a non-contrastive view of reasons.


Now I will argue that contrastivism about reasons is not sufficient for contrastivism about ‘ought’. So first, suppose that reasons are contrastive: a reason to ϕ is only such a reason relative to certain sets of alternatives. If we accept anything like the claim that you ought to ϕ iff you have most reason to ϕ , how can we maintain a non-contrastivist view of ‘ought’? Here’s one way. Adopt the plausible view that you ought to ϕ iff there’s more reason to ϕ than to not ϕ .²¹ Importing contrastivism about reasons now, even though reasons are only reasons relative to certain sets of alternatives, the particular reasons that we’re concerned with in trying to decide whether you ought to ϕ are the reasons relative to the set of alternatives, $\{\phi, \neg\phi\}$.

We could also make the set that matters for determining what you ought to do more interesting: perhaps it’s determined by the things that it’s possible for you to do. This doesn’t amount to accepting contrastivism about ‘ought’, since we can just say that what you ought to do is simply the best thing you can do. According to this view, the contrastive reasons we are concerned with are the ones that are reasons relative to the set of all the things that it’s possible for you to do. The central idea here is just that, when we are trying to determine what someone ought to do, we look at the reasons relative to some particular privileged set of alternatives. (This mirrors the strategy I used in the case of justification, when I introduced skeptical and everyday contrast classes.)

²¹This is essentially the principle that Schroeder (2007), pp. 130-131 adopts.



Contrastivism about reasons isn't useless on this picture, though. We can still appeal to the contrast-sensitivity of reasons to explain phenomena like the following, observed by Sinnott-Armstrong (2004, 2006, 2008). Today is your birthday. Suppose I'm trying to decide whether to bake you a chocolate cake, or no cake at all. Then intuitively, the fact that today is your birthday is a reason to bake you a chocolate cake. But now suppose I'm trying to decide whether to bake you a chocolate cake or a lemon cake. The fact that today is your birthday seems not to be a reason to bake you a chocolate cake in this case. The contrastivist has an easy explanation: it's a reason to bake you a chocolate cake out of {bake you a chocolate cake, don't bake you any cake}, but not out of {bake you a chocolate cake, bake you a lemon cake}. So there's still some motivation for adopting contrastivism about reasons, since it lets us explain these kinds of cases; it's just that when you're trying to decide if you ought to ϕ , you're only concerned with a certain set of your contrastive reasons: those relative to the privileged set.

I don't claim to have shown that Sinnott-Armstrong's argument to the best explanation cannot work. To salvage the argument, we would need to show that, even though contrastivism about 'ought'  justification neither requires nor is required by contrastivism about reasons, his explanation is nevertheless the *best* one. One challenge here is that the views I described which accept contrastivism about 'ought' or justification without accepting contrastivism about reasons seem very reasonable, and I'm not sure what evidence we could marshal in favor of the views which accept contrastivism about both 'ought'/justification and reasons over these, beyond giving direct arguments for contrastivism about reasons. Finally, of course, for the argument to actually be compelling, we would need to establish contrastivism about 'ought' and justification. Giving direct arguments for contrastivism about reasons, on the other hand, will let us avoid holding that view hostage to the fortune of these other contrastive theories.

1.2.3 Reasons and explanation

In the last two subsections, I discussed the relationship between two important normative concepts, 'ought' and justification, and normative reasons. In this subsection, I'm going to talk about the relationship between explanation, and *explanatory* reasons. Moral philosophers, understandably, focus mostly on normative reasons.

And even when they talk about explanatory reasons, they are usually most interested in explanatory reasons of the motivating variety: those that explain an agent's behavior by being considerations the agent took to be normative reasons. Even philosophers of science and epistemologists generally talk about explanations rather than explanatory reasons. So there's a lack of literature on the topic of explanatory reasons; everyone recognizes that there are such reasons, but few writers have had much to say about them. But it's important for my purposes to explore the topic, because Sinnott-Armstrong thinks that all reasons, normative and explanatory, are contrastive.

Explanatory reasons are clearly closely related to explanations. And a prominent idea in the philosophy of science is that explanation is contrastive.²² Contrastivism about explanation is the view that some fact *e* explains an event, or perhaps another fact, *f*, only relative to certain sets of alternatives. For example, the fact that it's humid outside explains why it's raining rather than not precipitating, but it does not explain why it's raining rather than snowing. And the fact that it's above the freezing point explains why it's raining rather than snowing, but not why it's raining rather than not precipitating. My sense of the literature is that contrastivism about explanation is more widely accepted than contrastivism about other concepts. As has been my policy, I'll assume that contrastivism about explanation is true.

So the next step is to figure out what the relationship is between explanation and explanatory reasons. As I said above, there has not been much written about this topic. The reason for this, I suggest, is that explanatory reasons just are explanations. Explanatory reasons are usually glossed as answers to 'Why?' questions: the reason why the carpet is wet is that the roof leaked; the reason why I'm sunburnt is that I forgot to put on sunscreen. And so on. But all of these things just look like explanations for the relevant facts. One idea is that explanatory reasons are *part* of the explanation for some fact. The explanation for why I'm sunburnt is that I forgot to put on sunscreen and I stayed in the direct sunlight for a long time and my skin has certain properties and Each of these individual parts of the explanation might be an explanatory reason; each explanatory reason contributes to the full explanation. The problem with this, though, is that when we cite explanatory reasons, we simply seem to be giving explanations: saying 'The reason why I'm sunburned is that I forgot to put on sunscreen' seems to be the

²²See van Fraassen (1980); Lipton (1990); Hitchcock (1996).

same as saying ‘The explanation for why I’m sunburnt is that I forgot to put on sunscreen’. It’s plausible that we’re just citing parts of explanations even when we use the word ‘explanation’, but the point is just that, whenever we purport to cite an explanation, using ‘explanation’, we could just as well have used ‘reason’.

If explanatory reasons just are explanations, or if to give an explanation for p is just to cite an explanatory reason why p , then obviously contrastivism about explanation entails contrastivism about explanatory reasons. So in that sense, Sinnott-Armstrong’s argument works in this case, though it doesn’t seem that he had anything quite so trivial in mind. The important point to consider now, though, is whether accepting contrastivism about explanatory reasons puts pressure on us to accept contrastivism about normative reasons. I’ll argue that it doesn’t, at least not without some further assumptions that can be rejected.

The first strategy, and the one that Sinnott-Armstrong seems to adopt, is to say that normative reasons and explanatory reasons are simply different species of the same genus, *reasons*. There are just reasons; normative reasons and explanatory reasons are just particular kinds. But since they’re part of the same general category, important structural features, like alternative-relativity, should be shared. So since explanatory reasons are contrastive, reasons are contrastive. And that means, on this view, that normative reasons are contrastive, as well. Some support for this view comes simply from the fact that we call both normative and explanatory reasons ‘reasons’. Surely this isn’t an accident, and one possible explanation is that the two categories are just species of the same genus.

One problem with this view is that it seems undermotivated. The simple fact that we call both explanatory and normative reasons ‘reasons’ is not enough to establish that there’s really just one thing, reasons, and two different kinds. This is especially so given the apparent differences—and seemingly structural differences—between explanatory and normative reasons. For example, there can be a normative reason for s to ϕ even if s does not ϕ , but there can’t be an explanatory reason why p if p is false. Another apparent structural difference is that explanatory reasons seem to take propositions as their complements: we give reasons why p for a proposition p . But normative reasons appear to take actions or attitudes as their complements: we have reasons to ϕ for an action ϕ , or reasons to form attitude A .²³ For these and

²³One promising route here might be to follow Hieronymi (2005) and note that both explanatory and normative reasons are considerations that bear on a question. This might be the underlying

other reasons, many philosophers think that there are big, important differences between normative reasons and explanatory reasons. Another way in which this thesis is controversial, if it is supposed to be different from the next strategy that I'll consider, is that many philosophers think that normative reasons are special kinds of explanatory reasons. On the view under consideration now, we have two different species of reasons, with neither being a special case of the other (although they are of the same genus).

A second strategy is to say that normative reasons are simply a special kind of explanatory reason; that is, normative reasons are explanatory reasons that explain a special kind of, usually normative, fact. Toulmin (1950); Finlay (2001, 2006, ms); Searle (2001); Broome (2004) all hold some variety of this view (with important differences, of course). If normative reasons are just a special kind of explanatory reasons, and explanatory reasons are contrastive, then normative reasons are contrastive, too.

A problem with this strategy is that it takes away the possibility of reducing the normative to reasons. This is obviously not something that everyone agrees can be done, but many people do think so. If we analyze normative reasons as considerations that explain some normative fact, like why you ought to ϕ , or why it would be good if you were to ϕ , or something like that, then we cannot analyze the notion of 'ought' or of goodness in terms of reasons, on pain of circularity. Since this strategy is incompatible with such a widely held view, it is controversial (which, again, is what I want to establish here).

A third strategy is to argue that the correct analysis of normative reasons must be in terms of explanation, or of explanatory reasons. For example, perhaps for r to be a reason for you to ϕ , it has to explain why your ϕ -ing would promote your desires, or why it would be good in some way if you were to ϕ . Then if explanation is contrastive, perhaps the contrastivity will be passed on to normative reasons.²⁴

But not all views of normative reasons have it that they are correctly analyzed in terms of explanation. Hampton (1998); Scanlon (1998); Parfit (2011), for example, simply take the reason relation as primitive. It's just a (normative) fact about

unifying feature in virtue of which both categories count as reasons. I won't pursue this here, though.

²⁴It is actually not at all straightforward to move from contrastivism about explanation, through one of these kinds of analyses, to contrastivism about reasons. But illustrating this is not necessary for the main point I want to make—that any analysis of reasons like this is controversial.

the world that some considerations favor some things, and explanation doesn't enter into it. And if the correct analysis of normative reasons is not in terms of explanation, then contrastivism about explanation will not support contrastivism about reasons.

I don't claim to have offered knockdown arguments against any of these three strategies for moving from contrastivism about explanation to contrastivism about normative reasons. But I hope I have shown that all three are controversial. Just as I do not want to rely on the truth of contrastivism about 'ought' or justification to establish contrastivism about reasons—which I think is a more defensible view, anyway—I do not want to rely on these controversial accounts of the relationship between normative and explanatory reasons. In the rest of the dissertation, I'm going to try to offer more direct evidence that reasons are contrastive.

1.3 The Plan

To close this chapter, I'll briefly preview the rest of the dissertation. In chapters two, three, and four, I aim to develop a detailed and attractive version of contrastivism in three steps, each of which is independently-motivated. After the full theory is on the table, I will illustrate other interesting applications of contrastivism in ethics, epistemology, and practical reasoning in the final two chapters.

In chapter two, I'll argue that our *talk* about reasons motivates a shallow kind of contrastivism—the view that the word 'reason', as used to ascribe normative reasons, expresses a relation with an argument place for sets of alternatives. The central argument will be that non-contrastive theories have trouble explaining puzzling features of explicitly contrastive reason ascriptions which employ the phrase 'rather than'. So the first step is to motivate and develop a contrastive account of our talk about reasons.

But this leaves open a view, which I call *shallow contrastivism*, on which contrastivity only goes "language-deep". According to this view, though the word 'reason' expresses a relation with an argument place for sets of alternatives, that relation is *not* the important normative relation for theorizing in ethics, epistemology, and practical reasoning. There is an underlying non-contrastive normative *favoring* relation that holds between actions or attitudes and considerations which favor them independently of the relevant alternatives, and in terms of which we can

analyze the contrastive relation expressed by ‘reason’. This is a more conservative view than *deep contrastivism*, according to which the contrastive relation expressed by ‘reason’ is the important normative relation, and is not to be analyzed in terms of a deeper non-contrastive normative relation. However, after sketching shallow contrastivism, I’ll argue that the normative favoring relation itself is contrastive, so that we should adopt deep contrastivism. This is the second step in developing contrastivism.

The theory, as it stands after chapter three, is that reasons to ϕ , out of a set of alternatives Q , are facts that favor ϕ -ing out of Q . This is just the contrastivist implementation of the ubiquitous, if uninformative, view that reasons are considerations that favor the things they are reasons for. The third step is to try to say something more substantive, by giving an analysis of the favoring, or reason, relation. In fact, I’ll show that we need to do this, since otherwise contrastivism faces a serious problem. I will argue, however, that we can solve this problem by appealing to the popular idea that reasons involve the *promotion* of various kinds of objectives, like desires or values. And moreover, I’ll argue that this idea itself provides strong independent support for contrastivism, since contrastivism provides the best way to understand how reasons could involve promotion. The result is an attractive, detailed version of contrastivism about reasons.

In chapter five, I’ll consider the possibility of the intransitivity of ‘more reason than’. Most people think this relation couldn’t possibly be intransitive, but others, like Temkin (1987, 2012); Rachels (1998, 2001); Friedman (2009), have given compelling arguments for the possibility of intransitivity. From this perspective, contrastivism might seem especially well-placed to accommodate and explain intransitivity, in virtue of the independence of reasons relative to different alternatives. I’ll explore ways we might do just that, and conclude that contrastivism does in fact provide an especially natural framework on which to accommodate intransitivity, but that doing so is ultimately optional for the contrastivist.

Finally, in chapter six, I’ll use contrastivism to provide an attractive account of rational withholding of both belief and intention—an important but underexplored topic in the theory of rationality. The upshot of these final two chapters is that contrastivism about reasons has interesting applications in normative domains, in which reasons are of fundamental importance.

Chapter 2

REASON CLAIMS

Now that I've gotten the preliminaries on the table, and argued that to establish contrastivism we need some *direct* evidence that reasons are contrastive, I'll start trying to provide that evidence. In this chapter, I'll argue that our talk about reasons—our use of what I'll call *reason claims*—supports contrastivism. There are puzzles with various sorts of claims which are best solved by contrastivism.

2.1 A Simple Argument

Sinnott-Armstrong (2004, 2006, 2008) points out that the following reason ascriptions are both intuitively true:

- (1) The fact that it's your birthday is a reason for me to bake you a chocolate cake rather than bake you no cake at all.
- (2) The fact that it's your birthday is not a reason for me to bake you a chocolate cake rather than bake you a lemon cake.

But this seems *prima facie* puzzling, if we think of reasons simply as considerations that favor some action. If the fact that it's your birthday is a reason for me to bake you a chocolate cake rather than bake you no cake at all, then it seems that it has to favor baking you a chocolate cake. But if it favors baking you a chocolate cake, why wouldn't it be a reason to bake you a chocolate cake rather than bake you a lemon cake? Sinnott-Armstrong concludes that reasons are always reasons for one thing *rather than* another: to know whether some fact is a reason for some action, we have to know 'Rather than what?'. That is, 'reason' expresses a relation

with an argument place for an alternative, or set of alternatives, relative to which the consideration is a reason for the action. This easily explains how (1) and (2) could both be true: the ‘rather than’ clause makes explicit the alternative relative to which the fact that it’s your birthday is or is not a reason for me to bake you a chocolate cake.

This argument for contrastivism, however, is too fast. First, contrastivism holds that *all* reasons are contrastive. But most reason ascriptions are not explicitly contrastive, unlike (1) and (2). So one way we may resist contrastivism, in the face of Sinnott-Armstrong’s observation, is to hold that ‘reason’ is ambiguous between a contrastive sense and a non-contrastive sense. The contrastive sense is used in claims like (1) and (2), while the non-contrastive sense is used in non-contrastive ascriptions, like ‘The fact that it’s your birthday is a reason to bake you a chocolate cake’.

I think this view is unattractive, because ‘reason’ fails a standard test for ambiguity, which Schaffer (2007) has called *coordination across conjunction*. ‘Reason’ fails this test for ambiguity because sentences like the following are perfectly appropriate:

- (3) The fact that you sprained your ankle is a reason to wear your brace, and to lift weights rather than run.

Notice that there is just one occurrence of ‘reason’ here; this suggests that ‘reason’ latches on to both conjuncts, ‘to wear your brace’ and ‘to lift weights rather than run’. The first conjunct, though, is non-contrastive, while the second is explicitly contrastive. If the single occurrence of ‘reason’ latches on to both conjuncts, that is evidence that ‘reason’ is univocal, whether it takes a contrastive or a non-contrastive complement. So sentences like (3) are a problem for the ambiguity view.¹

Ruling out the ambiguity view, though, does not yet establish contrastivism. We can still try to offer non-contrastive analyses of explicitly contrastive reason ascriptions like (1) and (2). If we can analyze these sentences in terms of non-contrastive reasons, then they will not cause a problem for the non-contrastivist.

A tempting, but ultimately problematic, idea is to give an analysis in terms of the strength of reasons, like one of the following:

¹See Schaffer (2007), p. 396.

RT-1: r is a reason to A rather than B iff r is a stronger reason to A than it is to B .

RT-2: r is a reason to A rather than B iff r is either (i) a stronger reason to A than it is to B , or (ii) r is a stronger reason not to B than it is not to A .

Either of these analyses would let us explain how (1) and (2) could both be true without resorting to contrastivism. Since the fact that it's your birthday is a stronger reason to bake you a chocolate cake than to bake you no cake at all, (1) is true, according to either analysis. But since this fact is an equally strong reason to bake you a chocolate cake and to bake you a lemon cake, (2) is also true, according to either analysis.

While these analyses do seem natural, I think both face a serious problem. The problem is that they require the phrase 'rather than' to function very differently in reason ascriptions than it functions in other contexts. In ordinary uses of 'rather than' like, 'I want pizza rather than salad', 'Bob went to the store rather than to the gym', and 'Two plus two equals four rather than five', 'rather than' seems to mean something along the lines of 'and not', perhaps with an implicature that the two things being contrasted are somehow especially relevant.² But according to **RT-1** and **RT-2**, 'rather than', as used in reason ascriptions, does not mean anything like 'and not', so they seem to rely on an ad hoc treatment of 'rather than' in reason ascriptions. This is a serious cost for these proposals.

Some writers, in objecting to contrastivism about explanation and about knowledge, have given non-contrastivist analyses of explicitly contrastive claims on which 'rather than' does mean 'and not'.³ Consider the following two proposals, which adopt this strategy, and differ only in the scope of the 'not':

RT-3: r is a reason to A rather than B iff r is a reason to A and r is not a reason to B .

²'Rather than' as used in preference ascriptions may seem to better fit with **RT-1** and **RT-2**: 'I prefer cake rather than pie' might seem to mean that I like cake more than I like pie. But I think a better analysis here is the following: I would choose cake and not choose pie, when those are the alternatives. In ordinary cases, I would make this choice because I like cake more than I like pie, but this need not be part of the analysis of the 'rather than' ascription.

³See Ruben (1987); Temple (1998) for this sort of strategy in resisting contrastivism about explanation. See Schaffer (2008) for arguments against this kind of analysis of 'rather than' knowledge ascriptions, and Rickless (fc) for further discussion.

RT-4: r is a reason to A rather than B iff r is a reason to A and r is a reason not to B .

Either of these proposals lets us explain how (1) and (2) could both be true, without resorting to contrastivism. Since the fact that it's your birthday is both a reason to bake you a chocolate cake and a reason to bake you a vanilla cake, (2) is true according to either proposal. And since this fact is not a reason to bake you no cake at all, and in fact a reason *not* to bake you no cake at all, (1) is also true according to either proposal.

Since there are reasonable, independently motivated non-contrastivist explanations of the data provided by (1) and (2), Sinnott-Armstrong is too quick in concluding on this basis that reasons are contrastive. What ascriptions like (1) and (2) do show, however, is that non-contrastivists need to say *something* about these explicitly contrastive ascriptions; I've argued in this section that they should take some version of the popular non-contrastivist strategy, treating 'rather than' as meaning 'and not'. But now that we've seen this, we're in a position to give a stronger argument for contrastivism.

2.2 A Stronger Argument

In this section I'll show that both **RT-3** and **RT-4** are problematic. This gives us good reason to reject non-contrastive accounts of reasons in favor of contrastivism. At the end of the section, I'll consider an objection to my argument, and argue that it fails.

2.2.1 The argument

Suppose that I need to get to campus sometime today, but live twenty miles away. Further, suppose I'm out of shape and don't want to wear myself out getting there. Now consider the following reason ascriptions:

- (3) The fact that campus is twenty miles away is a reason to drive to campus rather than bike there.
- (4) The fact that campus is twenty miles away is a reason to bike to campus rather than run there.

Both of these sentences are intuitively true, given the set-up of the case. But, as I'll now show, this causes problems for the non-contrastivist.

First, notice that it follows on both **RT-3** and **RT-4** that if ' r is a reason for A rather than B ' is true, then r is a reason for A .⁴ So on both proposed analyses, (5) follows from (4):

(5) The fact that campus is twenty miles away is a reason to bike to campus.

Since (5) follows from (4) on both proposals, and since (4) is true, the non-contrastivist is committed to the truth of (5).

Now consider **RT-3**, which says that we should analyze ' r is a reason to A rather than B ' as saying that r is a reason to A and r is not a reason to B . According to this view, (6) follows from (3):

(6) The fact that campus is twenty miles away is not a reason to bike to campus.

The problem, of course, is that (5) and (6), both of which the non-contrastivist who adopts **RT-3** is committed to, are inconsistent. So this analysis of 'rather than' ascriptions fails.

Now consider **RT-4**, on which ' r is a reason to A rather than B ' means that r is a reason to A and r is a reason not to B . This view also leads to a problem, though not to an outright contradiction. For on that proposal, (7)—instead of (6)—follows from (3):

(7) The fact that campus is twenty miles away is a reason not to bike to campus.

The non-contrastivist who adopts **RT-4** is thus committed to both (5) and (7); that is, she is committed to saying that the fact that campus is twenty miles away is both a reason *to* bike to campus and a reason *not to* bike to campus. But this is an implausible result. So the second proposal, while not leading to an outright contradiction, does give us an implausible result.

⁴This also plausibly follows on **RT-1**, since if r is a stronger reason for A than it is for B , it is surely a reason for A . As an anonymous referee points out, this does not follow on **RT-2**. But, as I argued in section 1, **RT-2** (as well as **RT-1**) relies on an idiosyncratic treatment of 'rather than' in reason ascriptions.

2.2.2 Response: Denying Exclusivity

The non-contrastivist may question the implausibility of the result, that the fact that campus is twenty miles away is both a reason to bike there and a reason not to bike there. If this result isn't really troubling, then the argument against **RT-4** is no good. So why think it is troubling?

One reason that many people will find this result implausible is that it violates the following principle:

Exclusivity: For all facts r , agents s , and actions A , if r is a reason for s to A , then it's not the case that r is also a reason for s not to A .

Philosophers who seem to accept this principle include Nagel (1970); Raz (1999); Crisp (2000). If **Exclusivity** is true, then the argument I've given shows that none of the non-contrastivist analyses of 'rather than' ascriptions that I've considered so far could be correct.

But in fact there's reason to question **Exclusivity**. Jonathan Dancy has argued that the principle is false by appealing to cases like the following.⁵ Suppose I love to talk to pretty girls, but hate to be snubbed. And suppose that, unfortunately, pretty girls are likely to snub guys like me. Then the fact that the girl across the bar is so pretty is a reason for me to go talk to her, since I love to talk to pretty girls. But it's also a reason for me not to go talk to her, since I hate to be snubbed, and since pretty girls are likely to snub guys like me. If this is the right way to describe this case (and I think it is), then **Exclusivity** is false.

Further, common theories of reasons seem committed to the falsity of **Exclusivity**. On a simple desire-based theory of reasons, for example, when you have a desire that A -ing would help promote, and r explains why this is so, r is a reason to A .⁶ So if you have *two* desires, such that r explains both why A -ing would promote one of them and why not A -ing would promote the other, this kind of theory entails that r is both a reason for you to A and a reason for you not to A . Similarly, on a standard value-based theory, all we need is a case in which there are two values such that r explains both why A -ing would promote or respect one of them and why not A -ing would promote or respect the other.⁷

⁵See Dancy (1993), p. 62 for discussion.

⁶See Schroeder (2007), for example.

⁷See Parfit (2011); Scanlon (1998); Moore (1912), for example.

This lets us see where **Exclusivity** goes wrong. It is common ground between desire-based and value-based theories that our reasons for action are provided or explained by various objectives—desires, on the desire-based theory and values on the value-based theory (and perhaps both on a hybrid theory). When a fact r helps explain why my A -ing would promote or respect one of these objectives, r is a reason for me to A , which is provided by that objective.⁸ The problem with **Exclusivity** is that on most theories, there are multiple objectives that can provide reasons—agents have multiple desires, and several kinds of values are worth promoting or respecting. In the pretty girl case above, for example, I have two relevant desires: first, a desire to talk to pretty girls, and second, a desire not to be snubbed. The first desire explains why the fact that she’s so pretty is a reason to go talk to her: this fact explains why doing so would promote my desire to talk to pretty girls. Similarly, this same fact explains why not going to talk to her would promote my desire to not be snubbed. It is easy to construct a similar case in which multiple values, rather than multiple desires, are involved.

What this shows is that popular theories of reasons allow for different objectives to explain why one and the same fact can be a reason to A and not to A . But they do not allow for one and the same objective to explain why one and the same fact is both a reason to do and not to do one and the same action. That’s because one fact r can’t explain both (i) why A -ing would promote or respect an objective o , and (ii) why not A -ing would also promote or respect o .⁹ So the following weaker principle is very plausibly true, even if **Exclusivity** is false:

Restricted Exclusivity: For all facts r , agents s , actions A , and objectives o , o cannot explain both why r is a reason for s to A and why r is a reason for s not to A .

⁸For discussion of objectives providing reasons, see Moore (1912) (though he talked about rightness rather than reasons); Nagel (1970); Anderson (1993); Scanlon (1998); Finlay (2001, 2006); Schroeder (2007); Wedgwood (2009); Parfit (2011). For analyses of reasons in terms of explanation, see Toulmin (1950); Finlay (2001, 2006); Searle (2001); Broome (2004); Schroeder (2007). I will return to this idea in Chapter 4.

⁹To illustrate this, consider Schroeder (2007)’s Hypotheticalism. He holds that r is a reason for s to A iff r explains why A -ing would promote p , where p is the object of one of s ’s desires, and where promoting p is just making p more probable. But it would be strange to think that r could explain both why A -ing would make p more probable, and why not A -ing would also make p more probable. Similarly, it doesn’t seem that one fact could explain both why A -ing would respect, say, the demands of justice, and why not A -ing would also respect the demands of justice.

What this principle rules out is that one and the same objective can explain why one and the same fact is both a reason to A and a reason not to A . This is overwhelmingly plausible. The cases that seem to violate **Exclusivity**, like the pretty girl case above, do not violate this principle. And, as I argued above, popular theories of reasons, though they allow for violations of **Exclusivity**, do not allow for violations of **Restricted Exclusivity**.

Crucially, this principle is enough for my argument, because the joint truth of (5) and (7) *does* violate **Restricted Exclusivity**. The objective that provides, or explains, both of them is my desire not to wear myself out getting to campus. This explains both why the fact that campus is twenty miles away is a reason to drive rather than bike, and why this fact is a reason to bike rather than run. And since these ‘rather than’ ascriptions, according to the non-contrastivist who adopts **RT-4**, are to be analyzed partly in terms of the non-contrastive reasons mentioned in (5) and (7), this single desire also provides the reasons ascribed in both (5) and (7). Thus, **RT-4** leads to an implausible violation of **Restricted Exclusivity**.

2.3 Contrastivism

In this section I’ll develop a contrastivist account of reasons, and then use that account to explain the problematic case from the last section.

2.3.1 A contrastive account reason claims

Contrastivism differs from traditional non-contrastive theories by claiming that the reason relation includes an argument place for sets of alternatives. So the relation holds between (at least) (i) some fact r which is the reason, (ii) the particular alternative A which the fact is said to be a reason for, and (iii) the relevant set of alternatives *out of which* r is said to be a reason for A .¹⁰

In an explicit ‘rather than’ ascription like ‘The fact that today is your birthday is a reason to bake you a chocolate cake rather than not baking you a cake’, the alternatives are provided explicitly: {bake you a chocolate cake, don’t bake you a cake}.¹¹ These kinds of ascriptions, according to contrastivism, provide the model

¹⁰I’m ignoring agents here.

¹¹The set of alternatives might be larger than just two members, but since I’ve been dealing with ‘rather than’ ascriptions here, I’ll mostly limit my discussion to two-member sets.

for all reason ascriptions. But most reason ascriptions aren't like this; most are what I'll call 'bare ascriptions', like 'The fact that it's your birthday is a reason to bake you a chocolate cake'. Since the relation that serves as the semantic value of reason ascriptions—say, $R(r, A, Q)$ —has an argument place for a set of alternatives Q , this set must be provided somehow before we can evaluate the ascription. The most natural way to develop the theory, and the way I'll adopt, is to let the set of alternatives be provided by the context of utterance.¹² In a context in which a reason ascription is made, some particular set will be relevant. One straightforward way in which a set of alternatives might count as the relevant one is by including the options under discussion, though there are likely other ways.¹³

Since an explicit 'rather than' ascription gives the set of alternatives explicitly, these ascriptions will have a stable content across contexts (assuming there are no other context-sensitive terms). But since bare ascriptions have the set of alternatives provided by context, their content will be shifty across contexts. Here are the semantic principles I propose:

Explicit Ascriptions: ' r is a reason to ϕ rather than ψ ' is true in context c ¹⁴ iff r is a reason to ϕ out of $\{\phi, \psi\}$.

Bare Ascriptions: ' r is a reason to ϕ ' is true in context c iff r is a reason to ϕ out of Q , where Q is the relevant set of alternatives in c .

Bare Against Ascriptions: ' r is a reason not to ϕ ' is true in context c iff r is a reason not to ϕ out of Q , where Q is the relevant set of alternatives in c .¹⁵

¹²We might develop a theory on which it's provided instead by the context of assessment, or perhaps by some features of the agent's situation (though to count as a contrastivist theory, it should not always simply be the set of alternatives which it is possible for the agent to perform). Sinnott-Armstrong (2004, 2006) defends a view on which there's no saying which set is relevant for evaluating any particular bare reason ascription, which leads him to adopt Pyrrhonian skepticism, and thus to refuse to evaluate any bare ascription as either true or false.

¹³One common contrastivist idea is that intonational stress can help fix the set of alternatives. If I say 'You have a reason to buy some *milk*', that suggests a set of alternatives like {buy milk, buy juice}, whereas if I say 'You have a reason to *buy* some milk', that suggests a set of alternatives like {buy milk, steal milk}. See Rooth (1992); Dretske (1970); Schaffer (2005b, 2008) for discussion of stress and the role of alternatives.

¹⁴Note that 'in context c ' doesn't show up on the right-hand side of the biconditional. That's because, as I said, explicit ascriptions are not context-sensitive.

¹⁵If we can think of reasons against an option A simply as reason for $\neg A$, then this clause is just a special case of **Bare Ascriptions**. I agree that this is an attractive view, though things are slightly more complicated for the contrastivist, since $\neg A$ will not necessarily be in the same set of alternatives as A —sets of alternatives need not be *exhaustive* in this sense. And it's not

Now I'll show how to use this account to explain the problematic case from the last section.

2.3.2 A contrastive solution

One way to think about the problem facing the non-contrastivist is that, no matter what analysis of 'rather than' ascriptions we considered, we always derived that r is a reason to A from ' r is a reason to A rather than B '.¹⁶ This result, when combined with claims about B that we could also derive (and that are independently plausible anyway) led either to outright contradiction or to otherwise implausible results. So the culprit seems to be this inference:

RT: If r is a reason to A rather than B , then r is a reason to A .

The trouble is that this inference seems very hard to reject for the non-contrastivist.

The contrastivist, on the other hand, can not only reject this inference, but can explain why it seems so natural. Note first that **RT** is, of course, a non-contrastivist principle—there is no mention of contrasts on the right-hand side. So I suggest that we replace it with this:

CRT: If r is a reason to A rather than B , then r is a reason to A out of $\{A, B\}$ and r is a reason not to B out of $\{A, B\}$.

The second half of the consequent is perhaps less obvious than the first, but I think it is nevertheless very compelling. Given that A and B are mutually exclusive (why else would we use 'rather than'?), B -ing precludes A -ing. So if r is a reason to A out of $\{A, B\}$, it is very plausible that it's also a reason not to B out of this set.¹⁷

Importantly, given **Bare Ascriptions**, we can now see why **RT** seems so natural: it's because the default interpretation of the consequent expresses a truth. The antecedent, ' r is a reason to A rather than B ', makes salient the set of alternatives

clear how to make sense of a reason to $\neg A$ out of a set of alternatives that doesn't contain $\neg A$. I return to this point in Chapter 4. Further, some philosophers have recently given accounts on which reasons against A cannot simply be thought of as reasons for $\neg A$. See Greenspan (2005), for example.

¹⁶Again, **RT-2** does not validate this inference, which is an advantage of that proposal. But it does rely on a problematic treatment of 'rather than' in reason ascriptions.

¹⁷This shows that my view does respect the fact that 'rather than' means something like 'and not': relative to the set of alternatives $\{A, B\}$, when ' r is a reason to A rather than B ' is true, r is a reason to A *and* a reason *not* to B .

$\{A, B\}$. And according to **Bare Ascriptions**, when we interpret the consequent, ‘ r is a reason to A ’ relative to this set, it means that r is a reason to A out of $\{A, B\}$. Finally, according to **CRT**, this will be true as long as ‘ r is a reason to A rather than B ’—the antecedent of **RT**—is true.

The mistake, though, is to move from the truth of ‘ r is a reason to A ’ in the context set up by the antecedent of **RT** to the truth of this claim in *every* context. If the relevant set of alternatives is not the one provided by the antecedent, $\{A, B\}$, then we have no guarantee that ‘ r is a reason to A ’ will be true, relative to this new set.

This puts us in a position to see how the contrastivist can solve the puzzle. First, here are the relevant sentences, re-written to make the contrasts explicit.

- (3*) The fact that campus is twenty miles away is a reason to drive rather than bike (out of $\{\text{drive, bike}\}$).
- (4*) The fact that campus is twenty miles away is a reason to bike rather than run (out of $\{\text{bike, run}\}$).
- (5*) The fact that campus is twenty miles away is a reason to bike (out of $\{\text{bike, run}\}$).
- (7*) The fact that campus is twenty miles away is a reason not to bike (out of $\{\text{drive, bike}\}$).

The truth of the ‘rather than’ claim (3), along with **CRT**, tells us that (7) is true in a context in which the relevant set of alternatives is $\{\text{drive, bike}\}$, so we interpret (7) as (7*), above. But in this context, by **Bare Ascriptions**, (5) would be true only if the fact that campus is twenty miles away is a reason to bike relative to this same set:

- (5’) The fact that campus is twenty miles away is a reason to bike out of $\{\text{drive, bike}\}$.

But this is false. So in the context in which (7) is true, (5) is false.

The truth of the ‘rather than’ claim (4), along with **CRT**, tells us that (5) is true in a context in which the relevant set of alternatives is $\{\text{bike, run}\}$, so we interpret (5) as (5*) above. But in this context, by **Bare Against Ascriptions**, (7) would be true only if the fact that campus is twenty miles away is a reason not to bike relative to this same set:

(7') The fact that campus is twenty miles away is a reason not to bike out of {bike, run}.

But this is false. So in the context in which (5) is true, (7) is false.

Thus, there's no one context in which both (5) and (7) are true—at least not one in which both reasons are provided by my desire not to wear myself out getting to campus.¹⁸ It's true, of course, that (3) and (4) are true in the same contexts, but that's not puzzling—no one should deny that these 'rather than' claims can be jointly true. The key to the contrastivist solution is, essentially, treating all reason ascriptions as (at least implicitly) 'rather than' ascriptions—as relative to sets of alternatives.

2.4 Negative Reason Existentials

So far I've shown that contrastivism can solve problems facing non-contrastive theories having to do with 'rather than' ascriptions. In this section, I'll show that contrastivism can solve a puzzle about a different kind of reason claim, *negative reason existentials*. These are claims like 'There's no reason to cry over spilled milk'. These sorts of claims are also very common, and play an important role in arguments against various theories in ethics, epistemology, and practical reasoning: some theory implies that there is a reason to ϕ , but intuitively there is no reason to ϕ , so the theory is false. Schroeder (2007) argues, though, that our intuitions about negative reason existentials are not to be trusted. After presenting the puzzle, I will argue that a contrastive solution is more attractive than Schroeder's own solution.

2.4.1 The puzzle

Consider the following famous case.

THE BOOK THIEF. You see Tom Grabit run out of the library, pull a book from under his coat, cackle gleefully, and then run away.¹⁹

¹⁸What is the status of **Restricted Exclusivity**? It's a non-contrastivist principle, so I need to reject it. But I relied on it in my argument against the non-contrastivist. I claim that the non-contrastivist is committed to this principle, which is why the case I presented in section 2 is problematic for her. The contrastivist should replace **Restricted Exclusivity** with a contrastive version, which says that one and the same objective can't explain why one and the same fact is a reason to do and not to do one and the same action *relative to* one and the same set of alternatives.

¹⁹Lehrer and Paxson (1969).

Intuitively, you have a reason to believe that Tom Grabit stole the book. But now consider this variation:

THE TWIN BOOK THIEF. You see Tom Grabit run out of the library, pull a book from under his coat, cackle gleefully, and then run away. Just as you begin to form the belief that Tom stole the book, your friend tells you about Tom's twin brother, Tim. Tom and Tim are indistinguishable, even to their mother.

Imagine that we've hauled Tom into the police station and charged him with stealing the book. He protests, "Wait! I have a twin brother! You have no reason to think I'm the one who stole the book!". His protest certainly seems warranted. So intuitively, in this case, you have no reason to believe that Tom stole the book. The reason provided by your visual evidence is *undercut*. And, it seems, when a reason is undercut, it's just not a reason anymore. But Schroeder argues that our intuitive judgment that there is no reason for you to believe that Tom stole the book in this case is mistaken. For consider the following variation:

THE TRIPLET BOOK THIEF. You see Tom Grabit run out of the library, pull a book from under his coat, cackle gleefully, and then run away. Just as you begin to form the belief that Tom stole the book, your friend tells you about Tom's two identical siblings, Tim and Tam. The triplets are indistinguishable, even to their mother.

Intuitively you have even *less* reason to believe that Tom stole the book in this case than you did in THE TWIN BOOK THIEF. But if you have even less reason in this case, you couldn't have had *no* reason in the previous case. So our intuition that you had no reason is mistaken. This is the puzzle about negative reason existentials: some of them are intuitively true, though we also seem to be committed to the existence of the very reasons that are said not to exist.²⁰

2.4.2 The pragmatic solution

Schroeder offers a pragmatic explanation of this puzzling data. The first observation is that some reasons are very weak, and very weak reasons are not generally

²⁰See Schroeder (2007), p. 93.

relevant in our deliberation about what to do. The second observation is that—following Grice (1989)—talking about irrelevant things is conversationally inappropriate. Thus, we are reluctant to assume that people who tell us about reasons are talking about very weak reasons, since this would be conversationally inappropriate. Instead, we make the presumption that they mean to be talking about relatively weighty reasons—their claim has this implicature. This presumption is reinforced by making a bare existential claim, which is less informative than a claim specifying what the reason is. For a bare existential claim to be sufficiently informative, the reason must be relatively weighty. But since the reasons in question are not weighty, the implicature is false, and so the claim seems false. Since the reason existential ‘You have a reason to believe that Tom stole the book’ seems false, its negation, the negative reason existential ‘You have no reason to believe that Tom stole the book’, in *THE TWIN BOOK THIEF*, sounds true.

This explanation predicts that saying what the reason is should make it sound less bad, since this will make the claim more informative. It also predicts that saying that the reason is a really weak one will make it sound less bad, since that would cancel the implicature that the reason is relatively weighty. And these predictions are plausibly correct. To illustrate the idea, Schroeder (2007) uses the example, ‘There is a reason to eat your car’. This seems obviously false. But if I say ‘The fact that eating your car will give you the daily recommended amount of iron is a reason to eat your car’, that sounds better. And it sounds better still if follow this up with ‘...but it’s a really weak reason’. This supports the hypothesis that there are pragmatic factors at work here. So Schroeder argues that there *are* in fact reasons in these cases, and thus that the negative reason existentials are false, but offers an explanation of why they sound true anyway.

But for this case to actually support Schroeder’s pragmatic solution, the explanation for why the negative reason existential initially seems true must be the reason’s low weight. But I think there is a better explanation. It’s not that we ignore the reason because it’s so weak; rather, I think that we simply *miss* the reason because it’s so non-obvious. Now, I agree that even once we acknowledge it, we’d do better to ignore it, since it’s so lightweight. But the point is that the weight of the reason is not the explanation for why the negative reason existential sounds true. The amount of iron you could get from eating your car is simply not

something we would ever think about, unless someone pointed it out to us. Confirmation for this explanation over Schroeder's comes from the fact that the difference in acceptability between (8) and (10) is more striking than the difference between (8) and (9):

(8) There's a reason to eat your car.

(9) There's a reason to eat your car, but it's a really weak one.

(10) The fact that it contains your daily recommended amount of iron is a reason to eat your car.

If Schroeder's explanation were correct, we should expect a much stronger difference between (8) and (9) than between (8) and (10), since (9) *explicitly* cancels the implicature that supposedly makes (8) sound false. But on my non-obviousness explanation, we should expect a stronger difference between (8) and (10). I think the data here confirms my explanation over Schroeder's.²¹

The important point for the Tom Grabit case is that the supposed reason, that you saw someone who looked just like Tom run out of the library with the book, isn't at all non-obvious. So we need a different explanation for the Tom Grabit case. I'll argue now that contrastivism offers an attractive explanation. The contrastivist explanation is also attractive because it holds that the intuitively true negative reason existential—Tom's compliant that we have no reason to think that he stole the book—is actually true, rather than trying to explain away that intuition.

2.4.3 A contrastivist solution

The puzzling data about negative reason existentials is that there are some claims like 'There is no reason to ϕ ' which are intuitively true, though we seem committed to there actually being reasons to ϕ . Thinking about THE TWIN BOOK THIEF, 'There is no reason to believe that Tom stole the book' seems true. But when we consider THE TRIPLET BOOK THIEF, it seems that there is even *less* reason to believe that Tom stole the book in this case than in the previous one. So there couldn't have been *no* reason before, contrary to our intuitions about the negative

²¹I agree that 'The fact that it contains your daily recommended amount of iron is a reason to eat your car, but it's a really weak one' sounds (a bit) better than (10). But I can explain that by pointing out that this claim is simply more informative than (10).

reason existential. Once we reflect on this fact, we seem to be committed to the claim that there *is* a reason to believe that Tom stole the book, after all.

Here's how to explain this phenomenon using contrastive reasons. In the context given by the THE TWIN BOOK THIEF, the most natural contrast to believing that Tom stole the book is believing that Tim, Tom's twin brother, stole the book, since this alternative has just been made relevant. And it's true that there is no reason to believe that Tom stole the book rather than believing that Tim stole the book. This is why it's true to say 'There's no reason to believe that Tom stole the book', when we're considering THE TWIN BOOK THIEF. But when we consider THE TRIPLET BOOK THIEF, and notice that there's even less reason to believe that Tom stole the book in this case than in the previous case, and so want to accept 'There is a reason to believe that Tom stole the book', the contrast to believing that Tom stole the book is most plausibly *not* believing (that is, failing to believe) that Tom stole the book.

Here's why. We need some reason that exists in both THE TWIN BOOK THIEF and in THE TRIPLET BOOK THIEF, whose weight we can compare across the cases. Since Tam wasn't in the picture in THE TWIN BOOK THIEF, we can't have in mind any reasons involving Tam. So the only two plausible candidates, on a contrastive framework, are a reason to believe that Tom stole the book rather than believing that Tim stole it and a reason to believe that Tom stole the book rather than not believing that Tom stole it. But there's no reason to believe that Tom stole the book rather than believing that Tim stole it in *either* case. So that only leaves a reason to believe that Tom stole it rather than not believing that Tom stole it. And there very plausibly is this reason in both cases, and it's very plausibly weaker in THE TRIPLET BOOK THIEF than it is in THE TWIN BOOK THIEF.

Any theory of reasons has to allow that some reasons are weightier than others, and that some consideration might be a weightier reason for one thing than it is for another thing. Contrastivism is no different—in THE TWIN BOOK THIEF, the fact that you saw (what looked like) Tom run out is a weightier reason to believe that he stole the book out of {believe that Tom stole the book, don't believe that Tom stole the book} than it is in THE TRIPLET BOOK THIEF.²² This solution

²²One important thing to keep in mind here is that according to contrastivism, the weight of a reason to ϕ out of Q does not depend on the weight of any non-contrastive reasons to ϕ and non-contrastive reasons for the other alternatives in Q . That's because, of course, there just are no non-contrastive reasons. This does raise the question of how best to think about the weight of

is available to the contrastivist because she holds that reasons are relativized to possibly non-exhaustive sets of alternatives. One fact can fail to be a reason to believe that p relative to one set, be a reason to believe that p relative to a second set, and be an even weightier reason to believe that p relative to a third set.

Contrastivism has a plausible explanation of why the negative reason existential is literally true, as well as why the positive reason existential that we accept after considering the comparative claim about the weight of reasons is true. The reason we're saying does not exist is a reason to believe that Tom stole the book out of {believe that Tom stole the book, believe that Tim stole the book}, while the reason whose weight we're comparing across cases is a reason to believe Tom stole the book out of {believe Tom stole the book, don't believe Tom stole the book}.²³ On a non-contrastive theory, on the other hand, these claims are incompatible. So this is an advantage for contrastivism: we need not bite the bullet and admit that one of our intuitive judgments is false. And moreover, contrastivism is independently well-motivated by the puzzle about reason against ascriptions.

2.5 Looking Forward

I've argued in this chapter that contrastivism about reasons gains some support from considerations of our talk about reasons. The solutions to the puzzles depend on the contrastive semantics that fit very naturally with contrastivism. But in the next chapter, I'll consider how far the argument in this chapter actually takes us.

a contrastive reason, but virtually any theory needs to give some account of weight, and few do. I'll rely on an intuitive understanding.

²³Some people deny that the negative reason existential even seems true. I think that framing things as I did above, by imagining Tom protesting when we accuse him of stealing the book, makes it even more intuitive that the claim is true. My hypothesis for their resistance is that they have in mind the set {believe that Tom stole the book, don't believe that Tom stole the book}. And contrastivism predicts that, relative to this set, there is a reason to believe that Tom stole the book, so the negative reason existential *should* seem false.

Chapter 3

FAVORING

In the last chapter, I argued that certain features of our talk about reasons supports contrastivism. Relativizing reason ascriptions to sets of alternatives solves puzzles that face non-contrastive theories. In this chapter, though, I'll consider how far that argument actually takes us. The puzzles concerned reason *ascriptions*, and the solutions came from a contrastive account of those ascriptions. It's very natural to think that this is strong support for the thesis that reasons are fundamentally contrastive. If reason ascriptions are contrast-sensitive, then it seems obvious that the reason relation, which presumably serves as the semantic content of 'reason', must be contrastive as well. But this is an apparently radical thesis, given that reasons are traditionally taken to be reasons for things *simpliciter*. We might hope, then, to avoid committing ourselves to this sort of theory prematurely.

3.1 Why Resist Contrastivism?

As I've emphasized, reasons are traditionally taken to be reasons for things *simpliciter*, independently of the contextually relevant alternatives. So the first reason why we might want to resist contrastivism is simply that it's *radical*. Nearly everyone who writes about reasons seems to assume a non-contrastive picture (usually implicitly, of course). So adopting contrastivism would require rejecting—or at least seriously reinterpreting—these writers' claims.

Second, contrastivism might seem mysterious. Scanlon (1998) says that reasons are considerations that count in favor of the things they are reasons for, and many other writers adopt this gloss. But if we accept this gloss, contrastivism might seem puzzling: how could the relevant set of alternatives matter for whether or not some

consideration counts in favor of an action or attitude? The *availability* of some alternative can arguably alter whether or not some consideration is a reason to do something. But contrastivism makes a more surprising claim than this. Reasons are only reasons *relative to* some set of alternatives, whether or not that set includes all of the available alternatives. So it's not simply the other available alternatives that matter in determining whether some consideration is a reason for an action: it also matters what we're comparing the action to.

Third, many theories in ethics and epistemology appeal to normative reasons, and some have even taken the notion of a reason to be *the* fundamental normative notion, reducing all other normative concepts to reasons.¹ But if we complicate the fundamental normative relation by adding an argument place for a set of alternatives, we might worry that this will complicate our analyses of other normative concepts. Even if, as I argued in Chapter 1, contrastivism about reasons doesn't require us to be contrastivists about these other concepts, the analyses of these other concepts might be more complicated if we adopt contrastivism about reasons.

3.2 Shallow Contrastivism

The arguments for contrastivism in Chapter 2 concerned reason ascriptions. So if we want to avoid contrastivism in the face of those arguments, a natural strategy is to develop a theory which can mimic the contrastivist's solution to the puzzles I presented in the last chapter, while retaining a more traditional non-contrastive metaphysics to avoid some of the complications I mentioned in the last section. To do this, we need a theory that offers a contrastive account of reason ascriptions but accepts a non-contrastive metaphysics of the important normative favoring relation that philosophers are concerned with when they theorize about reasons, and in terms of which we analyze other important normative concepts. Call this kind of theory *shallow contrastivism*; call the more radical theory *deep contrastivism*.

The basic idea is simple. Rather than tracking when some consideration does or does not favor an action *simpliciter*, our reason ascriptions track *the degree to which* the consideration favors the action as compared to the degree to which it favors the other relevant alternatives. So at the bottom we have a non-contrastive favoring

¹See, for example, Raz (1999); Schroeder (2007); Scanlon (1998). See, for example, Broome (2004); Väyrynen (2010) for skepticism about this project.

relation: facts favor certain actions (or attitudes) to a certain degree *simpliciter*. When a fact r favors ϕ -ing *more* than it favors ψ -ing, we can truly say ‘ r is a reason to ϕ rather than ψ ’. Or, if ψ -ing is the contextually relevant alternative to ϕ -ing, we can simply say ‘ r is a reason to ϕ ’, though what this means is that r favors ϕ -ing more than it favors the contextually relevant alternatives.

This view arguably has some precedent in Alastair Norcross’s work defending consequentialism.² He says: “[O]ur (moral) reasons for choosing between alternative actions, institutions, etc. are essentially comparative, and correspond to the comparative consequential value of the options”. This quote suggests that Norcross is actually a *deep* contrastivist about (moral) reasons. But he also says, in the next sentence, “I might have a better reason for choosing to do A than to do B, and better by a certain amount, but neither reason is either good or bad *simpliciter*”.³ This quote seems to fit better with a shallow contrastivist view, on which considerations favor alternatives with a certain weight, though we can only call them (good or bad) reasons relative to some comparison class. Moreover, Norcross understands his project as providing a contextualist *linguistic* theory for consequentialism, where context provides some appropriate alternatives, and we can only say something is good or bad *as compared* to these alternatives. But at bottom, he seems to assume a traditional view of favoring.

Now I’ll develop shallow contrastivism in more detail. First I’ll outline a traditional non-contrastive account of the favoring relation, then I’ll develop a contrastive account of reason claims and pair the two to end up with shallow contrastivism.

3.2.1 A traditional theory of favoring

A traditional, non-contrastive view of reasons says that reasons are considerations that favor certain actions with a certain weight or strength, some weightier than others. Most people agree that it’s a mistake to identify the weight of a reason with anything as precise as a number, but for simplicity that’s what I’ll do here (or at least *represent* it as a number). So we can think of the relation that a fact stands in to the thing it favors as something like this: $F(r, A, x)$. This means that fact r *favors* action A with weight x . This is a relation that holds between facts, actions, and weights independently of contextually relevant alternatives. We

²See Norcross (1997, 2005a,b).

³Norcross (2005a), p. 81.

can use this relation to obtain a *ranking* of alternatives based on how strongly r favors them. So if we have $F(r, A, 2)$ and $F(r, B, 5)$, we would rank B above A on the ranking of alternatives generated by r .⁴ We can understand disfavoring as favoring with a *negative* weight: $F(r, A, -5)$, for example.⁵ This favoring relation, according to shallow contrastivism, is the important normative relation that writers who talk about reasons are really concerned with, and the relation in terms of which we should try to analyze other normative concepts. That’s why nearly everyone glosses a reason as a consideration that *favors* some action or attitude (with a certain weight).

3.2.2 A contrastive account of reason claims

The next step is to pair this idea with a contrastive account of reason claims. The way deep contrastivism solves the puzzles is by making bare reason claims context-sensitive: as the contextually relevant set of alternatives varies, the content of the reason ascription shifts. The shallow contrastivist wants to mimic this semantics.

There’s an apparent tension here. What we want is for the underlying favoring relation F to hold independently of contextually relevant alternatives, but the content of the reason ascription to shift with those alternatives. So it looks like the relation which serves as the semantic value of reason claims—the semantic value of ‘reason’—has to include an argument place for sets of alternatives. But we want the important normative relation, the favoring relation, *not* to include such an argument place, so that it’s a non-contrastive relation.

So what we have to do is reject the natural idea that instances of the underlying favoring relation are the semantic content of reason ascriptions. Instead, I suggest we posit an “intermediate” relation, $R(r, A, Q)$, to serve as the semantic content of reason ascriptions. Q is the relevant set of alternatives. I’ll call R the ‘reason relation’, since it’s the content of reason ascriptions. But all this means is that the fundamental relation—the one we really care about in ethics, epistemology, and so

⁴We can take more facts into consideration to obtain a ranking of how strongly the balance of all the relevant considerations favor the different alternatives by performing some function on the weights with which each fact favors the various alternatives. (This is essentially the picture I sketched in Chapter 1 to show that contrastivism about ‘ought’ does not require contrastivism about reasons.)

⁵There might be complications here, but since I don’t ultimately want to defend shallow contrastivism, I won’t focus on those here.

on—is not (what I’m calling) the reason relation, but is rather the favoring relation F .

We can respect the fact that the favoring relation F is the fundamental relation by letting R depend on F in the following way:

Shallow Reasons: $R(r, A, Q)$ iff there’s some x such that $F(r, A, x)$, and for all B in Q not identical to A , there’s some y such that (i) $F(r, B, y)$, and (ii) $x > y$.

Intuitively what this says is that r is a reason for A out of Q iff r favors A more than it favors any other alternative in Q . And then we say that ‘ r is a reason for A ’ is true in context c iff $R(r, A, Q)$, for the relevant set of alternatives Q in c .⁶

So when philosophers say things like ‘Reasons are considerations that count in favor of actions or attitudes’, we have to do some mild reinterpretation: they mean that reasons are considerations that favor things *more* than they favor the relevant alternatives. This is not a very radical reinterpretation. When I say that r is a reason to ϕ , this can still be relevant for determining what I’m justified in doing, or what I ought to do. That’s because my claim can only be true if r really does favor (non-contrastively) ϕ -ing, and does so more than it favors other relevant actions. And our analyses of concepts like ‘ought’ should technically proceed in terms of favoring, not reasons. So instead of saying, for example, that you ought to ϕ when you have most reason to ϕ , we can say that you ought to ϕ when ϕ -ing is the most favored thing you can do.⁷

⁶There’s one complication I want to address here. Some considerations intuitively don’t stand in the favoring relation to some actions—not even with a negative weight. For example, the fact that $2+2=4$ doesn’t seem to favor or disfavor my drinking another cup of coffee: there’s no weight x such that $F(2+2=4, \text{drink another cup of coffee}, x)$. One initially attractive idea is to say that in these cases, the fact favors the action with weight zero. But this will give us odd results in some cases. Suppose that r neither favors nor disfavors A , and slightly disfavors B . Then it would be true, on the theory as I’ve developed it so far, to say ‘ r is a reason to A ’, in a context in which the set of alternatives is $\{A, B\}$, since zero is greater than whatever negative number we use for the weight with which r disfavors B . But it’s strange to say that r is a reason for A when it’s completely irrelevant to A .

I think there are a couple of ways to respond to this problem. We could say instead that when r is irrelevant to A , r is just doesn’t stand in the F relation to A at all. Then it wouldn’t be true to say that r is a reason to A . That would require complicating the analysis of R in terms of the F relation to accommodate this change. A second solution is to bite the bullet and admit that, if the salient set of alternatives is $\{A, B\}$, and r disfavors B , then it *is* a reason to A *rather than* B . I don’t think this is terribly implausible, and we could try to explain away the strangeness using pragmatic considerations. For now, though, I’ll just leave this issue aside.

⁷The particular analysis doesn’t really matter. The one I’ve mentioned allows us to avoid contrastivism about ‘ought’, since ϕ -ing has to be more favored than anything else you could do,

Shallow contrastivism can essentially adopt the semantic principles I gave in Chapter 2, but state things in terms of the R relation, which can then be cashed out in terms of the underlying F relation. Nothing about the solutions to the puzzles about reason claims depended on the relation which serves as the content of reason ascriptions being the fundamental relation. So the fact that deep contrastivism can solve puzzles is not (for all I've said so far) evidence for deep contrastivism over shallow contrastivism.

3.2.3 Shallow contrastivism and Exclusivity

Before I evaluate shallow contrastivism, I want to point out one interesting commitment of shallow contrastivism. It requires that **Exclusivity** is true, at least in cases in which the relevant set of alternatives is $\{\phi, \neg\phi\}$. Here again is that principle:

Exclusivity: If r is a reason for s to ϕ , then r is not also a reason for s not to ϕ .

Lots of people think that **Exclusivity** is true, so they won't see this as a problem for shallow contrastivism. But I gave a case in the last chapter which I think is plausibly a counterexample to this principle.⁸ The fact that the girl across the bar is so pretty is a reason to talk to her, since I love to talk to pretty girls. But it's also a reason not to talk to her, since I hate to be snubbed, and pretty girls are likely to snub guys like me. If this case, or one like it, is a counterexample to **Exclusivity**, then shallow contrastivism has a problem.

Here's why. In the case I presented, the set of relevant alternatives is clearly $\{\text{talk to her, don't talk to her}\}$. For the fact that she's so pretty to be a reason to talk to her out of this set, this fact has to favor talking to her more than it favors not talking to her. And for it to be a reason not to talk to her out of this set, it has to favor not talking to her more than it favors talking to her. But this is impossible. If the shallow contrastivist's favoring relation is going to be anything like the traditional understanding of the favoring relation, then at least this much has to hold: if r favors ϕ -ing more than ψ -ing, then it doesn't also favor ψ -ing

not just more than some (potentially non-exhaustive) set of alternatives. A contrastive view of 'ought' (see Sloman (1970); Jackson (1985); Finlay (2009); Cariani (2009, fc); Snedegar (2012)) could say instead that you ought to ϕ when ϕ -ing is more favored than any of the relevant alternatives.

⁸See Dancy (1993), p. 62 for another.

more than ϕ -ing. So again, if **Exclusivity** is false, then shallow contrastivism is in trouble.⁹ But since lots of people accept **Exclusivity**, I'll now move on to argue that the favoring relation itself is contrastive.

3.3 Favoring

In this section I'll assess the shallow contrastivist strategy for resisting deep contrastivism by discussing the favoring relation directly, rather than the reason relation.¹⁰ I'll argue that the favoring relation itself is contrastive—facts favor actions or attitudes only relative to sets of alternatives.¹¹

If the favoring relation is contrastive, that's strong support for deep contrastivism. First, the claim that reasons are considerations that favor things is ubiquitous.¹² If considerations only favor things relative to sets of alternatives, it follows from this claim that they're only reasons relative to sets of alternatives. Second, and particularly relevant for this chapter, of course, is that if the favoring relation is contrastive, then shallow contrastivism is false. We can't resist deep contrastivism by appealing to a non-contrastive favoring relation.

3.3.1 Against non-contrastive favoring

The argument that favoring is always favoring relative to some set of alternatives, rather than favoring *simpliciter*, begins from the following case. Suppose I have three dinner invitations: Invitation A is for Armenian with Ara. Invitation B is for burgers with Burt. Invitation C is for Chinese with Charlie. I love spending time

⁹The deep contrastivist doesn't have this problem. The fact that the girl across the bar is so pretty is a reason to talk to her out of this set, and a reason not to talk to her out of this same set. But since the explanations for these two reason claims are different, we don't have a violation of **Exclusivity***.

¹⁰Most people, including deep contrastivists, will think these are just the same relation, of course. But this is exactly what the shallow contrastivist has to deny.

¹¹The argument I'll give is a development of a discussion in Ross (2006), Chapter 9. I should note that Ross is careful to keep contrastivism about reasons at arm's length, claiming merely that talking in terms of contrastive reasons and contrastive favoring avoids some difficulties—difficulties that I'll develop into full-fledged objections—and is thus more convenient for his purposes.

¹²Scanlon (1998) is particularly explicit about this, but nearly everyone writing about reasons accepts it. Hieronymi (2005) argues that this is not the best way to think of reasons, because it leads to the "wrong kind of reasons" problem. She argues that it's better to think of reasons as considerations that bear on a question. But she also accepts that reasons favor the things they're reasons for *by* bearing on certain questions.

with Ara, really like spending time with Burt, but can barely tolerate spending time with Charlie. On the other hand, I love Chinese, really like burgers, but can barely stomach Armenian.

The following claims are clearly true. The fact that Invitation B is for dinner with Burt favors choosing Invitation B rather than Invitation C. And the fact that Invitation B is for burgers favors choosing Invitation B rather than Invitation A. But does the fact that Invitation B is for dinner with Burt favor choosing Invitation B *simpliciter*? Or does the fact that Invitation B is for burgers favor choosing Invitation B *simpliciter*? It might seem like the answer to both questions is clearly ‘yes’: after all, these two facts favor choosing Invitation B rather than Invitations C and A, respectively. And how could they do that if they didn’t favor choosing Invitation B *simpliciter*?

But the following claims are also true. The fact that Invitation B is for dinner with Burt does not favor choosing Invitation B rather than Invitation A—after all, I like spending time with Ara more than I like spending time with Burt. And similarly, the fact that Invitation B is for burgers does not favor choosing Invitation B rather than Invitation C—after all, I like Chinese food more than I like burgers. So perhaps these facts actually don’t favor choosing Invitation B *simpliciter*. I could do better with respect to my food preferences by not choosing Invitation B, and I could do better with respect to my dinner companion preferences by not choosing Invitation B.

But given the set up of the case, it’s very plausible that all things considered, I ought to choose Invitation B. If neither of the relevant facts—that Invitation B is for burgers, and that it’s for dinner with Burt—favor choosing B, then it looks like *nothing* favors choosing Invitation B. But now we have an odd result: I ought to choose Invitation B, though nothing favors choosing it. So it seems implausible, if you’re a non-contrastivist about favoring, to say that these facts don’t favor choosing Invitation B *simpliciter*.

At this point, we can conclude (with Ross (2006)) the following. There are cases, like the one above, in which it’s very hard to say whether some fact favors some action *simpliciter*, but very easy to say whether it favors the action *rather than* some other action. My suggestion is that this is because the favoring relation is fundamentally contrastive. Cases like this one make this clear: to answer the question ‘Does *r* favor *φ*-ing?’, we have to know ‘Rather than what?’. In lots of

cases that philosophers have been interested in, the relevant consideration favors the relevant action rather than any of the salient alternatives, so it's not so important to specify. But in cases like the one above, it *is* important.

Here's one natural response. We should understand 'rather than' claims as simple comparisons about how much the consideration favors the actions. So when I said that the fact that Invitation B is for burgers favors choosing it rather than Invitation A, that was true. But what this means is just that this fact favors choosing Invitation B *more* than it favors choosing Invitation A; we don't need a fundamentally contrastive favoring relation to explain this.

I think this response is problematic. The fact that Invitation B is for burgers does not favor choosing Invitation B rather than Invitation C—remember that I like Chinese food more than burgers. If we extend the explanation offered in the last paragraph in the natural way, we have to say that this is because the fact that Invitation B is for burgers does not favor choosing it more than it favors choosing Invitation C. But this fact doesn't seem to favor choosing Invitation C at all—it doesn't have anything to do with Invitation C. So if we insist that it does favor choosing Invitation B *simpliciter*, we have to say that despite this, it does not favor choosing Invitation B rather than Invitation C, even though it doesn't favor choosing Invitation C *at all*.

Here's a more sophisticated version of this response. I've actually misidentified the relevant facts that do the favoring in this case. The relevant facts are the following two more complicated facts:

- (i) I like Chinese more than burgers and burgers more than Armenian *and* Invitation A is for Armenian, Invitation B is for burgers, and Invitation C is for Chinese.
- (ii) I like spending time with Ara more than Burt and spending time with Burt more than Chris *and* Invitation A is for dinner with Ara, Invitation B is for dinner with Burt, and Invitation C is for dinner with Charlie.

Fact (i) favors choosing Invitation C to the highest degree, choosing B to a slightly lower degree, and choosing A to a much lower degree. Fact (ii) favors choosing Invitation A to the highest degree, choosing B to a slightly lower degree, and choosing C to a much lower degree. As suggested above, 'rather than' claims simply give us comparisons between the degree to which the relevant fact favors

the options. The reason why I ought to choose Invitation B, all things considered, is that it gets “the most” favoring, once we add up the amounts provided by facts (i) and (ii).

This view does seem to avoid the problems I’ve presented so far, but it has the counterintuitive consequence that *strictly speaking*, the fact that Invitation B is for burgers *doesn’t* favor choosing it rather than Invitation A. It’s only the more complicated fact (i) that favors choosing Invitation B rather than Invitation A. Contrastivism about favoring doesn’t have this consequence. We can avoid saying that our ordinary claims and intuitions about what favors what are actually false, or at least incomplete. The contrastivist can admit that these other facts about my preferences serve as background conditions which enable the fact that Invitation B is for burgers to favor choosing Invitation B rather than Invitation A.¹³

There is also a sharper problem with this response. Schroeder (2007) points out that a view analogous to this one (which he calls No Background Conditions) has implausible consequences when combined with a very plausible Deliberative Constraint. The Deliberative Constraint says that when we deliberate about what to do, we should think about our reasons, or about the considerations that favor the various options open to us. If the view under consideration now is correct, these considerations are (i) and (ii) above. But note that (i) and (ii) involve facts about me and my preferences, such as the fact that I like Chinese more than burgers and burgers more than Armenian. So if the Deliberative Constraint is true, I have to think about my own preferences in deciding which invitation to accept. But this seems objectionable: if I prefer Armenian to burgers, then just knowing that Invitation B is for burgers and Invitation A is for Armenian should be enough to move me (at least some way) toward accepting Invitation B rather than Invitation A. Thinking explicitly about my own preferences makes deliberation out to be, in Schroeder’s words, *objectionably self-regarding*.

There’s actually a more pressing related problem. Not only do I have to think about my own preferences in deliberating, according to this view. If what really favors Invitation B more than, say, Invitation C is fact (ii) above, then even if I’m only deciding between Invitations B and C, I have to think about how I feel about

¹³Again, compare Schroeder’s discussion of the No Background Conditions view in chapter 2 of Schroeder (2007). See also Dancy (2004)’s discussion of enablers. There Dancy argues explicitly that some facts play an important normative role by *enabling* other facts to favor certain alternatives, without themselves favoring the alternative.

Armenian food (if the very plausible Deliberative Constraint is true). But this is a bizarre result.

The defender of a fundamentally non-contrastive favoring relation (the shallow contrastivist, for example) could offer a more complicated story about these ‘rather than’ favoring claims, or just bite the bullet that most of our claims and intuitions about what favors what are incomplete, or strictly speaking mistaken, and try to explain away these odd results. But I do think I’ve said enough to show that whatever the non-contrastivist about favoring says is going to have to be somewhat complicated, or have some counterintuitive consequences. The contrastivist, on the other hand, offers a neat explanation. The fact that Invitation B is for burgers favors choosing Invitation B out of {choose Invitation B, choose Invitation A}, but not out of {choose Invitation B, choose Invitation C}. There’s just no question about whether some consideration favors choosing Invitation B *simpliciter* since all favoring is favoring out of a set of alternatives.

3.4 Contrastive Reasons and Favoring

So far, in describing contrastivism, I’ve remained at a pretty abstract level. In this section, I’ll develop contrastivism further and say more about what it takes for some consideration to be a reason to perform some action, or have some attitude, out of a set of alternatives. I’ll also consider the relationship between contrastive reasons and what an agent ought to do.

3.4.1 Contrastivism

I’ve already mentioned how the contrastivist about favoring can handle the case I’ve been considering in passing. We can say that the fact that Invitation B is for burgers favors choosing Invitation B out of {choose Invitation A, choose Invitation B}, but not out of {choose Invitation B, choose Invitation C}. Similarly, the fact that Invitation B is for dinner with Burt favors choosing Invitation B out of {choose Invitation B, choose Invitation C}, but not out of {choose Invitation A, choose Invitation B}. There’s just no question about whether it favors choosing Invitation B *simpliciter* according to contrastivism, which is good, since it’s such a hard question to answer.

Here is the contrastivist version of the idea that reasons are considerations that favor the things they are reasons for:

Contrastive Reasons: A reason to ϕ out of Q is a consideration that favors ϕ -ing out of Q

If we adopt **Contrastive Reasons**, we have two options. First, we can take the favoring relation, and so the reason relation, as primitive, like Scanlon (1998) and Parfit (2011). Second, we could try to give some analysis of this relation, saying what it takes for some consideration to be a reason for, or to favor, some alternative out of a set of alternatives. In this section, I'll just talk in terms of favoring, without offering an analysis. But in the next chapter I'll show that contrastivism fits nicely with a popular idea about how to analyze this relation, in terms of *promotion*.

3.4.2 A problem

So far, so good. But there is a problem. Which invitation should I choose? The relevant set is clearly {choose Invitation A, choose Invitation B, choose Invitation C}. And intuitively, I should choose Invitation B. But—just as it was hard to say whether anything favored choosing Invitation B *simpliciter*—it's not clear that anything favors choosing Invitation B relative to this set.

It isn't surprising that the same, or at least an analogous, issue should arise here, since I stipulated in this case that I have to choose exactly one of the three invitations. So the set {choose Invitation A, choose Invitation B, choose Invitation C} is, in this case, exhaustive. Contrastivism about favoring allows us to explain why the relevant considerations favor choosing Invitation B rather than Invitation A, on the one hand, and rather than Invitation C, on the other hand. But I haven't yet said anything to get around the problem of saying whether anything favors choosing Invitation B relative to the larger set.¹⁴ To solve this problem, I need to say something about the relationship between reasons and 'ought'.

¹⁴This problem actually depends on a certain understanding of what it takes for a fact to favor some action out of a set of alternatives. In the next chapter, I'll say more about this. For now, it's enough that intuitively, the fact that Invitation B is for burgers/dinner with Burt doesn't favor accepting Invitation B out of the larger set of alternatives, since I could do better with respect to both my food preferences and my dinner company preferences by not accepting it.

3.4.3 Reasons and ‘ought’

The problem is that it’s hard to see how we can say that I ought to accept Invitation B out of {accept Invitation A, accept Invitation B, accept Invitation C}, since it seems that neither the fact that Invitation B is for burgers nor the fact that it’s for dinner with Burt is a reason to accept it out of this set. To solve this problem, I’m going to give a (somewhat) non-obvious account of the relationship between contrastive reasons and ‘ought’.

Perhaps the most natural way to understand the relationship between reasons and ‘ought’, for a contrastivist, is the following (‘CRO’ for ‘Contrastive Reasons and Ought’, ‘*’ because I’m going to eventually reject it):

CRO*: s ought to ϕ out of Q iff s has most reason to ϕ out of Q

Note that this principle appeals to a contrastive ‘ought’, as well as contrastive reasons. I argued in chapter 1 that contrastivism about reasons neither requires nor is required by contrastivism about ‘ought’, but it’s hard to deny that the two do fit together nicely.¹⁵ **CRO*** is very natural for a contrastivist about reasons, but it generates the problem I’m trying to solve. If neither the fact that Invitation B is for burgers nor the fact that it’s for dinner with Burt is a reason to accept it out of the larger set, then it is hard to see how there could be *most* reason to accept Invitation B out of this set.

CRO* appeals to reasons out of arbitrarily large sets of alternatives. But as we’ve seen, it can be hard to say whether some fact is a reason out of larger sets of alternatives, though it’s generally easy to say whether it’s a reason out of two member sets.¹⁶ So I propose to make use of the facts about what’s a reason for what out of these two member sets in determining what an agent ought to do out of larger sets:

CRO: s ought to ϕ out of Q iff s has most reason to ϕ out of $\{\phi, \psi\}$ for all of the other alternatives ψ in Q .

When ϕ “wins” in all of the pairwise comparisons with other members in the set, you ought to ϕ . **CRO** may seem less natural than **CRO***, but it’s not *ad hoc*. I think

¹⁵For some arguments in favor of contrastivism about ‘ought’, see Sloman (1970); Jackson (1985); Cariani (2009, fc); Snedegar (2012).

¹⁶Of course, in many cases it’s easy to say whether some fact is a reason out of a larger set: in most situations, the fact that I’m hungry is a reason to eat lunch out of {eat lunch, throw all my food in the garbage, run a mile, watch a movie} (assuming these are all mutually exclusive).

it's one plausible way for the contrastivist to capture the intuitive idea that you ought to ϕ when ϕ -ing is best *on balance*, even if it's not the best in any particular category.

The dinner invitation case is exactly this sort of case: Invitation B isn't best with respect to my food preferences, nor is it the best with respect to my dinner companion preferences. But it's best *on balance*, and that's why I ought to accept it. There's most reason to accept Invitation B out of {accept Invitation A, accept Invitation B} because, though I slightly prefer Ara's company to Burt's, I greatly prefer burgers to Armenian food.¹⁷ Though the fact that Invitation A is for dinner with Ara is a reason to accept Invitation A out of this set, it's relatively weak. On the other hand, the fact that Invitation B is for burgers is a relatively strong reason to accept Invitation B out of this set. Similar reasoning shows that there's most reason to accept Invitation B out of {accept Invitation B, accept Invitation C}. Thus, according to **CRO**, I ought to accept Invitation B out of the larger set; this is the verdict we want.

CRO will be the principle I adopt going forward about the relationship between reasons and 'ought'. This principle allows the contrastivist to avoid the problem I raised for the shallow contrastivist in the previous section.

3.5 Looking Forward

In this chapter, I introduced a distinction between shallow contrastivism and deep contrastivism. The arguments from reason claims in the last chapter are neutral between these two versions. But I argued that there's good reason to prefer deep contrastivism. This is because the shallow contrastivist has to appeal to a non-contrastive *favoring* relation, which underlies the contrastive reason relation, but the favoring relation itself is contrastive. I began to develop a contrastivist account of favoring, and so of reasons, but ran into trouble with the relationship between 'ought' and reasons. To solve this problem, I adopted a certain view about this relationship.

¹⁷The idea here is that the strength of the reason corresponds in some way to the strength of the preference. I don't mean to endorse this as a perfectly general principle, but it does seem very plausible in cases like this one.

In the next chapter, I'll continue developing contrastivism. In particular, I'll show that the contrastivist can accommodate—and in fact is in a better position to capture—the widespread idea that reasons (and favoring) should be analyzed in terms of promotion. The account I'll develop there isn't a competitor to the account I sketched in this chapter in terms of favoring. Instead, it's one way of further spelling out the view from this chapter by giving an analysis of the favoring relation. I'll also continue to pay attention to just how radical contrastivism is.

Chapter 4

PROMOTION

In the last chapter, I argued that the favoring relation, which people writing about reasons generally identify with the reason relation, is contrastive. This closes off the shallow contrastivist strategy for dealing with the puzzles about reason claims from Chapter 2. I also began to develop contrastivism about reasons. The view, as I've stated it so far, says that some fact is a reason to ϕ out of Q when it favors ϕ -ing out of Q . This is just the contrastivist implementation of the ubiquitous idea that reasons are considerations that count in favor of the things they're reasons for.

I motivated shallow contrastivism by raising some worries about how radical deep contrastivism is. Now that I've argued against shallow contrastivism, these worries are live once again. In the first section of this chapter, I'll make this concern especially sharp by introducing a problem for contrastivism. According to contrastivism, reasons for and against an option can vary with the alternatives—this is the central feature of contrastivism, and the one that lets us solve problems facing non-contrastive theories. But this variation cannot be totally *unconstrained*—knowing what reasons there are relative to one set should tell us *something* about the reasons relative to certain other sets. The problem facing the contrastivist is simply that she seems unable to provide these constraints.

The remainder of the chapter will be concerned with developing a detailed version of contrastivism that solves this problem. This will essentially involve spelling out the view I developed in the last chapter in terms of *promotion*. This taps into a widespread idea about how to analyze the reason relation. I'll argue that, in fact, this idea itself supports contrastivism: the best way to make good on the idea that reasons are tied to promotion is on a contrastive framework. This argument motivates a particular version of contrastivism that straightforwardly entails just

the right kinds of constraints on the variation of reasons between different sets of alternatives.

4.1 A Problem for the Contrastivist

The independence of reasons relative to different sets of alternatives posited by the contrastivist seems like a good thing. But for all I've said so far, and for all other contrastivists have said, learning what reasons there are relative to one set doesn't tell us *anything* about what reasons there are relative to other sets. Reflection on this point shows that the contrastivist may have gone too far, and made reasons relative to different sets *objectionably* independent of one another. Intuitively, there should be constraints on this independence: reasons relative to one set should bear some kind of relation to reasons relative to certain other sets. The problem for the contrastivist is simply that there seems to be nothing about the theory to constrain the independence in the right ways. Further, the contrastivist has to be careful in trying to provide these constraints. Otherwise, she might lose the advantages over non-contrastive theories. In this section I'll make this problem sharp by describing a couple of different kinds of intuitive constraints that we should want, and that the contrastivist seems unable to provide, for all that's been said so far.

4.1.1 Intransitivity

Suppose I want to get in shape. Then the fact that physical exertion helps one get in shape is a reason to jog to Cafe 101 for my daily milkshake, rather than walk there. Similarly, it's a reason to walk there rather than drive. Is it a reason to jog rather than drive? Of course. And any plausible theory of what makes some consideration a reason for one thing rather than another will say so.

But the important thing to notice for my purposes here is that we shouldn't have to think about this at all, once we know it's a reason to jog rather than walk, and a reason to walk rather than drive. It should just *follow* that it's a reason to jog rather than drive. More generally, if r is a reason for A rather than B and a reason for B rather than C , it should simply follow that r is a reason for A rather than C . To deny this is to allow for a very troubling kind of intransitivity.¹

¹This isn't the same kind of intransitivity that receives a lot of attention in the literature, the intransitivity of 'better than' or 'more reason than'. (See especially Temkin (1987, 2012); Rachels

The problem for the contrastivist is that nothing about the theory, for all that's been said, seems to rule out this kind of intransitivity. Since reasons relative to different sets can be independent of one another—as they need to be for contrastivism to gain advantages over non-contrastive theories—we have no reason to think that the reason relation will guarantee that reasons for *A* rather than *B* and for *B* rather than *C* will be reasons for *A* rather than *C*. Ruling out this intransitivity is one important constraint.

4.1.2 Other entailment relations

The fact that Bill has high blood pressure is a reason for him to order fish out of {order fish, order pork, order beef}.² If the waiter comes back to the table and tells Bill that the restaurant is out of pork, that obviously shouldn't change this: the fact that he has high blood pressure is still a reason to order fish out of {order fish, order beef}. And again, any plausible theory about what makes a consideration a reason for an action out of a set of alternatives will say so.

But again, the important point is that we shouldn't have to think about this at all. It should just follow that, in general, when *r* is a reason for *A* out of {*A*, *B*, *C*}, it's also a reason for *A* out of {*A*, *B*}. Even more generally, reasons for an alternative out of one set should also be reasons for that alternative out of arbitrary subsets that include the alternative. If *r* is a reason for *A* when some other option *B* is relevant, making *B* irrelevant should not change this.

Similarly, suppose the fact that Bill has high blood pressure is a reason *against* ordering beef out of {order fish, order pork, order beef}. If the waiter comes back and informs Bill that the restaurant also has a salad, he shouldn't have to reconsider whether the fact that he has high blood pressure is a reason against ordering beef. Adding a new relevant alternative shouldn't change this. In general, reasons against

(1998, 2001); Friedman (2009) for advocates; see, for example, Broome (1991) for defense of the much more common view that intransitivity is impossible.) But if the considerations that are reasons for *A* rather than *B* and for *B* rather than *C* need not be reasons for *A* rather than *C*, then it seems that we might have a case in which there's more reason for *A* than *B* (out of {*A*, *B*}) and more reason for *B* than for *C* (out of {*B*, *C*}), but not more reason for *A* than for *C* (out of {*A*, *C*}). Further, this seems to be an independently troubling kind of intransitivity. I will return to this issue in Chapter 5.

²I'm assuming that these alternatives are mutually exclusive—Bill can order at most one of the dishes.

an alternative out of one set are intuitively reasons against that alternative out of *supersets* of that set.

These, and other, intuitive entailment relations should hold between reasons relative to different but related sets.³ In general, as we shift between contexts in which different sets of alternatives are relevant, we shouldn't have to start all over in figuring out what reasons there are. But for all that has been said, contrastivism cannot deliver these entailment relations, since reasons relative to different sets can be independent of one another—and they need to be for contrastivism to have an advantage over non-contrastive theories.

This is the problem facing the contrastivist that I'll attempt to solve. As a preview, my strategy will be to appeal to the popular idea that reasons involve the promotion of something—a desire or an objective value, for example. Once we see how to give the most plausible contrastivist analyses of reasons for and against in terms of promotion, I'll show that the right entailment relations straightforwardly follow. First, though, I'll argue that the idea that at least many reasons involve promotion itself provides independent support for contrastivism.

4.2 Promotion

Several writers with importantly different views about the nature of normative reasons accept something like the following schema:

Promote: There is a reason for s to ϕ iff s 's ϕ -ing would promote some X of the relevant kind.

Some writers, like Moore (1903, 1912); Wedgwood (2009); Parfit (2011) take the X s to be objective values, like happiness or justice.⁴ Others, like Schroeder (2007), take the relevant kind of X to be desires of the agent. And others, like Finlay (2006, ms), take the relevant kind of X to be contextually-specified ends. Nagel (1970) also analyzes reasons in terms of promotion, though his view is importantly different.

³I'll say more about entailments between sets that differ in resolution in section 5.

⁴Moore didn't talk about reasons explicitly, but rather rightness. But his view of rightness would extend straightforwardly to this kind of view of reasons.

Still, there is a plausible view in the neighborhood of Nagel's which explicitly fits **Promote**.⁵

We might follow writers like Scanlon (1998) and Anderson (1993), and draw a distinction between promoting values and respecting or honoring values. If we make this distinction, we might worry that some reasons will not be explained by, or even involve, the promotion of anything. For example, according to this view, the value of friendship primarily gives us reasons to respect or honor friendships, say by treating our friends well, rather than to promote friendship, say by matching up people we think would make good friends.⁶ The argument for contrastivism that I will give shortly is that the idea that reasons involve promotion is best captured on a contrastive theory. If not all reasons involve promotion, then this argument at best would seem to show that some reasons—those that involve promotion—are contrastive. But contrastivism, as I want to defend it, is the thesis that all reasons are contrastive.

I think we can allow that some reasons do not involve promotion, but rather involve respecting or honoring values, without limiting the scope of contrastivism in this way. Everyone should recognize that at least some reasons involve promotion.⁷ And if I am right that these reasons are contrastive, that is some pressure to adopt contrastivism about all reasons. It is theoretically unattractive to posit two separate reason relations, one of which is contrastive and one of which is non-contrastive. If there is a plausible unified view of reasons on offer, that would be preferable.

I will develop my contrastive theory in terms of promotion, since (i) everyone should agree that some reasons involve promotion, and (ii) doing so will simplify exposition of the view. At the end of the chapter, however, I will return to this point and show how to extend the view to cover non-promotional reasons.

⁵The complication with Nagel's view is that he thinks what we have reasons to do is to promote valuable states of affairs, whereas the other theories hold that the fact that performing certain actions would promote certain objectives is what explains why we have reasons to do those actions.

⁶See Pettit (1991) and Pettit's contribution to Baron et al. (1997) for a defense of the view that all reasons involve promotion.

⁷For example, instrumental reasons to take the means to our ends seem to be clearly tied to the promotion of those ends. See Kolodny (ms); Bedke (ms) for relevant discussion.

4.2.1 Promotion as probability-raising

The idea that reasons involve promotion is an attractive and widespread one. But what does it take for an action to promote some objective? A natural initial idea is that an agent's ϕ -ing promotes X when her ϕ -ing raises the probability of X . Of course, this is not yet very precise. Promoting a desire is making it more probable that the desire is satisfied—where the object of the desire is a proposition p , promoting the desire is raising the probability of p . It's harder to see what it would be to promote a value. I'll follow Wedgwood (2009) and say that to promote a value is to promote a state of affairs in which the value is instantiated.⁸ On the probability-raising view, to promote a value is to make it more probable that the value is instantiated.

Though this probability-raising view is a natural one, it is also controversial. One important problem concerns the promotion of values like justice. It is true that we have reasons to do things that would promote justice, in some sense. What is more problematic, though, is the claim that promoting justice involves making it more probable that justice is instantiated, because it seems that justice is *already* instantiated, at least to some degree. If that's right, then the probability that justice is instantiated, no matter what we do, is 1. So we can't make it *more* probable.⁹ I think there are things to be said in defense of the probability-raising view.¹⁰ But for now I want to put these worries aside, and work with this view. I do this in order to have a concrete version of the idea that reasons involve promotion to work with. Moreover, existing views about the relationship between reasons and objectives like desires, goals, or values are often framed in terms of probability-raising.¹¹ I show in

⁸Wedgwood does reject the probability-raising view of promotion, however.

⁹Thanks for Walter Sinnott-Armstrong and Jacob Ross for discussion of this point.

¹⁰For example, we may try specifying more fine-grained values, like 'I increase the amount of justice in the world' or 'I do the most just thing I can'. A further reason to be careful in specifying the objectives is the following kind of case, raised by Fabrizio Cariani in conversation. Suppose there are two objectives, X and Y , but that promoting X without promoting Y is worthless. Since promoting X when we also promote Y is valuable, it seems that X can provide reasons to do the things that would promote it. But it does *not* seem that it should provide reasons to do the things that would promote it in cases in which we will not also promote Y . In this kind of case, I want to say that the reason-providing objective is something like ' X when you also promote Y ', rather than just ' X '.

¹¹See, for example, Bedke (ms); Finlay (2001, 2006, ms); Kolodny (ms); Schroeder (2007). Bedke's and Kolodny's discussions concern instrumental reasons; Finlay's and Schroeder's concern all reasons.

section 4.5, though, that this probability-raising view of promotion is not necessary for my purposes.

So, according to this view, you have reasons to ϕ when your ϕ -ing would make some relevant objective more probable. But this leaves an important question: more probable than what? I'll consider two attempts to answer this question, and argue that both of them are problematic. Then I'll sketch a contrastivist version of the view, which I'll develop in the next section, that avoids the problems.

4.2.2 Doing nothing

The first view I'll consider is offered by Schroeder (2007), p. 113:

Doing Nothing: s 's ϕ -ing promotes X iff $P(X|s \phi s) > P(X|s \text{ does nothing})$

An important question about this view is the following: what is it for me to do nothing? Evers (2009), in a paper objecting to Schroeder's account, considers two ways to answer this question, but shows that both suffer from the same basic problem. On the first reading—which I think is pretty implausible even as an interpretation of Schroeder—doing nothing is simply sitting absolutely still. The second more plausible reading is that doing nothing is just doing what you're already doing. For example, if I'm walking to the store for milk, and realize I forgot my wallet, I have a reason to turn around and get it. That's because, on this account, the probability that I get milk is higher given that I turn around and get my wallet than the probability that I get milk given that I keep walking to the store without my wallet.

The problem Evers points out for both of these readings, and for **Doing Nothing** in general, is that I can have reasons to do nothing, however we understand it, but this account rules that out. Suppose doing nothing is sitting absolutely still. Well, suppose that I'm out in the woods and a predator has spotted me. Fortunately, the predator can only see me if I move. So I have a reason to sit absolutely still. Next, suppose that doing nothing is doing what I'm already doing. And suppose that I'm waiting at a restaurant for an important client who's just informed me that he'll be there in five minutes. Well, then I have a reason to keep doing what I'm doing—that is, to keep waiting. So I can have reasons to do nothing, no matter how we understand it. But for me to have a reason to do nothing on this account, the following must hold: $P(X|I \text{ do nothing}) > P(X|I \text{ do nothing})$. But

that's impossible. Since I can have reasons to do nothing, **Doing Nothing** is not the right way to understand promotion, in the sense relevant for **Promote**.

4.2.3 Not ϕ -ing

The second view I'll consider is offered by Finlay (2006, ms):

Not ϕ : s 's ϕ -ing promotes X iff $P(X|s \phi s) > P(X|s \text{ does not } \phi)$

This account avoids the problem for **Doing Nothing**: for you to have a reason to do nothing on this account, doing nothing just has to raise the probability of X more than not doing nothing would. But to evaluate this account, we need to know what not ϕ -ing is. I'll consider three different versions of this view and argue that all three are problematic.

4.2.3.1 Possibilism

The first proposal, call it *Possibilism*, is that ϕ -ing promotes X whenever ϕ -ing raises the probability of X more than anything else the agent could do that is incompatible with ϕ -ing. The problem with this suggestion is that it makes it too hard to have reasons to do things. For X to give you a reason to ϕ , ϕ -ing has to make X more likely than anything else you could do. Suppose I want to eat Thai food for dinner, and that there are two Thai places around, RCA and Thai Patio. Suppose that there is a very slight chance that RCA will be too crowded so that I'll have to go home without having Thai food if I go there. So I have a slightly better chance to get Thai food if I go to Thai Patio than if I go to RCA. Thus, according to this proposal, my desire for Thai food gives me no reason at all to go to RCA.

But this seems like a problem. After all, my desire for Thai food certainly gives me a reason to go to RCA *rather than* go to Zankou Chicken, which doesn't even serve Thai food. But how could this desire give me a reason to go to RCA rather than Zankou Chicken if it doesn't give me any reason *at all* to go to RCA? So this proposal makes it too hard to have reasons to do things.

4.2.3.2 Actualism

A second version of **Not** ϕ , call it *Actualism*, holds that not ϕ -ing is doing whatever it is you would actually do if you didn't ϕ .¹² This view would seem to make it easier to have reasons, since your ϕ -ing doesn't have to raise the probability of X more than anything else you could do, but only more than whatever it is you'll actually do, if you don't ϕ .

But now consider this case from Behrends and DiPaolo (2011). Debbie desires that p . She is given a choice to push one of three buttons, A, B, or C. If she pushes either A or B, p is guaranteed to be true. But if she pushes C, p will not be true. Debbie in fact pushes A, and if she hadn't pushed A, she would have pushed B. It turns out, according to the view under consideration, that she had no reason to push button A, since $P(p|\text{Debbie pushes A}) = P(p|\text{Debbie pushes B})$.¹³ But this is counterintuitive. After all, she certainly had a reason to push A rather than push C. So this proposal also seems to make it too hard to have reasons to do things.

4.2.3.3 Probabilism

Finally, consider the following *Probabilist* view. Rather than thinking of not ϕ -ing as doing *anything* incompatible with ϕ -ing, or with doing whatever it is you would actually do, if you didn't ϕ , we can just appeal to a probability distribution over the things you might do. So for example, if you don't go to RCA, there might be a 60% chance you go to Thai Patio, and a 10% chance you go to Zankou Chicken. This view will give us some value for $P(\text{you don't } \phi)$, and let us calculate one for $P(X|\text{you don't } \phi)$ by weighting $P(X|\text{you } \psi)$ by the probability that you ψ , for all different ways ψ of not ϕ -ing. What we end up with is something like the *expected value*, as far as X is concerned, of not ϕ -ing. Then to determine if you have a reason to ϕ , we can straightforwardly check to see if $P(X|\text{you } \phi) > P(X|\text{you don't } \phi)$.

This theory also makes it too hard to have reasons to do things. Suppose again that I want some Thai food, that the only two Thai places around are Thai Patio and RCA, and that the only other dinner option is Zankou Chicken, which doesn't

¹²The distinction between this version and the last version, and the corresponding names, are familiar from the debate between actualists and possibilists about obligation. See Jackson and Pargetter (1986).

¹³Behrends and DiPaolo point out a more basic problem. It turns out, according to this view, that Debbie didn't even promote her desire by pressing A.

serve Thai food. But now suppose that RCA is very likely to be so crowded that I'll have to go home without having Thai food if I go there. In particular, suppose that $P(\text{I get Thai food} \mid \text{I go to RCA})=.2$, while $P(\text{I get Thai food} \mid \text{I go to Thai Patio})=.9$, and of course $P(\text{I get Thai food} \mid \text{I go to Zankou})=0$. And finally, suppose $P(\text{I go to RCA})=.3$, $P(\text{I go to Thai Patio})=.6$, and $P(\text{I go to Zankou Chicken})=.1$. With some calculation, we end up with the result that $P(\text{I get Thai food} \mid \text{I don't go to RCA})$ is roughly .77 (where not going to RCA is just either going to Thai Patio or going to Zankou Chicken), which is greater than .2. So according to this view, my desire for Thai food doesn't give me any reason at all to go to RCA. But as before, surely this desire gives me a reason to go to RCA *rather than* Zankou Chicken. So it must give me some reason to go to RCA.¹⁴

It's also worth pointing out, I think, that if we're going to assign probabilities to the various *ways* of not ϕ -ing, we need some way of distinguishing one way from another. That is, we need some way of dividing up the *alternatives*. Different ways of doing this will lead to different answers about what you have reasons to do. And, as I pointed out in section 1, this kind of *resolution-sensitivity* is characteristically contrastivist. So spelling out this theory may require contrastivism after all.

4.2.4 Contrastive promotion

I think we can give a simpler contrastive view that avoids the problems I've discussed so far. This view does not rule out reasons to do nothing. And it does not need to come up with some way to understand not ϕ -ing. Instead, the idea is that X can give you a reason to ϕ , relative to a set of alternatives Q , when your ϕ -ing raises the probability of X more than all of the other alternatives in Q .¹⁵ Thus,

¹⁴In defending this sort of view, Finlay (2006, ms), makes the following assumption, which he calls the 'symmetry of choice': in making judgments about what reasons you have (or what's good, or what you ought to do, on his theory), we ignore your dispositions to take one option over another. So in figuring out which reasons I have, we would just assume that I am equally likely to go to each of the three restaurants. But we can still construct a case on which, according to this view, my desire for Thai food gives me no reason at all to go to RCA. Let $P(\text{I go to } x) = 1/3$, for all three restaurants, x , so that the symmetry of choice is satisfied. And, as above, let $P(\text{I get Thai food} \mid \text{I go to RCA})=.2$, $P(\text{I get Thai food} \mid \text{I go to Thai Patio})=.9$, and $P(\text{I get Thai food} \mid \text{I go to Zankou})=0$. Then by some simple calculations, $P(\text{I get Thai food} \mid \text{I don't go to RCA})$ is roughly .38, which is greater than $P(\text{I get Thai food} \mid \text{I go to RCA})=.2$. Thus, according to this modified Probabilist view, my desire for Thai food gives me no reason at all to go to RCA. But surely it gives me a reason to go to RCA rather than go to Zankou Chicken, and it's hard to see how it could do this unless it gave me some reason to go to RCA.

¹⁵I'll defend this claim in the next section.

in the Thai food case, my desire for Thai food gives me a reason to go to RCA rather than Zankou Chicken, even though it doesn't give me a reason to go to RCA rather than Thai Patio, since going to RCA raises the probability that I get Thai food more than going to Zankou, but less than going to Thai Patio. Similar remarks hold in the button-pressing case from Behrends and DiPaolo (2011). The key to these solutions to the problems facing non-contrastive theories is that, for the contrastivist, the sets of alternatives that reasons are relativized to need not be *exhaustive* of all the options open to the agent.

4.2.5 Motivating resolution-sensitivity

The arguments against the non-contrastive accounts of reasons in terms of promotion illustrate the advantages of non-exhaustivity. But when I introduced contrastivism in Chapter 1, I pointed out a second feature of the view, namely resolution-sensitivity. The arguments I've given so far don't motivate this feature of contrastivism.¹⁶ But I think we can motivate it. Once we relativize reasons to sets of alternatives, there are cases in which whether or not something is a reason for an action seems to depend, not just on what possibilities are covered by the alternatives in the set, but also how we split up those possibilities.

Consider the following version of the famous Professor Procrastinate case from Jackson (1985); Jackson and Pargetter (1986). Professor Procrastinate has been asked to write a review an important new book, because she is the most qualified person to write it. But she is a terrible procrastinator. If she accepts the invitation, she is very unlikely to actually write the review. And if that happens, the author's career and the field at large will suffer. If Procrastinate just declines the invitation, someone else—less qualified, but more reliable—will be asked to write, and will do so. In this case, the fact that Procrastinate is a procrastinator is a reason for her to decline the invitation rather than accept. That is, it's a reason to decline out of {decline, accept}, plausibly provided by the objective of doing what's best for the profession. But it is intuitively not a reason to decline out of the more fine-grained set, {decline, accept and write, accept and don't write}. So whether or not this fact is a reason to decline seems to depend on the resolution at which we divide up the

¹⁶I did mention that spelling out **Probabilism** may require resolution-sensitivity, but since I rejected that view—for reasons having to do with non-exhaustivity—I can't appeal to that argument in motivating resolution-sensitivity.

alternatives. Thus, once we relativize reasons to alternatives, there's good reason to allow the sets to vary in resolution, rather than, say, making them all maximally fine-grained.

4.3 Contrastive Reasons and Promotion

I've argued that promotion, in the sense relevant for **Promote**, is best understood contrastively: your ϕ -ing promotes X only relative to some set of alternatives. It does so when it raises the probability of X more than all the other alternatives in the set. In this section I'll begin to develop a contrastivist account of reasons based on this account of promotion. In developing the view, I will make the following two simplifying assumptions, both of which I have called attention to above: (i) all reasons involve promotion; and (ii) the relevant kind of promotion is probability-raising. Making these assumptions will let us see very clearly how the sort of view I want to defend can provide the constraints on how reasons can vary between sets of alternatives. But at the end of the chapter, I will show that we can actually relax both of them without sacrificing any of the attractions of the view.

Promote tells us what has to hold for you to have a reason to ϕ , or for there to be a reason for you to ϕ . But it would be nice to have an account of when some fact is a reason for you to ϕ . Philosophers who appeal to promotion to analyze the reason relation frequently appeal to *explanation* here. If you have a reason to ϕ because ϕ -ing promotes some objective, then some fact is a reason for you to ϕ when it *explains why* (or is part of the explanation for why) ϕ -ing promotes the objective.¹⁷ I have a reason to go to the store because doing so would promote the objective of my having breakfast in the morning. The fact that I'm out of milk is part of the explanation for why going to the store would promote this objective. So according to this idea, the fact that I'm out of milk is a reason to go to the store. I'll adopt this strategy here.

So what we want is a contrastivist analysis of reasons for and reasons against in terms of promotion. So far I've said that r is a reason for you to ϕ when it

¹⁷For analyses of reasons partly in terms of explanation, see Toulmin (1950); Broome (2004); Searle (2001); Finlay (2001, 2006, ms); Schroeder (2007). Like these other writers (with the possible exception of Finlay), I take the relevant sort of explanation here to be a non-epistemic notion. For r to explain why your ϕ -ing would promote X is for r to be part of what makes it true that your ϕ -ing would promote X .

explains why your ϕ -ing would raise the probability of some relevant objective X more than other relevant alternatives. Conversely, it's natural to think that r is a reason against ϕ -ing when it explains why your ϕ -ing would raise the probability of X *less* than other relevant alternatives. So what we need to know is, *how many* of the other alternatives does your ϕ -ing have to raise the probability of X more/less than, for X to give you a reason for/against ϕ -ing?

First a quick note about reasons against. It is often thought that a reason against some action ϕ is just a reason for not- ϕ .¹⁸ On this view, we can just have one kind of reason, reasons for. But the contrastivist cannot accept this view because, given non-exhaustivity, the sets of alternatives reasons are relativized to need not be closed under negation: just because ϕ is in a set, it does not follow that not- ϕ is also in the set. So if we want to talk about reasons against an alternative ϕ relative to a set Q , we can't just talk about the reasons for not- ϕ out of Q , since not- ϕ might not even be in Q . Thus, the contrastivist needs to define reasons against—reasons not to ϕ , instead of reasons to not- ϕ —separately.

So start (for expository purposes) with reasons against. I suggest the following analysis:

Against: r is a reason for you not to ϕ out of Q iff there's some X of the relevant kind¹⁹ and r explains why²⁰ $P(X|\text{you } \phi) < P(X|\text{you } \psi)$, for some $\psi \neq \phi$ in Q .

If you could do better, as far as X is concerned, by doing something other than ϕ -ing, then X very plausibly gives you a reason against ϕ -ing. The most natural alternative to **Against** would say that your ϕ -ing has to raise the probability of X less than *all* the other alternatives. But this misses out on some intuitive reasons against. For example, suppose I can travel the 15 miles to campus by jogging, biking, or driving. The fact that I want to get there in under an hour intuitively gives me a reason against biking, even though I have a better chance of getting

¹⁸See Nagel (1970), footnote 1 on p. 47, for example.

¹⁹Again, by 'some X of the relevant kind', I mean the following: there's some set of objectives that are the things to be promoted, according to our theory of reasons. A desire-based theory will say that this set contains desires of the agent; a value-based theory will say that it contains some independent values. X is an objective of the relevant kind when it is in this set. By existentially quantifying here, we allow that there are multiple objectives that can give you reasons in any individual case.

²⁰Or is part of the explanation of why—I'll ignore this qualification from now on.

there in under an hour if I bike than if I jog—it’s still quite unlikely that I get there in under an hour if I bike.

Once we accept **Against** as an analysis of reasons against, I think we should accept the following analysis of reasons for:

For: r is a reason for you to ϕ out of Q iff there’s some X of the relevant kind and r explains why $P(X|\text{you } \phi) > P(X|\text{you } \psi)$, for all $\psi \neq \phi$ in Q .²¹

This makes reasons for harder to come by than reasons against. But I think this is a very plausible analysis of reasons for. If you could do better, as far as X is concerned, by doing something other than ϕ -ing, it would be strange if X gave you a reason to ϕ out of a set of alternatives that includes that other thing. The most natural alternative would say that X can give you a reason to ϕ relative to some set when your ϕ -ing would raise the probability of X more than *some* other alternative in the set. But if we accept this in addition to **Against**, some consideration r could be both a reason for and a reason against all the alternatives in the middle—that do not do the best and that do not do the worst, as far as X is concerned. For example, in the case I introduced above, the fact that campus is fifteen miles away would be both a reason to bike and a reason not to bike, both provided by my desire to get to campus in under an hour. But this is a violation of the very plausible principle, **Restricted Exclusivity**, that I argued for in Chapter 2.²²

One potential objection to this view is that it gives too liberal a treatment of reasons against. A reason against some alternative, you might think, has to “imply a significant criticism” of that alternative, in the sense of Greenspan (2005, 2007), and merely being a reason for an incompatible alternative isn’t enough. But other philosophers seem to agree with me that a reason for one alternative is in fact a reason against other incompatible alternatives, for some relatively weak notion of incompatibility, like physical or psychological incompatibility.²³ I think this is very plausible: if some consideration r is a reason for A —if it highlights some benefit

²¹Interestingly, though I reject Finlay (ms)’s non-contrastive account of reasons, this contrastive account of reasons for closely mirrors his contrastive analysis of ‘ought’. According to that analysis, roughly, you ought to ϕ just in case your ϕ -ing makes the contextually salient end more likely than any other relevant alternative. By allowing multiple objectives—the analog of Finlay’s ends—to provide reasons in a single context, my analysis is better suited to the *pro tanto* notion of a reason.

²²This argument also tells against a revised version of **Possibilism**, according to which an objective provides a reason for you to ϕ when ϕ -ing raises the probability more than *something*—instead of *everything*—else you could do.

²³Compare the discussion of ‘General Substitutability’ in Sinnott-Armstrong (1992), section 1.

for doing A —and if doing B precludes doing A , then it is very plausible that r is a reason against B , since doing B would keep you from getting that benefit.²⁴

Note, though, that we can draw an interesting distinction between two types of reasons against on this view. A reason against some alternative that does imply a significant criticism of that alternative—like the fact that Zankou Chicken is really crowded at this time of day—often won’t turn out to be a reason for any other particular alternative. That’s because these sorts of facts are often irrelevant for differences between other alternatives in the set. For example, the fact that Zankou Chicken is really crowded doesn’t really tell us anything about the relationship between going to Thai Patio and going to Spicy Thai BBQ. So it doesn’t explain why either $P(\text{I don’t have to wait}|\text{I go to Thai Patio}) > P(\text{I don’t have to wait}|\text{I go to Spicy Thai BBQ})$, or vice versa. So by **For**, it won’t turn out to be a reason for either of these, though it’s still a reason against going to Zankou Chicken. Reasons against an alternative that do not imply a significant criticism, on the other hand, will often be reasons for other particular alternatives. So, even on this view, we can capture Greenspan’s insight: some considerations merely count against certain alternatives without counting in favor of others.

4.4 Providing the Constraints

In this section, I’ll show that the theory I developed in the last section, independently motivated by considerations about the relationship between reasons and promotion, straightforwardly provides the constraints on the independence of reasons relative to different sets of alternatives I motivated in section 1, as well as some other plausible constraints.

4.4.1 Transitivity

The first constraint we need on the independence of reasons relative to different sets is one that rules out intransitivity. That is, we want it to turn out that whenever

²⁴Sinnott-Armstrong (2008) says that a reason “favors one thing and disfavors others” (p. 258). Ruben (2009) says that when doing A is incompatible with doing B , “a reason to do an act of type B must also be a reason not to do an action of type A ” (p. 63). And Broome (ms) cites nice features of one alternative—which are reasons for taking that alternative—as reasons against incompatible alternatives. See the Montreux-Marrakesh example on p. 63.

r is a reason for A rather than B and for B rather than C , it is also a reason for A rather than C . In fact, we need an importantly but subtly different constraint, once we adopt a view on which we explain reasons in terms of the promotion of various kinds of objectives or ends. We don't just want reasons for A rather than B and for B rather than C to be reasons for A rather than C ; we want them to be provided by the same objective. A case in which my desire for Thai food gives me a reason r for A rather than B and for B rather than C , but not for A rather than C is counter to the transitivity intuition, even if, say, my desire to spend lots of money explains why r is a reason for A rather than C .

So what we need is the following kind of constraint:

Transitivity: If r is a reason for A rather than B , provided by X , and a reason for B rather than C , provided by X , then r is a reason for A rather than C , provided by X .

Fortunately, it's easy to see that this follows straightforwardly from the framework I developed in the last section. If r is a reason for A rather than B , and a reason for B rather than C , both provided by X , then we know that r explains why $P(X|A) > P(X|B)$ and why $P(X|B) > P(X|C)$. Since these facts about conditional probability don't depend on the comparisons we're making, and since the greater than relation is transitive, it follows that $P(X|A) > P(X|C)$. And given plausible assumptions about how explanation works in these cases—namely, that if r explains why $P(X|A) > P(X|B)$ and r also explains why $P(X|B) > P(X|C)$, then it will explain why $P(X|A) > P(X|C)$ (at least in cases in which this is true)— r explains this.²⁵ So, by **For**, r is a reason for A rather than C . Thus, by appealing to the idea that reasons are to be explained in terms of promotion, understood as probability-raising, we secure **Transitivity**.

4.4.2 Non-exhaustivity

The second constraint concerns entailment relations between reasons relative to certain kinds of related sets of alternatives. First I'll discuss entailments between

²⁵Remember that I'm using a non-epistemic notion of explanation here. Plausibly, the epistemic notion will not have this feature, but the sense I intend, on which x explains y when x is part of what makes it true that y , does.

reasons relative to sets that differ in the alternatives they include—that is, sets that differ along the dimension of *non-exhaustivity*.

First, consider reasons for an alternative relative to one set, and reasons for that alternative relative to a subset of that initial set. I said above that intuitively, facts which are reasons for A relative to Q should also be reasons for A relative to subsets of Q that contain A . Imagine that this were not true. Then there might be a case in which r is a reason for A rather than either of B or C , but not a reason for A rather than B . It's implausible that there could be such a case, so we should rule it out.

Fortunately, the theory I developed in the last section—in particular **For**—delivers precisely this result. If r is a reason for A out of Q , explained by end or objective X , then by **For**, $P(X|A) > P(X|B)$, for all the other alternatives B in Q . Since the value of $P(X|A)$ does not vary with the alternatives, $P(X|A)$ will obviously be higher than $P(X|B)$ for all the alternatives B in any subset of Q . And since r explains why $P(X|A) > P(X|B)$ for all of the other alternatives B in Q , r will also explain this for all of the other alternatives in a subset of Q . So r will be a reason for A out of any subset of Q that contains A .

Next, consider reasons against. I said above that reasons against an alternative out of one set should also be reasons against that alternative out of supersets of the initial set. If this entailment did not hold, then we might have a case like the following. r is a reason against A when the alternative is B alone, but not a reason against A when we consider C in addition to B . Again, this is implausible.

Fortunately, the theory delivers just the result we want. Suppose r is a reason against A out of Q , explained by X . Then, by **Against**, $P(X|A)$ must be lower than $P(X|B)$ for some other alternative B in the set. Merely adding some new alternatives to Q won't change this fact. Thus, r will also be a reason against A relative to the new superset of Q . So the independently motivated version of contrastivism that analyzes reasons in terms of promotion delivers precisely the constraints we need, on the independence of reasons relative to sets and their subsets and supersets.

4.4.3 Resolution-sensitivity

I also introduced a second feature of contrastivism, resolution-sensitivity. Two sets of alternatives may cover the same possibilities, but do so at different levels of detail. For example, we might have reasons to A relative to the more fine-grained set $\{A, B, C\}$ or relative to the more coarse-grained set $\{A, B \vee C\}$. Do we want a theory to deliver any relationships between reasons relative to sets related in this way?

I think there are some intuitive constraints here. Suppose the fact that I don't have any vacation days left is a reason to go to work out of $\{\text{go to work, stay home and clean, stay home and watch TV, stay home and do something else}\}$, explained by my desire not to get fired. It also seems to be a reason to go to work out of the more coarse-grained set $\{\text{go to work, stay home}\}$. The intuitive thought here is that if doing A is better, with respect to the objective X , than all of the (relevant) ways of doing B , then doing A is obviously better than doing B , with respect to X . And the theory delivers this result. It's a fact about conditional probability that $P(X|B \vee C)$ cannot be higher than both $P(X|B)$ and $P(X|C)$, and this secures the result: if $P(X|A)$ is higher than $P(X|B)$ for all the different fine-grained ways B of doing some more coarse-grained option D , then $P(X|A) > P(X|D)$. That is, reasons for A relative to more fine-grained sets entail reasons for A relative to more coarse-grained sets that just group together other alternatives.²⁶

Now consider reasons against. The fact that I'm out of vacation days is intuitively a reason against staying home out of the coarse-grained set $\{\text{go to work, stay home}\}$. And it also seems to be a reason against this alternative out of the more fine-grained set $\{\text{go to work by car, go to work by bus, go to work some other way, stay home}\}$. Again, the theory delivers the correct result here. Since $P(X|B \vee C)$ can't be higher than both $P(X|B)$ and $P(X|C)$ (it also cannot be *lower* than both) we know that, in general, reasons *against* an alternative A out of a coarse-grained set like $\{A, B \vee C\}$ entail reasons against that alternative out of more fine-grained

²⁶We might also wonder about the relationship between reasons for A relative to one set and reasons for more coarse-grained alternatives that subsume A (say, $A \vee B$) in more coarse-grained sets of alternatives. Similarly, we might wonder about the relationship between reasons for a more coarse-grained alternative (say, $A \vee B$) and reasons for the fine-grained alternatives that it subsumes (say, A) relative to more fine-grained sets. It turns out that the theory doesn't deliver any constraints in these cases. I think these are the intuitively correct results, though I won't argue for that here.

sets like $\{A, B, C\}$, according to the theory. If $P(X|A)$ is lower than $P(X|D)$, for some coarse-grained alternative D , then by this fact about conditional probability, $P(X|A)$ must be lower than *at least one* of the more fine-grained ways of doing D . And so by **Against**, X will provide a reason against A relative to the more fine-grained set.

4.4.4 Other entailments: unions and intersections

Now I will discuss some other entailment relations delivered by the theory that involve unions and intersections of sets of alternatives. It is independently interesting that the theory delivers entailments in these cases. But it also may have some interesting applications. For example, when two agents who have been deliberating separately, using two different sets of alternatives, come together to deliberate jointly, it's plausible that the relevant alternatives in this new context of joint deliberation will bear some relationship to the two separate sets the agents were using on their own. Either taking the union or the intersection of these sets would be reasonable.²⁷ Moreover, this illustrates that the theory delivers not only the most obvious constraints (e.g., the transitivity constraint, and the subset and superset constraints) but also less obvious constraints that seem intuitively correct, once the theory brings them into view. This provides additional confirmation for the theory.

First, note that the following two entailments follow as corollaries of the previous ones involving subsets and supersets. (i) Reasons for A out of both Q and Q' , both provided by X , will be reasons for A out of the intersection, $Q \cap Q'$, provided by X . This just follows from the result that reasons for A out of one set are reasons for A out of subsets of that set.²⁸ (ii) Reasons against A out of both Q and Q' will be reasons against A out of the union, $Q \cup Q'$. This just follows from the result that reasons against A out of one set are reasons against A out of supersets, as well.

We also get the following new result, which is not simply a corollary of a previous one. Suppose r is a reason to A out of Q , provided by X , and that r is also a reason to A out of Q' , provided by X . What about the *union* of these two sets, $Q \cup Q'$?

²⁷This will often be more complicated, of course, since the resolutions of the two separate sets may cross-cut one another.

²⁸The limiting case here is when $Q \cap Q' = \{A\}$. Technically, by **For**, everything will be a reason for A relative to this singleton set. But singleton sets will never be the relevant set—if there's only one relevant alternative, then there's no need for deliberation, and there won't be reasons for or against the single alternative.

By **For**, we can immediately see that r will be a reason for A relative to the union, as well. And this is plausible: if nothing in either set was better than A , as far as X is concerned, then putting them all together in the same set should not make a difference. Interestingly, we don't get the corresponding result for reasons against and intersections: it's not true in general that if r is a reason against A out of Q and out of Q' , then it's a reason against A out of $Q \cap Q'$, since we may lose all of the alternatives in either set that are ranked higher than A , when we intersect.

4.5 Promotion as Probability-Raising

Earlier I promised to reconsider my assumption of the probability-raising view of promotion in light of apparent problems with that view. I set those problems aside, since (i) existing discussions of these issues are often framed in terms of probability-raising, and (ii) the probability-raising view provided a concrete version of the idea that reasons involve promotion to work with. But now I want to show that this particular view of promotion is not crucial for my purposes here.

I have shown that the idea that reasons involve the promotion of objectives generates a version of contrastivism that provides important intuitive constraints on the independence of reasons relative to different sets of alternatives. The probability-raising view of promotion has two features that I used in proving that these constraints hold. First, it provides a *ranking* of options for each objective, in terms of how probable the actions make the objective. I used this ranking to generate the constraints involving non-exhaustivity: the transitivity constraint, as well as the constraints involving subsets and supersets. But any way of understanding promotion that provides a ranking of options for each objective in terms of how well the options promote the objective will generate these constraints, as long as the ranking itself is transitive.

The second important feature of the probability-raising view, which provides the resolution-sensitivity constraints, is that the rankings it provides have the following feature:

Disjunction Boundedness: For all actions A and B and for all objectives X , $A \vee B$ is ranked somewhere between (inclusive) A and B , in terms of how well it promotes X .

The fact about conditional probabilities, that when $P(X|A) \geq P(X|B)$, $P(X|A) \geq P(X|A \vee B) \geq P(X|B)$, secures this property for the probability-raising view. But I think that any plausible ranking of options in terms of how well they promote a given objective will have the **Disjunction Boundedness** property. Doing $A \vee B$ can't do worse than the worst of A and B , and can't do better than the best of A and B , as far as the relevant objective goes.

My arguments against **Doing Nothing**, **Actualism**, **Possibilism**, and **Probabilism** motivating a contrastive account of promotion were framed in terms of the probability-raising view. But analogous arguments can be run against more general versions of the first three views, stated just in terms of how well actions promote various objectives. For example, in Behrends and DiPaolo (2011)'s button-pressing case, Debbie surely has a reason to press button A rather than press button C. But according to **Actualism** and **Possibilism**, Debbie has no reason at all to press button A, since pressing button A does not better promote her desire than pressing button B. It is hard to see how to give a non-probabilistic version of **Probabilism**. But if we are worried about the probability-raising view of promotion, the difficulty with formulating a non-probabilistic version of **Probabilism** should make us worry about this view, as well.

We can also see now that the probabilistic principles, **For** and **Against**, were special cases, in the probabilistic framework, of the following more general principles:

For*: r is a reason for you to ϕ out of Q iff there is some X of the relevant kind such that r explains why your ϕ -ing would better promote X than any of the other alternatives in Q .

Against*: r is a reason for you not to ϕ out of Q iff there is some X of the relevant kind such that r explains why your ϕ -ing would not promote X as well as some other alternative in Q .

So the probability-raising view of promotion is not crucial for the theory I've developed here. Any view of promotion which provides rankings of options in terms of how well they promote the objectives and which has the **Disjunction Boundedness** property will work.

4.6 Non-Promotional Reasons

I've just argued that we need not understand the relevant kind of promotion as probability-raising: all we need is a ranking (with certain minimal properties) of alternatives in terms of how well they promote (in some sense) a given objective. In this section, I will relax the assumption that all reasons involve promotion.

In section 2, I argued that even if we think that some reasons do not involve promotion, but rather involve respecting or honoring values, my argument that the reasons that do involve promotion are contrastive puts rational pressure on us to adopt contrastivism about all reasons. But there is still a serious gap. I've spent the last few sections showing how a contrastive theory of reasons cashed out in terms of promotion straightforwardly delivers important constraints on how reasons can vary with sets of alternatives. In doing so, I made crucial use of rankings of alternatives in terms of how well they promote various objectives. But if some reasons do not involve the promotion of objectives at all, how can we provide the constraints for these reasons? Moreover, if we need to tell some other story for these reasons, there's a worry that it will generalize to explain the constraints on how the reasons that do involve promotion can vary. This would undermine my claim that the idea that reasons involve promotion and contrastivism about reasons provide an attractive mutually supporting package.

Fortunately, I think we can generalize the explanation of the constraints given in terms of promotion. The basic idea is that different alternatives can do better or worse at respecting or honoring values. For example, going into the burning building to save someone does better at respecting the value of human life than merely calling 911, which in turn does better than just walking by without even calling 911. There may be various ways of being better at respecting or honoring some value. For example, one action may respect a given value more than another action, or one action may be closer to the ideal way of honoring a given value than another action. The important point is that we will still be able to get a ranking of alternatives, not in terms of how well they promote a given value, but in terms of how well they respect or honor that value.²⁹ And this ranking will very plausibly

²⁹Some rankings will likely be pretty uninteresting, with one alternative at the top and all other alternatives tied at the bottom, as not respecting the value at all. For example, if I promise to *A*, then a ranking of the alternatives $\{A, B, C, D\}$ in terms of how well they respect the value of promise-keeping will be trivial in this way.

have the **Disjunction Boundedness** property. Thus, we will be able to appeal to these rankings in providing constraints on how non-promotional reasons can vary between sets of alternatives. To take just one example, if some fact r explains why A does a better job at respecting value V than B does, and why B does better than C , it will likewise explain why A does better than C . So we can see how the Transitivity constraint can be captured.

A potential difference between promotional and non-promotional rankings is that ties and perhaps even incommensurability may be more common in non-promotional rankings. This is just because, for many values, there are several different ways of respecting or honoring them, and these may be equally or incommensurably good. This will plausibly be less common in the promotional case (but probably not nonexistent). This is no problem for the view I am developing. Suppose it is your friend's birthday. And suppose that baking your friend a chocolate birthday cake is an equally good way of respecting the value of friendship as baking her a lemon birthday cake, while baking her no cake at all is not a good way of respecting this value. In a case like this, the fact that it is your friend's birthday is a reason to bake her a chocolate cake rather than bake her no cake, and a reason to bake her a lemon cake rather than bake her no cake, but not a reason to bake her a chocolate cake rather than bake her a lemon cake. And we can just explain this by appealing to the ranking of the alternatives 'bake a chocolate cake', 'bake a lemon cake', and 'bake no cake': the first two are tied, the third is ranked below them, in terms of how well they respect the value of friendship.

Let me sum up what I have done in the last two sections. In developing my contrastivist theory of reasons, it was convenient to make two assumptions. First, I assumed that all reasons involve promotion, and second, I assumed that the relevant kind of promotion is to be understood as probability-raising. Both of these assumptions are controversial. In the previous section I showed how to relax the second assumption—all we need are rankings of alternatives with certain minimal properties. In this section, I have shown that we can also relax the first assumption. This is because we can still appeal to rankings of alternatives, not in terms of how well they promote a given objective, but rather in terms of how well they respect or honor it. This is enough to secure the constraints on how reasons can vary between sets of alternatives.

4.7 Conclusion

In this chapter, I've shown that an independent source of motivation for contrastivism, the relationship between reasons and promotion, leads very naturally to a version of the theory that straightforwardly solves an important problem for the contrastivist. By appealing to this idea, we can capture intuitive relationships between reasons relative to different sets of alternatives without sacrificing the independence that lets the contrastivist solve problems facing non-contrastive theories. Reasons relative to different sets of alternatives are independent of each other, but the theory also has an important kind of *structure*.

The theory I developed here forms the heart of this dissertation. In the next two chapters, I'll apply the theory to some important issues in ethics, practical reasoning, and epistemology. I'll show, first, that contrastivism allows us to frame important questions in new and interesting ways, and second, that doing so lets us make some progress on them.

Before moving on, though, I want to take stock of where we are, in terms of the *depth* of contrastivism. In Chapter 3, I argued against a view I called *shallow contrastivism*, which holds that contrastivism only goes “language-deep”, and in favor of deep contrastivism, which holds that the important (possibly fundamental) normative reason—or favoring—relation is itself contrastive. Having developed the theory further in this chapter, however, we may wonder how deep contrastivism really goes. After all, the key to providing the constraints on how reasons can vary between sets of alternatives was to let things bottom out in something non-contrastive—perhaps conditional probabilities, perhaps something else that will give us the right kind of ranking.

This is an important point. I think what the need for constraints, which I motivated at the beginning of this chapter, shows is that we cannot adopt an extremely deep form of contrastivism, at least not without incurring significant costs. So the theory I am endorsing is a somewhat *moderate* form of contrastivism. A deeper version would hold that there is no non-contrastive foundation: perhaps we can rank a given set of alternatives, but how actions are ranked against each other itself would vary if the relevant alternatives changed. For example, *A* may be ranked higher than *B* relative to a set $\{A, B, C\}$, but once we introduce *D*, maybe *B* would be ranked higher than *A*. Perhaps there are some arguments in

favor of an extremely deep form of contrastivism like this, but I am not aware of any. And given the significant costs this theory would take on—in particular, the cost of losing the constraints I’ve been discussing in this chapter—we would have to think hard about how strong such arguments were.

But even though this theory is somewhat moderate, it is in opposition to nearly all existing work on reasons, which has traditionally assumed without argument that reasons are reasons for things *simpliciter*. Moreover, if we really can just appeal to something like conditional probabilities, then the non-contrastive foundation is not anything normative. (This may not be the case if we end up having to appeal to something like a ranking of alternatives in terms of how closely they approximate the ideal kind of honoring of some value.) So it may turn out that all the normative notions are contrastive, even if they bottom out in some non-contrastive, non-normative notion like conditional probability.

Chapter 5

INTRANSITIVITY

In Chapter 4, I developed a version of contrastivism (that we can call *moderate* contrastivism) that rules out a troubling kind of intransitivity of reasons, where r is a reason for A rather than B and for B rather than C , but not for A rather than C . This is an important result because (i) contrastivism, described abstractly as the view that reasons are only reasons relative to sets of alternatives, initially seemed to allow for intransitivity, and (ii) allowing for this kind of intransitivity is plausibly sufficient reason to reject a theory.

In this chapter, though, I introduce some pressure from the other direction: several writers have given powerful arguments that there really are cases of intransitivity in ethics and practical reasoning.¹ If they're correct—or even if we think they might be correct—a theory of reasons that simply rules out cases of intransitivity by definition looks problematic.

Importantly, from this perspective, an initial *attraction* of contrastivism, again described independently of the framework from Chapter 4, is that it seems better suited than non-contrastive theories to accommodate intransitivity. After all, a popular diagnosis of why intransitivity arises is that some moral factors are *essentially comparative*: whether or not a factor matters, or provides reasons, depends on the specific comparison being made.² If this thought is correct, then it seems that whether or not a fact is a reason for some action will depend on what you're comparing that action to, or to the relevant set of alternative actions, since the factors or ends that provide reasons only do so relative to some alternatives. And

¹See Temkin (1987, 1996, 2012); Rachels (1998, 2001); Friedman (2009).

²See in particular Temkin (1987, 1996, 2012).

that is the central contrastivist claim. So the objection that we don't even *want* to rule out intransitivity is especially sharp for the contrastivist.

A more general goal of this chapter is to give some support to the claim that contrastivism lets us make progress in important debates in normative philosophy. I hope to have shown in the first four chapters of the dissertation that the theory is independently well-motivated. In this chapter and the next, I want to show that it's also theoretically useful. If I'm right that the view I suggest in this chapter is a promising way to accommodate intransitivity, then we can use this view to frame important questions about that issue, since it will be a concrete example of a theory that allows for intransitivity. I want to emphasize right from the start, though, that it is not my goal to establish the possibility of intransitivity. Rather, I want to show that contrastivism sheds light on the debate, by showing us how to accommodate intransitivity, if we want to do so.

5.1 Transitivity and Reasons

First I'll say a bit about what kind of intransitivity I'll be concerned with here. It is different than the sort that is most often discussed in the literature, but the two sorts are connected.

The most frequently discussed normative relation in the literature on intransitivity is the 'better than' relation.³ Here is the principle at issue in those discussions:

Transitivity of 'Better Than': If A is better than B , and B is better than C , then A is better than C .

For the purposes of this chapter, I'll treat the things that are ranked (A , B , and C) as outcomes or states of affairs. Arguments against this principle tend to describe cases (often bizarre, yet possible, cases) in which the intuitive ranking of three (or more) states of affairs does not satisfy this principle. And frequently these arguments are based on cases devised by Parfit (1984), who does not, for what it's worth, think his cases show that 'better than' is intransitive. I'll present arguments in this vein in the next section.

³See, for example, Temkin (1987, 1996, 2012); Rachels (1998, 2001); Friedman (2009) for arguments in support of the intransitivity of 'better than'. Broome (1991) argues against this kind of intransitivity.

Since I am concerned with reasons, I'll need to amend these arguments for the possibility of intransitivity of 'better than' to apply to reasons. Fortunately, this is easy to do, given the following linking principle:

Reasons and Betterness: *A* is better than *B* iff there is more reason to choose *A* than to choose *B*, given the choice between *A* and *B*.

This principle should be amenable both to those who want to analyze value in terms of reasons (for example, Scanlon (1998)) and to those who want to analyze reasons in terms of value. Even if you don't think either of these sorts of analysis can succeed, the principle is very plausible.

Many writers have been attracted to *fitting-attitudes* accounts of value, on which to be good is just to be a fitting object of a pro-attitude, where fittingness is often, though not always, analyzed in terms of reasons.⁴ **Reasons and Betterness** is attractive for some of the same reasons that these accounts are attractive. In particular, as many writers have pointed out, there are important connections between reasons and value. For example, if some outcome is valuable, then it is very plausible that there's a reason to choose that outcome, if given the choice. Of course, there might be stronger reasons to choose some other outcome, but if that's so, then plausibly it is because the other outcome is more valuable—or better—in some way.⁵ **Reasons and Betterness** predicts that there should be such a connection.

But importantly, **Reasons and Betterness** does not require the truth of any fitting-attitudes account. We might instead take it as (part of) an analysis of reasons to choose in terms of value or betterness. Or we might think that there's some common third normative property that affects both the value and strength of the reasons to choose the outcome. So again, this principle is amenable to very different theories about the relationship between reasons and betterness.

The clause, 'given the choice between *A* and *B*' is important here. That's because reasons to choose, like all reasons, are reasons *for agents*. But it is controversial whether the value of an outcome is similarly connected to agents. We might have two outcomes, *A* and *B*, such that *A* is intuitively better than *B*, even

⁴See Scanlon (1998); Danielsson and Olson (2007); Rabinowicz and Rønnow-Rasmussen (2004); Suikkanen (2005, 2009); Reisner (2009); Schroeder (2010); Way (fc) for a sampling of this literature.

⁵See, in particular, Way (fc) on this point.

though no agent could have any reason to choose A because no agent would ever be in a position to make such a choice. The ‘given the choice’ clause blocks this. And moreover, in the cases I’ll discuss, I’ll simply stipulate that the agents involved face a choice between the relevant outcomes.

So while arguments in the intransitivity literature present cases that challenge **Transitivity of ‘Better Than’**, I’ll make use of these cases in challenging the following principle:

Transitivity of ‘More Reason’: If there is more reason to choose A than to choose B , and more reason to choose B than to choose C , then there is more reason to choose A than to choose C .

I think this principle is *very* plausible, and that nearly everyone (except advocates of intransitivity that I mentioned above) would accept it, even independently of **Reasons and Betterness**.

Often these kinds of transitivity principles are simply taken for granted. Temkin (1987) hypothesizes that Parfit (1984) took it for granted that ‘better than’ (all-things-considered) is transitive, and that this influenced his reaction to the cases he presents, which I’ll discuss in the next section. Broome (1991) claims that the transitivity of ‘better than’ is essentially a logical truth: “Whatever decisions we make, we shall always be guided by logic to preserve the transitivity of the comparative. [...] Logic requires betterness to be transitive”.⁶ The phrase ‘more reason’ suggests a view like this. It encourages us to think in terms of a linear scale, like the number line, so that to say that there is more reason to choose A than to choose B is like saying that choosing A falls to the right of choosing B on some line. And once we have this picture in mind, it’s easy to see why we would think that ‘more reason’ is transitive, since ‘to the right of’ is transitive. Nevertheless, I’ll rehearse arguments in the next section that call these transitivity principles into question.

Before moving on, I should note one qualification. Some people think that two outcomes can be *incommensurable* in value—that the value of the two outcomes cannot be compared. If A and C are two such outcomes, we might have that A is better than B and B is better than C without having that A is better than C . This could hold if B is commensurable with both A and C , though A and C

⁶Broome (1991), pp. 12, 136.

are incommensurable with one another. These kinds of cases, if they exist, are not generally treated as counterexamples to transitivity principles, but rather as exceptions. I could make this explicit by adding a clause to the principles like ‘and A and C are commensurable’, but I’ll stick with the simpler formulation, and simply assume that all the outcomes I’m discussing are commensurable.

5.2 Intransitivity

The most popular kind of argument for the possibility of intransitivity is based on cases from Parfit (1984). Temkin (1987, 1996); Rachels (1998, 2001); Friedman (2009) all give arguments based on Parfit’s cases. These cases are an obvious place to begin because (i) denying the transitivity assumption is one way to avoid the problematic results, and (ii) other proposed solutions are unsatisfactory. That’s not to say that many people find the solution of denying transitivity satisfactory, of course, but it does mean that we should at least give this solution a chance. The main argument I’ll present here is essentially just what I think is the most compelling way to present the common thread from these other arguments. So I make no claim to originality.

First I’ll show how one of Parfit’s cases combined with **Transitivity of ‘More Reason’** leads to what Parfit (1984) calls the *Repugnant Conclusion*. Parfit, and most other people, think this conclusion just can’t be true. But the common response is to reject a particular step in the argument other than the assumption of transitivity. But as advocates of intransitivity have shown, this is shortsighted: there are big problems with this strategy, and moreover, it leads us right back into a case—Parfit’s *Mere Addition Paradox*—which calls **Transitivity of ‘More Reason’** into question.

5.2.1 The Repugnant Conclusion

Consider the following case from Parfit (1984).⁷ Suppose a powerful being asks you to choose between twenty-six outcomes of the universe. Outcome A ⁸ contains

⁷See also Rachels (1998, 2001); Friedman (2009). As I said above, these authors are writing about ‘better than (all things considered)’, while I’m writing about ‘more reason than’.

⁸Non-italicized capital letters are just names of outcomes that I’ve described. Italicized capital letters are variables ranging over outcomes.

some relatively large population, say ten billion, all with a very high level of well-being. B contains a larger population—say for definiteness that it’s twice as large as A—all with a lower, but still quite high, level of well-being. This reduced level of well-being is not due to any sort of unfairness or other moral deficiency—roughly, all the moral factors except for total and average utility are the same in the two outcomes. I think that, intuitively, there is more reason to choose B than to choose A. Agree with me for now, and I’ll address worries about this first step below. Similarly, there is more reason to choose C than to choose B—it’s twice as large, and everyone has only a slightly lower level of well-being. So by **Transitivity of ‘More Reason’**, there is more reason to choose C than to choose A.

Continue this sequence until we get to Z. The population in Z is enormous—one billion times ²⁵. Everyone in Z has a level of well-being that makes their lives barely worth living (“muzak and potatoes”, as Parfit would say). Well, there’s more reason to choose B than to choose A, more reason to choose C than to choose B, more reason to choose D than to choose C, . . . , more reason to choose Y than to choose X, and more reason to choose Z than to choose Y. So if **Transitivity of ‘More Reason’** is true, it follows that there’s more reason to choose Z than to choose A. But this is what Parfit calls the *Repugnant Conclusion*: that for any outcome with a population of at least ten billion people, all of whom have unbelievably awesome lives, there is some possible outcome with a much larger population all of whose lives are barely worth living, such that there’s more reason to choose this second outcome than the first.

This result seems wrong. But every step in the sequence seemed correct—as we move through the options in alphabetical order, there seemed to be more reason to choose each option than to choose the one before it. If we want to retain both of these judgments—that the Repugnant Conclusion is false and that there’s more reason to choose each option in the sequence than the one before it—it looks like we have to give up on **Transitivity of ‘More Reason’**. That’s the first argument for intransitivity.^{9,10}

⁹Parfit (1984) considers and rejects several other kinds of responses, including an appeal to a *valueless level*: a level of well-being below which one’s life doesn’t contribute to the overall goodness of the outcome, or how much reason we would have to choose that outcome, though the life is still valuable to the one who lives it, since it is worth living. There are big problems with this view, as well as other attempts to avoid the Repugnant Conclusion.

¹⁰Rachels (1998) runs a similar argument by flipping the diagram upside down, and focusing on just one life for extremely long stretches of time. The idea is that in outcome A, you face very

5.2.2 Resisting the first step

One way to block this argument, of course, is to resist one of the steps along the chain. I'll focus on the first step, from A to B, but I don't see why this choice should make a significant difference. Parfit himself actually thinks the Repugnant Conclusion forces us to deny that there is more reason to choose B than to choose A, because he assumes that transitivity holds.

This strategy seems to require us to say that there is in fact more reason to choose A than to choose B. If we said merely that there is *not* more reason to choose B than to choose A, then we should say, by parity of reasoning, that there's not more reason to choose C than to choose B, and so on. Then we could say that there's not more reason to choose Z than to choose A. But if we don't say that there's more reason to choose A than to choose B, and so more reason to choose B than C, and so on, then transitivity-like considerations force us to say that there's not more reason to choose A than to choose Z. But this is also intuitively wrong—it isn't just that we think there's *not* more reason to choose Z than to choose A; we think there *is* more reason to choose A than to choose Z. The principle this argument relies on is the following.

Transitivity of 'Not More Reason': If there's not more reason to choose A than to choose B, and not more reason to choose B than to choose C, then there cannot be more reason to choose A than to choose C.

And this is very similar to **Transitivity of 'More Reason'**.

Some writers who accept **Transitivity of 'More Reason'** would be willing to deny this last principle. Cases of incommensurability provide the most compelling

intense pain for one year, while in B you face still intense, but slightly less intense, pain for one hundred years, and so on. In Z, you face very mild discomfort (with no redeeming qualities—you'd rather be temporarily unconscious, rather than experience the discomfort) for eons and eons. Each step—from A to B, from B to C, ..., from Y to Z—seems to be a step for the worse. So A is better than B is better than C ... is better than Y is better than Z. But then if 'better than' is transitive, it follows that A is better than Z. But intuitively, Rachels thinks, Z is actually better than A—experiencing the very mild pain in Z, even for eons and eons, is not as bad as experiencing the most intense pain imaginable in A, even for only one year (make it 100 years and increase everything else accordingly, if you aren't convinced). One potentially nice feature of this example, that I don't want to get into here, is that arguably the reasons involved are all agent-relative reasons: since you're the one that's going to suffer the pain, they're reasons for you to choose the relevant option. Of course, if I was given the choice between which of these fates you would suffer, I would have agent-neutral reasons to choose the same options. But that's compatible with you having agent-relative reasons.

counterexamples. Broome (1999) presents the following sort of case. Consider three possible careers: (i) a career as a lawyer, (ii) a career as a lawyer making \$1,000 less per year, but which is otherwise identical to career (i), and (iii) a career as an academic. If the careers as a lawyer and career as an academic are sufficiently satisfying, it's plausible that you don't have more reason to choose career (i) than career (iii), or more reason to choose career (iii) than career (ii). The intuition here is that there a career as a lawyer and a career as an academic are simply valuable in different ways, so that neither is better than the other. And \$1,000 per year is just not enough money to make a difference here. But if all of this is right, it would follow from the **Transitivity of 'Not More Reason'** that you don't have more reason to choose career (i) than career (ii). But that's clearly false—since the two careers are otherwise identical, an extra \$1,000 per year is exactly the kind of thing to give you more reason to choose (i) than to choose (ii). So there's some reason to think that this principle is false. This is far from uncontroversial, though.

But even if we deny the **Transitivity of 'Not More Reason'** while retaining **Transitivity of 'More Reason'**, we could just make B contain many more people with levels of well-being barely below the level in A. And that sort of modification clearly gives us some additional reason to choose B. If there wasn't more reason to choose A than to choose B to begin with, it seems like this sort of modification should give us more reason to choose B than to choose A. That doesn't follow immediately, of course. But if this sort of modification doesn't give us more reason to choose B than to choose A, we need some explanation for why.

So why think that there's more reason to choose A than to choose B? We're assuming that the outcomes are equivalent with respect to all the moral factors that they can be equivalent with respect to, while still exhibiting the relevant differences in total and average utility. So some moral factor with respect to which A is better than B must be making the difference, if there's more reason to choose A than to choose B. That is, the reasons provided by one of these moral factors must tip the scales in favor of A. The most obvious candidate is average utility.¹¹

Outcome A has a higher level of average utility than B, even though there's more total utility in B. If we want to appeal to average utility to resist the first step—to say that there's actually more reason to choose A than to choose B—we have to say that the difference in average utility is somehow more important than

¹¹See Temkin (1987); Friedman (2009) for more thorough discussions of this sort of strategy.

the difference in total utility. There are two ways to flesh out this response. First, we could say that average utility is one very important moral factor while admitting that total utility is also important, though less so, in determining what we have most reason to choose. Second, we could say that any increase in average utility is always better than any increase in total utility.

Suppose we take the first strategy. Then we admit that the fact that B has more total utility than A gives us some reason to choose B rather than A. But we just insist that the fact that A has higher average utility gives us *stronger reason* to choose A rather than B. Since A and B are identical with respect to other potentially relevant moral factors—we’re assuming, for instance, that any violations of duties that take place in B are matched in A—we can thus deny that there’s more reason to choose B than to choose A.¹² But as long as we admit that total utility gives us some reason to choose B rather than A, we can simply alter the case so that B has *much* more total utility than A and *very slightly* lower average utility. If we think total utility matters at all in this case, it’s implausible that this sort of modification could never flip the scales, so that the vast increase in total utility finally outweighs the tiny increase in level of average well-being.

This pushes us toward the second strategy: no increase in total utility can make up for any deficiency in average utility. So no matter how much more total utility B has, and no matter how small the difference in the level of average utility, there will always be more reason to choose A than to choose B. The problem with this strategy is that it’s simply implausible. Consider a version of the case in which outcome A has just one person with a very high level of well-being, and outcome B has tens of billions of people all with this same level of well-being as the person in outcome A, plus a single person with a very slightly lower level of well-being. We can even imagine that this single person lives lightyears away from the others, so that he’s not even aware of anyone with higher levels of well-being than himself.¹³ The advocate of the second strategy has to hold that we still have more reason to choose outcome A than outcome B. And that is obviously implausible.

¹²Temkin (1987) considers the moral factor of *perfection*. See that paper for arguments that that sort of strategy fails. See that paper, as well as Friedman (2009), for more discussion of the objections I raise here.

¹³Compare Parfit (1984)’s Mere Addition Paradox—more on this later.

5.2.3 The Mere Addition Paradox

I've been arguing that we have more reason to choose B than to choose A, given the choice. And this sort of reasoning, along with **Transitivity of 'More Reason'**, leads to the Repugnant Conclusion. This gives us some reason to question this transitivity principle. But I haven't considered all the possible ways one might resist the first step of the argument (or some other step along the chain). So in this subsection I'm going to present an argument that, even if we find some way to resist the first step, and claim that there's more reason to choose A than to choose B, we *still* have reason to question **Transitivity of 'More Reason'**.¹⁴

The argument turns on Parfit (1984)'s *Mere Addition Paradox*. Parfit was forced to face this paradox by denying that B is better than A, in order to avoid the Repugnant Conclusion. For my purposes, of course, we face the version I'll present by denying that there's more reason to choose B than to choose A. Imagine an outcome A+ that contains all the same people as A, plus some additional people whose lives are well worth living, but whose levels of well-being are lower than that of the original group. But the two populations are totally unaware of one another. Imagine you are forced to choose between these three outcomes. Parfit argues that there is definitely not more reason to choose A than to choose A+—we can not make an outcome worse by adding happy people to it. There is a drop in equality and in average well-being, but we haven't introduced any inequality via any sort of injustice. This sort of addition is what Parfit calls *mere addition*, and he claims that mere addition can never make an outcome worse. So we don't have more reason to choose A than to choose A+.¹⁵

A+ contains just as many people as B. Parfit argues that there's more reason to choose B than A+: I stipulate that B is better with respect to equality, total well-being, and average well-being. Thus, it's hard to see what would motivate us to think there's more reason to choose A+ than to choose B.

But remember that we're assuming that there's more reason to choose A than to choose B, in order to avoid the Repugnant Conclusion while retaining **Transitivity of 'More Reason'**. Note again that it's not easy to see how to avoid the Repugnant Conclusion if we say merely that there's not more reason to choose B than to

¹⁴See Friedman (2009).

¹⁵I think that we plausibly have more reason to choose A+ than to choose A. Parfit denies this. But nothing really turns on this for my purposes here.

choose A—we could just double the size of B and make the level of well-being of its inhabitants even closer to the level of the inhabitants of A. But if there’s more reason to choose A than to choose B, and more reason to choose B than to choose A+ then it follows from **Transitivity of ‘More Reason’** that there’s more reason to choose A than to choose A+. But Parfit argued, and it seems very plausible that, this is actually false. This is the Mere Addition Paradox. And we got ourselves into it by claiming that there’s more reason to choose A than to choose B—that is, by resisting the first step of the argument for the Repugnant Conclusion. And we could give a version of the Mere Addition Paradox for any adjacent members of the chain (for F and G, rather than for A and B, say). So once we try to avoid the Repugnant Conclusion, we walk right into the Mere Addition Paradox, which *also* gives us reason to question **Transitivity of ‘More Reason’**.

5.2.4 Intransitivity and contrastivism

In this section, I’ve argued that Parfit’s cases give us reason to question **Transitivity of ‘More Reason’**. I’ve considered some ways to resist the argument, and shown that they either lead to implausible results or else lead us right back to problems for transitivity. So we should take seriously the idea that ‘more reason than’ is intransitive.

This puts some pressure on the version of contrastivism I developed in Chapter 4, which seems to simply rule out intransitivity by definition. This is even more striking, given that an initial attraction of a contrastive theory of reasons might be that it seems better placed to deny **Transitivity of ‘More Reason’**. The idea that reasons relative to different sets of alternatives are independent of one another—that settling what the reasons are relative to one set does not tell us what the reasons are relative to a different set—might make us think contrastivism is well-placed to accommodate intransitivity. That’s because, as I emphasized in Chapter 4, transitivity seems to amount to a constraint on this independence. But I showed that the framework I developed there actually *does* constrain this independence. It cannot turn out, according to that framework, that *r* is a reason for *A* rather than *B* and for *B* rather than *C*, but not a reason for *A* rather than *C*. So if the arguments above are at least compelling enough that we don’t want to rule out the

possibility of intransitivity by definition, there seems to be a problem for the view I developed in the last chapter.

5.3 Two Kinds of Intransitivity

In this section I'll distinguish between two kinds of intransitivity. I'll argue that the sort of intransitivity that the view from Chapter 4 rules out is *not* the sort that the arguments from the last section purport to establish.

First, we might have what I'll call *intransitivity of reasons*. If this sort of intransitivity is possible, then there might be a case in which some fact r is a reason for A rather than B and for B rather than C , but not a reason for A rather than C . Recall that in Chapter 4, I pointed out that the intuitive transitivity requirement on reasons is even stronger, once we think of reasons in terms of promotion. It isn't just that whenever r is a reason for A rather than B and for B rather than C , it must be a reason for A rather than C . The end (whatever fills in for ' X ' in the conditional probabilities) that explains why r is a reason for A rather than B and for B rather than C (if there is a common end that explains both of these things) must also explain why r is a reason for A rather than C . The intuition behind a transitivity requirement wouldn't really be captured if *something else* explained why r is a reason for A rather than C . This is the variety of intransitivity that the framework from Chapter 4 straightforwardly rules out.

Second, we might have what I'll call *intransitivity of 'more reason'*. If this sort of intransitivity is possible, then there might be a case in which there is *more reason* for A than for B out of $\{A, B\}$ and more reason for B than for C out of $\{B, C\}$, but not more reason for A than for C out of $\{A, C\}$.

The arguments for intransitivity from the last section, if they're successful, only establish the intransitivity of 'more reason'. What seems wrong is that there's *more reason* to choose outcome Z than to choose outcome A . But it is not counterintuitive to say that there is *some reason* to choose outcome Z rather than outcome A . In particular, considerations of total utility provide a reason to choose outcome Z rather than outcome A . The fact that, say, more utility is better than less is a reason to choose B rather than A , a reason to choose C rather than B , ..., and a reason to choose Z rather than Y . It follows from the transitivity of reasons that it's a reason to choose Z rather than A . And this seems correct. It's just that the

reason is outweighed by the reasons to choose A rather than choosing Z, explained by some other factor, like average utility.

If this is right, then it's not simply that the cases don't support the intransitivity of reasons. They actually offer strong support for the *transitivity* of reasons. If this is the most compelling support for intransitivity—and the fact that advocates of intransitivity almost invariably appeal to these cases is some evidence for this—and it seems *wrong* to describe them as cases of intransitivity of reasons, that casts serious doubt on this kind of intransitivity.

5.4 Contrast-Sensitive Importance

Even once we recognize the difference between these two types of intransitivity, it's still a bit mysterious how we could have intransitivity of 'more reason' without the intransitivity of reasons. In this section I'll show how to pry them apart. I'll offer a way of accommodating the intransitivity of 'more reason' in the framework from Chapter 4, even though it clearly rules out the intransitivity of reasons. I'll take this in two steps. First I'll observe that the strength of reasons depends on the importance of the ends (desires, values, and so on) that provide those reasons. Then I'll suggest that the importance of a given end might *itself* be a contrast-sensitive matter: an end that matters a lot in making some comparisons might matter very little in making other comparisons. So the resulting view is that, even if a fact is a reason for some action relative to both of two different sets of alternatives, it might be a much weightier reason relative to one of them, due to the contrast-sensitivity of the importance of the end that provides that reason.

5.4.1 Strength of reasons and importance of ends

I argued in the last section that Parfit's cases don't push us towards the intransitivity of reasons, but only towards the intransitivity of 'more reason'. But this seems unstable. The intransitivity of reasons appears to be necessary for the intransitivity of 'more reason'. After all, if any fact that's a reason for A rather than B and for B rather than C is thereby a reason for A rather than C, the "amount" of reason, or the combined strength of the reasons, would also seem to be transitive. Thus, it

seems that since the view from Chapter 4 rules out the intransitivity of reasons, it also rules out the intransitivity of ‘more reason’.

But in fact the intransitivity of ‘more reason than’ does not require the intransitivity of reasons. An important first step is recognizing that the combined strength or weight of the reasons for A doesn’t depend only, or even most importantly, on *how many* reasons there are. Frequently there are more reasons to do A than B , but stronger reason to do B . The fact that I’d get my shoes muddy is a reason not to run into the pond to save the drowning child, as is the fact that I would be late for the movie. The fact that the child will die if I don’t is a reason to save him. So that’s two reasons not to save and only one reason to save. But clearly the reason to save is stronger than the combined strength of the reasons not to save. And the explanation for this is that the end that explains why I have this reason—saving the child’s life—is much more *important* than the ends that explain why I have the reasons not to save—keeping my shoes clean and seeing the movie. So how much reason there is to perform an action depends—in large part, at least—on the importance of the ends that provide those reasons.

An important upshot of this observation is that we cannot simply read off how strong a particular reason to A is, or what the combined strength of the reasons to A is, off of the facts about conditional probabilities that I used in Chapter 4. We need some account of the relative importance of the various reason-providing ends. I don’t have an account of this, but there are intuitively clear cases, like the one I mentioned above involving the drowning child.¹⁶

5.4.2 Contrast-sensitive importance

It’s still not clear, though, how to allow for the intransitivity of ‘more reason’ once we rule out the intransitivity of reasons. Even if the strength of an X -provided reason for A out of $\{A, B\}$ and for B out of $\{B, C\}$ depends on the importance of X , the strength of the X -provided reason to A out of $\{A, C\}$ will likewise depend on the importance of X . And if X is very important, then it’s going to provide

¹⁶It doesn’t follow, though, that the facts about conditional probabilities are irrelevant. First, and most importantly for my purposes, the conditional probabilities establish what reasons there are, relative to a given set of alternatives. This is very important, because it lets us retain the nice entailment relations I established in Chapter 4. I’ll have a bit more to say about this shortly. Second, facts about conditional probabilities can arguably make a difference when the ends that provide two reasons are equally important.

relatively strong reasons across all of these cases. So the transitivity of ‘more reason’ still seems to follow from the transitivity of reasons.

To see that this is wrong, it might be instructive to first see how we could allow for the intransitivity of reasons while retaining (something like) the probabilistic framework from Chapter 4. What we would have to do is to let the ends that provide reasons—the X s—*vary with sets of alternatives*. If X provides reasons relative to $\{A, B\}$ and $\{B, C\}$, but not relative to $\{A, C\}$, then we could have cases of intransitivity of reasons. The most important cost to adopting this kind of picture, though, is that we would lose the entailment relations between reasons relative to different, but related, sets that I established in Chapter 4. These are attractive results, so we should try to retain them if possible.¹⁷

The framework from Chapter 4 delivered relations between what reasons there are relative to different sets. Changing what ends can provide reasons between sets of alternatives leads to a difference in these facts. So to retain the entailment relations, we need to hold the ends fixed between sets of alternatives. No matter what the relevant set of alternatives is, there is one class of ends that provide reasons, though they’ll provide reasons for different things relative to different sets of alternatives, of course, depending on which alternatives in the set best promote them.

But as I argued above, the *strength* of the reasons for some action depends on more than what reasons there are: it depends on the importance of the ends that provide those reasons. So the way to accommodate the intransitivity of ‘more reason’ while retaining the transitivity of reasons (and so the entailment relations) is to let the relative importance of the ends vary between sets of alternatives. Some end X *always* gives us reasons, but these reasons are more or less weighty, depending on the set of alternatives, since the importance of the ends varies with sets of alternatives.¹⁸ Thus, even if every fact that is a reason for A rather than B and for B rather than C is thereby a reason for A rather than C , the weights of these

¹⁷One way to try to retain these entailment relations on this more radical picture would be to place some restrictions on how the reason-providing ends could vary between sets of alternatives. I won’t explore this option here, since as I argued above, the arguments for intransitivity don’t even purport to establish intransitivity of reasons.

¹⁸If there are some cases—though I haven’t seen any—that do seem to establish the intransitivity of reasons, we could even accommodate that, or something like it, on this picture. We would just let some ends provide reasons of *no weight*, relative to some sets of alternatives. So it would still be true to say that the fact is a reason, and we could retain the entailment relations. It’s just that the reasons have no weight.

reasons may vary depending on the comparison. We might have some very weighty reasons for *A* rather than *B* and for *B* rather than *C* which are much weaker reasons for *A* rather than *C*, because the relative importance of the ends that provide these reasons varies with specific comparisons being made. Then we can fail to have more reason for *A* than for *C* without giving up on the transitivity of reasons.

This is a way of spelling out the intuition, often appealed to in the anti-transitivity literature, that some values or ends are “essentially comparative”.¹⁹ Differences with respect to some value like equality, total utility, or average utility might make a big difference when comparing some alternatives, but matter much less when comparing other alternatives. We can allow that having more total utility, for example, is always a reason to choose an outcome rather than choosing one with less total utility. But depending on other features of the alternatives, this reason can be stronger or weaker. And that’s because considerations of total utility might be more or less important depending on the comparison.

Here’s an illustrative and particularly relevant example. When comparing two outcomes with relatively close levels of average utility, like any two adjacent outcomes in the sequence I used in the argument for the Repugnant Conclusion, total utility is more important than average utility. That’s why, when comparing any two adjacent outcomes, there’s always more reason to choose the one later in the sequence—more reason to choose *G* than to choose *F* out of {choose *F*, choose *G*}, for example. But when we have vast differences in average utility, total utility is much less important than average utility. If that’s right, it would explain why there’s more reason to choose *A* than to choose *Z*, even though there’s so much more total utility in *Z*. Even though considerations of total utility do provide a reason to choose *Z* out of {choose *A*, choose *Z*}, this reason is much weaker than the reason to choose *A* provided by average utility. That’s why there’s more reason to choose *A* than to choose *Z* out of {choose *A*, choose *Z*}. On the other hand, the reason provided by average utility to choose *A* out of {choose *A*, choose *B*} is apparently weaker than the reason provided by total utility to choose *B* out of this set—that’s why we think there’s more reason to choose *B*. The importance of total utility depends (probably among other things) on the relative levels of average utility, and vice versa.

¹⁹See especially Temkin (1987, 1996, 2012). See also Rachels (2001), p. 218: “[A] factor determining how *A*&*Z* compare differs dramatically in significance from how it figures in comparing *A*&*B*, *B*&*C*, *C*&*D*, and so on”.

5.4.3 Contrastivism about reasons and contrastivism about importance

Contrastivism about reasons is the view that whether or not some consideration is a reason for an action or attitude depends on the alternatives. The strategy for accommodating the intransitivity of ‘more reason’ that I’ve just outlined involves making the importance of reason-providing ends, and thus the strength of reasons, depend on the alternatives. What is the relationship between these two ideas?

First, we can be contrastivists about reasons without adopting a contrast-sensitive view of the importance of ends. That is, the strategy I’ve outlined in this section is *detachable* from the rest of the contrastivist theory I’ve developed in previous chapters. So if you do not find the arguments for intransitivity compelling, or just find the difficulties that come along with intransitivity (some of which I’ll discuss in the next section) too troubling, that’s no reason not to be a contrastivist. You can just hold the importance of ends fixed across sets of alternatives to rule out intransitivity.

Second, we can adopt a contrast-sensitive view of the importance of ends without being contrastivists about reasons. On this kind of view, whether or not some fact is a reason for *A* does not depend on the alternatives, though the weight of that reason does depend on the alternatives. For example, the total utility in *Z* always provides a reason to choose *Z* (*simpliciter*, of course). But when we compare choosing *Z* to choosing *A*, that reason is simply not very weighty, though it is quite weighty when we compare choosing *Z* with choosing *Y*. Note that this is precisely what the contrastivist says about these particular comparisons. The difference between the two views will be the following. The non-contrastivist about reasons says that the total utility in *A* provides a reason to choose *A* even when we compare it to choosing *Z*. But the contrastivist denies this: total utility doesn’t provide any reason to choose *A* when we compare it with choosing *Z* (or, as the contrastivist would rather say, no reason to choose *A* *rather than* choose *Z*). But the important point for now is that the non-contrastivist about reasons can still be a contrastivist about importance, and accommodate intransitivity in the way I’ve suggested here.

It’s worth pointing out, though, that accepting contrast-sensitive importance is a particularly natural thing for a contrastivist to do. The contrastivist about reasons is impressed by the role that relevant alternatives seem to have in our

judgments about what is a reason for what. So thinking that alternatives can also matter for *how strong* those reasons are is right up the contrastivist's alley.

5.5 Remaining Questions

In this section I'll discuss some remaining issues for the view I've developed here. Throughout the discussion, I'll flag issues that would need to be resolved before we should feel comfortable adopting the account in this chapter. Thus, the upshot of this discussion isn't that this account is the correct one. It's rather that it offers a plausible *start*, if we want to accommodate intransitivity.

5.5.1 What should I do?

An important question for any theory that allows for intransitivity of 'more reason' concerns how we should decide what to do. A natural view is that you should do what you have most reason to do. But consider the imagined choice you face in the Repugnant Conclusion case: {choose A, choose B, ..., choose Z}. What out of this set do you have most reason to do? In Chapter 3, I adopted **CRO**, which says that we have to look at the two-member subsets: the option you should take out of the larger set is the one that wins in all of its pairwise comparisons.

So suppose we know what to say for any two member subset of this larger set.²⁰ Well, part of the point of this case was that *nothing* wins in all of the pairwise comparisons. And this is the result that the account I developed here gives. So the account I've offered so far doesn't tell us what to do in this case. And a theory of reasons, we might hope, will tell us what to do, once we know the facts about the reasons and their weights. After all, what we should do is intuitively determined by our reasons.

I think that in fact a theory should *not* tell us that we should perform some particular alternative out of the larger set in this case. That's just because any answer the theory gives will seem intuitively wrong. If the theory gave us some answer here, there are two possibilities. First, it might say that we should choose Z. Well, that seems wrong, since there's more reason to choose A. Second, it might

²⁰This doesn't actually seem true to me—what should we say about {choose A, choose N}? More on this shortly. The important point, though, is that we do know what to say about any adjacent pair, and about {choose A, choose Z}.

say that we should choose some outcome besides Z. That also seems wrong, though, since there's intuitively more reason to perform the next outcome in the sequence. So I take it that it's actually a *virtue* of the account that it doesn't deliver an (inevitably unsatisfying) answer as to what we should do in a case of intransitivity.

Contrastivism (plus contrast-sensitive importance) gets all of the pairwise cases for adjacent alternatives, which we do have strong intuitions about, correct. But we don't have strong intuitions about pairwise comparisons between non-adjacent alternatives. For example, what is there most reason to do out of {choose D, choose P}? It's hard to say. This is going to depend on whether we think the additional total utility in P provides a strong enough reason to choose P to outweigh the reason provided by the higher average utility in D. These are hard questions, but I won't address them here.

So the view developed in this chapter gives the correct results in cases that we have strong intuitions about, and tells us exactly how to decide in pairwise cases we don't have strong intuitions about—we just have to compare the weights of the reasons provided by the various factors. A disadvantage is that, since the weights of these reasons can vary widely depending on the comparison being made, there doesn't seem to be a neat, systematic way to do that. A theory that did not allow the weight of the reasons to vary—or at least made them vary in a systematic way—would be neater, but I don't see how to give such a theory while still capturing the intuitively clear cases. This is an important open issue. If there's no good way to make progress on it, then we should be wary of accepting this account.

5.5.2 Resolution of pairwise intransitivity?

I've just argued that, if the two-member subsets of a set of alternatives exhibit intransitivity, as in the Repugnant Conclusion case, we want a theory that does *not* say that we ought to do some particular thing out of the larger set. This leaves us in a precarious practical position—there's just no answer about what we should do, out of the larger set—but any answer would seem incorrect.

But this raises a problem for the theory I've suggested in this chapter. The strategy for accommodating cases of intransitivity was to allow the importance of the reason-providing ends to vary with sets of alternatives, which generates variation in the weight of the reasons. Suppose the weights of the reasons to *A*, *B*, and *C*

vary in such a way that we have a case of pairwise intransitivity: more reason for A than B out of $\{A, B\}$, more reason for B than C out of $\{B, C\}$, but not more reason for A than C out of $\{A, C\}$. I've argued that we just can't say what there is most reason to do out of the larger set, $\{A, B, C\}$. And again, this seems to be the right thing to say about the Repugnant Conclusion case, since no answer is satisfying. But—and here's the problem—for all the theory I've developed so far says, pairwise intransitivities might be resolved when we consider a larger set. The weights of reasons for the different alternatives relative to the pairwise sets do nothing to constrain the weights of the reasons relative to the larger set. So it might be that, even though no single alternative wins out in all of its pairwise comparisons, one *does* win out relative to the larger set. And I've just argued that this is exactly what we *don't* want to happen, at least in many cases.

In a way, it would be comforting if intransitivity was always resolved in larger sets of alternatives. Out of any pairwise set, we'd know what to do (except in cases of ties or incommensurability). And we would know what to do out of larger sets, as well—there'd be a right answer to the question of what there is most reason to do. Moreover, we'd still be able to admit the possibility of pairwise intransitivities. But the problem is that this simply doesn't seem to be what happens in paradigm cases of intransitivity, like the case in the argument for the Repugnant Conclusion. Still, the framework seems to formally allow for it. And if it does, given the argument in the last subsection, that's a problem for the theory. In the rest of this subsection, I'll show how to block this result: we can guarantee that pairwise intransitivity will not be resolved in larger sets.

First, I need to say more precisely what it is for pairwise intransitivity to be resolved in a larger set. Roughly, it's just for there to be some thing that we should do out of the larger set, even though there is pairwise intransitivity between its members. But this is just rough. We might have pairwise intransitivity between three members, B , C , and D , of the larger four-member set $\{A, B, C, D\}$. But suppose that A wins in all of its pairwise comparisons with these three. Then intuitively there is something we ought to do out of the larger set, namely A , even though there is pairwise intransitivity. But this isn't really the interesting sort of resolution of pairwise intransitivity. The kind of case we're interested in, then, is one in which no option wins in all of the pairwise comparisons, but one does win out of the larger set.

We can guarantee the result we want by adopting **CRO**, which I argued for in Chapter 3:

CRO: s ought to A out of Q iff s has most reason to A out of $\{A, B\}$ for all of the other alternatives B in Q .

I argued that this principle is preferable to a somewhat more natural principle, called **CRO***, that says simply that what you ought to do out of a set is what you have most reason to do out of that set. This is because **CRO** lets us explain more easily why it is that you ought to A out of Q when A is, roughly, best *on balance*, even though it's not best with respect to any particular factor. **CRO** tells us to look at the reasons relative to the two-member subsets of Q ; when A wins in all of those pairwise comparisons, it's what you ought to do out of the larger set. **CRO***, on the other hand, tells us to look only at the reasons relative to Q . If we adopt **CRO**, and assume that intransitivity is resolved in the larger set only if there's something you ought to do out of that larger set, then **CRO** guarantees that pairwise intransitivity will never be resolved in a larger set, in the sense I described in the last paragraph. If no option wins in all of its pairwise comparisons, then there's nothing you ought to do out of the larger set.

One strange result that the view, even with the addition of **CRO** seems to allow for, is the following. No alternative wins in all of the pairwise comparisons, and so according to **CRO**, there is no alternative that we ought to do out of the larger set. But the weights of the reasons may still shift in such a way that there *is* most reason to do one of them out of the larger set. But since it doesn't win out in all of its pairwise comparisons, again, it's not the case that we ought to do it.

This highlights the fact that, once we adopt contrast-sensitive importance, we leave room for the possibility of a case in which what you ought to do out of a set is not what you have most reason to do out of that set. In Chapter 3, I argued against the principle, **CRO***, which I mentioned above. I rejected this principle in favor of **CRO** because there are cases in which you intuitively ought to do something even though it's hard to find *any* reason to do it out of the relevant larger set. But I claimed that this wasn't so counterintuitive, since in these cases, the thing that you ought to do, according to **CRO**, is what's best *on balance*. But the kind of case we're considering now isn't like that.

Here's a schematic description of a way to rule out this last result. This sort of strategy would actually make appealing to **CRO** unnecessary. To accommodate

intransitivity, I appealed to contrast-sensitive importance of reason-providing ends, and thus to contrast-sensitive *weight* of reasons. The issue is then that there seems to be nothing to predict that the weight won't always shift in such a way that pairwise intransitivity is resolved in a larger set. So we could let the importance of ends relative to subsets of a larger set influence the importance of those ends relative to the larger set. For example, we might say that when there is significant variation in how important an end X is relative to the different subsets of a set Q , there's simply no answer to the question of how important X is relative to Q . Then there would be no answer to the question of how weighty the X -provided reasons are relative to Q . And finally, there would be no answer to the question of what there is *most* reason to do relative to Q —we can't decide, since we can't say how weighty the X -provided reasons are. Obviously this view would require much more fleshing out to properly evaluate it. There might even be ways of developing the idea so that pairwise intransitivity was *sometimes* resolved in larger sets, though I don't see exactly how that would go. I just present it here as an alternative to the strategy of appealing to **CRO**, which can avoid the potentially troubling results.

5.5.3 Money pumps

A famous worry about intransitivity concerns *cycling*, sometimes presented as a “money pump”. The idea is that if we allow intransitivity, there might be situations in which doing what you have most reason to do would predictably lead you to a worse outcome. Suppose you have more reason to choose a chocolate ice cream cone and pay a dime than to choose vanilla, more reason to choose vanilla and pay a dime than to choose strawberry, but more reason to choose strawberry and pay a dime than to choose chocolate. If we allow intransitivity of ‘more reason’, this seems like a possible case. But it looks like you could be pumped for arbitrarily large amounts of money, simply by doing what you have most reason to do.

Suppose you have a vanilla ice cream cone. Then I might offer you a trade: your vanilla cone plus a dime for my chocolate cone. By hypothesis you have more reason to choose a chocolate cone plus paying a dime than to choose a vanilla cone. So if you do what you have most reason to do—as presumably you ought to do—you will make the trade. Now suppose I offer you a second trade: your chocolate cone plus another dime for my strawberry cone. Again, by hypothesis, you have more

reason to choose a strawberry cone plus paying a dime than to choose a chocolate cone. So you'll make the trade. Now I offer you a third trade: my vanilla cone (which, recall, was originally *your* vanilla cone) for your strawberry cone plus a dime. Again, if you do what, by hypothesis, you have most reason to do, you'll make the trade. So after this cycle, you'll be right back where you started, with a vanilla cone, only you'll have paid thirty cents to get there. So you're certainly in a worse position than you were to begin with. And we could obviously make things even worse by going around the loop again: I could offer you the chocolate cone for your vanilla cone plus a dime, then the strawberry cone for your chocolate cone plus a dime, and so on, at least until the ice cream melts. So you could end up paying an arbitrarily large amount of money to end up with the same vanilla cone you started with. But all you did was what you had most reason to do in each case.

This looks like a big problem. We might have hoped that doing what you have most reason to do could not—at least not *predictably*—land you in a worse position overall. And all we needed, it seems, to get this going is the assumption that intransitivity of 'more reason' is possible. Thus, according to this objection, we should conclude that intransitivity of 'more reason' is impossible, despite the powerful arguments for its possibility.

I think this would be hasty. There are two popular ways to respond to this type of concern, which correspond to two different ways to conceive of the money pump case. First, suppose you're minding your own business, getting ready to enjoy a vanilla ice cream cone, when I approach you with a chocolate cone. You don't suspect any bad intentions on my part. Then it's hard to deny that you should make the trade. Now suppose a different person comes up to you with a strawberry cone, and again you don't suspect any bad intentions. Again, it's hard to deny that you should make the trade. At this point you're down twenty cents and have a strawberry cone. But there's nothing irrational about ending up in this situation. Intuitively, it's only when you get back to where you started, but have paid thirty cents to get there, that you're in an obviously worse situation than you were to begin with. So if another person comes up to you with a vanilla cone, according to this line of response, you should refuse to make the trade, even though by hypothesis you have more reason to make the trade, at least synchronically. You just recognize that the facts about what reasons you have (plausibly due to your preferences) make you vulnerable to being money pumped, *if* you aren't careful.

Once you recognize this, you will not do what you recognize you have most reason to do, in the particular situation. That is, at least in cases of intransitivity, it's not true that you should do what you have most reason to do. Rachels (1998, 2001) seems to think of the money pump case in this way, and offers this sort of response. He says, "The money-pump objection assumes that a rational agent would always prefer what is better and act on those preferences, no matter what. But that assumption would be rejected along with Transitivity".²¹

It's straightforward to put this kind of response in a contrastivist framework. The relevant sets are just the obvious ones: first, {keep the vanilla cone, trade for the chocolate cone}; second, {keep the chocolate cone, trade for the strawberry cone}, etc. But we just deny the initially attractive idea that what you ought to do is what you have most reason to do. This is an unattractive consequence; but perhaps Rachels is right that once we deny that transitivity holds, some intuitive principles of rationality simply have to go, as well.

Here's the second way to think of the money pump objection, and the second way to respond to it. You're getting ready to enjoy your vanilla cone, when you see me approach with a chocolate cone and a strawberry cone, looking mischievous. If you think clearly about the situation, it seems that you are plausibly rationally required not to be pushed around the cycle—not to make the trades, at least not more than the first two. The idea is that rationality requires more than simply looking at the synchronic choices you have to make. If we think that intransitivity is possible, we should adopt a more global perspective. Darwall (1983) discusses this kind of response to intransitivity. Sometimes we have to make what looks *synchronically* like the choice we have less reason to make to avoid ending up in a worse situation overall—"take one step back to take two forward", or perhaps better, "stay put to avoid taking a big step back". If you see what I have in mind with my ice cream cones, you should rationally not make the trades.

One way to make this response more precise, in my contrastivist framework, is to say that if you see what I have in mind, the relevant set of alternatives for the first trade isn't actually {choose the chocolate cone and pay a dime, choose the vanilla cone}, but rather something like {choose the chocolate cone and pay a dime and then stay put, keep the vanilla cone, choose the chocolate cone and pay a dime and then accept the trade for the strawberry cone minus another dime

²¹Rachels (2001), p. 218.

and then stay put, choose the chocolate cone and pay a dime and then accept the trade for the strawberry cone minus another dime and then accept the trade for the vanilla cone minus another dime}. Spelling out exactly what the set is might get a bit complicated, but the idea is simply to include descriptions of what you'll do later—to look at more fine-grained alternatives. When we do this, staying put, or at least only trading once or twice, will very plausibly be what you actually have most reason to do, out of these fine-grained sets of alternatives. When you have a vanilla cone, an alternative that involves paying thirty cents to end up with that same vanilla cone will certainly be worse than the others. This is a very natural thing for a contrastivist to do, since (i) she thinks that reasons are always relativized to sets of alternatives, and (ii) one important way in which sets of alternatives can vary is by being more or less fine-grained. So this strategy appeals to what I've been calling, following Yalcin (2011) and Cariani (fc), the *resolution-sensitivity* of reasons.

I suspect that this strategy can be extended to stop the money pump, even when it's conceived of in the first way. The basic idea is that after you make the first two trades, you should recognize that you are vulnerable to the money pump, if you aren't careful. At that point, you should be considering a more fine-grained set of alternatives like the one I described in the last paragraph. And again, once you consider this fine-grained set, alternatives in which you get pushed around the loop will very plausibly be worse. So if you don't anticipate getting money pumped to begin with, you might end up losing some money. But you won't be forced around the loop twice. I think this response is more attractive than the idea that you should simply refuse to make the trades, and knowingly do what you have less reason to do, because it lets us retain the idea that you ought to do what you have most reason to do.

So there are ways to respond to the money pump objection to intransitivity. Moreover, an attractive way to respond to the problem can be made precise by appealing to the important contrastivist idea of resolution-sensitivity. However, what I've said here likely isn't completely satisfying. Given the importance of this objection, then, this is another place where more work would need to be done before we could be comfortable accepting an account like the one developed here.

5.5.4 Why so uncommon?

The last issue I want to consider is the following. Even if we think intransitivity is theoretically possible, and so want a theory that can accommodate it, it's plausibly a desideratum on any theory that allows for intransitivity to give some explanation for why intransitivity is so rare. After all, the cases that are most often appealed to as alleged examples of intransitivity are Parfit's cases I described above. And these are extraordinary cases, to say the least. In any ordinary case, transitivity seems to hold. But if there's no formal obstacle to intransitivity, it looks like every single situation any real person will ever face is just a special case. On the theory I've suggested here, in particular, cases in which transitivity holds look like *very* special cases. The importance of the reason-providing ends across different sets of alternatives just happen to work out conveniently. We need some explanation of this. If we can't offer such an explanation, then we should be very skeptical that the theory I've suggested here is correct.²²

To satisfy this desideratum, I need to give an explanation of why intransitivity will be rare, even though there's no formal obstacle. I don't see how to adjust the structural features of the account to give such an explanation. So the strategy has to be to appeal to substantive facts about particular cases. As long as the substantive facts that generate intransitivity are rare enough, we can get the explanation we want.

The crucial feature of the account is what I called contrast-sensitive importance of ends. That's the feature that lets the weight of reasons vary with sets of alternatives. So if the importance of ends only varies—or only varies *significantly*—when extraordinary sets of alternatives are involved, we could explain why intransitivity is so rare. If relative to all the sets of alternatives that are ordinarily relevant, the importance of a given end is relatively constant, we won't get intransitivity in ordinary cases. But we can still get intransitivity when we compare the strength of reasons between ordinary and extraordinary cases, or among extraordinary cases.

Consider the sequence of outcomes from the argument for the Repugnant Conclusion. Now, any given set of alternatives consisting of choosing members of this

²²It's worth reiterating here that the official view of this dissertation is the one developed in Chapter 4. So if you find intransitivity too problematic, you can simply reject the amendment I proposed in this chapter, namely tacking on contrast-sensitive importance of reason-providing ends.

sequence aren't really *ordinary*. But what is ordinary about sets like {choose A, choose B} or {choose Y, choose Z} is that we're choosing between outcomes in which the average levels of utility are relatively close. Usually when we choose between some alternatives, our choice won't make a drastic change in the average level of utility, at least not while leading to an increase in the total amount of utility. Of course, some people may face choices between one outcome and another which would lead to a much lower average level of utility for some part of the population. But the second choice will generally lead to a much lower total level of utility, as well.

The importance of the ends of total utility and average utility, I suggest, remains relatively stable between sets like {choose A, choose B} and {choose A, choose C}. It's only when we consider a set like {choose A, choose Z} that the importance of the ends can shift significantly. So the intransitivity of 'more reason' is only going to arise when some of the sets we're considering are extraordinary. That explains why intransitivity is so rare: we don't generally consider extraordinary sets of alternatives.

Some explanation along these lines has to work, if the theory I developed in this chapter is going to be acceptable. I think the prospects of developing this explanation into something plausible are promising. But it would obviously take more work than what I've done here. All I want to suggest is that the package I've developed in this chapter is one very promising way to begin accommodating intransitivity.

5.6 Conclusion

In Chapter 4, I developed a contrastivist view of reasons, and showed that it rules out a certain kind of troubling intransitivity. In this chapter, I considered an objection from the other direction: intransitivity actually is possible, and moreover, an initial attraction of contrastivism—ignoring the framework from Chapter 4—is that it seems especially well-placed to accommodate it. I argued that the sort of intransitivity that this objector has in mind actually is compatible with, though not required by, the framework from Chapter 4, and showed how to accommodate it. I want to remain neutral on whether intransitivity is possible; all I meant to do in this chapter was demonstrate that contrastivism can accommodate it.

Another more general goal of this chapter was to show that the view I developed in the first four chapters of the dissertation provides an interesting way to think about important debates in ethics and practical reasoning. We can reframe these debates in contrastivist terms and make progress by exploiting the extra resources that contrastivism gives us. In this chapter, I hope to have shown that appealing to natural contrastivist ideas—in particular, contrast-sensitive importance of reason-providing ends—lets us make some progress in the debate over the possibility of intransitivity. In the next chapter, I'll apply the theory to an interesting issue from epistemology, and an analogous issue in practical reasoning: the rationality of *withholding* belief, on the one hand, and intention, on the other.

Chapter 6

WITHHOLDING

My goal for both the previous chapter and this one is to illustrate some interesting applications of contrastivism about reasons. In the last chapter I showed how contrastivism lets us make progress in the debate over the possibility of the intransitivity of ‘more reason than’. In this chapter, I turn to an application from theoretical and practical rationality. An important, but often neglected, issue from these domains is the rationality of *withholding* an attitude—belief, in the theoretical case, and intention, in the practical case. Sometimes the most rational thing to do is to neither believe that p nor believe that $\neg p$, but to *not make up your mind* with respect to p , or *withhold* belief. Similarly, sometimes it’s more rational to withhold intention with respect to ϕ , rather than intend to ϕ or intend not to ϕ .

In this chapter I will develop an account of when an agent ought to withhold, which applies to both belief and intention, that follows very naturally from the contrastive account of reasons I have developed in this dissertation. This account has several attractive explanatory features, and avoids problems that face an initially attractive non-contrastive account of when agents ought to withhold.

6.1 Withholding Belief and Contrastive Reasons

I will develop the account for the epistemic case of withholding belief first, and then extend it to the practical case of withholding intention.

Since we are trying to provide an account of what doxastic attitude you *ought* to have, and what you ought to do depends on your reasons, we need to figure out what your reasons are in this kind of case. Given the contrastive theory of reasons I’ve developed in this dissertation, we need to look at your reasons relative to some set

of alternatives or other. If you are trying to decide what doxastic attitude to take towards p , the most natural set of alternatives is {believe p , believe $\neg p$, withhold belief with respect to p }.¹ That is, the relevant alternatives will be believing that p is true, believing that p is false (or *disbelieving* p), or not making up your mind one way or the other.

Recall the principle relating what an agent ought to do with her contrastive reasons that I have adopted:

CRO: s ought to A out of Q iff s has most reason to A out of $\{A, B\}$, for all of the other alternatives B in Q .

When A -ing wins in all the pairwise comparisons, you ought to A . Now we can just straightforwardly apply this to belief:

O(Bel)*: s ought to believe p out of {believe p , believe $\neg p$, withhold} iff s has most reason to believe p out of {believe p , believe $\neg p$ } and s has most reason to believe p out of {believe p , withhold}

This principle is very plausible: you ought to believe p just in case you have better reason to do so than to either disbelieve p or withhold belief.² More importantly for the purposes of this chapter, **CRO** gives us the following account of when an agent ought to withhold belief:

O(Wh-B)*: s ought to withhold belief in p out of {believe p , believe $\neg p$, withhold} iff s has most reason to withhold belief in p out of {believe p , withhold} and s has most reason to withhold out of {believe $\neg p$, withhold}.

This principle, like **O(Bel)***, is just the one that follows straightforwardly from the theory I've developed in this dissertation.

6.1.1 Contrastive epistemic reasons

So in providing an account of when an agent ought to withhold belief, we need to look at four kinds of reasons: (i) reasons to withhold rather than believe p , (ii)

¹I will ignore more fine-grained attitudes, like credences, here.

²Here is Chisholm (1976), p. 27: “ h is certain for S at $t =_d f$ (i) Accepting h is more reasonable for S at t than withholding h (i.e., not accepting h and not accepting $\neg h$)”. And here is Conee and Feldman (2004), p. 3: “a person is justified in believing a proposition when the person’s evidence better supports believing that proposition than it supports disbelieving it or suspending judgment about it”.

reasons to believe p rather than withhold, (iii) reasons to withhold rather than believe $\neg p$, and (iv) reasons to believe $\neg p$ rather than withhold. Also important, if we want to figure out what doxastic attitude an agent ought to take towards p will be (v) reasons to believe p rather than believe $\neg p$ and (vi) reasons to believe $\neg p$ rather than believe p . What we need next, then, is some account of what these reasons are, so that we know what to compare.

Before providing that account, though, I need to make one important clarification. I said above that I am giving a view of what doxastic attitude an agent ought to take toward a proposition, and more specifically, when an agent ought to withhold belief in a proposition. The ‘ought’ here needs to be read in a particular way, as the ‘ought’ of *epistemic rationality*. So, I will provide an account of when you ought, *epistemically*, to withhold belief in a proposition. This epistemic ‘ought’ is the one in play in intuitive claims like, ‘You ought to believe what your evidence supports’, or ‘You ought not believe contradictions’. Similarly, the reasons I will be talking about are *epistemic* reasons—reasons that bear on what you epistemically ought to do.³

The most useful way to characterize epistemic reasons for my purposes here is by distinguishing them from non-epistemic reasons to believe a proposition or withhold belief in a proposition. If a powerful demon threatens to kill you if you believe that there is a glass on the table, that is a (very strong!) reason not to believe that there is a glass on the table.⁴ But this reason is importantly different than the reasons provided by evidence that there is a glass on the table, like the fact that you seem to see a glass on the table. The fact that you seem to see a glass on the table is a reason to believe that there is a glass on the table.⁵ This is a paradigm example of what I mean by an epistemic reason—assuming that you don’t have any other evidence that there is not a glass on the table, or any reasons to disregard your visual evidence, the epistemically rational thing for you to do is believe that there

³Two clarifications: (i) I do not mean to claim that ‘ought’ is ambiguous, in any deep sense, between this use and a practical use; for all I have said, the difference can be traced to contextual elements, like Kratzer (1981)’s ordering source and modal base. (ii) There is another use of ‘epistemic ‘ought’ that differs from my use in this paragraph; this is sometimes also called the *predictive* use, and it is in play in the most natural interpretation of ‘The roast ought to be done by now’.

⁴For discussion of this “wrong kind of reasons” problem, see D’Arms and Jacobson (2000a,b); Rabinowicz and Rønnow-Rasmussen (2004); Schroeder (2012b). Pascal’s Wager provides another classic example of a practical reason for belief.

⁵I am glossing over issues about perceptual justification and evidence here.

is a glass on the table. But this may not be the *practically* rational thing to do in this case. In this practical sense, plausibly, you ought not believe that there is a glass on the table, since the demon's threat gives you very strong practical—but not epistemic—reason not to believe this. The demon's reason does not bear on what it is epistemically rational for you to do, though it does bear on what it is practically rational for you to do. There is a huge literature on these issues. But in my discussion of belief and withholding belief in this chapter, I will ignore them, and stipulate that I am only talking about epistemic reasons.

One final point before I move on to give a contrastivist account of epistemic reasons. This last example may tempt you to think that epistemic reasons are all evidence for or against the proposition, while non-epistemic reasons involve good or bad features of being in the doxastic state (belief, disbelief, or withholding) itself. In fact lots of people have been tempted to exactly this kind of view.⁶ But as Schroeder (2012b) argues, there are epistemic reasons besides evidence. For example, the fact that you have very little information and expect to get more soon is intuitively an epistemic reason to withhold belief. It bears on what doxastic attitude it is rational for you to have—in this case, withholding belief. This is, not surprisingly, a very important point for this chapter, since I will be talking mostly about reasons to withhold belief, and paradigm cases of reasons to withhold belief are *not* evidence.

Now I am ready to give a contrastivist view of the six types of epistemic reasons listed above. Begin with reasons to believe p rather than believe $\neg p$. This is probably the easiest case: a reason to believe p rather than believe $\neg p$ is simply *evidence* for p . Similarly, of course, a reason to believe $\neg p$ rather than believe p is evidence for $\neg p$ (or evidence against p). Evidence for p is also a reason to believe p rather than withhold—the fact that Tom's fingerprints are on the book is a reason to believe that Tom stole the book, rather than withhold belief in this proposition. Of course, this reason might not be weighty enough to justify belief in this proposition, but that is compatible with saying that it is a reason to believe it. Similarly, evidence for $\neg p$ is a reason to believe $\neg p$ rather than withhold. Let ' Ev_p ' stand for the set of evidence for p and ' $\text{Ev}_{\neg p}$ ' stand for the set of evidence for $\neg p$.

This leaves reasons to withhold rather than believe p and reasons to withhold rather than believe $\neg p$. Begin with considerations like the fact that you will receive

⁶For views along these lines (ignoring important differences between them), see Parfit (2001); Piller (2001, 2006); Olson (2004); Stratton-Lake (2005); Hieronymi (2005).

more evidence regarding p tomorrow. This is both a reason to withhold rather than believe p and a reason to withhold rather than believe $\neg p$. Similarly, the fact that p is completely irrelevant is plausibly a reason to withhold rather than believe that p and a reason to withhold rather than believe that $\neg p$. After all, making up your mind about p may take some cognitive effort and distract you from more pressing matters, and given the unimportance of p , it would be a waste to do so.⁷ These types of reasons are both reasons to withhold rather than believe that p and reasons to withhold rather than believe that $\neg p$. Call the set of such reasons for a proposition p ‘ W_p ’.

Sometimes, though, the reasons to withhold rather than believe that p and the reasons to withhold rather than believe that $\neg p$ can come apart. An advantage of contrastivism is that it gives us a very straightforward way to capture this. Stanley (2005) presents cases in which believing p is very risky, while believing $\neg p$ is not, or vice versa.⁸ In a case like this, the potential costs of believing p provide reasons to withhold rather than believe p , but not reasons to withhold rather than believe $\neg p$. Similarly, the potential costs of falsely believing $\neg p$ provide reasons to withhold rather than believe $\neg p$, but not reasons to withhold rather than believe p .⁹ Let ‘ C_p ’ stand for the set of reasons provided by potential costs of falsely believing p and ‘ $C_{\neg p}$ ’ stand for the reasons provided by potential costs of falsely believing $\neg p$.

Finally, I claim that *evidence* can also be reason to withhold belief. Initially, this claim seems completely wrong.¹⁰ I said above that the fact that Tom’s fingerprints are on the book is a reason to believe that Tom stole the book rather than withhold—it is a reason *not* to withhold, and to believe that Tom stole the book, instead. But remember that I am providing a contrastive account of epistemic reasons. So evidence that p , I claim, is reason to withhold rather than believe that $\neg p$. Similarly, evidence that $\neg p$ is reason to withhold rather than believe that p . Though the fact that Tom’s fingerprints are on the book is a reason to believe that Tom stole the book rather than withhold, it is also a reason to withhold belief rather than believe that Tom did not steal the book. As we will see below, this

⁷I am not claiming that making up your mind about propositions is always, or even usually, cognitively taxing. But sometimes, it is, even for unimportant propositions.

⁸The riskiness depends on the belief that p playing its characteristic role in guiding action—to suffer the potential consequences of believing that p , you have to actually act as if p is true.

⁹See Schroeder (2012a) for more discussion of these kinds of cases, which focuses on costs of believing and withholding.

¹⁰See Schroeder (2012a) for an argument that evidence cannot be reason to withhold.

is the feature of this contrastive account that gives it explanatory advantages over the rival non-contrastive picture.

So to sum up, I have offered the following account of epistemic reasons. Let ‘ \oplus ’ stand for the combining reasons operation (which is not simply additive):

- reasons to believe that p rather than believe that $\neg p$ = reasons to believe that p rather than withhold = Ev_p
- reasons to believe that $\neg p$ rather than believe that p = reasons to believe that $\neg p$ rather than withhold = $\text{Ev}_{\neg p}$
- reasons to withhold rather than believe that p = $\text{W}_p \oplus \text{C}_p \oplus \text{Ev}_{\neg p}$
- reasons to withhold rather than believe that $\neg p$ = $\text{W}_p \oplus \text{C}_{\neg p} \oplus \text{Ev}_p$

6.1.2 A contrastive account of rational withholding

Now we can plug this account of epistemic reasons into the principles about when an agent ought to believe a proposition, and when an agent ought to withhold belief in a proposition. Let ‘ $A > B$ ’ mean that the set of reasons A is weightier than the set of reasons B .

O(Bel): s ought to believe p iff (i) $\text{Ev}_p > \text{Ev}_{\neg p}$, and (ii) $\text{Ev}_p > \text{W}_p \oplus \text{C}_p \oplus \text{Ev}_{\neg p}$

O(Wh-B): s ought to withhold belief with respect to p iff (i) $\text{W}_p \oplus \text{C}_p \oplus \text{Ev}_{\neg p} > \text{Ev}_p$, and (ii) $\text{W}_p \oplus \text{C}_{\neg p} \oplus \text{Ev}_p > \text{Ev}_{\neg p}$

Note that condition (i) in **O(Bel)** is actually redundant, assuming W_p and C_p cannot have negative weight. That is, whenever an agent has more reason to believe that p than to withhold belief (out of these two options), she will also have more reason to believe that p than to believe that $\neg p$ (out of these two options). This is a direct result, of course, of the fact that this account treats evidence that $\neg p$ as reason to withhold rather than believe that p , along with the assumption that there are not epistemic reasons not to withhold belief.¹¹ As we will see, though, the analogous condition governing when agents ought to intend will not be redundant.

In the next subsection, I will discuss a particularly nice feature of this account that an otherwise promising existing account, developed by Schroeder (2012a), does

¹¹This assumption is not uncontroversial, but I will not defend it here.

not share. As we will see, contrastivism has this feature because of the claim that evidence for a proposition p is not only reason to believe p rather than either believe $\neg p$ or withhold, but also reason to withhold rather than believe $\neg p$. As I mentioned above, and will further argue below, this move seems proprietary to the contrastivist.

6.1.3 Ties

When your evidence for p and your evidence for $\neg p$ are balanced, or tied, you ought to withhold belief. This is a very plausible requirement on any theory of rational withholding. If we accept the appealing idea that what you ought to do is what you have most reason to do, this means that, whenever your evidence is tied, you must have more reason to withhold belief than to believe p or believe $\neg p$. So when your evidence provides very little reason to believe p and equally little reason to believe $\neg p$, you have to have more reason to withhold. And—crucially—when your evidence provides very strong reason to believe p , and equally strong reason to believe $\neg p$, you *still* have to have more reason to withhold.

What we need, then, is some explanation for why your reasons to withhold belief will always be weightier than your reasons to believe p and your reasons to believe $\neg p$, no matter how weighty they are, when your evidence is tied. If we only recognize as reasons to withhold (i) costs of falsely believing and (ii) considerations like the fact that you will get more evidence regarding p tomorrow, it is hard to see why this should be so. The weights of these kinds of reasons seem to have nothing to do with whether or not your evidence is tied. So we lack an explanation for why these reasons to withhold will always outweigh both the reasons to believe p and the reasons to believe $\neg p$, in cases of ties.

An account of rational withholding developed by Schroeder (2012a) faces exactly this problem. According to his account, you ought to withhold when the costs of falsely believing p outweigh the evidence for p and the costs of falsely believing $\neg p$ outweigh the evidence for $\neg p$. We can add to his picture reasons like the fact that you will get more evidence tomorrow. But as I said above, the weights of these reasons seem to have nothing whatsoever to do with the weight of the reasons provided by the evidence for and against p . So we have no reason to expect them to

be weighty enough to make it the case that you ought to withhold belief whenever your evidence for and against p is tied.

One response here is to claim that in cases of ties, there is a special kind of reason to withhold belief:

Meta-Reason: The fact that your evidence is tied is a reason to withhold belief.

Thus, in cases of ties, you get a new reason to withhold belief. But to predict that you always ought to withhold belief in the case of ties, we need to add the following stipulation:

Stipulation: This Meta-Reason is always sufficiently weighty that your reasons to withhold belief outweigh both your reasons to believe that p and your reasons to believe that $\neg p$.

With these two assumptions, a non-contrastive account like Schroeder's can predict that you ought to withhold belief whenever your evidence is tied.

But I think this strategy is problematic for two reasons. First, it is simply not very explanatory. While it is plausible that the fact that your evidence is tied is a reason to withhold, I do not think it is very explanatory to simply stipulate that it is sufficiently weighty to do the job. This comes very close to simply stipulating that in cases of ties, you ought to withhold belief. But it would be better if we could *explain* this fact, instead.

Second, these assumptions do not go far enough. In most cases, if your evidence is *almost* tied, you still ought to withhold belief. But if your evidence is merely *almost* tied, we cannot appeal to **Meta-Reason**, since you only have this reason when your evidence is in fact tied. We could try modifying **Meta-Reason** so that you had this sort of reason even when your evidence is merely *almost* tied. But this is starting to look very unexplanatory. Further, we need some account of *how close* your evidence must be, in order for this sort of meta-reason to apply. Complicating things even further, it will likely be a variable matter how close your evidence must be, for withholding to be what you ought to do.

The contrastivist account I've given here, on the other hand, explains this observation nicely. To see this, assume that your evidence for p and your evidence for $\neg p$ are tied; so set $Ev_p = Ev_{\neg p}$ in **O(Wh-B)** above. It is easy to see, then, that in any case like this, you can never have more reason to believe p than to

withhold or more reason to believe $\neg p$ than withhold. So as long as either (i) there is always *some* cost of falsely believing, e.g. C_p and $C_{\neg p}$ are not completely empty, or (ii) some reason to withhold rather than either believe p or believe $\neg p$, e.g. W_p is not empty, it will turn out that in every case in which your evidence is tied, you ought to withhold out of {believe p , believe $\neg p$, withhold}. It is overwhelmingly plausible that at least one of (i) and (ii) is true. I acknowledge, though, that for my account to make the prediction that you ought to withhold in cases of ties, I do need one of these to be true.¹² If there are cases of ties in which neither (i) nor (ii) is true, my account can at least explain why it is never the case that you ought to believe p and not the case that you ought to believe $\neg p$ when your evidence is tied. Non-contrastive accounts like Schroeder's do not explain this.

Moreover, the contrastive account can explain why you ought to withhold belief, even in many cases in which your evidence is not quite tied. As long as there are either (i) potential costs of falsely believing or (ii) reasons to withhold rather than either believe p or believe $\neg p$ that are sufficiently weighty to outweigh the difference in the weight of the evidence for p and for $\neg p$, you ought to withhold. So this account tells us how close to balanced your evidence must be, in order for withholding to be what you ought to do. The difference must be small enough for the potential costs of falsely believing and your reasons to withhold in W_p to outweigh it. This explains why we get the sort of variation I mentioned above: in cases in which you have very weighty reasons to withhold, in W_p , or in which the stakes are very high (as in Stanley (2005)'s bank cases), you ought to withhold belief even if your evidence is quite unbalanced. In lower stakes cases in which you don't have very weighty reasons to withhold in W_p , your evidence must be much closer to balanced, for withholding to be what you ought to do.

The key to this contrastivist account is that it allows us to treat evidence as reason to withhold. This is what allows us to explain why you should always withhold belief in cases of ties, and why you should often withhold in the case of near-ties. This move seems proprietary to the contrastivist. A non-contrastivist

¹²One way we could guarantee this is to accept a contrastivist version of **Meta-Reason** above: the fact that your evidence is tied is always a reason to withhold rather than believe p and a reason to withhold rather than believe $\neg p$. Then W_p will not be completely empty in cases of ties. Notice, that this strategy does not require accepting the unexplanatory **Stipulation** that this meta-reason is always sufficiently weighty. Rather, it would *explain* why it was sufficiently weighty: in cases of ties, all you need, according to **O(Wh-B)** is *some* reason in W_p , so the meta-reason will be weighty enough.

could not let evidence that p and evidence that $\neg p$ provide (non-contrastive) reasons to withhold, since in that case, there would always be at least as much reason to withhold belief as to believe p or believe $\neg p$, and ordinarily there would be most reason to withhold. Thus, if the non-contrastivist were to attempt to co-opt this strategy of treating evidence as reason to withhold, she would commit herself to the view that in almost every case, what you ought to do is withhold belief. This would be a strange route to skepticism.

6.2 Withholding Intention

Just as agents often ought to withhold belief in a given proposition, they often ought to withhold intention, with respect to a given action. For example, suppose I've been invited to a party next weekend. I'd like to go, but there's a chance an old friend will be in town the evening of the party, and if she is, then I should see her. I will know in the next couple of days if she will in fact make the trip. In this case, the thing for me to do is withhold intention about whether to go to the party. I shouldn't intend to go, since my friend may be in town, and I shouldn't intend *not* to go to the party, since she may not be. Rather, I should withhold intention until I get more information. Now that I have provided the contrastivist account of when agents ought to withhold belief, I want to generalize to an account of when agents ought to withhold attitudes that applies not only to belief but also to intention.

6.2.1 A unified account?

My goal is to give a unified account of rational withholding that applies to both the practical and theoretical cases. Besides the theoretical virtues, there are also striking structural similarities between withholding intention and withholding belief that make it reasonable to expect such a unified account. In both cases, there is what we might call the 'positive' attitude—believing p or intending to ϕ —the 'negative' attitude—believing $\neg p$ or intending not to ϕ —and the 'undecided' attitude—withholding belief or withholding intention.¹³

¹³By calling believing $\neg p$ and intending not to ϕ 'negative', I don't mean to suggest that there are any deep differences between these and what I called the 'positive' attitudes. This is especially obvious in the case of belief: believing that $\neg p$ is simply bearing the same attitude—believing—to

But Harman (2004) makes an observation that seems to raise a problem for this idea. He points out that there is, at least in many cases, an important difference between the theoretical and practical cases with respect to what you ought to do when your reasons for the positive and negative attitude are *tied*. In the practical case, it is generally permissible to simply decide. In fact, it's often *impermissible* to not decide, even if your reasons to intend to ϕ and reasons to intend not to ϕ are tied. Think of the case of Buridan's Ass, who starved to death trying to decide whether to take the bale of hay on the right or identical bale on the left. Since the reasons to intend to take the bale on the right were perfectly matched by reasons to intend to take the bale on the left, it withheld intention about whether to take the bale of hay on the right (as well as about whether to take the bale of hay on the left), and ended up starving to death. But things are different in the theoretical case. When your reasons to believe p and your reasons to believe $\neg p$ are tied, it is in general *not* rationally permissible to just pick one to believe. In this case, you ought to withhold belief. Since cases of ties are cases in which you ought to withhold belief but not, or not always, cases in which you ought to withhold intention, it looks like we can't give a unified account of when you ought to withhold that applies to both.

Harman's observation is correct—it is often (practically) permissible to simply form an intention to ϕ or an intention not to ϕ when your reasons for each are tied, and it is never (epistemically) permissible to simply form a belief in p or a belief in $\neg p$ when your reasons for each are tied. But it is a mistake to conclude that this must be due to any structural differences between the correct accounts of when you ought to withhold intention, on the one hand, and belief, on the other. Rather, we can explain the asymmetry Harman observes by noticing that though there are frequently weighty reasons *not* to withhold intention (e.g. you'll starve to death, if you do) there are never, or at least very rarely, reasons not to withhold belief (besides, of course, evidence). Thus, we should actually expect the asymmetry Harman observes, even if we have a unified account of rational withholding. It is just that in the practical case, it will often be much harder for the reasons to withhold intention to outweigh the reasons not to withhold intention, but this will not be true for withholding belief.¹⁴

a different proposition. *Perhaps* there's some more important difference between intending to ψ , where ψ is incompatible with ϕ , and intending not to ϕ . But I'll ignore this issue here.

¹⁴Schroeder (2012a) brings out essentially the same point, but in a different way.

6.2.2 Contrastive practical reasons

Now I will provide an account of what I'll call, for the purposes of this chapter, *practical* reasons. By this I simply mean reasons that bear on what practical attitude—intending to ϕ , intending not to ϕ , or withholding intention—it is rational for you to take toward a given action, ϕ . This is to use the term ‘practical reasons’ in a different way than it is often used. This term is often used to refer to reasons for *action* rather than reasons for intention. But again, since I am talking about intentions, I am just talking about reasons for intention.

Just as I restricted my discussion in the epistemic case to epistemic reasons, ignoring reasons like the fact that a demon has threatened to kill you if you believe that p , I will here restrict my discussion to practical reasons that bear on whether it is rational to intend to ϕ , intend not to ϕ , or withhold intention. So I will set aside reasons like those present in Kavka (1983)’s Toxin Puzzle. In this case, you are offered a lot of money to intend to drink a mild toxin, with the catch that you know you will receive the money before the time comes to drink the toxin. In this case, you have a reason to intend to drink the toxin that is importantly different than a reason like the fact that someone else will give you a lot of money if you actually drink the toxin, or the fact that the toxin is tasty. I do not have a good way to distinguish between these “right kind” and “wrong kind” of reasons in the practical case, just as I had no good way to distinguish between the right kind and wrong kind of reasons in the epistemic case. But rather than entering into this important but difficult debate, I will rely on an intuitive understanding of the difference, and only talk about reasons that, I hope, fall clearly on the “right kind” side.

I am concerned here with the question of what attitude an agent ought to take toward an action, ϕ . So the relevant alternatives will be intending to ϕ , intending not to ϕ , or withholding intention. As before, we can simply appeal to **CRO**. Applying **CRO** to the case of intention, we get the following principles:

O(Int)*: s ought to intend to ϕ iff (i) s has most reason to intend to ϕ out of {intend to ϕ , intend not to ϕ } and (ii) s has most reason to intend to ϕ out of {intend to ϕ , withhold intention}

O(Wh-I)*: s ought to withhold intention with respect to ϕ iff (i) s has most reason to withhold intention out of {intend to ϕ , withhold intention} and (ii) s has most reason to withhold intention out of {intend not to ϕ , withhold intention}

So what we need is some account of what these reasons are: (i) reasons to intend to ϕ rather than intend not to ϕ , (ii) reasons to intend to ϕ rather than withhold, (iii) reasons to intend not to ϕ rather than intend to ϕ , (iv) reasons to intend not to ϕ rather than withhold, (v) reasons to withhold rather than intend to ϕ , and (vi) reasons to withhold rather than intend not to ϕ .

Perhaps the easiest cases are reasons to intend to ϕ rather than intend not to ϕ , and reasons to intend not to ϕ rather than intend to ϕ . These are simply reasons for and against ϕ -ing, respectively. Similarly, reasons for and against ϕ -ing will be among the reasons to intend to ϕ rather than withhold, and to intend not to ϕ , rather than withhold, respectively. Call these sets of reasons ' R_ϕ ' and ' $R_{ag-\phi}$ '. As a contrastivist, though, I need to say what sets these reasons for and against ϕ -ing are relativized to. I conceive of things as follows. Deliberating agents face two subtly different kinds of questions. First, there is the kind of question I have been dealing with in most of the dissertation, the question of what to do. This question will ordinarily provide a set of alternative actions, like $\{\phi, \psi, \chi\}$. Second, the agent will face the kind of question I have been considering in this chapter, the question of whether to ϕ . This question will provide a set of alternatives like the one I am concerned with here, {intend to ϕ , intend not to ϕ , withhold intention}. So the reasons for and against ϕ -ing will be relativized to the contextually relevant set, like $\{\phi, \psi, \chi\}$. The reasons to ϕ relative to this set will then also be reasons to intend to ϕ rather than intend not to ϕ , and to intend to ϕ rather than withhold intention. The reasons not to ϕ , $R_{ag-\phi}$, will also be reasons to intend not to ϕ rather than intend to ϕ , and to intend not to ϕ rather than withhold intention.

But in addition, we need to account for some other reasons not to withhold—or, on this framework, reasons to intend to ϕ rather than withhold and to intend not to ϕ rather than withhold. These will include considerations like the fact that you will starve to death if you withhold intention. For a more realistic example of reasons not to withhold intention, notice that agents often simply have to decide because the time to decide has come. If you are trying to decide whether to go see the guest speaker today, at some point—say, shortly before the talk is scheduled to start—you simply have to make up your mind. Another related source of reasons

not to withhold intention come from the fact that intentions aid in coordination.¹⁵ Forming intentions can aid in your own planning and in making joint plans with others. You miss out on these benefits if you withhold intention. Thus, these considerations about coordination can provide reasons to intend to ϕ rather than withhold, and reasons to intend not to ϕ rather than withhold. Notice how common these kinds of reasons are. This reinforces the point above that there are frequently reasons not to withhold intention, and this lets us explain Harman's observation. Call the set of such reasons ' $\neg W_\phi$ '.

That just leaves reasons to withhold rather than intend to ϕ and reasons to withhold rather than intend not to ϕ . In the epistemic case, we had reasons in the set W_p , like the fact that you'll get more information regarding p tomorrow, or the fact that you simply don't have very much evidence one way or the other. These are reasons to withhold belief rather than believe p and to withhold belief rather than believe $\neg p$. Analogously, in the case of intention, there are reasons to withhold intention rather than intend to ϕ and to withhold intention rather than intend not to ϕ . For example, the fact that you'll get more relevant information tomorrow is such a reason. For another example, the fact that there's no pressing need to make up your mind yet (e.g., there's no one to coordinate with, not much advance planning required, and the time to ϕ , if you do it at all, isn't for several months) is also a reason to withhold intention rather than either intend to ϕ or intend not to ϕ . Call the set of reasons like this ' W_ϕ '.

Just as the reasons to withhold belief rather than believe p could come apart from the reasons to withhold belief rather than believe $\neg p$, the reasons to withhold intention rather than intend to ϕ may come apart from the reasons to withhold intention rather than intend not to ϕ . First, there may be particular costs associated with intending to ϕ that are not associated with intending not to ϕ , or vice versa. Suppose that intending to ϕ cuts off some other options that are desirable in certain ways, while intending not to ϕ does not. This is plausible, given the coordinating role of intention, especially in the interpersonal case. Forming intentions to do or not to do certain things will lead others to form intentions that may then disqualify certain options that would otherwise be open to you. These kinds of costs can provide reasons to withhold rather than intend to ϕ that are not reasons to withhold

¹⁵As emphasized by Bratman (1987).

rather than intend not to ϕ (or vice versa). Call the set of such reasons ‘ C_ϕ ’ and ‘ $C_{not-\phi}$ ’, respectively.¹⁶

Finally, just as I held that evidence for p provides reasons to withhold rather than believe that $\neg p$, I hold that reasons to ϕ , R_ϕ , are also reasons to withhold intention rather than intend not to ϕ , and that reasons not to ϕ , $R_{ag-\phi}$, are also reasons to withhold rather than intend to ϕ .

6.2.3 A contrastive account of rational withholding of intention

Now we can plug these reasons into the principles **O(Int)*** and **O(Wh-I)*** to get the account of when agents ought to intend and when they ought to withhold intention:

O(Int): s ought to intend to ϕ iff (i) $R_\phi > R_{ag-\phi}$ and (ii) $R_\phi \oplus \neg W_\phi > R_{ag-\phi} \oplus C_\phi \oplus W_\phi$

O(Wh-I): s ought to withhold intention with respect to ϕ iff (i) $W_\phi \oplus C_\phi \oplus R_{ag-\phi} > R_\phi \oplus \neg W_\phi$ and (ii) $W_\phi \oplus C_{not-\phi} \oplus R_\phi > R_{ag-\phi} \oplus \neg W_\phi$

As I said above, condition (i) in **O(Int)** is not redundant, as it was in the epistemic case, **O(Bel)**. In the epistemic case, as long as an agent has more reason to believe that p than to withhold (out of those two options), she was guaranteed to have more reason to believe that p rather than believe that $\neg p$ (out of those two options). This relied on the assumption that there are not epistemic reasons not to withhold belief. But as I said in defending the possibility of a unified account of rational withholding from Harman (2004)’s observation, there frequently are reasons not to withhold intention. On the contrastivist view I’m developing here, these are the reasons to intend to ϕ rather than withhold and to intend not to ϕ rather than withhold— $\neg W_\phi$. The upshot here is that an agent may have more reason to intend to ϕ than to withhold intention (out of those two options) without having more reason to intend to ϕ rather than intend not to ϕ . In fact, this kind of thing seems to happen frequently. I have to settle on an intention to go to the talk or an intention not to go to the talk—I can’t just keep deliberating forever—even if neither of these options is more attractive than the other.

In the epistemic case, I claimed that whenever your evidence for and against p is tied, you ought to withhold belief. The analogous claim is not true in the practical

¹⁶I use ‘ $C_{not-\phi}$ ’ simply for notational convenience, for the set of reasons associated with costs of intending not to ϕ .

case, as Harman observed. Sometimes you just have to decide. But it is still true that in *some* cases in which your reasons for and against ϕ -ing are tied, you ought to withhold intention. And, crucially, there are both cases in which you ought to withhold and have very little reason either way, and cases in which you ought to withhold and have strong reasons each way. So we still need an explanation for why your reasons to withhold intention are weightier than your reasons to intend to ϕ and your reasons to intend not to ϕ in these cases.

A non-contrastivist account, for example, a practical analog of Schroeder (2012a)’s view, will not have such an explanation. That is because the weights of the kinds of reasons to withhold that this kind of account could recognize—the fact that you’ll get more information before the time for action, costs of closing deliberation, and so on—seem to have nothing to do with the weight of your reasons to ϕ and the weight of your reasons not to ϕ .

But the account I’ve developed here, captured in **O(Wh-I)** can provide such an explanation of how your reasons to withhold intention could vary in the necessary ways with the weights of your reasons to ϕ and your reasons not to ϕ . And just as in the epistemic case, this feature comes from the move, proprietary to the contrastivist, that reasons to ϕ are also reasons to withhold intention rather than intend not to ϕ , and reasons not to ϕ are also reasons to withhold intention rather than intend to ϕ .

Finally, just as in the epistemic case, this account can explain why sometimes you ought to withhold intention even if your reasons to ϕ and reasons not to ϕ are merely *almost* tied. As long as they are close enough to be outweighed by (i) the reasons to withhold intention in W_ϕ and (ii) the costs of settling deliberation, and there aren’t weighty reasons not to withhold intention, you ought to withhold intention.

6.2.4 A unified account

So far I have shown how the following two principles make attractive predictions about when an agent ought to withhold belief and intention, respectively:

O(Wh-B): s ought to withhold belief with respect to p iff (i) $W_p \oplus C_p \oplus Ev_{\neg p} > Ev_p$,
and (ii) $W_p \oplus C_{\neg p} \oplus Ev_p > Ev_{\neg p}$

O(Wh-I): s ought to withhold intention with respect to ϕ iff (i) $W_\phi \oplus C_\phi \oplus R_{ag-\phi} > R_\phi \oplus \neg W_\phi$ and (ii) $W_\phi \oplus C_{not-\phi} \oplus R_\phi > R_{ag-\phi} \oplus \neg W_\phi$

These two principles are structurally identical, other than the addition of ‘ $\neg W_\phi$ ’—reasons not to withhold intention—in **O(Wh-I)**. Given the (somewhat controversial) assumption that there are not reasons not to withhold belief, we can simply add to **O(Wh-B)** an empty set, ‘ $\neg W_p$ ’, without changing anything substantive about the account. Then the two principles are structurally identical, giving us a unified account of when agents ought to withhold that applies to both belief and intention.

O(Wh): s ought to withhold attitude α with respect to content π iff (i) $W_\pi \oplus C_\pi \oplus R_{\neg\pi} > R_\pi \oplus \neg W_\pi$, and (ii) $W_\pi \oplus C_{\neg\pi} \oplus R_\pi > R_{\neg\pi} \oplus \neg W_\pi$

Moreover, this account is just a special case of the much more general principle relating an agent’s reasons with what she ought to do, **CRO**. Thus, the contrastivist view I have developed in this dissertation provides an attractive account of when agents ought to withhold that can solve problems facing otherwise attractive non-contrastive views. Crucially, as I’ve argued, the advantages of the contrastive account actually come from contrastivism: only the contrastivist, it seems, can allow (i) evidence to be reasons to withhold belief, rather than the appropriate alternative and (ii) reasons for action to be reasons to withhold intention, rather than the appropriate alternative.

6.3 A Non-Contrastive Alternative

So far I have presented a theory of rational withholding that falls out of the more general contrastivist theory of reasons and ‘ought’ that I have developed in the dissertation, along with a particular conception of epistemic and practical (in my restricted sense) reasons. A central attraction of this theory is that it can explain why agents ought to withhold in cases of ties—always, in the case of belief, and often, in the case of intention. The theory is able to explain this because reasons for the “positive” attitude (believing p , intending to ϕ) are also reasons to withhold rather than have the “negative” attitude, and reasons for the “negative” attitude are also reasons to withhold rather than have the “positive” attitude. I argued above that this kind of move is proprietary to the contrastivist. If a non-contrastive theory, like the one developed by Schroeder (2012a), allowed reasons to believe p

and reasons to believe $\neg p$, or reasons to intend to ϕ and reasons to intend not to ϕ , to also be reasons to withhold, then withholding would nearly always be what agents have most reason to do—what they ought to do.

But in this section I want to consider a non-contrastive alternative that seems to have this same advantage, without resorting to contrastive reasons. One way to think about the key to the contrastivist account's explanation of why you ought to withhold in the case of ties is the following. Non-contrastive accounts only let evidence that p , or reasons to intend to ϕ , count in favor of believing p , or intending to ϕ . When we're trying to determine whether your reasons to withhold outweigh your reasons to believe $\neg p$, or to intend not to ϕ , these reasons to believe p or intend to ϕ do not have any role to play. The contrastivist, on the other hand, gives these reasons for the positive attitude a role to play when comparing withholding with the negative attitude. The account I want to consider now takes this insight from the contrastivist, and builds it into a non-contrastive theory.¹⁷

6.3.1 The account

As before, I will begin with the epistemic case. Rather than conceiving of the alternatives as believing p , believing $\neg p$, and withholding belief, we think of them as believing p , not believing p , believing $\neg p$, and not believing $\neg p$. Withholding seems to have dropped out of the picture. But we can simply say that withholding belief with respect to p is neither believing p nor believing $\neg p$.¹⁸

Just like we have reason to ϕ and reasons not to ϕ for a given action ϕ , we have reasons to believe p and reasons not to believe p , for propositions p . Reasons to believe p are again just evidence that p . Reasons not to believe p include potential costs of falsely believing p , as before, but also evidence that $\neg p$. Similarly, of course, reasons to believe $\neg p$ are evidence for $\neg p$, and reasons not to believe $\neg p$ include both potential costs of falsely believing $\neg p$ and evidence that p .

We then adopt the following principles about when you ought to believe (listing both is actually redundant, but helps make things clear):

¹⁷Mark Schroeder has suggested this account in conversation.

¹⁸This is one place where the question of whether withholding is best thought of as a special kind of attitude or as simply the absence of belief is important. If withholding is a distinct kind of attitude, this account would need to be modified.

Believe p : You ought to believe that p iff your reasons to believe that p outweigh your reasons not to believe that p .

Believe $\neg p$: You ought to believe that $\neg p$ iff your reasons to believe that $\neg p$ outweigh your reasons not to believe that $\neg p$.

Here is how this account explains why you ought to withhold belief in cases of ties. In cases in which your evidence is tied, these principles will say that it's not the case that you ought to believe p and that it's not the case that you ought to believe $\neg p$. Again, this is because we let evidence for p provide reasons against believing $\neg p$, and vice versa. But to get the result that you ought to withhold whenever your evidence is tied, we need some principle about when you ought to withhold. A natural choice, given the two principles about belief above, is the following:

Withhold: You ought to withhold belief with respect to p iff (i) your reasons not to believe p outweigh your reasons to believe p , and (ii) your reasons not to believe $\neg p$ outweigh your reasons to believe $\neg p$.

When your evidence is tied, your reasons to believe p will never outweigh your reasons not to believe p (similarly for $\neg p$). If we adopt a version of **Meta-Reason**, and say that the fact that your evidence is tied is always some reason not to believe p and some reason not to believe $\neg p$, then we'll get the result that you ought to withhold in every case in which your evidence is tied.

This appeal to **Meta-Reason** raises an important question about this approach. We have reasons to believe p and reasons not to believe p , as well as reasons to believe $\neg p$ and reasons not to believe $\neg p$. What about reasons to *withhold*? Since withholding isn't a distinct alternative, on this approach, it's not surprising that we don't have a distinct category of reasons to withhold. But some considerations, like the fact that you'll get more evidence soon, and the fact that your evidence is tied, seem best thought of as reasons to withhold. On this picture, we can think of these kinds of considerations as *both* reasons not to believe p and reasons not to believe $\neg p$. So while we don't have a distinct category of reasons to withhold, we can just think of the set of reasons to withhold as the *intersection* of the sets of reasons not to believe p and reasons not to believe $\neg p$.

This alternative view gains the advantages of the contrastive view by letting evidence that p count against believing $\neg p$, in addition to counting in favor of

believing p . The contrastivist implementation of this idea is to treat evidence that p as providing not only reasons to believe that p rather than believe that $\neg p$ and to believe that p rather than withhold belief, but also reasons to withhold rather than believe that $\neg p$. This account, on the other hand, lets evidence that p provide both reasons to believe that p and reasons not to believe that $\neg p$.

Generalizing this account to the practical case is straightforward. The main difference is that we need to say what reasons not to withhold intention are. Again, withholding is not a separate attitude on this account, but instead must be reduced to not intending to ϕ and not intending not to ϕ . So just as we can think of reasons to withhold as both reasons not to intend to ϕ and not to intend not to ϕ , we can think of reasons *not* to withhold as both reasons to intend to ϕ and to intend not to ϕ .

6.3.2 Comparing

Which of these two accounts—the contrastivist view I developed in the last section, or the non-contrastivist alternative I’ve been developing in this section—should we prefer? Of course, the main difference between the two accounts is that one appeals to contrastive reasons and one does not. So if the arguments I’ve given in favor of contrastivism in the rest of the dissertation are compelling, we should prefer the contrastivist account. If not, perhaps we should prefer the non-contrastivist alternative. But besides this obvious difference, the two accounts are quite similar. In particular, both treat evidence for p as counting against believing $\neg p$ in addition to counting in favor of believing p , though they do so in different ways.

One striking difference between the two accounts is that the contrastivist account, but not the non-contrastive account, posits withholding belief and withholding intention as a third alternative, not necessarily reducible to believing and disbelieving. On the non-contrastive account, we reduce withholding belief with respect to p as not believing p and not believing $\neg p$, and withholding intention to not intending to ϕ and not intending not to ϕ .

Corresponding to this difference is a difference in what kinds of reasons each account recognizes. While the contrastivist account recognizes reasons for and against withholding belief or intention as a distinct class of reasons, the non-contrastive account reduces them to reasons for and against believing and reasons for and against

intending. What we would naturally call a reason to withhold belief with respect to p , like the fact that you'll get more information soon, the non-contrastive account understands as a reason both against believing p and against believing $\neg p$; similarly for reasons to withhold intention, like the fact that you'll get more relevant information soon. Finally, reasons against withholding intention, like the fact that you need to coordinate with a friend, are reasons to intend to ϕ and to intend not to ϕ . This is where I think the contrastivist account gains an advantage over the non-contrastive account.

In Chapter 4, I argued that reasons for one option are reasons against incompatible options. A reason r for an action A highlights some good feature of A —on the picture I've been working with, it explains why A -ing would promote some valuable objective. So doing something incompatible with A -ing would cause you to miss out on this good feature of A -ing, and r would be part of the explanation for why. Thus, it's very plausible to think that reasons for one option are reasons against incompatible options. This is just a way to capture the very plausible idea that the *opportunity costs* of an option—the things you miss out on by performing it—are reasons against it.¹⁹ So, while not totally uncontroversial, this thesis is attractive.

But the advocate of the non-contrastive account I'm discussing cannot accept it. I will focus on the case of withholding intention, since it should be uncontroversial that there are reasons not to withhold intention (beyond just reasons for and against ϕ -ing). Consider the fact that I need to coordinate my plans for tomorrow with my friend so that we can meet for lunch. This fact is a reason against withholding intention about whether to go in to campus tomorrow; call it r . On the non-contrastive account, we understand reasons against withholding like r as reasons both to intend to ϕ and to intend not to ϕ . But now add the attractive thesis that a reason for one option is a reason against incompatible options. Since r is a reason to intend to ϕ , given this thesis, it is a reason against intending not to ϕ . Similarly, since r is a reason to intend not to ϕ , given this thesis, it is a reason against intending to ϕ . So r is both a reason against intending to ϕ and against intending not to ϕ . But according to this account, that is just what it is for r to be a reason *for* withholding intention. So given (i) the non-contrastive account and (ii) the attractive thesis that reasons for one option are reasons against incompatible options, we get the absurd

¹⁹I cited Greenspan (2005) as someone who rejects this view, and Sinnott-Armstrong (2006); Broome (ms); Ruben (2009) as supporters.

result that all reasons against withholding are also reasons for withholding.²⁰ Thus, the advocate of this non-contrastive account must reject the attractive thesis that reasons for one option are reasons against incompatible options. This is a theoretical cost.

The contrastivist, on the other hand, can accept this thesis, since she does not reduce reasons for and against withholding to reasons for and against believing/intending. Moreover, the contrastivist, unlike the advocate of the non-contrastive account, does not need to claim that a reason against withholding intention with respect to ϕ is always a reason to intend to ϕ . This is an unintuitive claim, even when we add that it's also a reason to intend not to ϕ (similar remarks go in the belief case, of course). Some reasons against withholding intention, on the contrastivist account, will not be reasons for intending to ϕ or for intending not to ϕ , relative to $\{\text{withhold, intend to } \phi, \text{ intend not to } \phi\}$, since it could be that the objective X which provides r as a reason against withholding is such that $P(X|\text{you intend to } \phi) = P(X|\text{you intend not to } \phi)$, even though both are greater than $P(X|\text{you withhold intention})$.²¹ I conclude that the contrastivist account I developed has theoretical advantages over the non-contrastive account, even beyond the advantage of fitting into an independently attractive contrastivist view of reasons in general.

6.4 Wrap Up

In this chapter I've argued that the contrastivist can give an attractive account of when one ought to withhold both belief and intention. Given the importance of

²⁰This absurd result does not follow if we treat reasons against withholding as reasons to intend to ϕ *or* reasons to intend not to ϕ , instead of as *both* reasons to intend to ϕ and reasons to intend not to ϕ . And this treatment may actually seem more plausible. The problem, though, is that it is not clear how we could incorporate reasons not to withhold, thought of disjunctively, into this picture. In order to count against withholding intention, on this picture, the reasons must either count in favor of intending to ϕ or count in favor of intending not to ϕ , or both—counting in favor of doing one or the other will not work. Given this, the only reasonable choice is to let these reasons not to withhold both count in favor of intending to ϕ and count in favor of intending not to ϕ , since just picking one at the exclusion of the other would be arbitrary.

²¹This is one place where we can see the importance of being able to distinguish between reasons against an option which imply a significant criticism of the option (in Greenspan (2005)'s sense) and those which do not. I argued in Chapter 4 that the account I developed there can make that distinction.

withholding and the problems facing other accounts, if I'm right, then this is a compelling argument in favor of contrastivism about reasons. Moreover, this chapter and the last show that contrastivism has interesting applications in areas in which reasons are important, including ethics, practical reasoning, and epistemology.

I've argued in this dissertation for contrastivism about reasons. After introducing the project in Chapter 1, I began, in Chapter 2, by arguing that our talk about reasons supports a contrastive semantics for 'reason'. I then argued, in Chapter 3, that we shouldn't stop there: normative favoring itself is contrastive. And finally, in Chapter 4, I developed the theory in more detail by analyzing reasons in terms of promotion—a popular strategy which itself provides independent support for contrastivism. Moreover, the theory I developed in Chapter 4 provides some needed constraints on the independence of reasons relative to different sets of alternatives, giving the theory *structure*. In Chapters 5 and 6, I illustrated some applications of the theory I developed in Chapters 2 through 4. I think that, taken together, all of this gives us strong reason to accept contrastivism, rather than either reject it or withhold.

References

- Anderson, E. (1993). *Value in Ethics and Economics*. Harvard University Press.
- Baron, M., Pettit, P., and Slote, M. (1997). *Three Methods of Ethics: A Debate*. Blackwell Publishers.
- Bedke, M. (ms). Ends to means: Transmitted reasons and their weights. Unpublished manuscript, University of British Columbia.
- Behrends, J. and DiPaolo, J. (2011). Finlay and Schroeder on promoting a desire. *Journal of Ethics and Social Philosophy*.
- Blaauw, M. (2012). Contrastive belief. In Blaauw, M., editor, *Contrastivism in Philosophy*. Routledge.
- Bratman, M. (1987). *Intention, Plans, and Practical Reason*. Harvard University Press.
- Broome, J. (1991). *Weighing Goods*. Blackwell Publishers.
- Broome, J. (1999). Incommensurable value. In Crisp, R. and Hooker, B., editors, *Well-Being and Morality: Essays in Honour of James Griffin*, pages 21–38. Oxford University Press.
- Broome, J. (2004). Reasons. In Wallace, Pettit, Scheffler, and Smith, editors, *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*, pages 28–55. Oxford University Press.
- Broome, J. (ms). *Rationality Through Reasoning*. Manuscript.
- Cariani, F. (2009). *The Semantics of ‘Ought’ and the Unity of Modal Discourse*. PhD thesis, University of California, Berkeley.
- Cariani, F. (fc). ‘Ought’ and resolution semantics. *Noûs*.
- Chandler, J. (2007). Solving the tacking problem with contrast classes. *British Journal for the Philosophy of Science*, 58(3):489–502.

- Chandler, J. (2013). Contrastive confirmation: Some competing accounts. *Synthese*, 190(1):129–138.
- Chisholm, R. (1976). *Person and Object*. Open Court.
- Conee, E. and Feldman, R. (2004). *Evidentialism: Essays in epistemology*. Oxford University Press.
- Crisp, R. (2000). Particularizing particularism. In *Moral Particularism*, pages 23–47. Oxford University Press.
- Dancy, J. (1993). *Moral Reasons*. Blackwell Publishers.
- Dancy, J. (2004). *Ethics Without Principles*. Oxford University Press.
- Danielsson, S. and Olson, J. (2007). Brentano and the buck-passers. *Mind*, 116:511–522.
- D’Arms, J. and Jacobson, D. (2000a). The moralistic fallacy: On the ‘appropriateness’ of emotions. *Philosophy and Phenomenological Research*, 61(1):65–90.
- D’Arms, J. and Jacobson, D. (2000b). Sentiment and value. *Ethics*, 110(4):722–748.
- Darwall, S. (1983). *Impartial Reason*. Cornell University Press.
- Dretske, F. (1970). Epistemic operators. *Journal of Philosophy*, 67:1007–1023.
- Evers, D. (2009). Humean agent-neutral reasons? *Philosophical Explorations*, 12(1):55–67.
- Finlay, S. (2001). *What Does Value Matter? The Interest-Relational Theory of the Semantics and Metaphysics of Value*. PhD thesis, University of Illinois, Urbana-Champaign.
- Finlay, S. (2006). The reasons that matter. *Australasian Journal of Philosophy*, 84(1):1–20.
- Finlay, S. (2009). Oughts and ends. *Philosophical Studies*, (143):315–340.
- Finlay, S. (ms). *Confusion of Tongues*. Unpublished book manuscript.
- Finlay, S. and Snedegar, J. (fc). One ‘ought’ too many. *Philosophy and Phenomenological Research*.
- Fitelson, B. (2012). Contrastive Bayesianism. In Blaauw, M., editor, *Contrastivism in Philosophy*. Routledge.
- Friedman, A. (2009). Intransitive ethics. *Journal of Moral Philosophy*, 6:277–297.

- Greenspan, P. (2005). Asymmetrical practical reasons. In Reicher, M. and Marek, J., editors, *Experience and Analysis: Proceedings of the 27th International Wittgenstein Symposium*, pages 387–394. oebv&hpt.
- Greenspan, P. (2007). Practical reasons and moral ‘ought’. In Shafer-Landau, R., editor, *Oxford Studies in Metaethics*, volume 2, pages 172–194. Oxford University Press.
- Grice, H. (1989). Logic and conversation. In *Studies in the Way of Words*. Harvard University Press.
- Groenendijk, J. and Stokhof, M. (1997). Questions. In van Benthem, J. and ter Meulen, A., editors, *Handbook of Logic and Language*, pages 1055–1124. Elsevier Science Publishers.
- Hamblin, C. L. (1958). Questions. *Australasian Journal of Philosophy*, 36:159–168.
- Hampton, J. (1998). *The Authority of Reason*. Cambridge University of Press.
- Harman, G. (2004). Practical aspects of theoretical reasoning. In Mele, A. and Rawling, P., editors, *The Oxford Handbook of Rationality*. Oxford University Press.
- Hieronymi, P. (2005). The wrong kind of reason. *The Journal of Philosophy*, 102(9):437–457.
- Higginbotham, J. (1993). Interrogatives. In Hale, K. and Keyser, S. J., editors, *The View from Building 20: Essays in honor of Sylvain Bromberger*, pages 195–228. The MIT Press.
- Higginbotham, J. (1996). The semantics of questions. In Lappin, S., editor, *The Handbook of Contemporary Semantic Theory*, pages 361–383. Oxford University Press.
- Hitchcock, C. (1996). The role of contrast in causal and explanatory claims. *Synthese*, 107(3):395–419.
- Horty, J. F. (2007). Reasons as defaults. *Philosophers’ Imprint*, 7(3).
- Horty, J. F. (2012). *Reasons as Defaults*. Oxford University Press.
- Jackson, F. (1985). On the semantics and logic of obligation. *Mind*, 94(374):177–195.
- Jackson, F. and Pargetter, R. (1986). Oughts, options, and actualism. *Philosophical Review*, 95(2):233–255.

- Kavka, G. (1983). The toxin puzzle. *Analysis*, 43(1):33–36.
- Kierland, B. (ms). A new approach to the future dependence issue. Unpublished manuscript, Boise State University.
- Kolodny, N. (ms). Instrumental reasons. Unpublished manuscript.
- Kratzer, A. (1981). The notional category of modality. In Eikmeyer, H. J. and Rieser, H., editors, *Words, Worlds, and Contexts*, pages 38–74. de Gruyter, Berlin.
- Lehrer, K. and Paxson, T. (1969). Knowledge: Undefeated justified true belief. *Journal of Philosophy*, 66(4):225–237.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74:549–567.
- Lipton, P. (1990). Contrastive explanation. *Royal Institute for Philosophy Supplement*, 27:247–266.
- Moore, G. (1903). *Principia Ethica*. Cambridge University Press.
- Moore, G. (1912). *Ethics*. Williams & Norgate.
- Morton, A. and Karjalainen, A. (2003). Contrastive knowledge. *Philosophical Explorations*, 6(2):74–89.
- Nagel, T. (1970). *The Possibility of Altruism*. Princeton University Press.
- Norcross, A. (1997). Good and bad actions. *Philosophical Review*, 106(1):1–34.
- Norcross, A. (2005a). Contextualism for consequentialists. *Acta Analytica*, 20(2).
- Norcross, A. (2005b). Harming in context. *Philosophical Studies*, 123(1-2).
- Olson, J. (2004). Buck-passing and the wrong-kind of reasons. *The Philosophical Quarterly*, 54(3):295–300.
- Parfit, D. (1984). *Reasons and Persons*. Oxford University Press.
- Parfit, D. (2001). Rationality and reasons. In Egonsson, Josefsson, Petterson, and Rønnow-Rasmussen, editors, *Exploring Practical Philosophy: From action to values*, pages 17–39. Ashgate.
- Parfit, D. (2011). *On What Matters*. Oxford University Press.
- Pettit, P. (1991). Consequentialism. In Singer, P., editor, *A Companion to Ethics*, pages 230–237. Blackwell Publishers.

- Piller, C. (2001). Normative practical reasoning. *Proceedings of the Aristotelian Society*, 25:195–216.
- Piller, C. (2006). Content-related and attitude-related reasons for preferences. *Philosophy*, 59:155–181.
- Rabinowicz, W. and Rønnow-Rasmussen, T. (2004). The strike of the demon: On fitting pro-attitudes and value. *Ethics*, 114(3):391–423.
- Rachels, S. (1998). Counterexamples to the transitivity of ‘better than’. *Australasian Journal of Philosophy*, 76(1):71–83.
- Rachels, S. (2001). A set of solutions to Parfit’s problems. *Noûs*, 35:214–238.
- Raz, J. (1999). *Engaging Reason: On the theory of value and action*. Oxford University Press.
- Reisner, A. (2009). Abandoning the buck-passing analysis of final value. *Ethical Theory and Moral Practice*, 12(4):379–395.
- Rickless, S. (fc). The contrast-*insensitivity* of knowledge ascriptions. *Philosophy and Phenomenological Research*.
- Rooth, M. (1992). A theory of focus interpretation. *Natural Language Semantics*, 1:75–116.
- Ross, J. (2006). *Acceptance and Practical Reason*. PhD thesis, Rutgers University.
- Ruben, D.-H. (1987). Explaining contrastive facts. *Analysis*, 47(1):35–37.
- Ruben, D.-H. (2009). Con-reasons as causes. In Sandis, C., editor, *New Essays on the Explanation of Action*, pages 62–74. Palgrave Macmillan.
- Scanlon, T. (1998). *What We Owe to Each Other*. Harvard University Press.
- Schaffer, J. (2004). From contextualism to contrastivism. *Philosophical Studies*, 119(1-2):73–104.
- Schaffer, J. (2005a). Contrastive causation. *Philosophical Review*, 114(3):327–358.
- Schaffer, J. (2005b). Contrastive knowledge. In Gendler and Hawthorne, editors, *Oxford Studies in Epistemology*, volume 1, pages 235–271. Oxford University Press.
- Schaffer, J. (2007). Knowing the answer. *Philosophy and Phenomenological Research*, 75(2):383–403.

- Schaffer, J. (2008). The contrast-sensitivity of knowledge ascriptions. *Social Epistemology*, 22(3):235–245.
- Schaffer, J. (2012). Causal contextualisms. In Blaauw, M., editor, *Contrastivism in Philosophy*. Routledge.
- Schroeder, M. (2007). *Slaves of the Passions*. Oxford University Press.
- Schroeder, M. (2010). Value and the right kind of reason. In Shafer-Landau, R., editor, *Oxford Studies in Metaethics*, volume 5, pages 25–55. Oxford University Press.
- Schroeder, M. (2012a). Stakes, withholding, and pragmatic encroachment on knowledge. *Philosophical Studies*, 160(2):265–285.
- Schroeder, M. (2012b). The ubiquity of state-given reasons. *Ethics*, 122(3):457–488.
- Searle, J. (2001). *Rationality in Action*. MIT Press.
- Sinnott-Armstrong, W. (1992). An argument for consequentialism. *Philosophical Perspectives*, 6:399–421.
- Sinnott-Armstrong, W. (2004). Classy Pyrrhonism. In Sinnott-Armstrong, W., editor, *Pyrrhonian Skepticism*. Oxford University Press.
- Sinnott-Armstrong, W. (2006). *Moral Skepticisms*. Oxford University Press.
- Sinnott-Armstrong, W. (2008). A contrastivist manifesto. *Social Epistemology*, 22(3):257–270.
- Skorupski, J. (2010). *The Domain of Reasons*. Oxford University Press.
- Sloman, A. (1970). ‘Ought’ and ‘better’. *Mind*, 79(315):385–394.
- Snedegar, J. (2012). Contrastive semantics for deontic modals. In Blaauw, M., editor, *Contrastivism in Philosophy*. Routledge.
- Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford University Press.
- Stanley, J. (2011). *Know How*. Oxford University Press.
- Stratton-Lake, P. (2005). How to deal with evil demons: Comment on rabinowicz and rønnow-rasmussen. *Ethics*, 115(4):778–798.
- Suikkanen, J. (2005). Reasons and value – A defence of the buck-passing account. *Ethical Theory and Moral Practice*, 7(5):513–535.
- Suikkanen, J. (2009). Buck-passing accounts of value. *Philosophy Compass*, 4(5):768–779.

- Temkin, L. (1987). Intransitivity and the Mere Addition Paradox. *Philosophy and Public Affairs*, 16(2):138–187.
- Temkin, L. (1996). A continuum argument for intransitivity. *Philosophy and Public Affairs*, 25:175–210.
- Temkin, L. (2012). *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning*. Oxford University Press.
- Temple, D. (1998). The contrast theory of why-questions. *Philosophy of Science*, 55(1):141–151.
- Toulmin, S. (1950). *Reason in Ethics*. Cambridge University Press.
- van Fraassen, B. (1980). *The Scientific Image*. Oxford University Press.
- Väyrynen, P. (2010). A wrong turn to reasons? In Brady, M. S., editor, *New Waves in Metaethics*. Palgrave Macmillan.
- Way, J. (fc). Value and reasons to favor. In Shafer-Landau, R., editor, *Oxford Studies in Metaethics*, volume 8. Oxford University Press.
- Wedgwood, R. (2009). Intrinsic values and reasons for action. *Philosophical Issues*, 19:342–363.
- Yalcin, S. (2011). Nonfactualism about epistemic modality. In Egan, A. and Weatherson, B., editors, *Epistemic Modality*. Oxford University Press.