BPhil in Philosophy: Thesis

Candidate Number: 1038091

Thesis Title: **Ruling Reasons: A Defense of Constructivism about Reasons and Objectivism about Welfare**

Name: Zhicheng Fan

Word count: 29935

Abstract

I argue in this thesis in favor of a position that combines constructivism about reasons and objectivism about welfare. In Chapter 1, I define reasons constructivism, on the basis of the Kantian origin of the view. I argue that reasons constructivism is a claim about the perspective-relativity of normative reasons for action. Furthermore, I argue that reasons constructivism, as I define the view, is distinct from both reasons objectivism and reasons subjectivism. In Chapters 2 and 3, I offer an argument in favor of reasons constructivism. I argue that in interpersonal reasoning, one agent can only offer reasons to another by engaging in the "if I were you" mode of reasoning. I show that these considerations support reasons constructivism, rather than other meta-ethical views about reasons. Finally, in Chapter 4, I argue that reasons constructivism can be combined with welfare objectivism. I show that this approach allows constructivists to generate the right reasons for agents with fundamentally corrupt judgments. Furthermore, I argue that this mixed position is coherent.

Acknowledgements

I could not have written this thesis without the help of many different people. Some of the ideas here were first explored in my earlier BPhil essays as well as my PhD writing sample, and I benefitted from many stimulating conversations and helpful comments in the course of turning these ideas into this thesis. First and foremost, I would like to thank my thesis supervisor Jeremy Fix for being the best supervisor I could ask for and for always pushing me to refine every aspect of my view and my writing. I would also like to thank my former supervisor Ruth Chang for guiding me through the literature and seeing the potential in my then inchoate views. Both Jeremy and Ruth pushed me to become a better philosopher, and I am incredibly grateful for their mentorship and for teaching me the craft of philosophical writing. Moreover, I would like to thank the participants of Ruth's DPhil seminar for their insightful comments on the second and third chapters of my thesis. I am especially grateful to Cécile Degiovanni, Joshua Pike, Calvin Chan, and Wataru Inukai. I would also like to thank Matt Hitchens, Ethan Kerns, Karen Margrethe Nielsen, and two anonymous examiners for their comments on some of my earlier drafts and papers on these topics. As a prospective PhD student, I also received valuable feedback on my writing sample from Stephen Darwall, Sharon Street, and Julia Markovits. I am particularly grateful to my friend Kevin Zhang for reading almost everything I have written on the topics covered in this thesis. Our many spirited debates during our Port Meadow walks really helped me think through the different components of my theory. In addition, I could not have completed the thesis without the emotional and moral support of my parents. Last but not least, I am grateful for the Oxford-Kaifeng Scholarship, which enabled me to pursue the degree in the first place without any financial burden.

Table of Contents

Introduction 5
Chapter 1 What is Constructivism about Reasons? 8
1.1. The Kantian Origin of Reasons Constructivism 10
1.2. Alternative Definitions of Reasons Constructivism 13
1.3. My Proposal 18
1.4. Reasons Objectivism and Subjectivism 20
Chapter 2 The "If I Were You" Argument: Reasons Objectivism and Perspective-Relativity 25
2.1. The "If I Were You" Mode of Interpersonal Reasoning 26
2.2. Reasons Objectivism 34
Chapter 3 The "If I Were You" Argument: Humean Internalism, Subjectivism, and Constructivism 42
3.1. Humean Internalism 42
3.2. Street's Subjectivism 52
3.3. Reasons Constructivism 54
Chapter 4 Welfare and Reasons Constructivism 56
4.1. Constructivism and the Right Reasons Problem 58
4.2. Must Reasons Constructivists Adopt Welfare Constructivism? 64
4.3. Does Welfare Objectivism Entail Reasons Objectivism? 70
4.4. Getting the Right Reasons 79
Conclusion 88
Bibliography 90

Introduction

The aim of this thesis is to motivate and defend a meta-ethical view that combines constructivism about normative reasons for action and objectivism about welfare. Reasons constructivism usually starts from a distinctive interpretation of Kantian autonomy. While there is no consensus on how to define constructivism, constructivists such as Christine Korsgaard and Sharon Street believe normative reasons for action are constructed from an agent's exercise of the will or her subjective judgments. Welfare objectivism, on the other hand, has usually been associated with the Aristotelian tradition (Hurka 1996; Finnis 2011; Kraut 2009; Foot 2001), although the view has also been defended by others (Griffin 1988; Raz 1988). Welfare objectivists claim that facts about welfare are not grounded in the judgments and attitudes of an agent. I find both views plausible, and in this thesis I try to show how a combination of these two views is both coherent and attractive. Reasons constructivism and welfare objectivism fit together in my framework in roughly the following way. In Simon Blackburn's terminology (1998, 5), normative reasons for action are the *outputs* of non-defective practical reasoning from the perspective of an agent, whereas facts about welfare are among the *inputs*. Normative reasons are constructed from the agent's perspective, and welfare facts are among the materials of construction.

The overall argument will proceed as follows. In Chapter 1, I give a definition of reasons constructivism. I argue that the central claim is that normative reasons are *perspective-relative*, and the only way to determine an agent's reasons is to reason from her perspective. More precisely, normative reasons for action exist in virtue of the fact that they are the outputs of non-defective practical reasoning from the perspective of the agent. I also try to show how reasons constructivism can be distinguished from two other views about reasons, namely reasons

objectivism and reasons subjectivism. On the one hand, reasons objectivists claim that reasons are mind-independent, and the aim of practical reasoning is to respond to these reasons. Reasons constructivism is opposed to this view. On the other hand, reasons constructivism, unlike reasons subjectivism, does not claim that reasons are grounded solely in an agent's subjective judgments. In particular, I suggest that reasons constructivism is compatible with objectivism about rational norms as well as certain evaluative facts, such as facts about welfare. This paves the way for my arguments in the later chapters.

In Chapter 2, I give the "if I were you" argument for reasons constructivism. In interpersonal reasoning, I can only determine your reasons for action if I reason from your perspective in the "if I were you" mode of reasoning. I argue that reasons objectivists cannot explain this interpersonal requirement. Moreover, the argument shows that reasons for action are perspective-relative.

Chapter 3 considers whether the "if I were you" argument favors subjectivist theories of reasons. I argue that subjectivist theories unduly limit an agent's normative reasons by her motivations or judgments. Instead, I argue that the kind of non-defective perspective-relative reasoning that grounds reasons is governed by objective rational norms. One such norm is the norm of analogical consistency. Another norm is that a piece of non-defective reasoning should not contain mistakes about objective evaluative and descriptive facts.

Finally, Chapter 4 develops the suggestion that we can correct mistakes about objective evaluative facts in non-defective reasoning. Constructivism faces what Ruth Chang (2013a, 173–74) calls the "right reasons problem." The problem is that for agents with fundamentally corrupt subjective judgments, constructivism seems to deliver the intuitively wrong verdicts about their reasons. I argue that welfare objectivism, when combined with reasons constructivism, has the

right resources to tackle the problem. I also argue that it is not incoherent to adopt constructivism about normative reasons, but not about welfare. While normative reasons for action are claims on someone's agency, facts about welfare are not. The different roles that normative reasons and welfare play in our practical reasoning explain why it is coherent to adopt only constructivism about reasons.

There are a few crucial topics that I cannot cover in this thesis, and I will briefly discuss them in the conclusion. Nevertheless, my goal is to argue for what I think is an attractive meta-ethical framework, and to show that it has quite a lot of potential.

Chapter 1 What is Constructivism about Reasons?

Normative reasons for action are considerations that count in favor of or against a course of action (Scanlon 1998, 17). In this chapter, I try to define a meta-ethical position about such reasons that I call "constructivism about reasons," or "reasons constructivism." In particular, I give a definition that differentiates the view from two alternative views about reasons, which are "reasons objectivism" and "reasons subjectivism." Reasons objectivism is the claim that all normative reasons are mind-independent, where a normative property or fact is mind-independent if it is not metaphysically grounded in an agent's attitudes, desires, motivations, commitments, choices, judgments, etc.[1] Reasons subjectivists, on the other hand, insist that all normative reasons are ultimately mind-dependent.[2] Given this familiar division between objectivism and subjectivism, some meta-ethicists find it hard to make sense of meta-ethical constructivism, especially the kind of Kantian position associated with Christine Korsgaard. On the one hand, given her arguments against what she calls "substantive realism" or "dogmatic realism" (1996a, chap. 1; 2008a; 2008b), it seems natural to classify her view under the category of subjectivism. Indeed, this is how Derek Parfit (2011, vol. 1, 65, 94) and Allan Gibbard (1999, 9) interpret her view. On the other hand, the rather substantive norms of reasoning endorsed by Korsgaard, especially in her later defense of constitutivism (2009), are thought by critics to require the kind of mind-independent normative entities that she objects to in her earlier works (FitzPatrick 2013; Crisp 2006, 52; Hussain and Shah 2006, 291; Chang 2013a, 175–76; Wallace

[1] For the notion of metaphysical grounding in connection to meta-ethics, see Fine 2001; Chang 2013; 2015. My definition of mind-independence here is not meant to be rigorous, and there are difficult questions about how to define mind-independence. See Joyce (2015, sec. 5). I am only illustrating a common distinction between meta-ethical theories. Some well-known reasons objectivists by this definition include Derek Parfit (2011), T. M. Scanlon (1998; 2014), John Skorupski (2012; 1999), Jonathan Dancy (2018; 2002), John McDowell (1985), Joseph Raz (1999; 2011), and Roger Crisp (2006).

[2] Some well-known subjectivists include Bernard Williams (1981; 1995), Mark Schroeder (2007), David Sobel (2016a), Peter Railton (1986), and, as I shall discuss, Sharon Street (2006; 2008; 2012).

2012, 33–34). The Kantian norm of the categorical imperative is not grounded in the contingent attitudes and judgments of agents. It is therefore not obvious how to make sense of reasons constructivism as a coherent and distinctive meta-ethical position.

As I shall define the view, what makes reasons constructivism distinctive is the thesis that normative reasons are always *perspective-relative*. More specifically, the view is that reasons for an agent *x* can only be determined by reasoning from *x*'s perspective about how she in particular should exercise her own agency. As I shall argue, this view is distinct from reasons objectivism, since objective reasons are not perspective-relative. Yet reasons constructivism is also not just reasons subjectivism. While reasons constructivism can be developed into a subjectivist version, as Sharon Street has done, constructivists can also claim that reasons are constructed by perspective-relative reasoning that is constrained by mind-independent rational norms (such as Korsgaard's categorical imperative). I believe that my definition in terms of perspective-relativity not only captures the Kantian origin of reasons constructivism, but also shows how the view can be a significant alternative to reasons objectivism and subjectivism.

In §1.1, I discuss the Kantian origin of reasons constructivism, and what this tells us about the central motivation for the view. A suitable definition of reasons constructivism must be based on the distinctively Kantian conception of autonomy. §1.2 addresses two alternative definitions presented by Korsgaard and Street. I argue that neither definition represents reasons constructivism as a distinctive meta-ethical position. §1.3 presents my own definition of the view, which grounds reasons in perspective-relative reasoning. Finally, in §1.4, I compare reasons constructivism, as I define the view, with reasons objectivism and subjectivism.

My main goal in this chapter is to set the stage for my arguments in the rest of the thesis in favor of a specific version of reasons constructivism. Therefore, I cannot fill in many details of

my view until the arguments have been given. Nevertheless, I shall try to give a brief account of a view that I think is significant, attractive, but frequently misunderstood. I should emphasize that the constructivist view I define here is only constructivism about *normative reasons for action*. In the rest of the thesis, I will explain why I do not think constructivism generalizes to other normative entities and properties, such as reasons for belief (Chapter 2), the rational norms that govern practical reasoning (Chapter 3), and welfare (Chapter 4).

1.1. The Kantian Origin of Reasons Constructivism

It is worth examining first the Kantian origin of reasons constructivism in order to see its basic ideas. Note that I only present the following brief account as a Kant*ian* account, rather than as Kant's own view. I take no position on the vexed question of how to interpret Kant's metaethics.[3] Rather, I will only focus on how meta-ethical constructivists have interpreted certain Kantian ideas.

One idea that motivates reasons constructivism is that we are autonomous agents. Two prominent meta-ethical constructivists, Street (2012, 46; see also 2008, 244) and Korsgaard (1996a, 94), both quote from the following passage in Kant's *Groundwork of the Metaphysics of Morals*:

> […] one cannot possibly think of a reason that would self-consciously receive guidance from any other quarter with regard to its judgments, since the subject would not then attribute the determination of judgment to his reason, but to an impulse. Reason must view herself as the authoress of her principles, independently of alien influences […] (Kant 2012, 4:448)[4]

This Kantian conception of autonomy is defended by Korsgaard in the third lecture of *The Sources of Normativity*. After having considered realism in the first lecture and Humean views in

[3] For a brief discussion of the interpretive options here, see Johnson and Cureton (2016, sec. 14).
[4] Both Korsgaard and Street use a different translation. Their quote ends at "its judgments."

the second lecture, Korsgaard puts forward the Kantian claim that "autonomy is the source of obligation" or normativity more generally (91). This is so because the human mind is essentially reflective: as agents, we can question whether any external object or subjective desire gives us a *reason* to act (93). No normative reason from the outside (including from our desires) can be imposed on our will. Rather, we must live by the laws that the will makes for itself (98). Or as Kant puts it in the quote above, "[r]eason must view herself as the authoress of her principles, independently of alien influences."

While my presentation of the Kantian view is fairly basic, it serves as a good starting point. In my view, the Kantian conception of autonomy above suggests a rather appealing view about practical reasoning. The view, stated negatively, is that reasons for action cannot be imposed from the outside, and practical reasoning cannot be a matter of *responding to* objective reasons for action. The positive claim is that reasons for action only exist from the perspective of the agent who has these reasons. Let me discuss the negative claim first.

Many reasons objectivists claim that the aim of practical reasoning is to discover and respond to objective reasons. Moreover, rational agency is just the capacity to act in response to these reasons. I shall call this the "response model" of practical rationality. Joseph Raz, for example, writes that "rationality is the ability to realize the normative significance of the normative features of the world, and the ability to respond accordingly" (1999, 68). With respect to reasoning in particular, he writes that "the power of reasoning is essentially a power whose purpose is to identify and respond to reasons" (2011, 89). Scanlon claims that we rational agents "discover" normative reasons that are "independent of us" (2014, 70–71), and we then "respond […] to [our] judgments about reasons" (55). Jonathan Dancy argues that practical reasoning is a

matter of responding to "[f]eatures of our situation [that] favour or call for certain responses" (2018, 29). And finally, John Skorupski clarifies the response model in the following way:

> By 'responses' I do not mean responses to stimuli, but responses to what the actor takes to be reasons. Responses are 'intentional acts', in the broad sense that they all have an intentional content that is linked essentially to their answerability to reasons. (2012, 37)

While reasons objectivists disagree amongst themselves about the kind of response that practical reasoning leads to,[5] few of them dispute the response model of practical reasoning itself. The response model is so popular in the current literature on practical reason that Ruth Chang calls it the "orthodox view" of rationality (2020, 297–98).[6]

For constructivists influenced by Kant at least, the response model is too heteronomous. According to reasons objectivists, reasons are *metaphysically prior* to the activity of reasoning. The reasons that exist do not depend on how an agent reasons and determines her will. Rather, the aim of practical reasoning is to respond to the reasons that are imposed on the will from the outside. This, as I have argued, cannot be accepted by constructivists such as Korsgaard. So constructivism is opposed to the response model of practical reasoning.

Now the positive claim. Here it is not obvious what is the best way to develop a constructivist theory of practical reasoning from the Kantian notion of autonomy. Nevertheless, I shall venture a suggestion that I think is congenial to most meta-ethical constructivists. As I understand the view, the central commitment of reasons constructivism is that reasons are *perspective-relative*. Reasons for action do not exist from a "view from nowhere," to adapt Thomas Nagel's memorable phrase for my purposes (1989). Rather, reasons for action only exist

[5] See Dancy (2018, chaps. 8–9) for a good discussion of the main options in this debate.
[6] Chang thinks that the orthodox view treats agents as "*passive* automata" (297), but she still thinks that there are some objective reasons, in addition to the reasons created by our will (2013a). I think her rejection of the orthodox view does not go far enough, since our rationality in relation to these objective reasons is still too passive.

from the perspective of an agent who determines how she should exercise her own agency and live her own life. The aim of practical reasoning, on this view, is not to respond to the reasons that exist independently of perspective-relative reasoning. Rather, the aim of practical reasoning is to determine, from the agent's point of view, how her life in particular should go, and there are only reasons for action within this activity of perspective-relative practical reasoning. This view of practical reasoning is the opposite of the response model.[7]

The idea that reasons are relative to an agent's perspective fits quite well with the Kantian claim (at least as interpreted by constructivists) that the rational will cannot "self-consciously receive guidance from any other quarter with regard to its judgments."[8] If there is any guidance to be found, such guidance can only appear from within the agent's deliberative perspective. Therefore, I think perspective-relativity correctly identifies one of the central concerns of constructivists who reject reasons objectivism because of heteronomy. A suitable definition of reasons constructivism, I think, should give expression to this basic idea in the right way.

1.2. Alternative Definitions of Reasons Constructivism

In this section, I consider two alternative definitions of reasons constructivism put forward by Korsgaard and Street. I think neither definition sufficiently distinguishes reasons constructivism from alternative meta-ethical positions such as reasons objectivism and subjectivism. What it takes to distinguish the view from these alternatives is not the same in the two cases, however.

---

[7] As I shall explain in §2.2, many reasons objectivists believe that facts about reasons are relational facts, and in this weak sense reasons are always relative to some agent. However, for these reasons objectivists, reasons are still not *perspective-relative*, since responding to these relational facts does not require reasoning from the agent's perspective.

[8] However, I do not think Kant's remark here *requires* a constructivist interpretation or my interpretation in particular. Again, my focus is on how constructivists like Korsgaard and Street have interpreted Kant.

As we have seen, reasons constructivists explicitly reject reasons objectivism and the response model. A definition of reasons constructivism must rule out reasons objectivism as an option. But as I shall argue, a suitable definition of reasons constructivism only needs to leave open the possibility of a non-subjectivist version of constructivism.

Korsgaard endorses the following definition of constructivism:

> **Reasons Constructivism (Korsgaard):** A normative reason exists because it is the outcome of some constructive procedure.

Korsgaard also calls constructivism "procedural realism" (1996a, 35), which is the view that "there are answers to moral questions *because* there are correct procedures for arriving at them" (36). This "proceduralist" definition of constructivism is first adopted by Rawls in political philosophy (1980, 516), and it is usually considered the canonical definition of constructivism (Ridge 2012, 138; Enoch 2009, 319–20; Hussain and Shah 2006, 290).

Despite its canonical status, this definition of constructivism suffers from serious problems. Korsgaard's definition does not give an account of where the authority of the constructive procedure comes from (Wallace 2012, 33; Hussain and Shah 2006, 291; see also Darwall, Gibbard, and Railton 1992, 140–41). It seems that Kantian constructivists must believe that the validity of a particular constructive procedure does not depend on an agent's subjective attitudes and judgments. An agent who always rejects the authority of the categorical imperative is still bound by it. Some reasons objectivists, such as Roger Crisp, claim that for Korsgaard, there must be an antecedent objective reason to act in accordance with the outcome of the constructive procedure (2006, 53). Even if there is no such antecedent reason as suggested by Crisp, Korsgaard's definition is still compatible with reasons objectivism. For example, it could be that facts about which constructive procedures to adopt are objective normative facts, and facts about which normative reasons exist are all grounded in these facts about constructive

procedures. If so, facts about normative reasons are ultimately grounded in mind-independent facts, and reasons objectivism is true.

Let me clarify that I am not convinced that Korsgaard's definition *entails* reasons objectivism, as some of the philosophers quoted above believe. As I shall further clarify in §1.4, the existence of certain objective facts about rational norms does not entail the existence of objective normative reasons. The problem is that Korsgaard's definition does not *rule out* reasons objectivism, even though reasons constructivism, as a meta-ethical position, should not be compatible with reasons objectivism.[9] As I emphasized in the last section, reasons constructivism is based on the idea that normative reasons are perspective-relative. However, the existence of constructive procedures does not require reasoning from any particular perspective. In other words, there is nothing about the notion of a constructive procedure alone that rules out the existence of heteronomous reasons generated by objective constructive procedures. This is a problem, if Korsgaard's definition is meant to characterize what is distinctive about reasons constructivism.

Moreover, the notion of a procedure is not at all essential to the constructivist idea that reasons must be perspective-relative (see also Street 2010, 366). There may be no procedures that can determine which reasons exist for an agent. In fact, it seems to me that practical reasoning is a far more organic activity, governed by a few rational norms that do not by themselves constitute a procedure. Even the categorical imperative should be understood as a *test* for whether a maxim is acceptable, not a *procedure* for generating maxims.[10] Reasons are generated, under the Kantian framework, by the act of willing as constrained by rational norms,

[9] My claim is not that Korsgaard's constructivist *view* is compatible with reasons objectivism, even though some have made this accusation (Hussain and Shah 2006). I am only targeting her proceduralist *definition* here.
[10] Indeed, it is this "testing" model that Korsgaard uses to characterize the categorical imperative (2009, 51).

not by facts about constructive procedures. Therefore, the language of "constructive procedures" is ill-suited for the purpose of defining reasons constructivism.

Consider next Street's "Humean" constructivism. Here is how she defines the view (2010, 366):

> **Reasons Constructivism (Street):** "[T]he fact that *X* is a reason to *Y* for agent *A* is constituted by the fact that the judgment that *X* is a reason to *Y* (for *A*) withstands scrutiny from the standpoint of *A*'s other judgments about reasons." (Street 2008, 223)

This definition of constructivism does away with the notion of a constructive procedure. Instead, according to this definition, a consideration counts as a reason if it can be generated from the agent's subjective judgments about reasons. As Street puts it in multiple places, an agent's normative reasons are ultimately a *function* of her subjective normative judgments (2008, 208, 224, 235, 244). Importantly, she is a *thoroughgoing* constructivist: she thinks that *all* normative facts, including facts about value and rationality,[11] are grounded in the subjective normative judgments of an agent (209). Therefore, her test for whether one judgment withstands scrutiny from another does not appeal to any normative fact that is not grounded in the agent's subjective judgments.

Let me first note that I agree with Street's emphasis on the *standpoint* from which an agent makes normative judgments. Or as Street puts it, normative reasons are constructed from "the standpoint of a being who judges, whether at a reflective or unreflective level, that some things call for, demand, or provide reasons for others" (2010, 366). This is roughly equivalent to my claim that reasons only exist from the perspective of an agent who reasons about what to do.

[11] It should be noted that she thinks that the constitutive standards of valuing (and presumably judgments about reasons) are discovered through descriptive philosophical analysis rather than normative inquiries (2010, 374). These standards, such as the instrumental rule, are not grounded in subjective judgments. I am not entirely convinced that such standards are compatible with Street's overall subjectivism.

However, I think Street's definition of constructivism is too narrow. The idea that normative truth is ultimately grounded entirely in subjective normative judgments is closer to the usual motivation for *subjectivism*, rather than constructivism *per se* (Ridge 2012, 156). For Street, a subjective judgment only fails to generate a reason if it does not withstand scrutiny from another judgment. Street's notion of "withstanding scrutiny" will be discussed further in Chapter 3. In the meantime, I just want to emphasize that this notion is a very thin one, and it does not allow applying more robust rational norms to agents who do not endorse them. But for Kantians like Korsgaard, rational norms such as the categorical imperative are supposed to hold regardless of an agent's contingent subjective judgments.[12] As I shall explain in Chapter 3, I also endorse several norms that are, like the categorical imperative, not grounded in subjective judgments.

While Street can choose to define constructivism in such a way that only subjectivists can count as constructivists, I think there are good reasons to adopt a broader definition. One reason is that by Street's definition, Kantians (especially Kantian constitutivists)[13] who endorse categorical norms cannot be constructivists. Another reason is that reasons subjectivism suffers from many well-known counterexamples. For example, subjectivism has difficulty generating the right reasons for those with fundamentally evil, misguided, or imprudent aims (Cohen 1996, 183–84; Gibbard 1999, 145; Parfit 2011, vol. 1, 73–82). It could be that for these agents, all of their motivations are internally consistent by the standards of subjectivism, but these motivations are nonetheless deeply corrupt in some way. So if all normative facts, including facts about rational norms, are grounded in these agents' subjective judgments, then it seems like subjectivism cannot generate the right reasons. Regardless of whether reasons subjectivists can

[12] While Korsgaard is not a realist about these rational norms, she is not a subjectivist about them either. She thinks that the categorical imperative is rooted in features of rational agency (2009), rather than the subjective judgments of the agent. Thus, her "constitutivism" is not a form of subjectivism.
[13] See footnote 12.

tackle these counterexamples (and for reasons I discuss in §4.1, I am pessimistic), the idea of perspective-relativity alone does not require that we build *all* normative facts from the subjective judgments of agents. Perspective-relativity only requires that we situate reasons for action within the deliberative perspective of the agent who has these reasons, and as I shall argue, this idea can be developed in a way that does not entail subjectivism.

1.3. My Proposal

Here is how I propose to define the view that I shall defend in Chapters 2 and 3:

> **Reasons Constructivism:** There are normative reasons for *x* to φ (or not to φ) in virtue of the fact that these are the considerations that count in favor of (or against) φ-ing from *x*'s perspective, as determined through non-defective practical reasoning from that perspective.

This definition is rather complex, and let me try to explain each part and show why it is necessary.

First, reasons constructivism is a claim about what *grounds* normative reasons. Roughly speaking, the claim is that reasons exist in virtue of their role in non-defective perspective-relative practical reasoning. In contrast, the response model of reasons objectivism claims that what makes a piece of practical reasoning non-defective is that it responds to the objective reasons that exist independently. Therefore, reasons constructivism, as I define the view, rejects the order of explanation according to the response model of reasons objectivism.

Second, according to reasons constructivism, all reasons are *perspective-relative*. For reasons constructivists, there are no reasons from the view from nowhere, and all reasons must be firmly anchored within the deliberative perspectives of rational agents. I think it is not easy to characterize the idea of a deliberative perspective. But as I shall argue in Chapter 2, I think we have a clear enough pre-theoretical grasp of the idea when we reflect on the practice of "if I were

you" reasoning. At the moment, however, let me try to explain the idea that reasons are perspective-relative using Simon Blackburn's terminology of inputs and outputs.[14]

As Blackburn suggests, we can consider an agent as a "device whose function is to take certain inputs and deliver certain outputs" (1998, 5). The inputs, in my view, are descriptive facts about the world, but also certain objective evaluative facts (such as facts about welfare), as I shall argue in Chapter 3. But for our purposes now, it suffices to say that the inputs form a representation of the (perhaps value-laden) world. However, none of these facts alone grounds reasons directly for the agent. For an autonomous agent, facts about the world can only count in favor of or against an action to be performed by her when they are mediated through the agent's perspective. Otherwise, reasons would be imposed from the outside, and this would be heteronomy. Reasons are therefore among the outputs of an agent's perspective, and this is what it means for reasons to be perspective-relative.[15] I will use the language of inputs and outputs again in the rest of the thesis to characterize the relationships between objective facts (including objective evaluative facts) and reasons.

Third, reasons are the considerations that favor or disfavor actions as determined through *non-defective reasoning*. Actually, there are two points here: reasons are connected to reasoning, and the relevant kind of reasoning must be non-defective. There are familiar connections between normative reasons and the activity of reasoning. Many of our desires and impulses turn out to generate no reasons for us, when we reason or deliberate about what to do. Moreover, *normative* reasons in particular are connected to *non-defective* reasoning. A piece of practical

[14] While I think Blackburn's terminology here is helpful, I am not hereby endorsing any substantive meta-ethical claim made by Blackburn.

[15] Actually, it is perhaps more accurate to treat them as *intermediate* outputs. As I characterize normative reasons, there can be many different *pro tanto* reasons pulling in different directions. These reasons have to be combined through further reasoning to reach a conclusion about the agent's all-things-considered reason.

reasoning is non-defective if it does not violate any rational norm that governs the activity of practical reasoning. Setting aside the question of which norms exist for now, we can see that the "reasons" that are determined through particularly bad reasoning cannot count as normative reasons. If an agent who has an all-things-considered reason to do *x* still reasons that she has no reason to take the necessary means *y*, it does not follow that she does not have a reason to do *y*. Therefore, for reasons constructivists at least, the relevant kind of reasoning that can ground normative reasons must be non-defective reasoning.

I think my proposed definition coheres with the Kantian origin of reasons constructivism. Moreover, compared to Korsgaard's and Street's definitions, it does a better job at characterizing the central motivations for reasons constructivism discussed in §1.1. Recall that constructivists like Korsgaard do not accept the response model of practical reasoning, and reasons are perspective-relative for constructivists. If the view as I define it is correct, then the aim of practical reasoning is not to respond to objective reasons for action. Rather, the aim of practical reasoning is to determine, from an agent's particular deliberative perspective, which considerations count as reasons for her. Unlike Korsgaard's proceduralist definition, then, my definition emphasizes perspective-relativity and is not compatible with the view of practical reasoning endorsed by reasons objectivists. Moreover, unlike Street's definition, my proposed definition is also not too narrow. While reasons constructivism as I define it is compatible with reasons subjectivism, it could be that non-defective perspective-relative reasoning is constrained by rational norms that are not grounded in the agent's subjective judgments.

1.4. Reasons Objectivism and Subjectivism

How does reasons constructivism, as I define the position, differ from reasons objectivism and subjectivism? Consider subjectivism first. In order for reasons to be subjective, all the component facts that ground facts about reasons must themselves be acceptable by subjectivist standards. For example, as I have mentioned, a subjectivist like Street cannot accept that there is a mind-independent rational norm such as the categorical imperative, since for her reasons are ultimately entirely grounded in the subjective judgments of the agent. Moreover, there is reason to think that subjectivists must adopt subjectivism about *all* normative facts, in addition to those facts that ground facts about reasons. Some subjectivists, such as Mark Schroeder (2007, chap. 4) and Kate Manne (2013, 67–68; 2016, 124; 2017),[16] favor reducing normative facts to natural facts. Non-natural normative facts are metaphysically "queer," as J. L. Mackie famously argues (1977). For these subjectivists, any kind of normative fact that is mind-independent and non-natural is going to be "queer." Other subjectivists object to the existence of non-natural normative facts on epistemic grounds. For example, Street argues that even if there were such normative facts, we could not know them due to the pervasive contaminating effects of evolutionary forces (2006). Again, Street must endorse subjectivism about all normative facts, not just facts about reasons.

Although reasons constructivism, as I define the position, is compatible with this kind of thoroughgoing subjectivism, the former does not require the latter. According to my definition of reasons constructivism, reasons are grounded in perspective-relative practical reasoning that is constrained by rational norms, but rational norms may well be mind-independent normative entities. In fact, as I shall argue in §3.1, I think the norm of analogical consistency is not a norm that is grounded in the subjective judgments of agents. Furthermore, I think facts about welfare

[16] I shall count Manne's "democratic Humeanism" as a version of subjectivism (2016; 2017), even though for her an agent's reasons may be grounded in the desires of *other* agents.

are mind-independent facts that serve as inputs to practical reasoning. Here again, reasons constructivism is not in conflict with welfare objectivism, since my definition of reasons constructivism does not make any general metaphysical claim about all normative facts and entities.

But consider now the following objection. If my definition of reasons constructivism allows the existence of objective normative facts, isn't reasons constructivism just a version of reasons objectivism? The short answer is that for reasons constructivists, there are no objective *reasons for action*, even though reasons constructivists can be objectivists about other types of normative facts or entities. But let me develop this answer a bit more carefully.

To begin with, and this is a point I shall expand on in the next chapter, if there are objective reasons, then it cannot be the case that reasons are indexed to any particular perspective and only be "visible" from that perspective. Facts about normative reasons, like mathematical facts, should be available from any perspective and thus no perspective in particular. Reasons constructivists reject this, and thus reasons constructivists reject the existence of objective reasons.

Moreover, as I have already emphasized in the last section, reasons constructivism reverses the reasons objectivist's order of explanation. For the reasons objectivist, facts about reasons are not grounded in non-defective practical reasoning. In fact, objectivists like Parfit (2011, vol. 1, 62) claim that what makes a piece of reasoning non-defective is just the fact that it correctly responds to the right normative reasons. In contrast, for reasons constructivism, normative reasons exist in virtue of non-defective perspective-relative practical reasoning.

But it might be objected that if there are objective facts about rational norms, the violation of which makes reasoning defective, then these facts would ground reasons, and

reasons would be grounded purely in mind-independent facts. As I have mentioned, this is one of the objections against Korsgaard's proceduralist constructivism. But my definition of reasons constructivism does not entail that reasons are grounded in objective normative facts. Reasons constructivism, as I understand it, does not reduce reasons to other normative facts. It only claims that reasons exist in virtue of the *activity* of practical reasoning as performed by agents. The activity itself is not reducible to the norms that constrain it. To use a legal analogy, even though the constitution of a legal system supplies the framework for constructing law, the content of a legal system is not reducible to the content of the constitution. Clearly, it is the reasoning of judges and legislators that constructs the content of a legal system. I think likewise it is the activity of practical reasoning that constructs the reasons that apply to a specific agent. It is this activity that determines and grounds the normative reasons that an agent has, even though the activity is governed by certain objective rational norms.

There may be another worry here. If there are objective facts about rational norms, does it not follow that there are objective normative reasons to be, for example, analogically rational? The first thing I want to say is that I do not think the normativity of rational norms is grounded in an antecedent *reason* to be rational.[17] Nonetheless, it seems perfectly fine to say that "*x* has a reason to be rational," and reasons constructivists who endorse objectivism about rational norms need to explain this. There are two possible options available to reasons constructivists here. The first option is to claim that reasons to be analogically rational are not reasons for action. As I shall argue in Chapter 2, my argument for reasons constructivism targets explicitly reasons for action, and thus it does not apply to these reasons to be rational. I think constructivists about reasons for action can (and perhaps should) be objectivists about reasons for belief, for example,

[17] I am tentatively attracted to the idea that these norms are constitutive of the activity of practical reasoning, but I do not want to commit to this view here.

and similarly there may be objective reasons to be rational. The second option here is to claim that reasons to be analogically rational are only conditional reasons for action. They only apply to agents who have an antecedent reason to determine their reasons for action through practical reasoning in a given context. In split second decisions, perhaps an agent has no reason to engage in practical reasoning, and she should just follow her gut and hope that her action conforms to what she has most reason to do.[18] I am not committed to either option, although I slightly prefer the second one, since it fits better with my claim that practical reasoning is an activity.

In summary, my definition of reasons constructivism distinguishes the view from both reasons objectivism and reasons subjectivism. While reasons constructivism is compatible with objectivism about certain normative facts, it explicitly rejects the existence of objective reasons. While reasons constructivism is compatible with reasons subjectivism, the former position does not entail the latter.

[18] In such a case, it is still true that what she has most reason to do is determined by non-defective perspective-relative reasoning. However, she does not have the time to do this.

Chapter 2 The "If I Were You" Argument: Reasons Objectivism and Perspective-Relativity

I argued in Chapter 1 that reasons constructivism is best understood as the following claim:

> **Reasons Constructivism:** There are normative reasons for *x* to φ (or not to φ) in virtue of the fact that these are the considerations that count in favor of (or against) φ-ing from *x*'s perspective, as determined through non-defective practical reasoning from that perspective.

Reasons constructivism is opposed to reasons objectivism. Reasons objectivists claim that there are mind-independent (or objective) reasons, and we respond to them in practical reasoning. Reasons constructivists reject this picture of practical reasoning. I also argued that reasons constructivism does not entail reasons subjectivism, or the view that normative reasons are entirely grounded in an agent's subjective judgments and attitudes. This is because non-defective practical reasoning from the agent's perspective may be constrained by mind-independent rational norms (such as the categorical imperative). Christine Korsgaard seems to hold such a view, as I interpret her Kantian constructivism. In such cases, reasons as determined by such reasoning are not entirely grounded in mind-dependent facts, but they are still perspective-relative.

In this and the next chapter, I give an argument in favor of reasons constructivism. The argument is an inference to the best explanation of an interpersonal normative phenomenon which, following Bernard Williams, I call the "if I were you" mode of practical reasoning (1995, 36). I argue in this chapter that the necessity of "if I were you" reasoning in interpersonal contexts shows that normative reasons for action must be determined from the perspective of the agent who has the reasons. In other words, reasons for action are perspective-relative. In the next chapter, I will refine my constructivist theory of reasons on the basis of "if I were you" considerations.

I will argue for the perspective-relativity of reasons for action in this chapter by using reasons objectivism as a foil. First, in §2.1, I describe the "if I were you" mode of interpersonal reasoning. I argue that when I reason with you about what you should do, I must either implicitly or explicitly reason from your perspective in the "if I were you" mode. Otherwise, my attempt at interpersonal reasoning would be "infelicitous" (to adapt a term from speech act theory), and I would violate a condition that governs the activity of interpersonal reasoning. I then argue in §2.2 that if reasons objectivism were true, there would be no good explanation for why the "if I were you" mode of reasoning is mandatory in interpersonal practical reasoning. Moreover, I argue that reasons constructivism, in comparison, offers a much better explanation.

There are two reasons to take the argument I give here seriously. The first reason is that, while such "interpersonal" arguments are often made in support of some kind of internalism or constructivism (Williams 1981, 106; 1995, 39; Manne 2014; Korsgaard 1996a, 38; Smith 1995a, 110; 1995b, 151), they are often underdeveloped.[19] This is a shame for this second reason: it seems to me that the kind of interpersonal phenomenon that interests me here gets to the heart of a common source of dissatisfaction with reasons objectivism among Humeans and Kantians. My plan, then, is to develop in some detail an argument for constructivism that I think is dialectically significant and philosophically interesting.

2.1. The "If I Were You" Mode of Interpersonal Reasoning

[19] Manne's defense of Humean internalism (2014) is a notable exception, although she focuses on different aspects of interpersonal reasoning. Another notable critical discussion of the "if I were you" argument for Humean internalism is given by Samuel Asarnow (2019). As I shall discuss in §3.1, I am sympathetic with some of Asarnow's objections. As the citations above reveal, when such arguments are given, they are often developed quickly within single pages.

In this section, I describe what Williams calls the "if I were you" mode of interpersonal practical reasoning (1995, 36). I argue that when I reason with you about what you should do, I can only determine your reasons for action if I reason from your perspective in the "if I were you" mode of reasoning.

To see the kind of normative phenomenon that I have in mind, consider first Williams's Owen Wingrave example, based on Henry James's eponymous story:

> **Owen Wingrave:** Owen Wingrave is a military student. Family: "You will stay in the military for the sake of family honor. You are a Wingrave and a soldier." Owen: "No, the profession of arms is beneath me." (Williams 1981, 106; James 1893, 167)[20]

In the original story, James emphasizes that Owen is a thoughtful person who has his own ideas about how his life should go. From his perspective, we can imagine as readers, there is no place in his life for a military career. Unfortunately, Owen's family consistently disregard Owen's independence of thought. They only care about what they want, which is the continuation of the family military tradition. They tell Owen that, as a member of the Wingrave family, he must do what it takes to preserve family honor. They are *imposing* the family tradition on Owen.

Focus on *how* Owen's family interact with Owen. For the moment, let's set aside questions about whether the *consideration* of family honor is suitable as a reason for action. Let's set aside whether the consideration links up with any of Owen's motivations (Williams 1981, 106), or whether it still counts in favor, perhaps to a small degree, of staying in the military despite Owen's motivations (Callard 2016, 146). I am interested in how the *activity* of reasoning about what Owen should do governs the *individuals* involved. What explains the intuition that some kind of "browbeating" occurs in this interaction (Millgram 1996, 205)? What explains our unease toward the way Owen's family interact with Owen?

[20] My discussion of the example extends my previous discussions in (Fan 2020a; 2020b).

Here is my first conjecture: what Owen's family try to do with Owen is better described as browbeating, because they are violating some condition that governs the activity of interpersonal reasoning. We think Owen's family are not doing what it takes to reason with an autonomous agent like Owen. Owen is trying to figure out how to live his own life, but his family are telling him that a military career is what the family tradition requires. There is a mismatch between what each side is doing: Owen's family are not engaging in the right kind of activity to help Owen figure out how to live his own life. It seems to me that we can only judge that there is such a mismatch if we judge that Owen's family have violated some *condition* of interpersonal reasoning.

What can this condition be? My second conjecture is that reasoning about what *x* should do must proceed from *x*'s perspective. Or as Williams puts it, interpersonal practical reasoning must implicitly or explicitly adopt the "if I were you" mode of reasoning (1995, 36). I think this second conjecture, if true, best explains our intuitive discomfort with how Owen's family interact with Owen. Owen's family are not thinking from *Owen*'s perspective at all. Their reasoning is entirely situated within the perspective of the Wingrave family. As a result, their reasoning fails to determine how *Owen* should live his own life.

Thus, I claim that we think there is something wrong with how Owen's family interact with Owen, because they are not fulfilling the condition that governs the activity of interpersonal practical reasoning. The condition is that interpersonal practical reasoning with *x* requires adopting *x*'s perspective. This condition is analogous to the felicity conditions that govern speech acts. The speech act of promising with no intention of keeping the promise is infelicitous (Austin 1962, 16). Similarly, adapting Stephen Darwall's terminology (2009, 4), we might say that

claiming that Owen has a reason to stay in the military without reasoning from Owen's perspective is *normatively* infelicitous.

So far I have argued that if the "if I were you" mode of reasoning is mandatory in interpersonal reasoning, then this fact can explain what goes wrong in the interaction between Owen and Owen's family. I want to emphasize how familiar this mode of reasoning is. If I want to help a first-year student decide which classes to take, it is very natural for me to say "if I were you, I would take *X* rather than *Y*." This is a case of explicit "if I were you" reasoning. Perhaps from my own perspective, I would never consider taking a theater class. But this obviously has no direct relevance for what advice I should give the student.[21] A good advisor is someone who can easily switch to the perspective of the student and reason from her perspective. Even though I cannot yet offer an analysis of the "if I were you" mode of reasoning,[22] I think we have an intuitive, pre-theoretical grasp of what it is like to reason from someone's perspective.

Again, my claim so far is not about the *content* of "if I were you" reasoning. For example, I do not suggest that reasoning from the student's perspective is just *instrumental* reasoning about how to achieve her academic goals, or reasoning about what she is motivated to do. So far, I am only claiming that the *mode* of "if I were you" reasoning is necessary in such interpersonal contexts. The necessity of "if I were you" reasoning can explain why Owen's family are only browbeating Owen, rather than reasoning with him. It can also explain why it is so natural to explicitly state that we are reasoning from the other agent's perspective when we give advice.

[21] I am considering here the US context, where many undergraduates take a variety of classes in different disciplines.
[22] For reasons I shall discuss in §3.3, I am skeptical about reductive analyses of the "if I were you" mode of reasoning.

From Williams's Owen Wingrave, we are led to the "if I were you" mode of practical reasoning. I want to suggest now that we can reach the same point from Korsgaard's well-known "normative question" argument against what she calls "substantive realism" as well. In other words, I argue that the normative question argument can be framed entirely in the interpersonal terms I give above. Korsgaard in fact imagines a conversation between a realist and someone who has fallen into doubt about whether she really has reason to sacrifice her life in order to do the moral thing (1996a, 38). The person in doubt asks the normative question "why should I be moral?" (Or rather, "why should *I* be moral?") Korsgaard thinks that a realist who only rehearses the first-order criteria that determine which actions are morally required is not answering the right question. The realist, at the end of the day, can only say that the action *just is* normative for the person in doubt. As Scanlon nicely summarizes this argument (even though he is not convinced in the end), the realist's reply to the person is mere "foot-stomping" (2014, 67).

While I think Korsgaard is ultimately correct that the realist cannot answer the normative question without such foot-stomping, many realists are not convinced by the argument as formulated. Derek Parfit, for example, does not see why the realist cannot answer the normative question "why should I be moral?" by pointing out the objective reasons that apply to the person (2011, vol. 2, 420). Here I offer an interpersonal reinterpretation of the normative question argument, and my argument against reasons objectivism proceeds in two stages. In the rest of this section, I argue that Korsgaard's normative question can only be answered by "if I were you" reasoning. In the next section, I argue that if reasons objectivism were true, then "if I were you" reasoning would not be necessary. I think this interpersonal argument captures the spirit of Korsgaard's argument, because foot-stomping, like browbeating, seems to be a violation of a normative felicity condition of interpersonal reasoning.

Imagine a conversation between "Anna" and "Brian," where Brian plays the role that the realist plays in Korsgaard's example. In my example, however, Brian *refuses* to reason from Anna's perspective. Suppose that Anna wants to know why she really should, for example, sacrifice her life in order to save five other people. Given that Brian refuses to reason from Anna's perspective, Brian might point out that Anna should sacrifice her life because this action produces the best outcome impartially considered. Now, it could be that Brian is right about what morality requires, given that impartiality seems to be a key aspect of morality. But it seems to me that it is legitimate for Anna to complain that Brian has not answered her original question. After all, Anna is not trying to figure out what impartial morality requires in this case, but rather why she should govern her own life as morality dictates.

So Anna might ask the question now with a different emphasis: she is asking why *she* should live her own life as impartial morality requires. The urgency of this question is amplified by the fact that she is required to sacrifice her life, although the question should be intelligible even when less is at stake.[23] When the question is posed in this way, it seems to me that Brian has to adopt Anna's perspective and answer her question in the "if I were you" mode. After all, the question only makes sense from Anna's own perspective of planning how her life should go. If Brian still refuses to engage in this kind of reasoning, then he can only emphasize again that sacrificing her life just is what morality requires. Brian can only stomp his foot, and Anna can again complain that he is not answering her question. I hope you agree that Anna's question is a

[23] The point here is not that moral reasons have to be justified on the ground of self-interest, as H. A. Prichard warns against (2002a; 2002b). If Anna wants to donate to a charity that addresses an issue that she cares a lot about, but Brian tells her that impartial morality requires that she donate to a more effective charity that addresses a different issue, it seems to me that Anna can still ask why she should govern her life as morality requires. Anna is not asking whether Brian's answer can be justified on the basis of self-interest, since both courses of action are altruistic.

perfectly intelligible (and perhaps familiar) one, and that Brian fails to answer Anna's question because of his refusal to adopt Anna's perspective.

It may be helpful to compare the situation to interpersonal *theoretical* reasoning. If Chris asks Danielle whether he has reason to believe that climate change is real, of course Danielle can say something like "if I were you, I would believe that…". But this is not so natural in the theoretical case. Most likely, Danielle would answer Chris's question by directly explaining why the proposition "climate change is real" is true, and she can do this without reasoning from Chris's perspective. In general, given that belief is governed by truth (Velleman 2000; Shah 2003), an explanation for why *p* is true is also an explanation for why someone has reason to believe that *p*. It is not intelligible for Chris to ask, "I understand why *p* is true, but why do I have reason to believe that *p*?"[24] Therefore, it is not intelligible for Chris to complain that Danielle fails to answer his original question, even though Danielle does not reason from his perspective.

Thus, there are key differences between interpersonal practical and theoretical reasoning. In interpersonal practical reasoning, an agent asks why she should live her own life in a certain way when she asks what reason she has to φ. In interpersonal theoretical reasoning, an individual asks why some proposition *p* is true when she asks why she has reason to believe that *p*. Practical decisions are importantly personal in a way that theoretical inquiries are often not.[25]

It is worth emphasizing that the claim that Korsgaard and I want to make is not that moral requirements are not binding for someone who is skeptical about morality. Nor is it the claim

[24] But see Street (2009a) for an attack against objective epistemic reasons. While I think there are objective epistemic reasons, here I just need to compare practical reasons with an objectivist *conception* of epistemic reasons.
[25] This point is also emphasized by Williams (1985a, 67–69; 1985b). His mistake, I think, is to conclude that reasons for action must therefore be constrained by the agent's motivations. This Humean claim can be separated from the claim about the perspective-relativity of practical reasoning, as I shall argue in §3.1.

that morality, like one's family military tradition, only generates reasons if the agent has a relevant desire.[26] Moral reasons may well be categorical, if they are conclusions that can be drawn from every deliberative perspective.[27] The point is that as long as the realist in Korsgaard's example does not reason from the perspective of the person in doubt, the realist's attempt at interpersonal reasoning is still infelicitous. The point is again not just about *which* considerations are put forward in interpersonal contexts, but about *how* they are put forward. It is by engaging in reasoning from an agent's perspective that any reason offered as a result can have normative "grip" on the agent (Korsgaard 2008b, 53–54, 56).

I think the fact that Korsgaard's argument can be expressed in interpersonal terms is notable, because Korsgaard's basic point should not be expressed in terms of motivation at all.[28] Although she often frames the normative question in terms of how reasons can motivate rational agents (2008a, 315), she notes in "The Normativity of Instrumental Reason" that she no longer thinks that her argument should be understood in motivational terms (2008b, 54). The basic point, I believe, is that reasons must have the right kind of *rational authority* over autonomous agents.[29] The rational authority that needs to be explained is distinctively practical, and it is not merely a matter of motivational efficacy. I shall come back to this point in §3.1. As I shall argue

[26] Cf. Foot (1972).
[27] I wish to develop an account of moral reasons in future work.
[28] Here I differ from Jamie Dreier's interpretation of Korsgaard's "normative question" (2015), which centers on connections between reasons and motivation. My interpretation of the normative question argument also differs from Shamik Dasgupta's "external normative argument" (2017), which claims that objectivists fail to explain how any purported *sui generis* action-guiding property can be action-guiding (n. 11). Markovits also gives an interpersonal argument against reasons objectivism. She writes that "[f]aced with such insistence [that there is an objective reason to be moral], the amoralist may well respond, '*whose* reasons? *Which* rationality?' An assertion that rules of some old-fashioned religion—a religion that in no way reflects what I care about—belong to the requirements of (substantive) rationality and are reason-giving for me would do little to convince me of the normative import of those rules. What makes the externalist response to the amoralist any better?" (2014, 68). This is notable, because Markovits's internalism is also not expressed in motivational terms.
[29] Scanlon also interprets Korsgaard in this way (2014, 10).

there, I think this point recommends reasons constructivism, rather than the kind of Humean internalism defended by Bernard Williams and Kate Manne.

In summary, I have argued that in each of the examples of Owen Wingrave, academic advising, and the conversation between Anna and Brian, an agent can only determine another agent's reasons for action by engaging in the "if I were you" mode of reasoning. As I shall argue in the next section, I think this normative phenomenon suggests that normative reasons are essentially perspective-relative, as suggested by reasons constructivism.

2.2. Reasons Objectivism

Can reasons objectivists explain why the "if I were you" mode of practical reasoning is necessary in interpersonal contexts? I shall argue that they cannot. I shall primarily focus on "reasons fundamentalists" such as Parfit and Scanlon, although as I shall suggest, the problem seems more serious for objectivists who think that reasons are grounded in objective values.

According to reasons fundamentalists, reasons are the irreducible, fundamental normative entities of the normative domain. Scanlon argues that values pass the buck to normative reasons, and the normativity of reasons is just an objective and irreducible fact about them (1998, chaps. 1–2). Furthermore, he claims that "'is a reason for' is a four-place relation, R(*p, x, c, a*), holding between a fact *p*, an agent *x*, a set of conditions *c*, and an action or attitude *a*" (2014, 31).[30] All claims about practical normative reasons together form the "normative domain," as he calls it. The normative domain is just like other theoretical domains, such as the mathematical domain, in that truths within each domain are discovered by the domain-specific type of reasoning (19–20).

[30] John Skorupski thinks that R is a five-place relation between "facts; times; degrees of strength of reason; actors; and acts" (2012, 37).

Just as mathematical reasoning responds to mathematical truths, normative reasoning responds to normative truths.

Reasons fundamentalists such as Scanlon, then, believe that to reason about what Owen has reason to do is just to reason about all the normative reason relations where the relatum *x* is Owen Wingrave. Or in other words, reasoning about what Owen should do is just reasoning about normative claims *about* Owen. One mathematical analogue is that doing group theory is just reasoning about mathematical claims *about* groups. I think this analogy is not unfair to Scanlon, who emphasizes the similarities between practical and mathematical reasoning.

If the "if I were you" mode of reasoning really is mandatory in interpersonal reasoning, then I think the above comparison suggests a problem for reasons fundamentalism. No mathematician thinks that we can only do group theory by reasoning from the perspectives of groups. If moral reasoning is just like mathematical reasoning, except that moral reasoning is about agents rather than mathematical objects like groups, then it seems like we can identify Owen's reasons directly without reasoning from Owen's perspective. We just need to respond to those reason relations in which Owen is one of the relata. If this is the case, why is it necessary to adopt the "if I were you" mode of reasoning in interpersonal practical reasoning? After all, interpersonal theoretical reasoning does not require this. If Scanlon is right about the normative domain, then there should be no differences between the modes of interpersonal practical and theoretical reasoning.

Is it perhaps possible that practical reasons are those mind-independent truths that are only discernible when we reason from the perspective of the agent who has those reasons? Can reasons fundamentalists adopt such an *epistemic* explanation of the necessity of "if I were you" reasoning? Again, it is not clear how such an explanation might work. If there is a normative

domain of all normative reasons, why can't we directly track these reasons in interpersonal practical reasoning as we would in interpersonal theoretical reasoning? It is true that in many theoretical domains such as science, we non-experts cannot always track the truths directly. But in such cases we can only *trust* the opinions of those who have better access. We can't overcome our own epistemic limits just by reasoning from the perspectives of the experts. So there is no general epistemic analogue of the "if I were you" mode of reasoning.

The explanatory challenge seems even harder to meet for those who think that objective values ground objective normative reasons for action.[31] If there are objective values that generate practical reasons directly, why does it matter which perspective we reason from? If Owen's reason is directly determined by the value of the military career alone, then reasoning about what Owen should do only needs to respond to that value. But as I have argued, it would be objectionable browbeating for Owen's family to then claim that Owen has a reason without adopting his perspective. It is open for Owen to complain that his family have not answered the important question, which is whether the military career *belongs in* his life.[32]

In contrast, reasons constructivism can easily explain why the "if I were you" mode of reasoning is necessary. If it is in the nature of reasons for action that your reasons can only be determined by practical reasoning from your perspective, then it follows directly that I am required in interpersonal reasoning to reason from your perspective. Otherwise, I cannot determine what reasons you have.

---

[31] Barry Maguire (2016) offers the clearest statement of the kind of value-based reasons objectivism I discuss here, and Roger Crisp (2006) endorses a similar view. While Joseph Raz (1999; 2011) is commonly considered a value-based reasons objectivist, as I shall suggest in footnote 72, it is far from clear that this is the case.

[32] As I will argue in Chapter 4, I do think there are objective facts about welfare. Unlike value-based reasons objectivists, I do not think these facts alone settle the question of what reasons for action Owen has.

I shall now consider two replies (as I understand them) that Scanlon offers to such "interpersonal" arguments against reasons objectivism. I shall consider in particular whether his replies undermine my "if I were you" argument.

The first reply can be stated roughly as follows. Perhaps the necessity of "if I were you" reasoning is not to be explained by the nature of normative reasons, but by some specific features about the interpersonal situation that the agents are in. As Scanlon puts it, "what it takes for a claim to be correct need not be the same as what it takes for the claim to be one that one's opponent in argument cannot consistently deny" (2014, 13). The question about the nature of normative reasons is only a question about "what it takes for a claim to be correct." If so, Scanlon would probably say, against my arguments, that interpersonal considerations about what "one's opponent in argument [can] consistently deny" cannot reveal anything about the nature of normative reasons.

Let me call this activity of showing a claim in a way that cannot be consistently denied "interpersonal persuasion." His reply, in a nutshell, is that even if interpersonal persuasion requires the "if I were you" mode of persuasion, it does not follow that normative reasons are perspective-relative. There can still be objective reasons for action that apply to you even if I cannot persuade you in a way that you cannot consistently deny.

However, I do not think this reply succeeds, because the activity of interpersonal reasoning, as I have described, is not the same as interpersonal persuasion. The two activities have different *aims*. Interpersonal reasoning aims to answer an agent's question about what reasons she *really* has. Interpersonal reasoning does not require in addition that the person be persuaded in a way that she cannot consistently deny. To see how the two activities come apart, consider interpersonal theoretical reasoning first. Danielle tries to explain to Chris why climate

change is real, in order to explain why he has reason to believe that it is real. It could be that given how limited Chris's knowledge is about the issue, it would be consistent for Chris to deny what Danielle is claiming. In such a case, Danielle has failed to persuade Chris. But crucially, she has not failed in interpersonal reasoning, since she has indeed answered Chris's question about what he has reason to believe.

Similarly, in interpersonal practical reasoning, Frank might reason with Elizabeth about whether she has reason to do something about climate change. If my arguments in the last section are sound, Frank can only reason to this conclusion by reasoning from Elizabeth's perspective. However, this piece of practical reasoning depends on certain descriptive facts (a point that I come back to in the §3.1), such as the fact that climate change is real. Again, it could be the case that Elizabeth can consistently deny the proposition that "climate change is real," given her lack of knowledge about the matter. If so, Elizabeth can also consistently deny the practical conclusion that she has reason to do something about climate change. But again, crucially, Frank has not failed in interpersonal reasoning, since he has indeed answered Elizabeth's question about what she has reason to do.

Unlike interpersonal persuasion, interpersonal reasoning does not fail just because the individual is not receptive to the conclusion of such reasoning. The aim of interpersonal reasoning is to determine her reasons in an interpersonal setting. When Anna complains that Brian has not answered her question about what reasons she has to sacrifice her life, she is not just saying that she is not persuaded. Rather, she is claiming that Brian is not answering the right question (or answering it in the right way), and as a result he fails to determine her reasons for action. Given this tight conceptual connection between interpersonal reasoning and normative reasons, it is hard to see how the necessity of "if I were you" reasoning in the practical case does

not reveal anything about the nature of practical normative reasons. And as I have argued, if reasons for action are mind-independent, rather than perspective-relative, we have no explanation for the necessity of "if I were you" reasoning.

The second reply that Scanlon gives, as I reconstruct it, is an objection against the very idea that interpersonal practical reasoning can have some *sui generis* mode. Constructivists, according to Scanlon, believe that reasons given in interpersonal reasoning "must be grounded in something that is already true of that agent" (2014, 11). However, when an agent reasons about what she herself should do, "the question of how reasons 'get a grip on one' properly disappears" (14). Scanlon calls interpersonal reasoning "external reasoning," and the kind of reasoning that an agent does for herself "internal reasoning." In internal reasoning at least, an agent is not reasoning about what follows from subjective facts about herself.[33] In fact, from the internal point of view, it seems that the aim of her reasoning is just to discover the objective reasons that apply to her, rather than the reasons that are grounded in facts about herself. How can external reasoning be so different from internal reasoning according to constructivists, if both kinds of reasoning aim to determine an agent's reasons?

I do not think constructivists like myself posit two different forms of reasoning between internal and external reasoning. The claim that Scanlon ascribes to his opponents—that external reasoning "must be grounded in something that is already true of that agent"—is most naturally understood as a claim about the *content* of reasoning. But I have only discussed the *mode* of external reasoning. I do not think in the external case, when Brian, for example, reasons about what reasons Anna has, Brian is merely reasoning about which normative facts are grounded in facts about Anna, such as facts about Anna's subjective motivations. Rather, I think that Brian

[33] In his earlier work, Scanlon makes a similar argument for the claim that desires do not generate reasons for action directly (1998, 43–49). I shall not examine that argument in detail in this paper.

must reason from Anna's perspective, and this form of reasoning is the same as the form of Anna's internal reasoning. Trivially, when Anna reasons about what she has reason to do, she is reasoning from her own perspective.[34] Therefore, I do not posit a distinctive form of reasoning for external reasoning in particular.

Now I happen to think that the perspective-relative nature of reasoning can be seen in internal reasoning as well (Fan 2020a). When I reason about what to do, I am reasoning about what *I* in particular should do, how *I* should live my life, etc. When I reason about what to believe about climate change, I am reasoning about whether this fact about the world is true. But as Scanlon illustrates, the perspective-relativity of practical reasoning tends to get obscured when we only focus on what the right answer is in internal reasoning. While the phenomenology of internal reasoning might be ambiguous, I have appealed to another set of intuitions about the kind of interpersonal reasoning necessary for answering an agent's question about her reasons. I think we have intuitions about whether Owen's family are addressing Owen's question about what reasons *he* has to stay in the military, whether an advisor is addressing a first-year student's question about what reasons *she* has to take a class, and whether Brian is addressing Anna's question about what reasons *she* has to sacrifice her life. It is these intuitions, I claim, that reveal the perspective-relative nature of practical reasoning.[35]

We have thus not only discovered an argument against reasons objectivism, but also an (at least partial) argument for reasons constructivism. As I have argued, if reasons objectivism

---

[34] Those familiar with the "backgrounding desires" (Pettit and Smith 1990) literature might notice some similarities here. Mark Schroeder argues that while normative reasons are grounded in desires, desires are not part of the content of normative reasons (2007, chap. 2). This maneuver is controversial. See Hieronymi (2021, 351). In my case, however, I think it should be less controversial that the *mode* of reasoning cannot be part of the foreground, or the content, of reasoning.

[35] My method here thus has some similarities to Allan Gibbard's method (2003), despite our different conclusions. Gibbard presents G. E. Moore's open question argument (1993) by considering how different agents might express their normative disagreement interpersonally (23–24). In my case, the interpersonal context reveals a gap between conclusions that are not drawn from an agent's perspective, and the agent's reasons for action.

were true, we would have no explanation for the necessity of the "if I were you" mode of reasoning. In contrast, reasons constructivism offers a much better explanation. Recall that, in Chapter 1, I explained reasons constructivism in terms of its three components: (1) reasons exist *in virtue of* a kind of reasoning, (2) the kind of reasoning is *perspective-relative,* (3) the reasoning must also be *non-defective*. The perspective-relativity of reasons (component (2)) naturally explains why the "if I were you" mode of reasoning is necessary. Moreover, the claim that reasons exist in virtue of perspective-relative reasoning (component (1)) explains why there aren't any objective reasons that we can directly respond to in interpersonal reasoning. Reasons constructivism thus offers a metaphysical explanation for the interpersonal phenomenon described in §2.1.

It still remains to be discussed what kind of reasons constructivism we should accept on the basis of the "if I were you" argument. Specifically, in Chapter 3, I shall give an account of the norms that govern the kind of non-defective reasoning (component (3) above) that grounds reasons for action.

Chapter 3 The "If I Were You" Argument: Humean Internalism, Subjectivism, and Constructivism

In Chapter 2, I discussed the *mode* of interpersonal practical reasoning, which is the "if I were you" mode of reasoning. In this chapter, I engage in an "intramural" debate about the *content* of "if I were you" reasoning. More specifically, I discuss what kind of perspective-relative reasoning grounds normative reasons for action. Humean internalists are my primary target in this chapter. Humean internalists, such as Bernard Williams and Kate Manne, agree that interpersonal considerations rule out reasons objectivism, but they have developed very different theories of reasons from mine. In the course of arguing against Humean internalism, I shall further develop my alternative constructivist account of reasons.

In §3.1, I argue that "if I were you" considerations support reasons constructivism, rather than Humean internalism. In §3.2, building on my critique of Humean internalism, I argue that the "if I were you" mode of reasoning does not support Sharon Street's reasons subjectivism either. I conclude in §3.3 that reasons constructivism offers the best explanation for the kind of reasoning that we engage in interpersonal contexts. Furthermore, I clarify the kind of constructivist view that is best supported by the interpersonal considerations already discussed.

3.1. Humean Internalism

I have argued so far that the "if I were you" mode of reasoning is mandatory in interpersonal practical reasoning, and reasons objectivism cannot explain this. The essential idea is that an agent's perspective sets a limit on interpersonal reasoning, and it is the job of a meta-ethical theory of reasons to explain this.

Some philosophers have developed this idea into an argument for a view that I call "Humean internalism." In this section, I argue against this influential line of argument. This argument (and indeed the locution "if I were you" itself) was first presented by Williams, and recently Manne has revived the argument and developed it in more detail. I shall use "Humean internalism" to refer to any view derived from the following intuition:

> **The Core Intuition of Humean Internalism:** Normative reasons for an agent *x* are *constrained* by *x*'s motivational profile.[36]

Humean internalism, especially the version defended by Williams, is perhaps the most influential version of reasons subjectivism in the literature.[37] Both Williams and Manne are Humean internalists who share the above intuition. Williams writes that "[a]n internal reason statement is falsified by the absence of some appropriate element from *S*" (1981, 102), where *S* is the agent's subjective motivational set. Manne similarly writes that "there may be real limitations on what can be said and done, which we might have hoped could be said and done, within the interpersonal activity of reasoning with an agent" (2014, 91), and she goes on to suggest that these limits are set by the motivational propensities of the agent.

To see what they have in mind, consider Williams's following example:

> Suppose, for instance, I think someone (I use 'ought' in an unspecific way here) ought to be nicer to his wife. I say, 'You have a reason to be nicer to her'. He says, 'What reason?' I say, 'Because she is your wife.' He says—and he is a very hard case—'I don't care. Don't you understand? I really do not care.' I try various

[36] As I characterize the view here, some other prominent internalists would not count as Humean internalists. Indeed, while Williams himself is a Humean internalist (as I argue below), his definition of internalism ties reasons to "sound deliberative route[s]" (1995, 35), and a robust conception of sound deliberation can be used to support non-Humean internalisms as well as reasons constructivism. Michael Smith, for example, defends an "anti-Humean" internalist view according to which normative reasons are not relativized to an agent's motivations (1995a; 1995b; 2012). Christine Korsgaard's theory of reasons is also not a version of Humean internalism (1986), and neither is Julia Markovits's (2011; 2014). For Korsgaard, the internalism requirement does not show that the reasons we have are *limited* by our actual motivations (such as our ends or desires). Finally, Markovits's internalist account rejects that reasons for action have to potentially *motivate* the agent who has the reasons (2014, 52). I think her view is closer to the kind of constructivism I defend here, rather than Humean internalism.

[37] Some subjectivists, however, do not see their views as versions of Humean internalism. See, for example, David Sobel (2016b) and Mark Schroeder (2007, 7).

> things on him, and try to involve him in this business; and I find that he really is a hard case: there is nothing in his motivational set that gives him a reason to be nicer to his wife as things are. (Williams 1995, 39)

Williams concludes that in this case the man simply does not have a reason to be nicer to his wife. Williams thinks that while we can call him "inconsiderate," "ungrateful," or "brutal" (39), we cannot say that he has a reason to improve his behavior.

From a very similar starting point to mine, Manne reaches the same conclusion as Williams does. Her argument for Humean internalism is based only on the conditions that govern interpersonal reasoning. With respect to the husband in the example above, Manne thinks that if there is simply no way to get him to be motivated, no matter what we try, then "something has now changed in the normative and dialogical space between us" (2014, 103). As long as we are still trying to reason with him, we must aim to motivate him and get him to act out of his own recognition that there is a reason (107–8). Once we recognize, however, that this is not possible with him, we can no longer reason with him. We can at most treat him as some kind of "human object to be managed, cured, or navigated around," rather than as a "sovereign creature" over whom we have "no particular practical (as opposed to epistemic) authority" (95). Thus, she thinks that normative reasons are constrained by the agent's motivations, no matter how bad or cruel they are. As the title of her paper suggests, Manne thinks that Humean internalism is "sad but true."

Here is a quote that nicely summarizes her reasoning, and I will use it as the starting point of my discussion of her view:

> Perhaps the issue is that we ignore someone's receptivity in this sense [i.e. the agent's recognitional capacities and motivational profile] at the expense of no longer envisaging her as a sovereign creature, whose enduring unwillingness to do something is in a certain sense definitive. For, a reason belongs to the practice of calling each other to account. And a call requires the realistic psychological possibility of its being heard and answered, I think. Otherwise it does not address

> the agent as she is but rather says something more impersonal *about* what she is doing—or addresses an idealized counterpart of the person who is not actually on the scene. And we might also point out that we ignore someone's receptivity in this sense at the expense of no longer treating them as someone over whom we have no *prima facie* authority. (Manne 2014, 110)

I am sympathetic to Manne's worry that we might end up "say[ing] something more impersonal *about* what she is doing" if we violate what I have called the "normative felicity conditions" of interpersonal reasoning. As I have emphasized, interpersonal reasoning is not browbeating or foot-stomping. But I do not think that our reasoning must therefore be constrained by the agent's motivational profile, as a result. It seems to me that in cases like Williams's example of the cruel husband, his bad motivations are not "definitive" for whether he really has a reason to treat his wife better.[38] It is in our concept of normativity, it seems to me, that normative reasons should be able to transcend an agent's motivations.[39] How can we, then, avoid both the Scylla of browbeating and thus no longer reasoning with an agent, and the Charybdis of having our interpersonal reasoning unduly constrained by an agent's bad motivations?

My suggestion is to focus instead on how reasons count as *rationally* authoritative (rather than *motivationally* efficacious) from an agent's perspective. In the rest of this section, I consider the "moves" that we can make in "if I were you" reasoning. These moves reveal the "rules" that govern practical reasoning. When we reason from other agents' perspectives, we often engage in a kind of analogical reasoning that extends an agent's prior commitments to new cases, even when the agent lacks the relevant motivation in such new cases. Moreover, I think we often appeal to objective facts, even when the agent has the wrong beliefs about such facts and is not

[38] Whether he in fact has such a reason has to be established by reasoning from his perspective, and this question will be addressed in Chapter 4. My present point is just that his bad motivations are not definitive.
[39] For critiques of the idea that reasons are tied to motivations, see Johnson (1999, 61) and Markovits (2014, chap. 2).

motivated to change her mind. I argue that Humean internalists cannot explain either of these moves that we make in interpersonal reasoning, whereas reasons constructivists can. In fact, I argue that the kind of non-defective perspective-relative reasoning that grounds normative reasons according to constructivism is governed by rules that allow these moves.

Consider analogical reasoning first. Here is a very simplified example.[40] Suppose I know that you, from your past experience facing racial discrimination, judge that you have a weighty reason to fight against racism. However, you see nothing wrong with gender discrimination, and you are not motivated at all to do something about your own sexist beliefs or the sexist beliefs of those around you. Suppose that your motivations toward both sets of issues are relatively stable. Furthermore, while you are motivated and see a reason to do something about racism, you are not motivated and do not see a reason to follow more general principles like equality or fairness. You are only concerned about racism. It still seems like when I reason from your perspective, I can conclude that you have a reason to fight against sexism as well. I can point out to you that the relevant factors in both cases are similar, that the harm that people suffer from gender inequality are analogous (even if not completely the same) to the harm that people suffer from racial inequality. I can point out that your motivations are rationally *inconsistent*, because you do not treat analogous cases alike in your practical reasoning. I am still reasoning from your perspective, even though you are completely unmotivated to do something about sexism, because I am showing to you what reasons you have on the basis of your own rational commitments. I am not reasoning from a "view from nowhere" (Nagel 1989) when I conclude that you in particular have a reason to be against sexism.

[40] It is simplified because the dynamics behind racism and sexism may be very different. For the purpose of getting my point across, however, this example is sufficient.

The norm of consistency that I appeal to here is not just *logical* consistency.[41] It is logically consistent to be against racism but not sexism. The kind of norm I appeal to here is similar to, but significantly weaker than, what Michael Smith (1995b, 159; 1995a, 114–16) and Julia Markovits (2014, 132–34) call "systematic justifiability." Both philosophers believe that from a Kantian internalist (rather than Humean internalist) foundation, we can generate a norm according to which a more unified set of desires (or ends) is *prima facie* better than a set of disconnected desires (or ends). With respect to my example above, both would think that your current attitudes toward racism and sexism fail to be systematically justifiable. I am not convinced that reasoning that fails to follow the norm of systematically justifiability must always be defective.[42] But like systematic justifiability, what I have called the norm of consistency that governs analogical practical reasoning is also a *rational* norm. It governs how prior commitments in one case can generate new rational commitments in analogous cases. It seems to me that we reason in this analogical way all the time, and reasoning that fails to draw the right conclusions in analogous cases is rationally defective.[43]

While the norm just described seems to be one that we actually follow both in our own reasoning and in reasoning from other agents' perspectives, it does not seem to me that Humean internalists can accommodate such a norm. First, it seems like agents are not always motivated to

[41] As I shall discuss in the next section, the similar norms that Street appeals to are conceptual or analytic (2008, 228). The norm of consistency I appeal to here is considerably stronger than hers.
[42] I hope to justify this claim in future work.
[43] In further support of this claim, it seems to me that *legal reasoning* is often governed by analogical norms. When the U.S. Supreme Court ruled in *Obergefell v. Hodges* that restrictions on same-sex marriage were unconstitutional, they relied on the precedent in *Loving v. Virginia*, where the Court ruled that restrictions on interracial marriage were unconstitutional (2015). In addition, David Wong has defended the importance of analogical reasoning in connection with Mencius's theory of practical reasoning (2002; 2020, 115–17). In one excerpt, Mencius urges King Xuan to extend his compassion for an ox that he spares from ritual killing to his own people (*Mengzi* 1A7). Wong argues that Mencius's theory of analogical reasoning involves "inference of particular from particular" (2002, 212), rather than identification of general principles, as in the Kantian and utilitarian traditions. As Wong argues, this model of reasoning receives considerable support from work in psychology about ordinary reasoning (213–214).

be rational in the way I described. An agent might lack the right element in her motivational profile to be receptive to analogical arguments. It is natural to think that if Humean internalism is true, all rational norms that we are allowed to appeal to must gain their credentials from motivational facts about the agent. If so, it is not clear how Humean internalists can accommodate the norm of analogical consistency. Williams claims at one point that we are allowed to correct an agent's factual mistakes and mistakes in instrumental reasoning in the "if I were you" mode, only because the agent is likely to have a general motivation to be factually informed and instrumentally rational (1995, 37). (His point about factual mistakes will be discussed very soon.) But it is hardly obvious that a general motivation to be "analogically rational" exists for every agent, especially given that analogical rationality often requires us to do things that we are completely unmotivated to do.

Second, even if a given agent is generally motivated to be "analogically rational," it is still not true that there is always a "realistic psychological possibility" (as Manne puts it) that she will be motivated by the conclusion of analogical reasoning. Back to the example about anti-racism and anti-sexism: maybe you are motivated to fight against racism and motivated to follow analogical norms in your reasoning, but when you see that these commitments require that you fight against sexism, you suddenly lose the motivation to be committed to this. Manne, I think, believes that your "enduring unwillingness" to be motivated by the conclusion of analogical reasoning should be "definitive." But it just seems counter-intuitive to me to say that I have not shown to you that you have a *reason* to be against sexism, despite your lack of motivation. Note that Williams and Manne cannot appeal to some kind of deeper commitment to equality either. As I constructed the example, you are, at the end of the day, only motivated to be against racism in particular. It seems to me that Humean internalism does not have any resources to generate

reasons that go beyond your particular motivations in this case. As far as fidelity to our practice of interpersonal reasoning is concerned, this counts against Humean internalism as an account of "if I were you" reasoning.

I shall now consider appealing to objective facts in "if I were you" reasoning. Williams believes that when I reason from your perspective, the factual information that I can appeal to is not limited to what you know. Therefore, if you falsely believe that the liquid in the bottle is gin rather than its true content, gasoline, then you do not have a reason to make gin and tonic with this liquid (1981, 102). I think, and I believe that Williams agrees, that if you are not motivated to change your mind about what is in the bottle (perhaps you are very certain about your belief), it is still the case that you have no reason to make a cocktail with this liquid.[44] However, Williams does not think that I can correct your evaluative mistakes in a similar way. Suppose the husband in Williams's example stubbornly believes that his behavior toward his wife is not bad for her at all, and he is not motivated to change his mind about this belief.[45] When we reason with him, we are not allowed to correct his belief according to Williams.

But as Samuel Asarnow argues, it is not clear what rationale there is for correcting only factual mistakes but not mistakes about externalist evaluative facts when we reason from someone's perspective (2019). As Asarnow points out, neither Williams nor Manne denies that there are evaluative facts that do not depend on the agent's motivations or attitudes (51). Here is one such fact: the fact that the cruel husband's behavior is bad for his wife does not depend on

[44] Manne does not seem to believe that I can just correct your mistake if you are completely not receptive to this piece of information (2014, 105, 107). I think this would be an implausible consequence of her view.
[45] Crucially, Williams's example does not have to be interpreted in this way. For example, the husband may be aware that his action is bad for his wife, but he sees no reason to improve his behavior. I explore how to interpret examples like Williams's in the next chapter.

his motivations or attitudes.[46] Both Williams and Manne want to say that his behavior is cruel,[47] even if they think that he has no reason to improve his behavior. Moreover, Williams denies that internalism depends on a distinction between fact and value (1995, 36). So why should we only correct descriptive facts in "if I were you" reasoning? It seems to me that reasoning that employs false premises about evaluative facts is just as *rationally* defective as reasoning that employs false descriptive premises. The cruel husband's reasoning is not any better just because he makes a mistake about his wife's welfare. In fact, it seems like we do correct such evaluative premises in "if I were you" reasoning. In the advising example I give in §2.1, it seems natural for me, when I reason from the student's perspective, to correct her beliefs about whether a given class is objectively good for her academic development. I can do so even though it is a further question whether she has any *reason* to take such a class. As I will argue in §4.3, while it might be good for her to learn Bulgarian history, she might nevertheless have very little or even no reason to do so (Scheffler 2011, 26–27), and this further reason is determined by perspective-relative reasoning. To put this in Simon Blackburn's terminology introduced in Chapter 1, in "if I were you" reasoning, I can directly correct your inputs (both evaluative and descriptive) but not your outputs (normative reasons).

Thus, if Humean internalists think we can always correct descriptive mistakes in an agent's reasoning, we should also be allowed to correct mistakes about objective evaluative

[46] There are technically two ways for this fact to be "externalist." This fact could be a mind-independent (or objective) fact that is not grounded in anyone's motivations or subjective judgments, or this fact could be grounded in someone else's (such as his wife's) motivations or judgments. As I argue in the next chapter, I think the former is the case when it comes to facts about welfare, although Manne probably thinks that the latter is the case. See Manne (2017). In either case, I think Asarnow is right that we should be able to correct the cruel husband's belief in "if I were you" reasoning. I will primarily talk about objective facts, however, given my own view.

[47] For reasons I discuss below, my own view is that only beliefs about welfare should be corrected in such cases, rather than beliefs about which actions are cruel. I am inclined to agree with Scanlon that calling an action cruel is to say that the agent has a reason (1998, 367). Thus, for me at least, facts about which actions are cruel cannot be purely objective facts.

facts. Both kinds of mistakes make one's reasoning defective. But it should now be clear why Humean internalism cannot accommodate a blanket rule for correcting the agent's inputs. Williams, as I have mentioned, thinks that we can correct factual mistakes because of our general motivation to be factually informed. He thinks, quite reasonably, that no such motivation exists with respect to our evaluative mistakes. But such motivation might not even exist with respect to factual mistakes. If you think that you have reason to sacrifice your children because of your mistaken religious views, I can clearly correct your factual mistakes when I reason from your perspective about what you have reason to do, even if you are not motivated at all to change your religious views.[48] This problem is perhaps even more serious when it comes to mistakes about objective evaluative facts. It should now be obvious that, as in the case of analogical reasoning, an agent can lack both the general motivation to correct mistakes about objective descriptive and evaluative facts, and the motivation to endorse a specific objective descriptive or evaluative fact. Since reasons constructivism does not tie reasons to motivations, constructivists can claim that non-defective perspective-relative must only use correct premises about objective facts.

Let me briefly clarify my view here, which I develop further in Chapter 4. I do not think *all* evaluative mistakes should be corrected. Rather, I think "if I were you" reasoning can only appeal to objective evaluative facts that do not entail or ground normative reasons for action. Otherwise, we would be directly correcting some normative reasons without engaging in perspective-relative reasoning. As I shall argue in the next chapter, I think welfare facts are objective, and they do not entail reasons directly. Thus, we should correct mistakes about welfare facts in "if I were you" reasoning.

[48] Williams is therefore too optimistic when he thinks that agents have a general motivation to be rationally informed. Manne agrees (2014, n. 43), but I think she chooses the less plausible option of disallowing even corrections of descriptive mistakes when an agent is not motivated to accept new information (107).

In summary, I have argued that there are certain moves that we are allowed to make when we reason from other agents' perspectives. Furthermore, Humean internalism cannot accommodate these moves. If so, I think we should conclude that Humean internalism does not correctly describe the rules of interpersonal reasoning. I have instead described two constructivist-friendly rules that govern non-defective perspective-relative practical reasoning. If my arguments in the last chapter are sound, this kind of non-defective reasoning grounds normative reasons for action.

3.2. Street's Subjectivism

I shall now briefly discuss Street's subjectivist version of constructivism. Street, in fact, thinks that her view is just a version of Williams's Humean internalism (2008, 244). Indeed, most of the objections above against Humean internalism also apply to her view.

As I mentioned in §1.2, Street holds the view that an agent *x*'s reasons are a *function* of *x*'s subjective normative judgments about her reasons (2008, 208, 224, 235, 244). To identify *x*'s reasons, we only need to check which of *x*'s subjective normative judgments *withstand scrutiny* from the point of view of her other judgments (223). The notion of withstanding scrutiny in Street's account is a rather minimal one. Street claims that one normative judgment can set a standard for another judgment when the agent cannot make the former judgment without the latter, and the force of the "cannot" here is not *rational*, but rather *analytic* or *conceptual* (228). For example, it is in the concept of a reason that one cannot take oneself to have a reason to pursue the end, without taking oneself to have a reason to pursue the necessary means (228).[49] A

[49] Street emphasizes that normative judgments do not merely set *instrumental* standards for other judgments. Here is another example: "if Lois Lane judges that the fact that Superman is back is a reason for celebration, then she thereby legislates a standard according to which she's making a mistake if she thinks that the fact that Clark Kent is back is no reason for celebration" (2008, n. 38).

normative judgment fails to withstand scrutiny from another judgment, when the former judgment is mistaken according to the standards legislated by the latter judgment in combination with non-normative facts (230).

It should be clear that Street's "thin" notion of what it takes for one judgment to withstand scrutiny from another cannot generate the "thick" rational norms that I describe in the last section. Under Street's subjectivism, one judgment cannot set analogical standards for another. Moreover, Street's subjectivism clearly does not allow us to correct an agent's evaluative beliefs directly, if we judge that these beliefs are mistaken by the standards of objective evaluative facts. Like Humean internalism, then, practical reasoning according to Street's subjectivism is also unduly constrained by the agent's subjective judgments. Street is of course very open about this aspect of her view, and she is willing to bite the bullet to a certain extent (2009b). It could be that we have other good philosophical reasons to accept Street's subjectivism (for example, for evolutionary reasons, Street 2006). I just want to emphasize that as far as fidelity to our practice of interpersonal reasoning is concerned, Street's view does not achieve a high score. In interpersonal reasoning, it seems commonplace to employ analogical reasoning and correct evaluative mistakes when we reason from other agents' perspectives.

Indeed, it seems that for reasons very similar to the ones I discussed in the last section, subjectivists in general cannot accept the two norms that I have identified. It seems possible that there can be agents whose subjective judgments cannot ground a general norm to be analogically rational and take into account the right objective facts. Thus, I think it has to be an objective fact that perspective-relative practical reasoning is governed by these two norms. And as I have emphasized in Chapter 1, there is no reason to think that constructivists cannot endorse objective rational norms, as long as these norms describe the rules of perspective-relative reasoning. I do

not impose reasons on you from the outside, if I extend your commitments in an analogically consistent way from your own perspective and correct your mistaken inputs.

3.3. Reasons Constructivism

In the course of examining other theories of normative reasons, especially reasons objectivism, I have also argued for the various components of reasons constructivism. In Chapter 2, I argued that normative reasons for action are essentially perspective-relative. Moreover, there are no objective reasons that we respond to in practical reasoning. Instead, reasons for action are determined through perspective-relative reasoning. In this chapter, I have discussed various proposals for the kind of perspective-relative reasoning that grounds reasons for actions. I have argued that non-defective perspective-relative reasoning, which grounds reasons, is governed by at least two norms: analogical rationality, and correct inputs of objective facts. To use Blackburn's terminology, normative reasons are the outputs of perspective-relative reasoning that is analogically consistent and based on the correct inputs. I have thus given an argument for reasons constructivism, which, as we recall, is the following claim:

> **Reasons Constructivism:** There are normative reasons for *x* to φ (or not to φ) in virtue of the fact that these are the considerations that count in favor of (or against) φ-ing from *x*'s perspective, as determined through non-defective practical reasoning from that perspective.

Unlike reasons objectivism, reasons constructivism explains the necessity of "if I were you" reasoning in interpersonal contexts. And unlike Humean internalism and reasons subjectivism, reasons constructivism does not unduly limit an agent's reasons by her motivations and subjective judgments. I think it is therefore an attractive intermediate position.

One might still wonder what it means to reason from someone's perspective. So far I have only relied on our commonsensical judgments about whether someone is or is not reasoning

in an "if I were you" manner, or putting themselves in the shoes of other agents. For example, I think most people would judge that Owen's family are not doing that when they tell Owen to stay in the military. I suspect that a reductive analysis of this mode of reasoning cannot be given. Williams seems to believe that the "if I were you" mode of reasoning can be analyzed as reasoning that is constrained (in a way specified by Humean internalists) by motivational facts about you. As I have argued, I do not think this analysis is correct. Moreover, while I have given my account of the *content* of perspective-relative reasoning in this chapter, I am not offering a reductive analysis of the *mode* of "if I were you" reasoning. I do not think *how* we engage in interpersonal reasoning can be analyzed in terms of *what* is involved in this kind of reasoning. At the end of the day, I think the perspective-relativity of practical reasoning is just a fundamental feature that cannot be further analyzed. Fortunately, as I argued in §2.1, I think we do have an intuitive, pre-theoretical understanding of what it is to reason from someone's perspective.[50]

I also argued in §1.1 that reasons constructivism is often motivated by a distinctive conception of Kantian autonomy. The "if I were you" argument brings this conception to the fore. The necessity of reasoning from the agent's perspective in order to identify her reasons, as I argued in §2.1, is explained by the fact that reasons must fit in the agent's own life in the right way. Reasons cannot be imposed from the outside, as Owen's family and Brian try to do. Autonomy matters, not because reasons must be grounded entirely in an agent's motivations or subjective judgments. Rather, autonomy matters, because an agent who plans her life has her own perspective from which she engages in practical reasoning. If so, other agents can only offer her reasons when they reason from this perspective.

[50] There are other features of "if I were you" reasoning that I have not discussed. For example, it seems like when I reason from your perspective, I owe a certain kind of deference to your own reasoning. If I engage in analogical reasoning to a conclusion that you do not accept, as long as your alternative reasoning is rationally beyond reproach, then I should defer to your own conclusion.

Chapter 4 Welfare and Reasons Constructivism

In Chapters 2 and 3, I gave what I called the "if I were you" argument in favor of reasons constructivism. The "if I were you" argument gives expression to the core ideas in some prominent Humean and Kantian attacks against reasons objectivism. In particular, I argued that "if I were you" considerations explain what goes wrong in Bernard Williams's Owen Wingrave example, as well as in Christine Korsgaard's imagined conversation between a realist and a person in doubt about the demands of morality. I also argued in §3.1 that we should be allowed to correct an agent's mistakes about objective facts when we reason from her perspective. Following Samuel Asarnow's line of argument, I argued that if there are objective evaluative facts, then we should be allowed to correct mistakes about these facts as well.

Meta-ethical constructivism, however, is often rejected because of what Ruth Chang calls the "right reasons problem" (2013a, 174): it fails to generate the right normative reasons for all agents, especially those with coherent but fundamentally corrupt or imprudent attitudes and judgments.[51] In this chapter, I argue that reasons constructivists should be welfare objectivists, and this combination of meta-ethical positions can plausibly tackle the right reasons problem. I argue that the right reasons problem only seems fatal to constructivism because most constructivists embrace *thoroughgoing meta-ethical constructivism*, or the view that all normative facts are constructed.[52] However, by accepting welfare objectivism, reasons constructivists can correct an agent's mistakes about facts about welfare as well. I argue that this

[51] Chang's objection is actually that constructivists (or what she calls "source internalists") cannot generate the right *all-things-considered* reasons for these agents. However, I shall focus on *pro tanto* reasons in particular, since if my account can generate the right *pro tanto* reasons of the right weight, it should be able to generate the right all-things-considered reasons as well.

[52] I adopt this term from Street (2008). Examples of thoroughgoing meta-ethical constructivists include Korsgaard (1996a; 2013) and Street (2006; 2008).

version of *restricted meta-ethical constructivism* is less vulnerable to the right reasons problem, and it is an attractive alternative to thoroughgoing constructivism.

In §4.1, I explain the differences between restricted and thoroughgoing meta-ethical constructivism. I also explain why the right reasons problem seems intractable for thoroughgoing meta-ethical constructivists, and why I think endorsing welfare objectivism is the right first step to solving the problem. Then I address two sets of worries regarding whether welfare objectivism and reasons constructivism are compatible. In §4.2, I address the worry that since facts about reasons and facts about welfare are both normative, any meta-ethical argument for reasons constructivism would also support welfare constructivism. While this may be true for certain arguments for constructivism (for example, Street 2006), my "if I were you" argument does not support welfare constructivism. In §4.3, I consider the worry that welfare objectivism entails a version of value-based reasons objectivism. If facts about welfare either conceptually entail, or metaphysically ground, facts about normative reasons, then it seems that welfare objectivism entails the existence of objective reasons. However, I argue that we have reasons to doubt the antecedent of this conditional. After addressing these two sets of worries and showing that restricted meta-ethical constructivism is a coherent position, I argue in §4.4 that my view has the right resources to tackle the right reasons problem. While for certain normative reasons (such as moral reasons) I can only sketch an account how they can be generated, my main claim is that compared to thoroughgoing constructivism, my view is far more promising as a solution to the right reasons problem.

My overall dialectical aim is twofold. With respect to fellow reasons constructivists, I want to show that restricted meta-ethical constructivism both has the resources to tackle the right

reasons problem, and preserves what makes reasons constructivism an attractive view.[53] Against the opponents of reasons constructivism, I want to show that the right reasons problem is only devastating for thoroughgoing constructivism, rather than reasons constructivism itself. Given that the right reasons problem is one of the main objections against constructivism in the literature, this conclusion is a significant one for the viability of reasons constructivism *per se*.

4.1. Constructivism and the Right Reasons Problem

According to Sharon Street, there are two kinds of constructivism: *restricted constructivism in ethics*, and *meta-ethical constructivism* (2008). According to restricted constructivism in ethics, a subset of normative facts are constructed via a justifying procedure, but the materials of construction may include other normative facts (209–10). For example, Rawls's Kantian constructivism constructs reasons of justice from reasons to adopt a given political arrangement from the original position (1971; 1980). Scanlon's contractualism, on the other hand, constructs reasons about what we owe to each other from reasons to reject any given moral principle (1998).[54] These forms of restricted constructivism are not *meta-ethical* positions, since Rawls and Scanlon adopt constructivism to solve first-order normative problems. In contrast, meta-ethical constructivism, in Street's view, is *thoroughgoing*: all normative facts are constructed,

[53] It is notable that the Kantian internalist Julia Markovits also considers the possibility of combining value externalism with reasons internalism in a footnote of her book (2014, chap. 6, n. 5). She motivates value externalism on the ground that there might be agents (such as the severely depressed Maria von Herbert, as considered by Rae Langton 2007) who lack the right attitudes and judgments to ground certain values. While there are many significant differences between her proposal and mine, I suspect that the kind of welfare objectivism I defend here is congenial to many philosophers who have constructivist and internalist intuitions about reasons in particular. I want to thank Julia Markovits for directing my attention to this helpful footnote in her book.

[54] For Scanlon's discussion of the constructivism that underlies his and Rawls's views, see 2014, 96–97.

and no objective normative fact counts among the materials of construction (2008, 223).[55] This kind of thoroughgoing constructivism, as I argued in Chapter 1, is just a version of subjectivism.

This taxonomy is rather curious. Street implies that any restricted version of constructivism can only be a first-order normative theory, rather than a meta-ethical one. Why does Street think that meta-ethical constructivists must also be thoroughgoing constructivists? This is because, as I argued in §1.4, Street's constructivism is based on her general meta-ethical anti-realism. However, recall my own definition of reasons constructivism:

> **Reasons Constructivism:** There are normative reasons for *x* to φ (or not to φ) in virtue of the fact that these are the considerations that count in favor of (or against) φ-ing from *x*'s perspective, as determined through non-defective practical reasoning from that perspective.

This kind of reasons constructivism can be combined with, for example, welfare objectivism, as long as facts about welfare do not exist because they are the conclusions of perspective-relative practical reasoning about welfare. This combination *is* a version of restricted *meta-ethical* constructivism. According to my view, on a fundamental meta-ethical level, facts about normative reasons are constructed, even though facts about welfare are not. This is important, since those (like myself) who are attracted to constructivism for meta-ethical reasons but find thoroughgoing constructivism unappealing may wish to consider this alternative instead.

Therefore, I argue that we should take seriously the possibility of *restricted meta-ethical constructivism*. According to the version of restricted meta-ethical constructivism that I intend to defend, normative reasons exist because they are the conclusions of perspective-relative practical reasoning, but non-defective practical reasoning of this kind takes into account objective facts about welfare.

---

[55] Technically, Street's definition of thoroughgoing constructivism concerns only normative reasons, but I think it is clear that her constructivism applies to *value* as well (2008, 207; 2006, 110).

I find restricted meta-ethical constructivism attractive because it has more resources to tackle what Chang calls the "right reasons problem" (2013a, 174). According to Chang, the purely formal rational norms that constructivists can adopt are too weak to generate the normative reasons that we think apply to agents,[56] especially those agents who have attitudes that are fundamentally "bad" or "wrong" as judged by our commonsense intuitions.

There are a number of examples in which constructivism is thought to have difficulty generating the right normative reasons. In the last chapter, I have already considered the example of the husband who is nasty and mean toward his wife. The husband may be formally rational in his attitudes, but he does not judge that he has a reason to treat his wife better.

In addition, an idealized Mafioso can count as formally rational by the standards of constructivism and still choose to live by a code of honor that requires him to maim and kill others (Cohen 1996, 183–84). An idealized anorexic might "prefer[] […] starvation to life with a figure plump enough to sustain life" (Gibbard 1990, 166). Lastly, an idealized Caligula might "aim[] solely to maximize the suffering of others" (Gibbard 1999, 145). All of these agents are stipulated to be formally rational.

Lastly, Parfit thinks that subjectivists in general (a group that includes Korsgaard by his classification) cannot even generate the reason to avoid future agony. According to Parfit, if the agent currently discounts the badness of future agony simply because the agony still seems quite far in the future, then formal rationality alone cannot guarantee that the agent would recognize the reason to avoid future agony. He calls this the "agony argument" (Parfit 2011, vol. 1, 73–82).

In each example above, it is claimed that constructivists cannot account for the right normative reasons. I want to argue that the problem is especially serious for thoroughgoing meta-

[56] There is no consensus on what makes a rational norm "formal," but I think one basic requirement is that any criterion for formal rationality cannot already appeal to objective normative reasons. See Parfit (2011, vol. 1, 62).

ethical constructivists, who cannot appeal to any normative fact that is not constructed from the agent's subjective judgments. Moreover, I think if we understand why the problem seems intractable for thoroughgoing constructivists, we will see that endorsing welfare objectivism is an attractive first step to solving the problem.

The first thing to notice is that all examples above involve welfare in some way.[57] In some of these examples, such as the idealized anorexic and the agent in Parfit's example, it is the agent's own welfare that is relevant. In the other examples, such as the husband, the idealized Mafioso, and the idealized Caligula, it is other people's welfare that is relevant for practical reasoning. It is natural, then, to wonder whether the right reasons problem reveals some special difficulty with facts about welfare in particular. And if there is a special difficulty here (as I shall argue there is), then a natural strategy to approach the right reasons problem would be to abandon welfare constructivism.

Second, it seems that the examples considered above can be divided into two categories. There are those examples in which the agent does not use the right welfare facts in her reasoning. The idealized anorexic, for example, does not seem to fully recognize how bad her actions are for her in her reasoning, and the agent in Parfit's example reasons as if her future agony is less bad for her. I shall call this the *first category* of examples. And then there are examples in which the agent uses the right welfare facts, but reaches the intuitively wrong conclusions about reasons. The idealized Caligula, notably, knows that his actions are bad for his victims, but he sees no reason not to torture them. In fact, it is *because* the idealized Caligula recognizes how bad torture is for his victims that he keeps torturing them. Call this the *second category* of

[57] The claim is not that welfare is the only relevant factor. Idealized Mafioso and Caligula also violate *rights*, but the topic is beyond the scope of my chapter. I think an account of rights has to be based on a full account of moral reasons.

examples. It should be noted that in many cases, it will be difficult to tell which kind of mistake an agent is making. Moreover, many of the examples above may fall under either category, depending on how they are interpreted. The idealized anorexic likely reasons from the wrong welfare beliefs, but it is also possible that she recognizes how bad starvation is for her but still judges that she has more reason to stay anorexic. Parfit's example can also be interpreted in these two ways. Nonetheless, I think the two categories are conceptually distinct.

In order to tackle the right reasons problem, I think constructivists should tackle the two categories of examples separately. Moreover, I think thoroughgoing constructivists cannot address the first category of examples, because I do not think thoroughgoing constructivists can construct facts about welfare from the perspective of every idealized agent. John Rawls's grass counter, who is formally rational but who prefers counting blades of grass "in various geometrically shaped areas" to developing his mathematical talents, still seems to me like a good counterexample (1971, 432).[58] Even if we set aside this example, it seems possible that an idealized anorexic might come to agree with the supermodel Kate Moss that "nothing tastes as good as skinny feels" (Costello 2009).[59]

Constructivism about welfare facts also suffers from another class of counterexamples, which I shall call cases of *absent judgments*. Even though I have emphasized in previous chapters that meta-ethical constructivism is different from subjectivism, constructing normative facts about X from an agent's perspective requires that the agent have *some* relevant judgments

[58] As Chris Heathwood argues in a related context, the thoroughgoing constructivist must either argue that an idealized agent would never prefer grass counting, or simply bite the bullet. But he thinks there is no reason to think idealizing would get rid of the preference (2015, 140–41; see also Crisp 2006, sec. 4.2). Heathwood prefers the second option, and he argues that while the grass counter's life may be lacking in excellence, it is not lacking in well-being (2005, 499–500). I am not convinced that his life is not also lacking in well-being. Just as it benefits children to learn to engage in worthwhile activities, it would benefit the grass counter if he does something else.

[59] It should be noted, however, that actual people who suffer from anorexia nervosa are quite far from being ideally rational, as Street emphasizes (2009b, 276–77).

about X that we can correct or extend. But it seems possible for things to be good and bad for welfare subjects who have no relevant judgments about their own welfare whatsoever. Eden Lin (2017), for example, argues that things can be good and bad for newborns, even though they lack the relevant judgments that can ground these welfare facts.[60] While newborns may have desires for certain experiences, these desires are not enough to ground the fact that a baby cared for by parents is better off than a baby with the same experience in an experience machine (362). Unlike adults, they do not have the capacity to form the judgment (or even the desire) that it is better to experience the love of actual parents. They lack the mental capacity to form fine-grained mental states about reality and simulation. Similarly, Ian Tully (2017) considers examples of individuals who suffer from extreme depression such that they no longer have any desire to pursue or avoid anything. In particular, he argues that there are people (such as the victims of the Nazi concentration camps) who suffer from what he calls "complete conative collapse" as a result of severe maltreatment (6). These individuals become completely apathetic. Intuitively, these people seem to be doing terribly. But given their lack of any evaluative judgment or attitude, it seems that the badness of such a state cannot be grounded in their judgments or attitudes at all.

My goal in the preceding discussion is not to offer decisive arguments against constructivism about welfare, although I hope I have highlighted some difficulties. Rather, I have argued that to a significant extent, the intractability of the right reasons problem may be explained by the difficulty of constructing welfare facts. Moreover, I have argued that we should treat the first category of examples differently from the second category. We cannot begin to address these agents' reasons related to welfare, if we cannot get the right welfare facts in the

[60] Similar neo-natal examples are also presented in Richard Kraut (1994, 42–43).

first place. My proposal, then, is that we should adopt welfare objectivism to deal with the first category, and then show that restricted meta-ethical constructivism can also deal with the second category.

Before I try to show, in §4.4, that restricted meta-ethical constructivism has the right resources to solve the right reasons problem, I need to deal with some potential worries first in the next two sections about whether welfare objectivism is even compatible with reasons constructivism. First, it might be thought that there is no meta-ethical motivation for adopting constructivism about only part of the normative domain. I address this worry in §4.2. Second, it might be thought that if welfare directly gives us certain normative reasons, then welfare objectivism would entail the existence of some objective normative reasons. I address this in §4.3. I hope by the end of §4.3, I will have shown that restricted meta-ethical constructivism is a coherent meta-ethical position.

4.2. Must Reasons Constructivists Adopt Welfare Constructivism?

I think there is probably a natural suspicion toward any meta-ethical view that adopts two different positions about two different sets of normative facts. One worry, I suspect, is that any meta-ethical argument in favor of reasons constructivism will also support constructivism about other normative facts. For example, Street's argument for anti-realism, based on evolutionary considerations, targets *all* realist views in meta-ethics (2006). According to Street, realists have no evolutionarily acceptable explanations for how we can have any kind of normative knowledge. This argument, if sound, applies to both reasons and welfare. Similarly, metaphysical arguments against non-natural objective normative facts are also thoroughgoing. It

seems to me that objective facts about welfare would not be any less metaphysically "queer" (to use J. L. Mackie's term, 1977) than objective facts about normative reasons.

However, I do not think these epistemological or metaphysical considerations constitute the only arguments for constructivism, and in any case they are not what convinced me to be a constructivist.[61] But it is easy to see why most constructivists are thoroughgoing constructivists. Street's and Mackie's arguments do not distinguish between normative properties, since their target is realism about *normativity* in general. In contrast, my argument for constructivism emphasizes that *normative reasons for action* in particular have to be identified from an agent's perspective. This is because an autonomous agent always has her own perspective from which she deliberates about what she should do, and considerations can only favor or count against actions from this perspective.

I shall argue that, unlike normative reasons for action, facts about welfare are not similarly connected to an agent's deliberative perspective. Therefore, there are relevant differences between reasons and welfare, and my motivations for adopting reasons constructivism do not carry over to welfare. I shall offer two arguments in support of this claim.

First, it seems that when we try to identify certain facts about welfare, we do not reason from the perspective of the welfare subject. Since, as I have argued, we determine reasons for action by engaging in perspective-relative reasoning, there is at least one relevant difference between welfare and reasons.

[61] It is obviously beyond the scope of this paper to address these arguments. But for a useful discussion of Street's evolutionary debunking argument, see William FitzPatrick (2020, sec. 4.1). For a discussion of the metaphysical commitments of realism, see Scanlon (2014, chap. 2).

I shall use Korsgaard's constructivist analysis of welfare as my foil here.[62] Korsgaard believes that "we use the concept of the good to mark out the solution to the problem of what things to aim at" (2013, 23–24). In particular, she thinks that what is good for *x* can only be determined from *x*'s perspective and constructed from *x*'s own valuing attitudes (2013, 23–24; 2018, 27, 135). It should be noted that the concept of relational good, for Korsgaard, is an action-guiding one, just like the concept of normative reasons. What is good for *x*, when *x* is a being with a valuing perspective, marks out the ends that *x* should pursue (2018, 17).

However, it seems like there are such welfare subjects that the things that are good and bad for them can outstrip what they can value from their perspectives. One class of such subjects, as I have already mentioned, are *newborns*.[63] Consider

> **Neo-Natal Surgery:** Kevin is born with a physical defect. If he is not immediately treated, his mental and physical capacities will be severely diminished. Fortunately, the procedure to fix the defect is completely safe and effective. However, from Kevin's perspective, he has no pro-attitude or favoring judgment whatsoever toward the procedure and its benefits. He experiences no pain and derives no pleasure from the procedure. He also has no judgment or belief about the value of the procedure. In fact, he completely lacks any capacity to respond to the value of the procedure.

Is the procedure *good for* Kevin? Intuitively, yes. It is simply good for Kevin to have his mental and physical capacities intact. But Kevin does not have any attitude, belief, or judgment regarding the capacities that allow him to respond to his parents' love, move around, apply to universities when he is older, become a professional athlete, etc.[64] While a full account of the

---

[62] Technically, the object of Korsgaard's analysis is "relational value." She offers a constructivist analysis for things that are *good for x*, in terms of *x*'s valuing attitudes. It should be clear that the value of welfare I describe here is at least one element of relational value.

[63] Again, the example here is inspired by Lin's discussions (2017).

[64] Welfare constructivists might try to account for the latter benefits on this list by appealing to the future attitudes or judgments of the agent, when the agent acquires the right capacities to respond to the value of the procedure. Lin (2020) explores the prospects of such a futurist approach in the context of reasons, whereas Jennifer Hawkins (2019) proposes a version of internalism about realized goods that is indexed by time. The first point is that such accounts still do not explain why the surgery is good for Kevin before he develops such capacities to a sufficient degree. Moreover, it is possible that the surgery still leaves Kevin's mental capacities severely diminished, even though the

value of the surgery would include all of these benefits, Kevin is in no position to grasp them or value them.

More importantly, it seems that *we* can directly know that the surgery is good for Kevin without trying to adopt Kevin's perspective. If we try to determine whether the surgery is good for Kevin from his perspective, we will see that Kevin has no relevant judgment or attitude that indicates how the surgery might be good for him. There is simply nothing there for us to work with when we try to engage in perspective-relative reasoning.

Let me try to explain why I do not find welfare constructivism attractive here. Roughly speaking, both Korsgaard and I think that autonomous agents have their own perspectives on how to live their lives, and reasons must have the right "grip" on their agency as a result (2008b, 53–54, 56). It is because of such connections between reasons and agency that reasons constructivism is true. But if the procedure can be good for Kevin even though Kevin has no perspective from which he reasons about the value of the procedure, then the status of being a welfare subject is very different from the status of being a rational agent.[65] It seems that welfare facts do not need to have any "grip" on the welfare subject. If so, the motivations for reasons constructivism do not carry over to welfare constructivism.

More concretely, as I mentioned above, Korsgaard thinks that "we use the concept of the good to mark out the solution to the problem of what things to aim at." But it is not plausible to think what is good for Kevin refers to what Kevin should aim at. Kevin has no rational agency,

---

surgery prevents even greater harm to him. Perhaps he still cannot grasp the value of the surgery for the rest of his life, even though his life is still better than his life would be without the surgery.

[65] Connie Rosati (2003) thinks that facts about welfare must engage our agency, and she presumably thinks that the status of being a welfare subject is connected to (or perhaps even identical to) the status of being an agent. I do not think this is right. In one sense, I agree that agency *makes a difference* to which things can be good for someone. The fact that human beings possess rational agency makes it good for us to make plans with friends, study philosophy, or manage a *polis*. But I do not think what is good for *x* must *engage x*'s agency in the internalist or constructivist sense. As I have argued, what is good for Kevin does not engage his agency.

so he cannot aim at what is good for him.[66] In contrast, Korsgaard's remark fits normative reasons much better. Indeed, it seems wrong to say that *Kevin* has a reason to ask doctors to perform surgery on him, and this is exactly because Kevin cannot aim at things himself or even ask anyone to help him.

Let me briefly address one possible objection here, although the objection perhaps deserves more in-depth treatment than I can give here. One might worry that the conclusions I reach with respect to Kevin might not generalize. Could it be, for example, that welfare constructivism is false for newborns like Kevin, but true for agents who have acquired rational capacities? I think such a view needs to be motivated. Why would constructivism about the same normative property be true only for a subset of welfare subjects? Moreover, as Lin (2017, 359–60) emphasizes, such a disjunctive view would have paradoxical consequences, given that infants usually become rational agents eventually. If this disjunctive view is true, it should be the case that at some point constructivism becomes true for welfare facts about Kevin. But this would mean that some of the things that are good for Kevin before the transition point might suddenly lose value, if Kevin does not have the right valuing attitudes to sustain their value after the transition point.

I shall now move on to my second argument. If welfare constructivism is true, then I must engage in "if I were you" reasoning to determine what is good for you in interpersonal contexts. If so, the normative felicity condition that constrains interpersonal reasoning about normative reasons should also constrain interpersonal reasoning about welfare. As I shall argue, however, no such condition exists in the case of welfare.

---

[66] Actually, Korsgaard also does not think that only *rational* agents can have a good. Things can be good for non-rational animals as long as animals can aim at things through instincts (2018, sec. 3.2). But my point with the example of Kevin is that there are things that can be good or bad for Kevin beyond what he can even aim at instinctively. The world of welfare is much richer than what Kevin can respond to with any of his current capacities.

To adapt an example from Stephen Darwall (2009, 128), a mother can tell her five-year-old daughter that broccoli is good for her, in order to teach her the health benefits of vegetables. It seems like the mother cannot possibly browbeat in this case, since she is not trying to make any claim on the daughter's agency. (For example, it could be the case that *the mother* has a reason to make her daughter eat broccoli, or trick her into doing so.) In this case, the mother is just giving her daughter a piece of information. When her now middle-aged daughter refuses to eat broccoli, however, the mother needs to offer the daughter a *reason* to eat broccoli. Intuitively, the mother can now violate the normative felicity conditions of interpersonal reasoning by browbeating her daughter into eating broccoli. This is because offering a reason, unlike offering a fact about welfare, is making a claim on the daughter's agency.

Actually, offering a welfare consideration to *x* cannot violate the normative felicity conditions of interpersonal reasoning, *as long as the speaker does not conversationally implicate that x has a reason*.[67] Suppose that Joseph plans to be a professional skateboarder. His academic mentor, Sarah, tells him that an academic job would be good for him given his talent. This act cannot count as browbeating, if she is merely drawing his attention to a fact about welfare that may be relevant to his practical reasoning. She is introducing the fact only as an input for his reasoning, and she does not tell Joseph what reasons he has. But if Sarah is conversationally implicating that Joseph has a reason, perhaps even a decisive reason, to become a professor,[68] she *could* violate normative felicity conditions if she does not reason from Joseph's perspective. This is especially the case if Sarah does not consider Joseph's plan to be a skateboarder at all,

[67] For the notion of conversational implicature, see Grice (1991).
[68] The kind of example I have in mind is the following. Joseph: "I want to be a skateboarder." Sarah: "You haven't thought about how good your life would be as an academic, have you?" It is natural to interpret Sarah as implicating that Joseph has a stronger reason to be an academic.

perhaps because she thinks that skateboarding is too undignified to even be considered as a career choice.[69]

Given that my argument for reasons constructivism in previous chapters is based on these interpersonal considerations, that argument does not support welfare constructivism. So it seems like we can have meta-ethical reasons for endorsing reasons constructivism but not welfare constructivism.

4.3. Does Welfare Objectivism Entail Reasons Objectivism?

The second worry I shall address is connected to the last one. As I mentioned in the last section, it may be objected that we cannot adopt different meta-ethical positions with respect to reasons and welfare, because both facts about reasons and facts about welfare are *normative* facts. In this section, I explore another way to develop this worry. It might be thought that if certain connections between welfare and reasons hold, then welfare objectivism would entail the existence of objective reasons. The connections I have in mind can roughly be described by the claim that *welfare directly gives me reasons*. Indeed, it might be thought that this is what it means for welfare facts to be normative.[70] If welfare facts are objective, and these objective facts directly give me reasons, then I have objective reasons, and reasons constructivism is false.

Whether this objection is sound depends on how we should understand the claim that welfare directly gives me reasons. Here I shall primarily focus on how those who endorse the

[69] Similarly, if Owen's family only tell Owen that given his talents, a military career would be good for him, it seems to me that this act does not count as browbeating either. They only browbeat when they make claims on Owen's agency without reasoning from Owen's perspective.

[70] I think this is essentially Rosati's (2009, 217–18) complaint against Kraut's very similar version of welfare objectivism (2009). Kraut, like me, believes that facts about welfare leave open what reasons we have for promoting welfare (2009, 75). Rosati objects that if this is so, "Whether anyone has reason to promote what is good for S, and under what circumstances, would […] be a matter left to normative theory" (2009, 218). As I shall explain, I think it is entirely right that our reasons to promote welfare should be left to further normative theorizing.

value-based theory of reasons might develop the objection.[71] I can see two ways to understand the claim that welfare directly gives me reasons. The first way to understand it is as a conceptual claim: it is a conceptual fact that I have a reason to act in ways that promote my (or someone else's) welfare. The second way to understand it is as a claim about metaphysical grounding: I have a reason to act in ways that promote my (or someone else's) welfare in virtue of the relevant fact about welfare. Actually, since reasons always have weights, the claim that welfare gives me a reason must be understood as the claim that welfare gives me a reason *of a certain weight w*. This will be important later when I examine how the weight is determined.[72] The worry would then be that if every welfare fact conceptually entails or metaphysically grounds a reason of weight *w* for me, then welfare objectivism entails the existence of objective reasons. To anticipate my reply later, I will argue that there is a "gap" between welfare facts and reasons, and welfare facts do not directly give me reasons in either of the two ways above.

The grounding claim seems to be more popular among those who develop a value-based theory of reasons recently. One of the clearest statements of the value-based theory of reasons is provided by Barry Maguire (2016).[73] According to Maguire, "Some fact of the form [φ would promote S] is a reason to φ if and only if and if so *in virtue of* the facts that φ would promote S

[71] An exactly parallel objection can be made if the buck-passing account of welfare is true. That is, if reasons are explanatorily prior to welfare. However, as Guy Fletcher (2012) argues, no buck-passing account of welfare in the current literature is both extensionally adequate and non-circular.

[72] Joseph Raz would object to my claim that welfare can only give me a reason if the weight of the reason is also determined by the value. Raz (2016, 144) thinks that the question of whether value directly gives me reasons can be separated from the question of how weighty these reasons are. I find this implausible. First, the two questions are at least connected in the following sense: a reason that has *zero* weight is no reason at all. Second, it is not clear how value-based reasons objectivism does not already commit one to a theory of the weight of reasons. While in a similar Humean context, Mark Schroeder does not think that the weight of a desire-based reason is proportional to the strength of the desire promoted (2007, 100–101), this is because he does not think the strength of one's desire is a normative property. But for value-based reasons objectivists at least, value, unlike the strength of desires, is clearly a normative property.

[73] However, Maguire's definition of the value-based theory focuses on *valuable states of affairs*. His theory of reasons seems to be intended as an intermediate step in an argument for consequentialism (233–234). For my purposes here, I shall consider a version of *welfare*-based theory of reasons.

and that S is valuable" (237, emphasis mine). Furthermore, the more valuable S is, the weightier the reason is (242). This is meant to be an *analysis* of reasons, rather than a *substantive account* of reasons (237). (This point will come up again very soon.) Maguire does not commit to a view about the nature of analysis, although he seems to prefer (as revealed both by the above definition and a footnote) characterizing the view in terms of metaphysical grounding (n. 7).

Similarly, Roger Crisp develops a version of value-based reasons objectivism. On the first-order level, he thinks that I have self-interested reasons whose strength is proportional to the degree that my welfare is promoted (2006, 73). I also have weaker reasons to promote the welfare of others (132). It is not clear what meta-ethical view Crisp holds with respect to the connections between reasons and welfare, although he explicitly rejects conceptual or analytic connections between them (62). However, according to his dualism of practical reason, facts about welfare and facts about which values are instantiated in *my* life in particular *ground* self-interested reasons (139), and facts about other people's welfare (but notably *not* facts about my relationships to others) ground weaker reasons to promote their welfare (133, 143).[74] I think it is natural to interpret these claims as claims about metaphysical grounding.[75] If this interpretation is correct, such claims along with welfare objectivism would also undermine reasons constructivism. This is because, again, all the facts that ground welfare-based reasons here are mind-independent facts.

[74] The relevance of relationships to welfare-based reasons is important, since as I shall suggest, restricted meta-ethical constructivism offers a good explanation for how our normative reasons (according to commonsense morality at least) seem to depend on relationships. It is interesting that one of Crisp's arguments against the relevance of relationship facts takes the form of a partial evolutionary debunking argument (143). Our intuitions about the relevance of relationships, he argues, can be entirely explained by evolutionary influences. For a better developed version of the argument, see de Lazari-Radek and Singer (2012). This argument seems to presuppose reasons objectivism. For a reasons constructivist, this evolutionary explanation should not undermine relationship-based reasons.

[75] It is likely that the grounding relation here for him is only a normative one. As I will argue later, this claim poses no problem for my view. Nonetheless, I shall interpret these claims as claims about metaphysical grounding, since only such claims pose problems for me.

Is it enough, as an objection to my view, for it to be a *first-order normative* truth that welfare directly gives me reasons? No. Restricted meta-ethical constructivism is compatible with the existence of self-interested reasons or even categorical reasons to promote the welfare of others. For example, it might be the case that when we reason from each agent's perspective, we conclude that the agent has reason to promote her own and other people's welfare. Similarly, if it is only a normative truth that a reason given by welfare has weight *w*, this fact may be explained via perspective-relative reasoning. In fact, restricted meta-ethical constructivism might even be compatible with utilitarianism, as long as it can be shown from each agent's perspective that there is equal reason to promote everyone's welfare. But the normative truth of utilitarianism (if it is true) is not directly grounded in objective welfare facts, but rather in non-defective perspective-relative reasoning. Therefore, the existence of such a normative fact about the relationships between welfare and reasons does not by itself undermine restricted meta-ethical constructivism. The objection against my view must instead appeal to conceptual or metaphysical connections between welfare and reasons, as I have argued.[76]

Thus, in the rest of the section, I shall address whether it is a conceptual truth or a metaphysical truth that welfare gives me reasons. Roughly, I argue that in the presence of a fact about welfare, I can always ask further questions about whether it gives me a reason and how weighty the reason is. I argue that these further questions are substantive normative questions, rather than conceptual or metaphysical ones. If so, claims about whether welfare gives me

[76] I assume here that there is an intuitive distinction between what counts as a normative claim and what counts as a metaphysical claim. There is thus an intuitive difference between Maguire's metaphysical claim that value grounds reasons, and the first-order normative claim (usually favored by consequentialists) that there is reason to promote value. Moreover, one can accept consequentialism without accepting the value-based theory of reasons. For those who do not accept the distinction between normative and metaphysical claims, it suffices for my purposes that value does not *directly* ground reasons. I just need to show that we have no reason to rule out a constructivist explanation of reasons that are ostensibly generated by welfare.

reasons are not conceptual claims, nor are they metaphysical ones.[77] Note that in order to defend my view against the objection, I only need to shift the burden of proof to the objectors. In order to show that my view is coherent, I only need to show that it is plausible to deny direct conceptual or metaphysical connections between welfare and reasons. Therefore, I shall not (and cannot) discuss every possibility of how welfare might give me reasons in a way that undermines my reasons constructivism. Given that it seems to make a difference whose welfare is at stake, I shall first consider one's own personal welfare, and then the welfare of others.

4.3.1. Personal Welfare

Consider:

> **Bulgarian History:** I believe that learning Bulgarian history would be good for me, but do I have a reason to learn Bulgarian history, and, if so, how weighty is my reason? (Scheffler 2011, 26–27)[78]

As Samuel Scheffler argues, there may be *many* activities that I believe are good for me (philosophy, ballet, Mongolian throat singing, learning Bulgarian history…), but I actually *value* very few of these activities. In addition to my belief (as well as certain affective dispositions), the attitude of valuing requires that I judge there to be a *reason* for me to engage in the activity. Valuing the activity requires that I treat the activity as having deliberative significance (26–31). Setting aside whether Scheffler's analysis of the valuing attitude is correct, I think he is right that regarding an activity as valuable (or specifically, good for me) is different from judging there to be a reason for me to pursue the activity. I can fully believe that knowledge of Bulgarian history

[77] There is an interesting historical connection here between the argument that I am about to give and Bertrand Russell's objection against G. E. Moore's value-based objectivism about practical normativity. As Russell points out, even with a conception of intrinsic value in place, it is still an open question whether someone *ought* to promote the best results (1904; see also Darwall 2009, 283).

[78] My discussion of this example extends my previous discussion in (Fan 2020b).

is good for me. But I might still see little or even no reason to learn Bulgarian history. It seems to me that it is one question whether an activity is good for me, but another question entirely whether this fact generates a reason for me.

When I think about this example, my own intuition is that I have *no* reason to learn Bulgarian history. But the intuition here is not necessarily reliable: I also do not think learning Bulgarian history has much value for me either. As Mark Schroeder points out, negative existential intuitions (such as my intuition that there is *no* reason to learn Bulgarian history) can be unreliable (2007, 92–97). It could be that the reason is actually very small, and for pragmatic reasons we often do not communicate these small reasons.[79] So perhaps the small value of learning Bulgarian history grounds or conceptually entails a very small reason for me.

But now consider a variation on this example. Suppose that, as an undergraduate student, there are many different classes I can take, and one class I can take is a theater class. I judge that a theater class would be quite good for me, because it would be good for me to be able to express myself in new ways. Now if this fact alone gives me a reason to take the class, it seems that I have to accept that there is quite a weighty reason to take it. Nonetheless, the matter is not so simple, because the prospect of acting leaves me cold.[80] Perhaps I care more about computer science or mathematics, even though I think the skills I can acquire from these classes are

[79] Schroeder appeals to Paul Grice's maxim of quantity, according to which I should only tell you that there is a reason when this is informative, i.e. the reason is relatively weighty (Schroeder 2007, 95; Grice 1991, 26).

[80] Does the fact that the prospect of doing something leaves me cold show that the activity cannot be very good for me? Peter Railton suggests that "It would be an intolerably alienated conception of someone's good to imagine that it might fail in any such way to engage him" (1986, 9). However, it should be emphasized that the notion analyzed by Railton is not the notion of welfare, but rather a matter of what ends one should seek (n. 9). In fact, I think the intuition about alienation here tracks reasons, rather than welfare. It seems to me that the skills I can acquire from learning how to act cannot be any less good for me, just because the class is lower on my priority list. If the class does not engage my interest, it can still be very good for me, even though I have less reason to take it.

equally as good for me as the skills I can acquire from a theater class.[81] If welfare directly gives me reasons, it seems that I am provided with the same reason to take each class. But intuitively, the deliberative question of how much reason I have to take the theater class, compared to my other options, remains open. If so, it does seem to be a further question how much reason I have to do what is good for me.

Maybe my intuitions about my reasons and their weight are wrong. But even if I am wrong, what kind of mistake am I making? I do not think I am making a *conceptual* mistake. I do not lack the concept of welfare just because I do not think facts about welfare directly give me reasons. Nor do I think that I am making a mistake about whether welfare facts *metaphysically ground* reasons for action. I do not see how it can be a metaphysical truth that the welfare contribution of the class alone grounds a reason, even when the class leaves me cold.

I could be making a *normative* mistake, however. It seems to me that there is a gap between what is good for me and what I have a reason to do, because not everything that is good for me is a personal *project*. But it may be a normative truth, as Crisp argues, that the weight of my self-interested reasons is strictly proportional to how good the option is for me. It may be a normative truth that my personal projects have no relevance to the weight of my reasons. My intuition at least is that my personal projects do make a normative difference. Restricted meta-ethical constructivism can explain this, since facts about which project I should take on seem to be the kind of facts that can be constructed from my deliberative perspective. But even if my intuition is wrong, it is still possible that Crisp's proportionality claim about the weight of self-interested reasons can be constructed from my deliberative perspective. Perhaps if I judge that a

[81] It may be true that there is a fourth comparative relation here, as Chang (2017) suggests, and these classes are "on a par" in terms of how good they are for me. But the parity relation helps me here, since I might have more reason to take one of the classes if I value the relevant skills more.

theater class and a computer science class are equally good for me, then I should see the two cases as exactly analogous, and I have reasons of the same weight to take them. If I am only making a normative mistake, then the further question whether I have a reason of a certain weight to do what is good for me is just a first-order normative question. If so, welfare objectivism does not undermine reasons constructivism, since it is neither a conceptual nor a metaphysical truth that welfare gives me reasons. Instead, it could well be that faced with certain objective welfare facts, I need further perspective-relative reasoning to determine which reasons exist and how weighty they are.

4.3.2. The Welfare of Others

If you already accept that personal welfare does not directly generate normative reasons, the case for the analogous claim about the welfare of others should be easier to accept. Consider the following example:

> **Cancer Research:** I believe that cancer research would be good for many other people, but do I have a reason to choose cancer research as my career, and, if so, how weighty is my reason?[82]

I do not claim that the answer to this question is that I have *no* reason. Rather, I am only claiming that beyond the question of what is good for other people, there does seem to be a further question whether I have a reason (and, if so, how weighty my reason is) to promote other people's welfare. This further question is not an empty one. For example, at least according to commonsensical intuitions, whether I have a reason to promote the welfare of others, and how weighty my reason is, seems to depend on my *relationships* to these other people. While a warm bowl of soup may be equally good for my child as it is for my neighbor's child, it is a further

[82] My discussion of this example extends my previous discussion of a similar example in (Fan 2020b).

normative question what reasons I have to feed my own child and my neighbor's child. Similarly, in the case of cancer research, the fact that cancer research would be good for many other people does not close the deliberative question of how much this fact gives me a reason to engage in cancer research. To answer this further question, I need to reflect on my relationships to those people who would benefit from the research.

Again, I might be making a normative mistake here, but I do not think I am making a conceptual or metaphysical mistake. Restricted meta-ethical constructivism is not undermined if the mistake I am making here is a normative one. Perhaps by reflecting on Peter Singer's pond analogy about our reason to save a drowning child (1972, 231), we see that we have an exactly analogous reason to donate to charities that prevent famines. Perhaps as we keep reflecting on these cases, we see that utilitarianism is correct, and our relationships make no normative difference on a fundamental level. Again, it seems like restricted meta-ethical constructivism can accommodate this, since (as I emphasized in §3.1) analogical reasoning and correction of mistakes about objective evaluative facts are two moves that we can make in perspective-relative reasoning.[83] Moreover, it seems like Singer is making a normative claim about our welfare-given reasons, rather than a conceptual claim or a grounding claim about welfare and reasons. Welfare objectivism can only undermine reasons constructivism if it is a conceptual or metaphysical truth that welfare gives us reasons, and we have seen no compelling reason to believe this.

[83] Actually, Singer's argument fits my view surprisingly well. The first premise in his argument is a claim about the badness of suffering, which, according to my view, is an objective fact. His second premise is about whether we have a duty to prevent harm if we incur no comparable moral cost, and his argument for the premise is an analogical argument that extends our judgment about our duty to save a drowning child to the case of preventing famines (231). I shall not get into the complicated debate about whether this analogy is sound.

So far, I have merely tried to show that it is far from obvious that welfare objectivism undermines reasons constructivism. In fact, if restricted meta-ethical constructivism is correct, we have a good explanation for why it often seems like it is a further question whether we have a reason to promote our own welfare and the welfare of others. The idea that our reasons to promote welfare depend on our projects and relationships is not new. Scheffler (2004; 2010) has emphasized this idea in numerous places. Williams's integrity-based argument against utilitarianism also makes a similar point (1973).[84] Maybe Scheffler and Williams are wrong. But it seems to me that if they are wrong, they are only making normative mistakes, rather than conceptual or metaphysical ones. It does seem to me that they are right though, and restricted meta-ethical constructivism, far from being undermined by reflections on whether welfare gives us reasons, actually offers a good meta-ethical explanation for how our reasons are affected by personal projects and relationships. Facts about welfare do not directly give us reasons. Instead, in order to determine our reasons, we have to engage in perspective-relative practical reasoning that takes welfare facts as inputs.

4.4. Getting the Right Reasons

I argued in §4.1 that welfare objectivism seems to be the right first step to tackling the right reasons problem for reasons constructivists. I then argued in §§4.2 and 4.3 that restricted meta-ethical constructivism is a coherent position. In this section, I come back to the right reasons problem. It should be noted, however, that I cannot possibly try to generate *all* the normative reasons that there are. I cannot solve normative ethics in one section. So the goal is simply to

[84] As does Scheffler's (1995). For other similar notions, Korsgaard's "practical identity" as a source of reasons makes a more obvious connection to constructivism (1996a, chap. 3). In addition, Chang (2013b) argues that our commitments directly shape our reasons, although she is not a constructivist.

show that restricted meta-ethical constructivism has the right *resources* to generate the right normative reasons, or at least the reasons that thoroughgoing constructivists have difficulty generating.

Let me first mention that one benefit of restricted meta-ethical constructivism, which I have already discussed, is that it gives a good account of the agent-relative reasons that vary with our projects, relationships, and personal commitments in general. I might judge that both ballet and philosophy greatly contribute to a flourishing life. But if I choose to be a philosopher, then I have much weightier reasons to study philosophy rather than train for ballet.[85] Like thoroughgoing constructivism, restricted meta-ethical constructivism gives a good explanation for such reasons. Therefore, restricted meta-ethical constructivism preserves one key advantage of thoroughgoing constructivism.

Next, I shall consider the first category of examples, as described in §4.1. Recall that the first category includes agents who do not use the right welfare facts in their reasoning. Both the idealized anorexic example and Parfit's agony argument can be interpreted as cases that belong in this category. The category also includes the example of the husband who cares about his wife's welfare, but mistakenly believes that it is better for her (in the welfare sense) to stay at home than to engage in work that she finds meaningful.[86]

Restricted meta-ethical constructivism can deal with this category of examples quite easily, whether it is the agent's own welfare or the welfare of someone else that is at stake. I argued in the last chapter that if constructivists are allowed to correct descriptive mistakes in an agent's reasoning, there is no reason why constructivists are not allowed to correct mistakes

[85] This point has been emphasized by Ruth Chang in numerous places. See Chang (2017; 2013a; 2013b).
[86] This is, again, an idealized example, since in most cases it is not just the welfare belief that explains this kind of expectation. Very likely, the husband's reasoning is also defective in other ways.

about objective evaluative facts. The idealized anorexic is perhaps mistaken about how good it really is for her to have good health. Given that non-defective practical reasoning, according to my view, should not contain mistakes about objective evaluative facts, the idealized anorexic's reasoning is not so "idealized" after all. Similarly, the husband in the example above makes a mistake in his reasoning about what is good for his wife. Therefore, his reasoning is also defective. In contrast, as I argued in Section I, thoroughgoing constructivists cannot find fault with the judgments of these agents at all as long as they are formally rational, even though they are mistaken about facts about welfare.

As for Parfit's agony argument, like a kind of optical illusion, future pain often seems less vivid and somehow less bad from my present perspective. But just like optical illusions, this judgment does not track how things really are. The same pain in the future will still be the same pain. So if the agent does not reason from the welfare fact that the future pain is equally bad for her, then her reasoning is defective.

Consider now the second category of examples. This category includes examples in which agents use the right welfare facts but reach the wrong conclusions about reasons. Examples that belong to this category include the idealized Mafioso, the idealized Caligula, and Williams's original example of the husband who does not care about how much she is harmed by his cruelty. In addition, both the agony argument and the idealized anorexic can also be interpreted as examples in the second category. I am going to use the norm of rational consistency in analogical reasoning, as discussed in the last chapter. Recall that according to the norm of rational consistency, if two actions are similar, including (but not limited to) the value of welfare that the actions promote, then the reasons to perform these actions would also be

similarly weighty to the extent that the actions are similar. Using this norm, we can extend the agent's judgments in one set of cases to another set of relevantly similar set of cases.

In Parfit's example, suppose that the agent recognizes that future agony is equally serious and equally bad for her. However, she does not see an equally weighty reason to avoid future agony. Her reasoning is defective. It will still be the agent herself who suffers from the pain in the future. The agent's relations to her present self and to her future self are exactly analogous, since her present self and her future self are the same person.[87] The norm of analogical rationality would then require that she recognize an equally weighty reason to avoid future agony.

It could also be that the idealized anorexic discussed by Gibbard belongs to my second category rather than my first. Such an agent recognizes that starvation and death can stand in the way of many of her projects (including the project of keeping a slim figure). She also recognizes that good health is *very* good for her, but achieving a slim figure is *much* less good for her. Nevertheless, it may be the case that she cares so significantly more about having a slim figure and so little about her health that she thinks there is more reason to pursue a project with so little value. I do think this is a possibility, because I think we can affect the weight of our reasons by caring and not caring about things. But I think this kind of agent is *psychologically* highly unlikely. It is far more likely that someone with her priorities does not fully recognize the value of good health, and such an agent would belong to the first category. (It might be helpful to imagine asking such a person if she thinks having a healthier figure would be very good for her.) Nonetheless, if her reasoning is not defective in any way, then I think she does have more reason to stay anorexic.

[87] There are important metaphysical questions about personal identity here, which I cannot delve into. For a similar response to the agony argument, see David Sobel (2016c).

It may be objected that, just like the thoroughgoing constructivist, I also need to bite the bullet here. But I think thoroughgoing constructivism has much less plausible consequences, since a thoroughgoing constructivist has to say that someone who does not even have the intuitively right welfare beliefs might nevertheless make the right reason judgments. In contrast, the kind of idealized anorexic I discuss here has full command of all relevant descriptive and evaluative facts. I do not see how we can fault her judgments without engaging in the kind of browbeating that I critique in Chapter 2. After all, we do not fault the judgments of someone who cares a lot more about computer science and thus sees a much weightier reason to take the computer science class rather than the theater class. So I think my restricted meta-ethical constructivism actually reaches the right conclusion here.[88]

Before I try to tackle the rest of the examples, let me note that there is a missing piece that I have not been able to develop here. I have not given an account of the *moral reasons* that apply to the idealized Mafioso, Caligula, and the cruel husband. I think morality has important social and intersubjective dimensions that I cannot fully discuss here. I will nevertheless try to give a sketch of an account toward the end of this section. I think the reasons that apply to the cruel husband in Williams's example are entirely moral reasons. Given that he does not recognize at all his relationship-based reasons, we can only appeal to moral judgments about his

[88] At the end of the day, we have to choose the most plausible theory as a whole. I think thoroughgoing constructivism incurs a much more significant cost if it cannot generate the right welfare facts for infants, as discussed by Lin, and people who suffer from severe depression, as discussed by Tully. In contrast, I think it counts much less (if at all) against a theory that says autonomous agents with full command of the relevant welfare and descriptive facts might have reasons that the rest of us do not have. This is because many philosophers find it independently plausible that autonomous agents might have different reasons if they have different projects, commitments, and identities (Chang 2017; 2013b; Scheffler 2004; 2010; Williams 1973; Korsgaard 1996a, chap. 3).

duties as someone who is bound by the institution of marriage.[89] Of course, one might wonder what kind of relationship he still has with his wife at this point.

Besides moral reasons, restricted meta-ethical constructivism can actually say a bit more about the idealized Mafioso and Caligula. First, the idealized Mafioso and Caligula likely do not think that they have reasons to torture their own loved ones. We can show that these strangers are relevantly similar to their own loved ones. Perhaps they have no reason to be partial to these strangers, but nevertheless they have reasons not to harm those who, like their own loved ones, can feel pain and suffer.[90] After all, I think most people's reason for not torturing their loved ones is based not on partiality, but rather on what it is like to suffer as a human being.[91] Thus, if the idealized Mafioso and Caligula do not extend their judgments in an analogically rational way, their reasoning is defective.

Second, and this is perhaps controversial, I think while a Mafioso may have *moral* reasons not to kill or maim others, it would be wrong to say that he does not even have *pro tanto* reasons to kill or maim others. If someone from a rival criminal organization threatens to harm his friends, I think he in fact has a reason to kill preemptively. Maybe the reason is outweighed by other moral reasons, but I think it would be wrong to say that the Mafioso has no reason whatsoever to kill his enemies. The Mafioso naturally has different reasons from the reasons of law-abiding citizens. Of course, he should recognize his moral reasons and see how much his

---

[89] A lot more needs to be said here about the institution of marriage, but I think a full discussion of this point far exceeds the scope of this paper. Again, I just want to show that restricted meta-ethical constructivism has the right *resources*.

[90] Recall that in footnote 43 in the last chapter, I described how Mencius reasoned with King Xuan in order to get him to extend his compassion for the ox to his own people. The kind of reasoning we have to do with the idealized Mafioso and Caligula is very similar.

[91] This last claim is again not the claim that pain or suffering conceptually entails or metaphysically grounds reasons. Rather, it is a substantive normative conviction. It is also possible that this is a conviction not shared by the idealized Mafioso and Caligula, in which case the only reasons there are for them not to torture their victims are moral reasons.

reasons are in conflict with each other, but as long as he stays in the Mafia, his identity would be a source of reasons for him.[92] It does seem like browbeating if we insist to the Mafioso that he has no reason whatsoever to kill his enemies. This conclusion may seem less counterintuitive if we keep in mind that the moral reasons are *very* weighty, as we have very good reasons to be morally outraged toward his behavior. However, I acknowledge that it is not a widely shared intuition that the Mafioso has these *pro tanto* reasons,[93] and this may be because it is never analogically consistent for the Mafioso to even see *pro tanto* reasons to kill his enemies. So I do not wish to press the point that these reasons really exist.

Let me now briefly sketch the promised account of moral reasons. I think moral reasons are constructed when each person reasons from the plural "we" perspective of a moral community.[94] Moreover, this moral "we" perspective is embedded in the personal perspective of every agent. For reasons discussed by P. F. Strawson (2008), nobody can simply "opt out" of morality. When other people wrong us, we cannot help but feel that we are *morally* wronged, and as social creatures, we cannot help but hold each other *accountable* for transgressions. Moreover, we hold each other accountable, not just from our particular perspectives, but from the perspective of the moral community as a whole (Strawson 2008, 18; Darwall 2009, 9). Without the perspective of the moral "we," there are only individual grievances, but no moral

[92] Korsgaard (1996a, 257) shares my intuition about the Mafioso case. Admittedly, my intuition is different if the Mafioso is not even trying to protect the honor of his crime family, and like the idealized Caligula, he is only inflicting senseless pain on innocent people. Perhaps in such a case his reason is completely silenced by the moral reason that forbids harming innocent people.

[93] However, I interpret Cohen's objection against Korsgaard to be that her constructivism cannot construct an *all-things-considered* reason to "do the specifically *moral* thing" (1996, 184), so these *pro tanto* reasons may be acceptable.

[94] For a similar view, see Darwall (2009, 9, 102). For a version of moral constructivism that is society-centered, see Copp (1995). For versions of moral constructivism that emphasize the intersubjective dimensions of morality, see Habermas (1990) and Korsgaard (1996b).

complaints.[95] As long as the idealized Mafioso, Caligula, and the husband in Williams's example are participants in our moral community, they have moral reasons as well.[96]

If this account can be justified (and a lot more work needs to be done), then we can explain the moral reasons that apply to these blameworthy agents. From the perspective of the moral community, we endorse rules that regulate the duties associated with marriage. As someone who voluntarily chooses marriage, the cruel husband in Williams's example is bound by the moral duties that come along. We also endorse moral rules that protect people from harm, and these rules likewise bind the idealized Mafioso and Caligula as long as they are participants in the moral community.[97]

Again, my goal here is not to show that this account generates all the right moral reasons. Rather, the goal is to show that, compared to throughgoing constructivism, restricted meta-ethical constructivism has the right resources to generate the right reasons. One worry with the kind of account I propose here is that it leads to a kind of moral relativism between moral communities, both at present and in the past. I am prepared to accept this, although I do not think the kind of relativism that my account leads to is worrying. But compared to a parallel

---

[95] The norm of analogical rationality plays an important role here: an agent cannot only *partially* opt in morality. Take the golden rule for example: if the agent judges from the perspective of the moral community that it is wrong for someone else to wrong her by φ-ing, then she herself would also have a moral reason not to φ. Both reasons are exactly analogous from the perspective of the moral community.

[96] I leave open the possibility that moral reasons might not apply to some individuals, who are so completely beyond the reach of the moral community that they do not even feel morally wronged when other people wrong them. But like Aristotle, I think only beasts or gods can live outside a moral community (*Politics* 1253a29). See, however, Gary Watson's subtle and thoughtful discussion of the limits of such a Strawsonian approach in cases of extreme evil (2004, 234–39). My sense is that, with respect to someone completely beyond the reach of morality such as the cold-blooded murderer Robert Harris, little is lost if we have to say that he has no moral reasons. After all, *we* have moral reasons to defend people within our community (and our community itself) against people like Harris.

[97] It may be thought that when reasoning from the "we" perspective of the moral community, the idealized Mafioso and Caligula will still reach conclusions very different from those reached by law-abiding citizens. While genuine moral disagreements are possible, and an account of moral reasons must include structural norms that help us resolve such disagreements, it is far from a foregone conclusion that their non-defective reasoning from the "we" perspective of the moral community would reach different conclusions. The problem with them is that they reject the authority of the moral community in the first place, and the Strawsonian account above is intended as a solution. See also footnote 96 above.

thoroughgoing constructivist account, my view has one important advantage. Given that agents in most (and likely all) moral communities in human history have reasoned from many false beliefs about what is good for individuals, restricted meta-ethical constructivism allows us to correct all such beliefs when we reason from the perspectives of these communities. As I argued in Section I, thoroughgoing constructivism has difficulty generating the right welfare facts. Given the centrality of welfare beliefs in the moral reasoning of any moral community, this is a serious problem for thoroughgoing constructivism.[98]

In conclusion, I have argued that restricted meta-ethical constructivism allows us to not only correct welfare mistakes in perspective-relative reasoning, but also extend an agent's reactions to welfare in some situations to other analogous situations. I have only sketched a few ways to address the right reasons problem here. Further normative theorizing is needed for a full account of moral reasons, as well as the reasons of agents who have bad projects and commitments. However, I hope I have shown that there are reasons to be optimistic. Thoroughgoing constructivists who are worried about the right reasons problem should consider whether my view can better tackle the problem and still preserve what makes reasons constructivism attractive.

[98] Another thing to keep in mind about moral relativism (especially cultural relativism) is that moral communities are not always delineated along cultural lines. As Williams argues, as long as there are "real confrontations" between cultures and societies, the boundary of our moral community must be extended in order to resolve such confrontations (1985a, 162–63; 2005). I think this argument rules out very strong forms of moral relativism in the modern world.

Conclusion

In this thesis, I have defended reasons constructivism and welfare objectivism. I argued in Chapter 1 that it is perspective-relativity, rather than mind-dependence, that characterizes the central motivation for reasons constructivism. I then argued in Chapters 2 and 3 that reasons constructivism alone explains the necessity of the "if I were you" mode of reasoning in interpersonal contexts. Lastly, in Chapter 4, I argued that welfare facts are among the facts that can be corrected in perspective-relative reasoning.

Among the topics that still need to be addressed in future work are *analogical reasoning* and *morality*. First, I need to show that analogical consistency, rather than stronger norms such as the categorical imperative and the norm of systematic justifiability, governs perspective-relative reasoning. Second, I need to give a precise formulation of the norm of analogical rationality, as well as an account of what makes different cases analogous.

With respect to morality, I think a satisfactory account has to fully grapple with the different perspectives from which we make moral judgments. As I have argued, all reasons have to be located within the perspective of the agent who has these reasons. But as Stephen Darwall (2009) has argued, we not only make moral judgments from the second-personal perspective, but also from the perspective of the entire moral community. A theory of moral normativity needs to show how these different levels are in harmony. The moral standpoint requires a kind of ascent from the particular perspectives of the individuals to the plural "we" perspective of the moral community.

As I have indicated in Chapter 4, the challenge of explaining moral normativity is not easy to meet. But this is exactly what we should expect. "Why should I be moral?" is one of the most central questions of moral philosophy, and *pace* H. A. Prichard (2002a), it is not a mistake

to ask this question. It is not a question about what self-interested or desire-based reasons we have to be moral. Rather, it is an invitation to demonstrate that moral reasons can be situated within each agent's deliberative perspective. I have not been able to show how this can be done. Nonetheless, I hope I have at least described the basic components of what I think is a promising constructivist theory.

Bibliography

Aristotle. 1998. *Politics*. Translated by C. D. C. Reeve. Indianapolis, Ind: Hackett Publishing.

Asarnow, Samuel. 2019. "Internal Reasons and the Boy Who Cried Wolf." *Ethics* 130 (1): 32–58.

Austin, J. L. 1962. *How to Do Things with Words*. Oxford: Oxford University Press.

Fan, Zhicheng. 2020a. Objective Reasons and the First-Personal Nature of Practical Reasoning. BPhil Essay. University of Oxford.

———. 2020b. Value and Reasons. BPhil Essay. University of Oxford.

Blackburn, Simon. 1998. *Ruling Passions: A Theory of Practical Reasoning*. Oxford; New York: Oxford University Press.

Callard, Agnes. 2016. "Proleptic Reasons." In *Oxford Studies in Metaethics, Volume 11*, edited by Russ Shafer-Landau, 129–54. Oxford: Oxford University Press.

Chang, Ruth. 2013a. "Grounding Practical Normativity: Going Hybrid." *Philosophical Studies* 164 (1): 163–87.

———. 2013b. "Commitments, Reasons, and the Will." In *Oxford Studies in Metaethics, Volume 8*, edited by Russ Shafer-Landau, 74–113. Oxford: Oxford University Press.

———. 2015. "Practical Reasons: The Problem of Gridlock." In *The Bloomsbury Companion to Analytic Philosophy*, edited by Barry Dainton and Howard Robinson, 474–99. London: Bloomsbury Academic.

———. 2017. "Hard Choices." *Journal of the American Philosophical Association* 3 (1): 1–21.

———. 2020. "Do We Have Normative Powers?" *Aristotelian Society Supplementary Volume* 94 (1): 275–300.

Cohen, G. A. 1996. "Reason, Humanity, and the Moral Law." In *The Sources of Normativity*, 167–88. Cambridge; New York: Cambridge University Press.

Copp, David. 1995. *Morality, Normativity, and Society*. New York: Oxford University Press.

Costello, Brid. 2009. "Kate Moss: The Waif That Roared." WWD. 2009. https://wwd.com/beauty-industry-news/beauty-features/kate-moss-the-waif-that-roared-2367932/.

Crisp, Roger. 2006. *Reasons and the Good*. Oxford; New York: Oxford University Press.

Dancy, Jonathan. 2002. *Practical Reality*. Oxford: Oxford University Press.

———. 2018. *Practical Shape: A Theory of Practical Reasoning*. Oxford: Oxford University Press.

Darwall, Stephen, Allan Gibbard, and Peter Railton. 1992. “Toward *Fin de Siècle* Ethics: Some Trends.” *The Philosophical Review* 101 (1): 115–89.

Darwall, Stephen L. 2009. *The Second-Person Standpoint: Morality, Respect, and Accountability*. Cambridge, Mass: Harvard University Press.

Dasgupta, Shamik. 2017. “XV–Normative Non-Naturalism and the Problem of Authority.” *Proceedings of the Aristotelian Society* 117 (3): 297–319.

Dreier, Jamie. 2015. “Can Reasons Fundamentalism Answer the Normative Question?” In *Motivational Internalism*, edited by Gunnar Björnsson, Caj Strandberg, Ragnar Francén Olinder, John Eriksson, and Fredrik Björklund, 167–81. Oxford: Oxford University Press.

Enoch, David. 2009. “Can There Be a Global, Interesting, Coherent Constructivism about Practical Reason?” *Philosophical Explorations* 12 (3): 319–39.

Fine, Kit. 2001. “The Question of Realism.” *Philosophers’ Imprint* 1 (1): 1–30.

Finnis, John. 2011. *Natural Law and Natural Rights*. Oxford; New York: Oxford University Press.

FitzPatrick, William J. 2013. “How Not to Be An Ethical Constructivist: A Critique of Korsgaard’s Neo-Kantian Constitutivism.” In *Constructivism in Ethics*, edited by Carla Bagnoli, 41–62. Cambridge: Cambridge University Press.

———. 2020. “Morality and Evolutionary Biology.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. https://plato.stanford.edu/archives/spr2021/entries/morality-biology/.

Fletcher, Guy. 2012. “Resisting Buck-Passing Accounts of Prudential Value.” *Philosophical Studies* 157: 77–91.

Foot, Philippa. 1972. “Morality as a System of Hypothetical Imperatives.” *The Philosophical Review* 81 (3): 305–16.

———. 2001. *Natural Goodness*. Oxford; New York: Oxford University Press.

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

———. 1999. “Morality as Consistency in Living: Korsgaard’s Kantian Lectures.” *Ethics* 110 (1): 140–64.

———. 2003. *Thinking How to Live*. Cambridge, Mass: Harvard University Press.

Grice, Paul. 1991. “Logic and Conversation.” In *Studies in the Way of Words*, 22–40. Cambridge, Mass: Harvard University Press.

Griffin, James. 1988. *Well-Being*. Oxford: Oxford University Press.

Habermas, Jürgen. 1990. “Discourse Ethics: Notes on a Program of Philosophical Justification.” In *Moral Consciousness and Communicative Action*, 43–115. Cambridge, Mas.: MIT Press.

Hawkins, Jennifer. 2019. “Internalism and Prudential Value.” In *Oxford Studies in Metaethics Volume 14*, edited by Russ Shafer-Landau, 95–120. Oxford: Oxford University Press.

Heathwood, Chris. 2005. “The Problem of Defective Desires.” *Australasian Journal of Philosophy* 83 (4): 487–504.

———. 2015. “Desire-Fulfillment Theory.” In *The Routledge Handbook of the Philosophy of Well-Being*, edited by Guy Fletcher, 135–47. Abingdon: Routledge.

Hieronymi, Pamela. 2021. “Reasoning First.” In *The Routledge Handbook of Practical Reason*, edited by Kurt Sylvan and Ruth Chang. 349-365. Abingdon, Oxon; New York, NY: Routledge.

Hurka, Thomas. 1996. *Perfectionism*. Oxford: Oxford University Press.

Hussain, Nadeem, and Nishi Shah. 2006. “Misunderstanding Metaethics: Korsgaard’s Rejection of Realism.” In *Oxford Studies in Metaethics Volume 1*, edited by Russ Shafer-Landau, 265–94. Oxford: Oxford University Press.

James, Henry. 1893. *The Wheel of Time; Collaboration; Owen Wingrave*. New York: Harper & Brothers.

Johnson, Robert, and Adam Cureton. 2016. “Kant’s Moral Philosophy.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. https://plato.stanford.edu/entries/kant-moral/.

Johnson, Robert N. 1999. “Internal Reasons and the Conditional Fallacy.” *The Philosophical Quarterly* 49 (194): 53–72.

Joyce, Richard. 2015. “Moral Anti-Realism.” Stanford Encyclopedia of Philosophy. 2015. plato.stanford.edu/entries/moral-anti-realism/.

Kant, Immanuel. 2012. *Groundwork of the Metaphysics of Morals*. Edited by Mary J. Gregor and Jens Timmermann. Cambridge: Cambridge University Press.

Korsgaard, Christine M. 1986. “Skepticism about Practical Reason.” *The Journal of Philosophy* 83 (1): 5–25.

———. 1996a. *The Sources of Normativity*. Cambridge; New York: Cambridge University Press.

———. 1996b. “The Reasons We Can Share: An Attack on the Distinction between Agent-Relative and Agent-Neutral Values.” In *Creating the Kingdom of Ends*, 275–310. Cambridge: Cambridge University Press.

———. 2008a. “Realism and Constructivism in Twentieth-Century Moral Philosophy.” In *The Constitution of Agency: Essays on Practical Reason and Moral Psychology*, 302–26. Oxford; New York: Oxford University Press.

———. 2008b. “The Normativity of Instrumental Reason.” In *The Constitution of Agency*, 27–68. Oxford: Oxford University Press.

———. 2009. *Self-Constitution: Agency, Identity, and Integrity*. Oxford; New York: Oxford University Press.

———. 2013. “The Relational Nature of the Good.” In *Oxford Studies in Metaethics, Volume 8*, edited by Russ Shafer-Landau, 1–26. Oxford: Oxford University Press.

———. 2018. *Fellow Creatures: Our Obligations to the Other Animals*. Oxford: Oxford University Press.

Kraut, Richard. 1994. “Desire and the Human Good.” *Proceedings and Addresses of the American Philosophical Association* 68 (2): 39–54.

———. 2009. *What Is Good and Why: The Ethics of Well-Being*. Cambridge, Mass: Harvard University Press.

Langton, R. 2007. “Objective and Unconditioned Value.” *Philosophical Review* 116 (2): 157–85.

Lazari-Radek, Katarzyna de, and Peter Singer. 2012. “The Objectivity of Ethics and the Unity of Practical Reason.” *Ethics* 123 (1): 9–31.

Lin, Eden. 2017. “Against Welfare Subjectivism.” *Noûs* 51 (2): 354–77.

———. 2020. “Future Desires, the Agony Argument, and Subjectivism about Reasons.” *The Philosophical Review* 129 (1): 95–130.

Mackie, J. L. 1977. *Ethics: Inventing Right and Wrong*. Harmondsworth; New York: Penguin.

Maguire, Barry. 2016. “The Value-Based Theory of Reasons.” *Ergo, an Open Access Journal of Philosophy* 3 (9): 233–62.

Manne, Kate. 2013. “On Being Social in Metaethics.” In *Oxford Studies in Metaethics, Volume 8*, edited by Russ Shafer-Landau, 50–73. Oxford: Oxford University Press.

———. 2014. “Internalism about Reasons: Sad but True?” *Philosophical Studies* 167 (1): 89–117.

———. 2016. “Democratizing Humeanism.” In *Weighing Reasons*, edited by Errol Lord and Barry Maguire, 123–40. Oxford: Oxford University Press.

———. 2017. “Locating Morality: Moral Imperatives as Bodily Imperatives.” In *Oxford Studies in Metaethics, Volume 12*, edited by Russ Shafer-Landau, 1–26. Oxford: Oxford University Press.

Markovits, Julia. 2011. “Why Be an Internalist about Reasons?” n *Oxford Studies in Metaethics, Volume 6*, edited by Russ Shafer-Landau, 255–79. Oxford: Oxford University Press.

———. 2014. *Moral Reason*. Oxford: Oxford University Press.

McDowell, John. 1985. “Values and Secondary Qualities.” In *Morality and Objectivity*, edited by Ted Honderich, 110–29. London: Routledge.

Mencius. 2021. “Liang Hui Wang I.” Chinese Text Project, edited by Donald Sturgeon. ctext.org/mengzi/liang-hui-wang-i. Accessed June 11, 2021.

Millgram, Elijah. 1996. “Williams’ Argument Against External Reasons.” *Noûs* 30 (2): 197–220.

Moore, G. E. 1993. *Principia Ethica*. Edited by Thomas Baldwin. Cambridge; New York, NY: Cambridge University Press.

Nagel, Thomas. 1989. *The View from Nowhere*. New York, NY: Oxford University Press.

Parfit, Derek. 2011. *On What Matters*. 2 vols. The Berkeley Tanner Lectures. Oxford; New York: Oxford University Press.

Pettit, Philip, and Michael Smith. 1990. “Backgrounding Desire.” *The Philosophical Review* 99 (4): 565–92.

Prichard, H. A. 2002a. “Does Moral Philosophy Rest on a Mistake ?” In *Moral Writings*, by H. A. Prichard, edited by Jim MacAdam, 7–20. Oxford: Oxford University Press.

———. 2002b. “Duty and Interest.” In *Moral Writings*, by H. A. Prichard, edited by Jim MacAdam, 21–49. Oxford: Oxford University Press.

Railton, Peter. 1986. “Facts and Values.” *Philosophical Topics* 14 (2,): 5–31.

Rawls, John. 1971. *A Theory of Justice*. Cambridge, Mass: Harvard University Press.

———. 1980. “Kantian Constructivism in Moral Theory.” *The Journal of Philosophy* 77 (9): 515–72.

Raz, Joseph. 1988. *The Morality of Freedom*. Oxford: Oxford University Press.

———. 1999. *Engaging Reason: On the Theory of Value and Action*. Oxford; New York: Oxford University Press.

———. 2011. *From Normativity to Responsibility*. Oxford: Oxford University Press.

———. 2016. "Value and the Weight of Practical Reasons." In *Weighing Reasons*, edited by Errol Lord, 141–56. New York, NY: Oxford University Press.

Ridge, Michael. 2012. "Kantian Constructivism: Something Old, Something New." In *Constructivism in Practical Philosophy*, edited by James Lenman and Yonatan Shemmer, 138–58. Oxford: Oxford University Press.

Rosati, Connie S. 2009. "Relational Good and the Multiplicity Problem." *Philosophical Issues* 19 (1): 205–34.

Rosati, Connie S. 2003. "Agency and the Open Question Argument." *Ethics* 113 (3): 490–527.

Russell, Bertrand. 1904. "The Meaning of Good." The Bertrand Russell Society. https://users.drew.edu/~jlenz/br-moore-review2.html.

Scanlon, T. M. 2014. *Being Realistic about Reasons*. Oxford: Oxford University Press.

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

Scheffler, Samuel. 1995. *The Rejection of Consequentialism: A Philosophical Investigation of the Considerations Underlying Rival Moral Conceptions*. Oxford: Clarendon Press.

———. 2004. "Projects, Relationships, and Reasons." In *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*, edited by R. Jay Wallace, Philip Pettit, Samuel Scheffler, and Michael Smith, 247–69. Oxford: Oxford University Press.

———. 2010. "Morality and Reasonable Partiality." In *Partiality and Impartiality*, edited by Brian Feltham and John Cottingham, 98–130. Oxford: Oxford University Press.

———. 2011. "Valuing." In *Reasons and Recognition: Essays on the Philosophy of T.M. Scanlon*, edited by R. Jay Wallace, Rahul Kumar, and Samuel Freeman, 23–39. Oxford: Oxford University Press.

Schroeder, Mark. 2007. *Slaves of the Passions*. Oxford; New York: Oxford University Press.

Shah, Nishi. 2003. "How Truth Governs Belief." *The Philosophical Review* 112 (4): 447–82.

Singer, Peter. 1972. "Famine, Affluence, and Morality." *Philosophy and Public Affairs* 1 (3): 229–43.

Skorupski, John. 1999. “Irrealist Cognitivism.” *Ratio* 12 (4): 436–59.

———. 2012. *The Domain of Reasons*. Oxford: Oxford University Press.

Smith, Michael. 1995a. “Internal Reasons.” *Philosophy and Phenomenological Research* 55 (1): 109–31.

———. 1995b. *The Moral Problem*. Oxford; Cambridge, Mass: Blackwell.

———. 2012. “XV–Agents and Patients, or: What We Learn About Reasons for Action by Reflecting on Our Choices in Process-of-Thought Cases.” *Proceedings of the Aristotelian Society (Hardback)* 112 (3): 309–31.

Sobel, David. 2016a. *From Valuing to Value: A Defense of Subjectivism*. Oxford, United Kingdom: Oxford University Press.

———. 2016b. “Explanation, Internalism, and Reasons for Action.” In *From Valuing to Value: Towards a Defense of Subjectivism*, 144–60. Oxford: Oxford University Press.

———. 2016c. “Parfit’s Case against Subjectivism.” In *From Valuing to Value: Towards a Defense of Subjectivism*, 275–97. Oxford: Oxford University Press.

Strawson, P. F. 2008. “Freedom and Resentment.” In *Freedom and Resentment and Other Essays*, 1–28. London; New York: Routledge.

Street, Sharon. 2006. “A Darwinian Dilemma for Realist Theories of Value.” *Philosophical Studies* 127 (1): 109–66.

———. 2008. “Constructivism about Reasons.” In *Oxford Studies in Metaethics, Volume 3*, edited by Russ Shafer-Landau, 207–45. Oxford: Oxford University Press.

———. 2009a. “Evolution and the Normativity of Epistemic Reasons.” *Canadian Journal of Philosophy Supplementary Volume* 35: 213–48.

———. 2009b. “In Defense of Future Tuesday Indifference: Ideally Coherent Eccentrics and the Contingency of What Matters.” *Philosophical Issues* 19 (1): 273–98.

———. 2010. “What Is Constructivism in Ethics and Metaethics?” *Philosophy Compass* 5 (5): 363–84.

———. 2012. “Coming to Terms with Contingency: Humean Constructivism about Practical Reason.” In *Constructivism in Practical Philosophy*, edited by James Lenman and Yonatan Shemmer, 40–59. Oxford: Oxford University Press.

Supreme Court of the United States. 2015. Obergefell v. Hodges. 576 U.S. 644.

Tully, Ian. 2017. “Depression and the Problem of Absent Desires.” *Journal of Ethics and Social Philosophy* 11 (2): 1–16.

Velleman, James David. 2000. “On the Aim of Belief.” In *The Possibility of Practical Reason*, 244–81. Oxford: New York: Oxford University Press.

Wallace, R. Jay. 2012. “Constructivism about Normativity: Some Pitfalls.” In *Constructivism in Practical Philosophy*, edited by James Lenman and Yonatan Shemmer, 18–39. Oxford: Oxford University Press.

Watson, Gary. 2004. “Responsibility and the Limits of Evil: Variations on a Strawsonian Theme.” In *Agency and Answerability*, 219–59. Oxford: Oxford University Press.

Williams, Bernard. 1973. “A Critique of Utilitarianism.” In *Utilitarianism: For and Against*, 77–150. Cambridge: Cambridge University Press.

———. 1981. “Internal and External Reasons.” In *Moral Luck: Philosophical Papers 1973–1980*, 101–13. Cambridge: Cambridge University Press.

———. 1985a. *Ethics and the Limits of Philosophy*. Cambridge, Mass: Harvard University Press.

———. 1985b. “Formal and Substantial Individualism.” *Proceedings of the Aristotelian Society* 85: 119–32.

———. 1995. “Internal Reasons and the Obscurity of Blame.” In *Making Sense of Humanity, and Other Philosophical Papers 1982-1993*, 35–45. Cambridge: Cambridge University Press.

———. 2005. “Human Rights and Relativism.” In *In the Beginning Was the Deed: Realism and Moralism in Political Argument*, 362–74. Princeton, N.J.: Princeton University Press.

Wong, David. 2002. “Reasons and Analogical Reasoning in Mengzi.” In *Essays on the Moral Philosophy of Mengzi*, edited by Xiusheng Liu and P. J. Ivanhoe, 187–220. Indianapolis, Ind: Hackett Publishing.

———. 2020. “Practical Reasoning in Early Chinese Philosophy.” In *The Routledge Handbook of Practical Reason*, edited by Ruth Chang and Kurt Sylvan, 113–25. Abingdon, Oxon; New York, NY: Routledge.