

THESIS

CONSTRUCTIVISM AND THE PHENOMENOLOGY OF MORAL DELIBERATION

Submitted by

Alexander Ponzo

Department of Philosophy

In partial fulfillment of the requirements

For the Degree of Master of Arts

Colorado State University

Fort Collins, Colorado

Fall 2017

Master's Committee:

Advisor: Elizabeth Tropman

Moti Gorin

Anne Cleary

Copyright by Alexander Ponzo 2017

All Rights Reserved

ABSTRACT

CONSTRUCTIVISM AND THE PHENOMENOLOGY OF MORAL DELIBERATION

In this thesis I claim that constructivism is an attractive position when it comes to explaining the nature of moral facts. The central tenet of constructivism is that moral facts are a function of what actual (or hypothetical) people think. According to constructivism, if it is true (for example) that murder is morally wrong, it is because some actual (or hypothetical) person or group of people under certain conditions thinks that murder is morally wrong. A primary competitor to constructivism is realism. Realism denies the central tenet of constructivism and instead holds that moral facts are objective; this means that the truth of moral facts does not depend on any person's or group of persons' attitudes regarding the object of moral evaluation. For realists, then, if murder is morally wrong it is not because any (actual or hypothetical) person or group of people thinks that it is wrong; for realists, this fact obtains regardless of what anyone thinks about it. I argue in Chapter One that there is a significant epistemological problem with the claim that moral facts are objective; namely, that it is hard to see how we could have moral knowledge if the truth of moral facts is not a function of what we think. I argue in Chapter Two that constructivism does not share the epistemological problem that realism has. This problem is significant and it should be enough, I submit, to push us into accepting constructivism over realism.

While I devote the first chapter (and some of the second) to explaining what constructivism is and how it compares to other competing theories, this is not the only task of this thesis. My other central task is to examine two specific versions of constructivism—namely,

the respective theories of Christine Korsgaard and Sharon Street—and find out whether each theory can maintain a plausible phenomenology of moral deliberation. In other words, I am interested in finding out whether each version of constructivism can give a plausible account of *what it is like* for people to engage in moral deliberation. Since it's possible to study the phenomenology of moral deliberation empirically (or by plausible speculation), we can compare how moral deliberation is actually experienced—or at least how it seems to be experienced—with how it would likely be experienced if the theories in question were true. Ultimately, I argue that the phenomenology of moral deliberation occasioned by Korsgaard's position is more plausible than the one occasioned by Street's because Korsgaard's position does a better job of preserving what we tend to think the experience of moral deliberation is like.

A summation of this thesis can be given in the following way: I make a presumptive case for constructivism and, specifically, I argue that Korsgaard's constructivism is more attractive than Street's because, among other things, the phenomenology of moral deliberation occasioned by Korsgaard's position is more plausible than the one occasioned by Street's.

TABLE OF CONTENTS

ABSTRACT.....	ii
TABLE OF CONTENTS.....	iv
INTRODUCTION	1
1. CHAPTER ONE – CONSTRUCTIVISM AND THE EPISTEMOLOGICAL PROBLEM FOR REALISM	4
1.1 CONSTRUCTIVISM’S PLACE IN METAETHICS	4
1.2 MOTIVATIONS FOR GOING CONSTRUCTIVIST	13
1.3 THE EPISTEMOLOGICAL PROBLEM FOR REALISM	19
2. CHAPTER TWO – THE METAETHICAL CONSTRUCTIVISMS OF KORSGAARD AND STREET.....	25
2.1 KORSGAARD’S CONSTRUCTIVISM	25
2.2 A PROBLEM FOR KORSGAARD’S CONSTRUCTIVISM	33
2.3 STREET’S CONSTRUCTIVISM	35
2.4 COMPARING AND CONTRASTING KORSGAARD AND STREET	37
2.4.1 HUMEAN CONSTRUCTIVISM VS KANTIAN CONSTRUCTIVISM	38
2.4.2 STREET’S CONSTRUCTIVISM AND INTUITIVELY WRONG ACTS	41
2.4.3 CONSTRUCTIVISM AND THE EPISTEMOLOGICAL CHALLENGE	48
2.5 CONCLUSION.....	50
3. CHAPTER THREE – CONSTRUCTIVISM AND DELIBERATION	52
3.1 ENOCH’S OBJECTION: CAN GLOBAL CONSTRUCTIVISM MAKE SENSE OF DELIBERATION?.....	54
3.2 ENOCH’S CASE FOR WHY ROBUST REALISM CAPTURES DELIBERATION.....	56
3.3 THE OBJECTION REGARDING DELIBERATION AS TAILORED TO STREET’S CONSTRUCTIVISM	58
3.4 A REPLY TO THE OBJECTION ON BEHALF OF STREET’S CONSRUCTIVISM	62
3.5 WHY MORAL DELIBERATION ACCORDING TO STREET’S POSITION IS FLAWED ..	65
3.6 MORAL DELIBERATION: WHY KORSGAARD DOES IT BETTER THAN STREET ..	71
3.7 CONCLUSION.....	74
4. REFERENCES	77

INTRODUCTION

My first task in this thesis is to make a presumptive case for metaethical constructivism. The central tenet of metaethical constructivism is that moral facts are a function of what actual (or hypothetical) people think. According to constructivism, if it is true, for example, that murder is morally wrong, it is because some actual (or hypothetical) person or group of people under certain conditions thinks that murder is morally wrong.

As you probably guessed from the name ‘metaethical constructivism’, the position in question is a *metaethical* position. Accordingly, this is a thesis in metaethics. Roughly put, metaethics is the area of philosophy in which we try to understand the metaphysical, epistemological, semantic, and psychological commitments of our thoughts and practices regarding moral matters.¹ For example, metaethicists ask whether there really are moral facts and, if so, whether they are natural or non-natural. And they ask, if there are moral facts, how do we come to know them? With respect to moral practices, metaethicists take great care in examining how we use moral language—are we attributing properties to things as it seems we are doing when we utter a standard sentence like “The banana is yellow” or are we expressing emotions, preferences, or prescriptions? If we are ascribing moral properties to things, then it seems as if moral statements are capable of being true or false. If this is true then metaethicists can ask another question: are moral utterances true or false relative to what an individual or a society thinks, or is their veracity objective? While these inquiries do not exhaust all the possible questions metaethicists might investigate, it should be enough to give a fairly clear picture of what metaethicists are up to.

¹ Sayre-McCord (2014)

In Chapter One, I will examine a few of these questions in an effort to create a metaethical map. Each theory's answers to the above questions (or a subset of them) will serve to distinguish it from other metaethical positions. In Chapter One I will offer several reasons to prefer constructivism over other metaethical theories. Constructivism's main competitor is realism. Realists deny the central tenet of constructivism and instead hold that moral facts are objective; this means that the truth of every moral fact does not depend on any person's or group of persons' attitudes regarding the object of moral evaluation. For realists, then, if murder is morally wrong it is not because any (actual or hypothetical) person or group of people thinks that it is wrong; for realists, this fact (and every moral fact) obtains regardless of what anyone thinks about it. I argue in Chapter One that there is a significant epistemological problem with the claim that moral facts are objective; namely, that it is hard to see how we could have moral knowledge if the truth of moral facts is not a function of what we think. This problem is significant because constructivism and realism share so many of the same advantages over other theories—it is something that constructivists can point to as an advantage over realism.

In Chapter Two, I will zoom into the metaethical map and focus on two distinct versions of metaethical constructivism. These are the positions of Christine Korsgaard and Sharon Street. In addition to giving a detailed explanation of what each of these theories is and which claims each theory does and does not maintain, I will examine central problems that each of these theories face. Here I will explicitly argue that both Korsgaard and Street avoid the epistemological problem that faces realism. It is in this chapter (and Chapter Three) that I argue for my other central claim in this thesis; that is, I argue that Korsgaard's constructivism is more attractive than Street's because, among other things, the phenomenology of moral deliberation occasioned by Korsgaard's position is more plausible than the one occasioned by Street's. When

I use the word ‘phenomenology’ I mean to indicate the experience of what it is like for someone to do something. In this case, we will look at the experience of what it is like to deliberate about moral questions

In Chapter Three we’ll see what the phenomenology of moral deliberation would likely be for the respective theories of Korsgaard and Street. Here I will appeal to a plausible phenomenology of moral deliberation given by Enoch (2011). This characterization claims that moral deliberation is similar to straightforward deliberation about non-moral facts (such as whether Denver is farther north than San Francisco). During such deliberations it feels as if we are trying to find a correct answer and that the answer is correct regardless of whatever answer we come to at the end of deliberation. I will provide more justification for why I think this is a plausible phenomenology of moral deliberation in Chapter Three.

We can think of this document as having the shape of a funnel. In Chapter One we start off with a wide overview of the metaethical terrain and where exactly constructivism is supposed to fit in. In Chapter Two we get a bit narrower and focus on the advantages and weaknesses of the respective theories of Korsgaard and Street. Then in Chapter Three we get even narrower and focus on one issue facing these two views; that is, the challenge of maintaining a plausible phenomenology of moral deliberation.

This thesis can be boiled down into the following sentence: I make a presumptive case for constructivism and, specifically, I argue that Korsgaard’s constructivism is more attractive than Street’s because, among other things, the phenomenology of moral deliberation occasioned by Korsgaard’s position is more plausible than the one occasioned by Street’s.

CHAPTER ONE: CONSTRUCTIVISM AND THE EPISTEMOLOGICAL PROBLEM FOR REALISM

The purpose of this chapter is to present a map of several popular positions in the field of metaethics. The central focus will be on a group of positions called metaethical constructivism. Before I present those theories it will be helpful to get clear on where exactly constructivism is supposed to fit on the metaethical map. This will be the task of section 1.1. Section 1.2 will be a discussion of the purported benefits of constructivism over other theories. Many of the benefits that metaethical constructivism is thought to have over other theories are shared by metaethical realism. The goal of section 1.3 will be to present the epistemological problem for realism as it is characterized by Street.² If it is true that constructivism and realism are preferable to other metaethical theories and that realism has a significant problem that constructivism does not share, then, barring any severe problems unique to constructivism, constructivists can claim to have an edge over realism and other competing theories.

1.1 Constructivism's place in metaethics

What is constructivism in metaethics? For now I would like to offer just a general characterization of it. The specific versions I will discuss in Chapter Two have their own unique bells and whistles, but what I want to do here is to show where constructivism is generally put on the metaethical map; in other words, I want to show which borders it does and doesn't share with other positions.

First, metaethical constructivism is a type of cognitivist position. Cognitivists treat moral utterances—such as “Stealing is wrong”—as capable of being true or false. Non-cognitivists

² Street, 2006.

treat moral language differently. For them, moral utterances are not capable of being true or false; instead, what's really going on is that when we utter moral statements we express states such as non-cognitive attitudes, emotions, or prescriptions. For example, emotivism—a popular non-cognitivist position from the middle of the twentieth century—treats the statement “Stealing is wrong” as something akin to “Boo stealing!”³ Of course “Boo” is not an emotion, but it is supposed to stand for an emotional attitude internal to the subject who utters the moral statement. This understanding of “Stealing is wrong” is not something that is true or false because it's just an expression of emotion. Similarly, an emotivist position might render a claim of moral approbation, such as “Donating is the right thing to do”, in the following way: “Hurray donating!”. Like “Boo stealing!”, this utterance is incapable of being true or false.

Some non-cognitivist positions get confused with a cognitivist position called subjectivism. Subjectivists hold that the rightness or wrongness (or goodness or badness) of an action depends on whether an individual thinks it is right or wrong (or good or bad). According to this view, the phrase “Stealing is wrong” is really more akin to “I disapprove of stealing”. Another way of phrasing this point is that the sentence “Stealing is wrong” is made true or false by whether or not I disapprove of stealing. My opinion about the issue is the truth-maker. While this subjectivist move might seem similar to turning “Stealing is wrong” into “Boo Stealing!”, they are crucially different in that the subjectivist's claim is truth-apt while the emotivist's is not. The truth of the subjectivist's claim depends on whether or not in fact I disapprove of stealing. “Boo Stealing!” is not apt for the assignment of truth or falsehood.

The cognitivist view it is most important to distinguish constructivism from is realism. Realism describes any view which holds that:

- (1) there are moral properties (such as right and wrong)

³Ayer (1952) offers a classic version of emotivism.

(2) some of them are instantiated, and

(3) they are objective.

By describing moral properties as ‘objective’, realists mean that an instance of a moral property must not depend on any agent’s or group of agents’ attitudes regarding the object of moral evaluation.⁴ If a theory holds that instances of moral properties depend on a particular agent’s attitudes regarding the object of moral evaluation then that theory is subjective in the way I sketched above. It’s important to note, though, that for realists the objective moral facts they claim exist might be *about* an agent’s attitudes—for example, it might be a fact that someone’s hateful attitudes towards certain ethnic groups is morally wrong—but the truth of objective moral facts does not *depend* on any agent’s attitudes.

Constructivists agree that (1) and (2) are true, but they disagree about (3). But denying that moral properties are objective is not so simple. One might think that in denying (3), constructivists must accept that moral properties are subjective. If this were true it might seem like constructivism would have no room to be its own distinct position because subjectivism already occupies the role of a cognitivist position which holds that there are instantiated subjective moral properties. This is a tricky matter and one that I will hold off on giving a full response to until the positions of Korsgaard and Street are presented fully in Chapter Two. However, there are two initial things to point out in response to the subjectivism/objectivism dichotomy regarding moral properties that should make the prospect of constructivism achieving its own space on the metaethical map seem more likely.

⁴ At times I will just speak about moral *properties* and at other times I will speak about *facts* regarding those properties. When an action or object has a property we say that there is a fact about that action or object. For example, if stealing has the property of wrongness, we would say that it is a fact that stealing is wrong. I will trade off which terms I use because sometimes talk of properties fits the sentence better than talk of facts (and vice versa).

First, it is plausibly a false dichotomy to think that moral properties have to be either subjective or objective. They might instead be intersubjective in the sense specified by Sayre-McCord: “Intersubjectivism grants (with subjectivism) that people figure in the truth-conditions [of moral claims], but it holds that the truth of moral claims doesn’t turn on facts about *particular* [emphasis added] individuals” (1986, p. 20). Thus, intersubjectivism positions itself as a middle ground between objectivism and subjectivism. It affirms that moral properties depend on agents’ attitudes (viz., it denies (3)), but it rejects that moral properties depend on the attitudes of *particular* agents.

There are several ways in which an intersubjective position might be developed and I would like to focus on two of them. The first is cultural relativism. This takes the truth and falsity of moral claims to be dependent on the actual practices and conventions of particular societies. For example, if a society condones honor killings then this practice is morally right or permissible. Constructivists want to distinguish themselves from this kind of intersubjectivism for a few reasons. The obvious one is that they want to have a unique place on the metaethical map. The other reasons have to do with the shortcomings of cultural relativism. One shortcoming is that this picture of morality makes it hard to see how we could ever make sense of change in a society’s practices and conventions. Just ten or fifteen years ago in the United States gay marriage was not practiced and was not allowed by law. But in 2015 the Supreme Court ruled that it was legal and many people now think that marriage should be a legal right for gay couples. Are we to say that gay marriage was wrong a few years ago but right now? The cultural relativist has to say that the allowance of gay marriage is morally right in 2017 because it has become conventional—but the shift in convention is better explained by the likelihood that people reconsidered the reasons they thought gay marriage was wrong in the first place and then

changed their minds. This point leans against another problem for cultural relativism which Sayre-McCord points out (1986, p. 20): it seems that cultural relativism can't explain how we could intelligibly make moral objections to the habits approved of by our society. Under such a position, members of society can't say that a social convention is wrong because *by definition* it is right.

But there is another way to incorporate intersubjectivism into a metaethical theory. This way is idealistic in the sense that it requires that we abstract away from actual practices and people and instead use the hypothetical practices and conventions of hypothetical people as the agents who determine the truth of moral claims (Sayre-McCord, 1986, pp.7-8). This seems to avoid the two problems of cultural relativism discussed above. This idealistic intersubjectivism can take the case of gay marriage in the U.S. and say that a hypothetical group of people with full access to legal and moral scholarship would decide that gay marriage has been right all along. Even if one argues that the same group of hypothetical people would decide that gay marriage has always been wrong, we dodge the problem of a seemingly arbitrary shift in the moral permissiveness of gay marriage. Notice that this also solves cultural relativism's problem of not being able to explain how we could intelligibly make moral objections to the habits approved by our society—idealistic intersubjectivism allows dissenters to argue that a hypothetical group of people given all the facts relevant to whatever issue is at hand would decide that the current practice in question is wrong.

The second point I want to make in regard to the line constructivists walk between subjective and objective moral properties is that positions which claim that there are subjective moral properties are generally presented as a crude subjectivism of the sort I described above. It's at least possible, on first blush, that a more sophisticated metaethical theory that holds moral

properties to be subjective could ultimately be an attractive theory. This point will become of interest in Chapter Two when I discuss how Korsgaard and Street characterize moral properties.

For now, I want to give the following general definition of constructivism bearing in mind that many of the details will have to be filled in by the respective theories of Korsgaard and Street. Here's the general formulation: for constructivists, moral properties are constructed from the practical point of view. There are at least three questions that, by my lights, need to be answered in response to this definition:

Question 1: What does 'constructed' mean?

Question 2: What is the practical point of view?

Question 3: Does this definition of constructivism necessarily commit constructivists to objectivism, intersubjectivism, or subjectivism about moral facts?

Question 1: When I say that moral properties or facts are constructed, I have in mind Enoch's definition of metaethical constructivism:

According to metaethical constructivism, there are no actions we ought to perform or to avoid performing independently of some (actual or hypothetical) procedure – say, that of reaching consensus in an open discussion about principles of conduct; rather, these procedures determine the moral status of actions (2009, p. 322).

On this definition, there are no moral facts about what we ought to do in the moral realm without performing some actual or hypothetical procedure. This is what I mean by 'construct'—that is, I mean that moral facts are created by a procedure that is either actual or hypothetical. As this might still be unclear, let me give an example of one of the most famous and, I think, easiest to understand examples of constructivism. Consider a certain reading of John Rawls's views about justice in liberal democratic societies. In *A Theory of Justice* (1971), Rawls can be read as a

constructivist about social and political justice.⁵ On this reading, facts about social and political justice consist in their being constructed or created from the point of view of the original position. By ‘original position’ I am referring to the hypothetical situation we are supposed to imagine in which hypothetical deliberators who are free and equal persons commit themselves to coming up with principles of social and political justice (Rawls, 1971, pp. 17-21). When we imagine these deliberators in the original position, we assume that they take for granted some facts regarding the nature of fair bargaining conditions and about how traits such as sex, race, and class shouldn’t make a difference from a moral point of view. On this reading, the principles of justice that these hypothetical people would choose are Rawls’s principles of justice.

Metaethical constructivists claim that all moral facts, including the ones assumed in the original position (if one is a Rawlsian), are constructed from the practical point of view. Metaethical constructivism, while not entirely global, casts a much wider net than the above reading of Rawls (on that reading, only facts about social and political justice are constructed). While there is plenty of discussion regarding views which are said to be ‘globally’ constructivist (Enoch, 2009; Street, 2010), in this thesis I am limiting myself to the discussion of metaethical constructivism (except when brief detours into global constructivism are needed to make a point about metaethical constructivism). The main difference between the two is that global constructivism claims that *all* normative facts are constructed, whereas metaethical constructivism is only committed to the claim that *moral* facts are constructed. Moral facts are necessarily normative, but there are often thought to be other, non-moral kinds of normative fact. To illustrate, consider the following example of a claim that a global constructivist will say is constructed, but that a metaethical constructivist need not have a stance on: “You ought to

⁵ See Street (2010, pp. 364-366) for such a reading.

proportion your beliefs to the evidence you have for them”. Such a claim is normative, but not—at least not obviously—moral.

Question 2: What is the practical point of view? Put simply, it is the position we occupy when we perform practical reasoning. A helpful way to understand practical reasoning is to distinguish it from theoretical reasoning which is the kind of reasoning concerned with matters of explanation and prediction. Following Wallace (2014), we can say about theoretical reasoning that, “looking backward to events that have already taken place, it asks why they have occurred; looking forward, it attempts to determine what is going to happen in the future” (Practical and Theoretical Reason, para. 2). In other words, theoretical reasoning is largely concerned with matters of fact and their explanation. Practical reasoning, on the other hand, is concerned with our decisions about which actions to take when we face a problem. When we reason practically, we are always in a situation which requires us to choose a path of action among several options. In choosing, we must consider what reasons we have for acting, the implications of each path of action, and, ultimately, which action would be best to commit.

Rawls uses the practical point of view in his theory of justice. The problem in his case is one of distribution. Liberal democratic societies need a way to distribute the benefits and burdens of living in a cooperative society (Korsgaard, 1996b, p. 113). Here, once we take into account the moral constraints mentioned above and have our hypothetical deliberators enter the original position, the distributive problems of liberal democratic societies are supposedly solved.

Question 3: When an actual or hypothetical agent (or group of agents) enters the practical point of view and is ready to construct moral facts, will the facts created be subjective, intersubjective, or objective? Recall that in order for an instance of a moral property to be objective it must not depend on any agent’s or group of agents’ attitudes regarding the object of

moral evaluation. We can rule out that constructed moral properties will be objective since they will always be constructed from the practical point of view of some actual or hypothetical observers. Ultimately, whether a constructivist metaethical position considers moral facts subjective or intersubjective will depend on the details of that position. Thus constructivism as given a general characterization is only committed to the negative claim that moral facts are not objective.

The blueprint of metaethical constructivism I've given here will be filled in by the details of the respective positions of Korsgaard and Street in Chapter Two. I have delayed that task until Chapter Two so that I can preserve a narrative in this chapter starring the two main characters of this chapter; i.e., metaethical constructivism and metaethical realism. Section 1.2 tells the story of the benefits of adopting metaethical constructivism—metaethical realism is thought to share all of these benefits. In section 1.3 I'll introduce a problem that is thought to be unique to realism.

1.2 Motivations for going Constructivist

There is a particular frame I would like to give this section. Metaethical constructivism is a relatively new stance that has gained popularity in recent years often because it is thought to capture the insights of realism without some of its purported costs (Enoch, 2009). I will begin with a discussion of why I think the cognitivism of realism and constructivism is strength—this largely amounts to a discussion of why non-cognitivism is problematic. I will then note a few of the other strengths realism and constructivism are thought to share.

A thorough discussion of the pros and cons of cognitivism and non-cognitivism would require more than I can do here, but I would like to present an important objection to non-cognitivism that some take to be a significant reason to prefer cognitivism to non-cognitivism. I

will use emotivism as the exemplar non-cognitivist position just to avoid the redundancy of having to make the same argument against each different non-cognitivist position. If we were to use a different form of non-cognitivism here in place of emotivism, the same objections would still ultimately apply. Indeed, these objections are aimed at any and every position classified as non-cognitivist.

Many take the linguistic evidence against non-cognitivism and in favor of cognitivism as the best reason to prefer cognitivism.⁶ We typically utter statements with grammatical forms similar to the following statements: “Stealing is wrong”, “Donating to Doctors without Borders is the right thing to do”, and “Lying is bad”. Each of these appears to be a straightforward sentence in which the property of rightness or wrongness (or goodness or badness) is attributed to an action. This is the same way we analyze uncontroversial sentences like “Bananas are yellow” or “That guitar is made of mahogany”. In all of these statements—the moral and the non-moral alike—a property is predicated upon an object which makes the statement capable of being true or false. Furthermore, this is the most intuitive way to understand the statement.⁷ But non-cognitivist positions have to give a different account of the meaning of moral statements—one that involves rejecting the *prima facie* way of understanding statements. Return again to the emotivist understanding of the sentence “Stealing is wrong”. The emotivist renders this sentence an expression of one’s emotions represented in the following way: “Boo Stealing!”. This expression is not capable of being true or false and thus it fails to keep with the way we generally analyze language.

⁶ Glassen (1959); Huemer (2005).

⁷ By “intuitive” I mean something like “the most common way of understanding the statement for competent speakers of the (in this case) English language.”

Following Huemer (2005, pp. 18-25), I think it's plausible to say that attempts to make sense of moral language by cognitivists and non-cognitivists alike is a task with its goal being to accurately portray the current meaning and function of moral language; neither side is suggesting that we change the way we speak of moral matters. This makes it so that we can check each theory's claims about how we use language. If a theory were to claim that 'chair' means 'thing with a tail' we could check to see if competent users of a language actually mean this. This strategy has the desirable feature of being such that we can just check uncontroversial cases of how competent speakers use the language (Huemer, 2005, p. 19).

Huemer provides an extensive list of linguistic examples which favor cognitivism (2005, pp. 18-25). I will present four of them. The general strategy of my argument here will be to show that since only cognitivism can successfully capture the way we use evaluative language, we have strong reason to reject non-cognitivism.

- (1) Evaluative statements are declarative statements. We say things like "Helping is the right thing to do". This has the same form as "This banana is yellow". These kinds of sentences are usually how we assert facts. Statements like "Boo Stealing!" and "Hurray donating!" are not declarative statements
- (2) Moral predicates can be transformed into abstract nouns. This suggests that they refer to properties. We can change, for example, 'good' into 'goodness' and 'wrong' into 'wrongness'. That is, it makes sense to say, "It's not that I don't enjoy some aspects of eating meat, it's that I worry about its wrongness."
- (3) We can describe evaluative statements as having the same properties as other, non-controversial statements. For example, we can add "It is true that", "It is false that", and "It is possible that" in front of any evaluative statement, as in "It is possible that I

was wrong to act in such a way”. We cannot, however, say “It is possible that hurray donating!” or “It is true that boo murder!”

- (4) We can transform evaluative statements into yes or no questions. I can assert that “Bananas are yellow” and I can also ask “Are bananas yellow?”, just as I can assert that “Murder is wrong” and ask “Is murder wrong?” However, it doesn’t seem like we can make sense of an emotional expression like “Boo murder?” or “Hurray a woman’s right to choose an abortion?”

Alone, each of these observations is not enough to craft an argument against non-cognitivism. However, the fact that non-cognitivism doesn’t easily capture all four of those common ways we use language is a reason to suggest that non-cognitivism might be false. Of course the cognitivist/non-cognitivist debate has volumes of its own, but my discussion of it must end here.

Cognitivism is not the only strength that realism and constructivism share—they each can allow the possibility of moral disagreement, moral error, and moral progress (Tropman, 2014). Subjectivism and non-cognitivism cannot hope to leave open all of these possibilities.

How is either subjectivism or non-cognitivism supposed to make sense of moral disagreement? Recall that subjectivism is the view which holds that the rightness or wrongness (or goodness or badness) of an action depends on whether an individual thinks it is right or wrong (or good or bad). Consider the case of abortion; it is a highly charged topic the moral permissibility of which is debated with much verve. Some folks are convinced that abortion is always morally wrong and others are convinced that it is permissible at least in some cases. The subjectivist is committed to the claim that the moral rightness or wrongness of abortion depends on whatever an individual thinks is the case. But this should strike us as odd—people who are pro-life think that abortion is wrong no matter what anyone thinks about it! Similarly, people

who are pro-choice will think abortion is permissible in certain cases regardless of what anyone thinks about it. (Of course they will care what the pregnant woman thinks about the matter, but this doesn't change the fact that they would think that a woman should be able to choose an abortion in many cases even if she decides to go ahead with the birth).⁸

The non-cognitivist—again using emotivism as the model non-cognitivist position—will say that claims about the rightness or wrongness of abortion are just expressions of emotion. Thus, when people claim that “abortion is wrong” they are really just expressing something more akin to “Boo abortion!”, where “Boo” represents a non-cognitive emotional attitude internal to the agent who utters the statement. Similarly, when people in the pro-choice camp claim that “abortion is morally permissible” they will be expressing an internal emotional attitude more akin to “Hurray a woman’s right to choose!”. Importantly, “Boo abortion!” and “Hurray a woman’s right to choose!” are not capable of being true or false on a non-cognitivist framework. We cannot genuinely disagree if we are not even making claims that are being capable of being true or false.

When it comes to allowing for the possibility of moral error—something that seems to happen frequently—subjectivism and non-cognitivism fail. Imagine a person who as a child thought that stealing wasn't morally problematic. After acquiring some formal education and accumulating life experience, she comes to think that stealing is wrong. What happened in this example is that the woman came to see that she made an error in her moral evaluation when she was a child. But this story doesn't make sense according to subjectivism. Since subjectivists are committed to the claim that the moral rightness and wrongness of an action is determined by whatever an individual thinks is the case, an individual cannot error because *by definition*

⁸ Of course those who are pro-choice may disagree about term limits and other factors, but there are cases where virtually all of them will agree; for example, nearly all pro-choice folks will agree on cases of incest or rape where young women and girls face the unwelcome prospect of being an extremely young mother.

whatever judgement she makes about the rightness or wrongness of an action is true. The woman's younger self was not in error—she merely made a different moral appraisal. This should strike us as problematic because there are many apparent cases of moral error.

Non-cognitivists do no better here. The child version of the woman from the prior paragraph would say that stealing is permissible, while the adult version would say that stealing is wrong. The emotivist non-cognitivist will say that the child really expresses an emotional attitude akin to “Hurray Stealing!” while the adult expresses an emotional attitude akin to “Boo Stealing!”. As I'm sure you've anticipated, these two expressions are not capable of being true or false. This is problematic for non-cognitivists if they hope to be able to explain how we make moral errors, because in order for us to make an error in moral judgement there has to be a fact about the matter at hand to be wrong about. But non-cognitivists can only say that the woman's emotional attitude toward stealing has changed over time, rather than that she was making an error in judgement when she was a child.

This example involving the woman and her past self brings to light another problem for subjectivism and non-cognitivism—namely, that they cannot account for moral progress. Most of us like to think that throughout our lives we grow and come to see things more accurately than we used to; presumably, some evidence of personal growth can be found in our judgements about morality. But it's hard to see how subjectivists can make sense of this. The woman who changes her opinion about the moral permissibility of stealing throughout her life cannot be said to be progressing in the way we tend to think we do. Why is this? It is because the individual determines what is right and wrong. Thus the woman cannot say that she was wrong when she was younger to think that stealing was permissible—individuals determine what is right and wrong, so there is no way that they can be mistaken.

Non-cognitivists again face the problem that in order for us to claim that we were wrong about certain moral judgements, we must have made a claim that we could have been wrong about—but when we are only expressing emotional attitudes we are not making claims capable of being false.

Thus far, realists and constructivists are on even ground. They each get to claim the benefits of cognitivism while eschewing the many problems that plague subjectivism and non-cognitivism. The goal of the next section is to present an objection that is supposed to apply uniquely to realism.

1.3 The Epistemological Problem for Realism

Realism's biggest problem has to do with epistemology (Enoch, 2009, p. 324). The classic problem should be credited to Mackie (1977, p. 38). According to Mackie's characterization of realism, moral facts are non-natural and intrinsically motivational,⁹ rendering them extremely strange. Because of this, he thinks that the only way we could come to know these facts would be through a mysterious faculty dedicated to the detection of such strange facts—yet there is no reason to think we have such a faculty. If there is no way we could come to know these non-natural moral facts, then we can't have moral knowledge. To accept this conclusion would amount to accepting skepticism about moral knowledge—this is something that realists are keen on avoiding. Since Mackie presented this classic problem, realists have developed several ways to respond. For example, Brink argues that objective moral properties need not be intrinsically motivational or non-natural (1989). Additionally, even if moral

⁹ While there are several ways to characterize non-natural moral facts, it will suffice to say that for Mackie they are facts that are in no way reducible to or identical to natural properties; a property is natural just in case it is studied by the natural or social sciences (Brink, 1989). By describing moral facts as intrinsically motivational we mean that once an agent makes a moral judgement, that judgement will at least somewhat motivate the agent to act in accordance with it. So for example, if I judge that it is wrong to steal, I will be at least somewhat motivated to not steal.

properties are non-natural, it may not be as objectionable as it first seems—other areas such as mathematics and other parts philosophy sometimes appeal to the existence non-natural properties and facts (Audi, 1997; Shafer-Landau, 2003).

Recent epistemological objections to realism do away with discussion of strange moral properties and perceptual faculties. The new objections instead allow that there could be objective moral facts, but they then argue that these would not explain why we form the moral beliefs we do. If these objective moral facts don't explain why we form the moral beliefs we do, then we can't have moral knowledge—this will become clear in a moment.

Sometimes epistemological objections to realism appeal to evolutionary theory. The thought is that the real reason we currently judge certain things right (or good) and certain things wrong (or bad) is that it was adaptive for our ancestors to react in such ways when faced with certain stimuli. There are several ways to develop this objection, but I would like to focus on Street's (2006) version of the argument. I do so for two reasons. The first is that it does a nice job of articulating the objection and the second is that she takes this objection to be a reason to prefer her brand of metaethical constructivism—I agree with her on this point.

Street begins by noting that, evolutionarily speaking, there were some motivational tendencies or proto-judgments that were clearly better than others for the reproductive success of our ancestors (2006, pp. 13-21). By 'proto-judgement' I mean to indicate "an unreflective, non-linguistic, motivational tendency to experience something as 'called for' or 'demanded' in itself, or to experience one thing as 'calling for' or 'counting in favor of' something else" (Street, 2006, p. 119). The creature that values putting itself in harm's way is clearly going to be less reproductively successful than the creature that avoids harm and instead values acquiring sufficient nutrition for itself and its kin. Having motivational tendencies to avoid harm, to seek

out proper nutrition, and to help one's kin accomplish these goals would essentially be to enjoy a higher level of fitness than one's contemporaries. These motivational tendencies were heritable and over time became frequent in the population while the motivational tendencies of less fit creatures became more uncommon. According to Street, what is important about this is that these motivational tendencies, these proto-judgements, influenced the capacity of our more recent ancestors to form full-fledged cognitive judgements about morality. In fact, Street argues, the selection of certain motivational tendencies also affected the *content* of our current moral judgements. For example, Street would say that the reason we now generally think that stealing from our kin is wrong is because having the motivational tendency to respect and help one's kin contributed to the reproductive success of our ancestors.¹⁰

This puts a stick in the spokes of realism. Because now, if Street's explanation of how we come to form the moral beliefs we have is correct, even if objective moral properties exist, it's highly unlikely that our moral judgements are true. This is because our moral judgement making capacities are tainted by evolutionary forces in such a way that we are likely to be skewed towards making moral judgements that are reproductively advantageous and not necessarily objectively true; and since truth is a necessary condition for knowledge, we cannot say that what we come to believe via our moral judgements counts as knowledge. While it's true that this argument doesn't establish that there can't be objective moral properties, its conclusion that we don't have moral knowledge should not be taken lightly. To endorse a metaethical theory that denies the possibility of moral knowledge is a move in the direction of skepticism about morality, something that most, if not all, realists want to avoid.

¹⁰ Street acknowledges that we can sometimes use rational reflection to change our moral beliefs and judgements. However, she thinks that even rational reflection has been influenced by evolutionary forces since rational reflection ultimately—she argues—requires the weighing of some evaluative or moral judgements against others (2006, pp. 123-124).

One might offer the following response on behalf of realists (Street, 2006, pp. 125-135). We haven't yet ruled out the possibility that what made our ancestors more reproductively fit than their competitors was that they were sensitive to moral features of the world. For example, if this claim is true, we can say that our ancestors survived because they were sensitive to objective facts such as that helping kin is morally right while harming those who have helped you is morally wrong. Street rejects this view, however, because it is much more plausible to suppose that our ancestors acted in reproductively advantageous ways for the reason that acting in such ways made it more likely that they would survive and pass on their genes—an appeal to objective moral facts here is uncalled for. Ultimately, Street thinks it's more plausible to suppose that evolutionary forces distorted our ability to gain knowledge of objective moral facts than it is to suppose that they led us perfectly to these facts.

Eventually Street argues that her objection to realism should point us toward accepting her metaethical constructivism—we'll get to that in Chapter Two. But it's worth noting here that there are several ways to object to realism on similar epistemological grounds. So even if there is a problem with Street's particular objection to realism or if realists can offer compelling responses to it, realism's epistemological worries may not be solved.

Just to give one more example of such an objection of realism, critics point to strong influence from cultural indoctrination as a plausible explanation of why we form the moral beliefs that we do. Huemer, for example, asks us to consider common moral judgments about polygamy, infanticide, and slavery (2008, p. 374). While it's true that most of us in contemporary western societies (and many other places) refrain from these practices, there have been plenty of societies in the past, and, indeed, plenty of societies which exist today, that did or currently do partake in these practices. While the frequency of said practices does not necessarily

reveal what members of any culture thinks of them, it seems likely that the moral judgements of people in societies which practice (or did practice) polygamy, infanticide and/or slavery are different than our own. If they did reject those practices and had moral judgements more aligned with our own, we would expect that those societies would cease to practice such activities. Thus, it seems clear that one's moral judgements are strongly influenced by one's culture.

The upshot, as Huemer puts it, seems to be that,

The variation in moral [judgements] across cultures strongly suggests that many or most [judgements]—including, of course, one's own—are explained more by historical accident than by objective ethical truths, even if such truths exist (2008, p. 375).

Like Street's argument, this argument leaves room for the existence of objective moral facts, but it is similarly crippling for realism in that it seems to entail that we don't have knowledge of those facts; except this time it is because our moral judgements are heavily influenced by our culture. One might want to insist that some particular culture is good at finding out the objective moral facts and that people in that culture have moral knowledge, but such a claim would require a case for why that culture is so much better at figuring out the objective moral facts than other cultures. Making such a case would require using your own culturally indoctrinated moral attitudes to explain why your culture's moral judgements are the objectively correct ones and why other culture's moral judgements are wrong. The problem with making such a case is that one's own starting premises are likely just culturally indoctrinated moral attitudes, so the initial objection still applies.

We could make more epistemological objections to realism by replacing the influence of evolution and cultural indoctrination on our moral judgements with the influence of emotions

and personal biases on our moral judgements,¹¹ but I think that would be redundant. The point of mentioning the various epistemological objections to realism was to show that even if one objection fails, realists will not be finished with refutation.¹²

With this chapter I have situated constructivism into its place on the metaethical map. In section 1.2 we saw the advantages that constructivism and realism share over other competing metaethical theories. In section 1.3 we saw Sharon Street's epistemological objection to realism. In Chapters Two and Three we'll see some objections to constructivism, but for now, realism is the only one of our starring characters that faces any serious problems.

¹¹ See Huemer (2008, p. 377) for a presentation of Walter Sinnott-Armstrong's argument that emotion and personal biases often influence our moral judgements.

¹² Of course, realists might try to craft a response that dismisses all such epistemological arguments against realism in one fell swoop. See Huemer (2008) for such an attempt.

CHAPTER TWO: THE METAETHICAL CONSTRUCTIVISMS OF KORSGAARD AND STREET

The purpose of this chapter is to present the constructivist positions of Korsgaard and Street.¹³ In doing this, I will also point out a problem unique to each of the positions. However, these problems aren't as severe as the epistemological problem facing realism.

In section 2.1 I'll present Korsgaard's position. It is the oldest of the currently popular strands of metaethical constructivism, so starting there is appropriate. In section 2.2 I'll discuss a central problem facing Korsgaard's position. Section 2.3 will be dedicated to presenting Street's position. Section 2.4 is a comparison of the views—it is pretty lengthy. In this section we will see that the Kantian nature of Korsgaard's view secures her position a few benefits usually held by objectivist metaethical theories. In subsection 2.4.2 I will argue that Street's position has a shot at gaining some of these advantages. In subsection 2.4.3 I will explain how constructivism avoids the epistemological worry that we saw in Chapter One.

2.1 Korsgaard's Constructivism

The version of Korsgaard's position I will present is one that is common in the literature.¹⁴ However, I should note that her work does not come assembled neatly in a metaethical package with the ingredients listed in the explicitly metaethical language of the sort I've been using—thus some interpretation and rephrasing are necessary. For example, I don't think she once uses the word 'constructivism' in her book *The Sources of Normativity* (and this is

¹³ Korsgaard, 1996b; Street 2010

¹⁴ Bagnoli, 2017; Street, 2010.

the text of Korsgaard's from which I draw most heavily). In any case, I take my presentation of her view to be in line with other popular interpretations of her view.

Korsgaard describes her position as a kind of procedural realism. This position holds that "there are answers to moral questions *because* there are correct procedures for arriving at them" (Korsgaard, 1996b, p. 36). It is helpful to define this position in contrast to what she calls substantive realism. This position holds that there are correct and incorrect procedures for answering moral questions "*because* there are moral truths or facts which exist independently of those procedures, and which those procedures track" (Korsgaard, 1996b, pp. 36-37).

Notice the ways in which those definitions fit into the metaethical map I presented in the last chapter. Substantive realism is similar to what I have been calling realism since it holds that there are moral facts which our moral judgements (procedures) track rather than construct. While I haven't made it entirely clear yet, what Korsgaard calls procedural realism is a form of constructivism. Recall from Chapter One the discussion of what it meant for a moral fact to be constructed—there I said that constructed moral facts are those that are created by agents occupying the practical point of view using either hypothetical or actual procedures. So when Korsgaard says that "there are answers to moral questions *because* there are correct procedures for arriving at them" we should get a strong whiff of constructivism about her position (1996b, p. 36).

What is still unclear is whether the facts that agents construct are subjective or intersubjective (as I said in Chapter One, no constructivist position holds that moral facts are objective). The goal of the next several pages will be to present Korsgaard's position and argue that according to her view moral facts are subjective.

Korsgaard ultimately thinks that moral obligations arise out a central feature of human agency (Bagnoli, 2017).¹⁵ This central feature is that we are reflective creatures who can act rationally. In fact, Korsgaard defines ‘humanity’ as the ability to rationally reflect (1996b, pp. 92-93). Since it is far from obvious that moral obligations arise from this feature of human agency, it will require explanation. We can take this claim and break it down into two steps.¹⁶ The goal of the first step is to establish that our humanity is the condition for the possibility of valuing anything at all. The goal of the second step is to show that through our shared valuing of humanity, we get moral obligations.

To secure the goal of the first step, Korsgaard argues that because of our humanity we require reasons for our actions. Consider the following example: after finishing my lunch at a picnic table I desire to rid my trash as soon as possible due to the flies that are swarming around it; but I also feel compelled not to litter. Here I have two competing impulses—how do I decide on which to act? According to Korsgaard, “we endorse or reject our impulses by determining whether they are consistent with the ways in which we define ourselves” (1996b, p. 120); and, furthermore, we define ourselves in terms of our practical identities. Ultimately, I have to choose whether to litter or toss my trash in a waste basket. I will make this decision according to a reason that comes from one of my practical identities. In this case, my identity as an environmentalist will provide me reason not to litter.

According to Korsgaard, most of our practical identities are contingent (1996b, pp. 120-122). We have identities as friends of certain people, lovers of others, and as citizens of a certain countries. While some of these identities often stay with us for life—such as the identity of being

¹⁵ Agents are beings with the ability to act. Human agents are beings who have the ability to act and who, according to Korsgaard, are rational and reflective.

¹⁶ Throughout this chapter I will speak of moral facts and obligations interchangeably.

a citizen of a certain country—they are contingent and thus may change (and may have been different in the first place). Sometimes these identities clash and we may end up shedding an identity. One’s identity as a scientist may lead her to shed her religious identity, or one’s newly acquired identity as a father may lead him to shed his identity as captain of the tri-weekly bowling team.

While many of these practical identities may grow or fade away, “what is not contingent is that you must be governed by *some* conception of your practical identity” (Korsgaard, 1996b, p. 120). If we didn’t have practical identities, we would essentially be sacks full of impulses and desires pulling every which way. A sack full of impulses is not a human—we need to rationally reflect and endorse or reject certain impulses inside of us in order to be human (recall that, for Korsgaard, rational reflection is what makes us human). Thus, we conform to at least some of our practical identities because we would lose our humanity if we didn’t. Notice that this reason for conforming to at least some of your practical identities is not a reason that comes from one of those practical identities: “It is a reason that *springs from your humanity itself* [emphasis added], from your identity simply as a human being, a reflective animal who needs reasons to act” (Korsgaard, 1996b, p. 121). It is here that the goal of the first step is obtained. That is, we need our humanity in order to value anything at all (since when we value we do so for reasons). In our reflections about what we should value or how we should act, we necessarily endorse and reject our options for valuing (or acting) due to various *reasons*. Thus, when we value something, our rational capabilities (i.e., our humanity) are always in use.

Now for the second step of Korsgaard’s claim that moral obligations arise out of our rational abilities—the goal here is to show that from our valuing of humanity we get moral obligations. As we just said, we must value our humanity if we are to value anything else. The

interesting thing about our humanity is that it is not unique to any single human. Remember, it is that quality of rational reflection that all human agents share. Thus for Korsgaard, there is no relevant distinction between my humanity and your humanity. And given that I value my humanity, I must value yours. If I did not, I would contradict myself; and as a human being—a *rational* reflective creature—I would compromise my own humanity if I did not value yours (Bagnoli 2017).

The moral obligations we have to one another stem from our shared humanity. Interestingly, Korsgaard does not say much about any of the particular obligations we have to each other. Here's one that is probably a safe bet, however: everyone is obligated not to kill each other. Why? Because doing so would not be consistent with valuing your humanity. It would be the opposite of valuing your humanity.¹⁷ Another might be that we are all obligated to help those people in need of food or shelter because if they don't get those things they will likely die. It seems possible that there could be other moral obligations we have to each other based on shared practical identities that are extremely important to us, but I don't want to speculate here—the point of this discussion is not to find out the particular obligations we have to one another, but instead to show Korsgaard's argument for the claim that we can derive moral obligations from our rational capabilities (i.e., our humanity).

Now we can sort out whether moral facts on Korsgaard's view are subjective, intersubjective or objective. The moral obligations we get from Korsgaard's position depend on the attitudes of agents (since they stem from our attitude of value toward humanity). Bearing in mind the definition of objective moral facts I've provided—that is, the one that says that in order

¹⁷ There seems to be room for Korsgaard's position to accommodate complexity, too. For example, killing may be permissible if it is in self-defense. Even though this may seem like it is an act which violates the humanity of another person (which is not relevantly different than your own), it is an act which will at least preserve your own humanity.

for a moral fact to be objective, its truth must not depend on any agent's or group of agents' attitudes regarding the object of moral evaluation—we can conclude that the moral facts constructed according to Korsgaard's theory are not objective. Now the question is whether they are subjective or intersubjective. At first it might seem tempting to latch onto the fact that it is our shared humanity that gets us moral obligations according to Korsgaard's view. If we go down this path it may seem like moral obligations come from the group of beings that are committed to valuing humanity. If it's a group of people that determine the truth of moral claims then we are pushed towards accepting intersubjectivism about moral facts—recall the following characterization from Sayre-McCord: “Intersubjectivism grants (with subjectivism) that people figure in the truth-conditions [of moral claims], but it holds that the truth of moral claims doesn't turn on facts about *particular* [emphasis added] individuals” (1986, p. 20). If it's true that it is not a particular individual who determines the truth of moral claims, and given that we've already eliminated the possibility that moral facts are objective on Korsgaard's view, then the only option we are left with is intersubjectivism.

This would be a mistake, however; for it is not true that groups of individuals determine the truth of moral claims according to Korsgaard's position. Moral obligations are those commitments we are forced to accept on pain of contradicting our own commitment to the value of humanity. So when I judge that ‘It is wrong to murder my neighbor’, it will be true because I am committed to it in virtue of the fact that I value my own humanity. My obligation to my neighbor comes about because it is the case on Korsgaard's view my humanity is the same as the humanity of everyone else; thus, to murder my neighbor would be to devalue humanity and, thus, to contradict myself. You'll notice here that the moral obligation not to murder my neighbor comes from the features of *my* agency. And while it's true that everyone else will share

a commitment to the value of humanity, ultimately, my obligation not to murder my neighbor comes from the features of an individual's agency (that is, my own). If this is the case—that is, if it is true that moral obligations come from a *particular* individual's agency rather than from judgements of a group of people—then I have established that according to Korsgaard's position moral facts are subjective, not intersubjective.¹⁸

There several ways one can be a subjectivist about moral facts. We saw one of them in Chapter One. I consider it to be a crude form of subjectivism since it doesn't allow for moral error. According to this crude subjectivism, rightness and wrongness and goodness and badness are determined by whatever an individual judges to be the case. Thus, if Mr. Stabbs is in a particularly bad mood and judges it to be the case that slashing the tires of his foes in the name of vengeance is the right thing to do, then that action is morally right. But, let's say that Mr. Stabbs undergoes a change of heart in a few minutes after having a little snack and decides that such behavior is despicable and obviously morally wrong. According to crude subjectivism, slashing the tires of his foes is now a morally wrong thing to do.

Luckily for proponents of Korsgaard's position, her view is more complex than this. If Mr. Stabbs were to be in a particularly bad mood and decide that a bunch of violent and destructive acts were morally acceptable, Mr. Stabbs would be mistaken—his actions would be wrong. This is because the violent and destructive acts would disrespect the humanity of other people. The people who would be affected by Mr. Stabbs' actions would have all sorts of reasonable concerns about not wanting destruction and violence to occur, and because Mr. Stabbs' humanity is the same as theirs, he would have to respect their reasonably held desires not

¹⁸ This is a good place remind us why Korsgaard's theory is not objective. The correctness of moral judgements *does* depend on an agent's attitudes. For every moral evaluation, an individual's valuing of humanity (which is an attitude) is involved in the moral evaluation (in evaluating, I am relying on my rational reflective abilities). If I am evaluating whether or not taking home a bottle of hot sauce from a restaurant is morally permissible, this answer depends in part on my attitude towards humanity. Thus, the correctness of my moral judgement is subjective.

to want violence and destruction to occur.¹⁹ Thus those actions would be wrong according to Korsgaard's position.

There's a kind of subjectivist *ideal observer* theory of moral facts which, on first blush, seems to capture Korsgaard's characterization of moral facts (Sayre-McCord, 1986, p. 19). An ideal observer theory agrees with the crude subjectivism about moral facts from above in that they both hold that the truth of moral claims depends on what a particular individual judges to be true. However, the ideal observer theory claims that the truth of moral claims depends on an idealized individual, rather than an actual individual. For example, an ideal observer might be characterized as an individual who has perfect non-moral information as well as a dedication to treating everyone as fairly as possible.

However, as we've seen above, moral obligations come from a feature of the agency of *actual* human agents. It is actual people that are committed to valuing their own humanity (and thus the humanity of others). This may seem worrisome to some because it may seem that if moral obligations stem from the actual features of human agency, then those actual humans couldn't ever make errors in their moral judgements. But as we saw above, this is not the case. People make errors when they fail to act in accordance with the value of humanity. Thus, while the subjectivity of moral facts according to Korsgaard's view should not be classed as an ideal observer theory, it does not entail that every moral judgement an actual human agent makes is correct.

If what I've said here is correct, Korsgaard's theory should be considered as one which holds that moral facts are subjective. We should add to the last statement that her theory is not an ideal observer theory and that the subjectivism to which her position is committed is more

¹⁹ There will surely be situations where one person's reasonably held desires clash with another person's reasonably held desires. These are tough situations, but they don't pose any unique problem for Korsgaard. Most moral theories have a tough time giving totally confident answers to tough moral decisions and dilemmas.

complex than the crude subjectivism we saw in Chapter One. With her position on the table, let's now examine a flaw.

2.2 A Problem for Korsgaard's Constructivism

The problem has to do with the fact that that Korsgaard assumes that it is entailed by our valuing of ourselves (that is, all of our various practical identities) that we must also value humanity. Recall that, according to Korsgaard, we must value humanity because our humanity is what allows us to endorse and reject (based on reasons) different acts and values that issue from our practical identities; and we all have practical identities. We endorse and reject various possible actions and values, and in doing so—in using our reflective rational capacities—we necessarily value humanity. However, according to Bagnoli, “humanity may be the condition of the possibility of value and yet lack value itself” (2017, *Constructivism as Procedural Realism*, para. 7). It's worth slowing down to understand this claim. First, let's restate that humanity according to Korsgaard is the ability to rationally reflect. I restate this here because it's tempting to read ‘humanity’ in its most common usage which refers to human beings in general—this would be a mistake here. With this in mind we can now read Bagnoli's (slightly altered) claim to be that ‘the ability to rationally reflect may be the condition of the possibility of value and yet lack value itself’.

An obvious way to answer this objection would be to say that rational reflection—that is, humanity—is unconditionally valuable. In fact, this is exactly what Korsgaard does in her characterization of Kant's formula of humanity (Korsgaard, 1996a, pp. 106-32). This is relevant because in *The Sources of Normativity* Korsgaard claims that her argument is just a “fancy new” version of Kant's formula of humanity (Korsgaard, 1996b, p. 122). Thus, I think it's likely that she holds this view as well. Moreover, she says the following about the value of humanity:

We find that the unconditioned condition of the goodness of anything is rational nature, or the power of rational choice. To play this role, however, rational nature must itself be something of unconditional value – an end in itself (Korsgaard, 1996a, p. 123).

Bagnoli argues that this seems to be a step toward accepting a kind of realism (2017). This is because when our humanity is considered an unconditioned value it then seems to count as an objective moral fact. Recall that for a moral fact to be objective its truth must not depend on any agent's or group of agents' attitudes about the object of moral evaluation. The worry is that if humanity is valuable regardless of what any agent or group of agents thinks about it (i.e., if it is unconditionally valuable), then it seems to be an objective moral fact that humanity is morally good and that actions which respect it are morally right.

I think that there is more to be said on behalf of Korsgaard's position though. It might not be the case that just because humanity is an unconditioned condition that it is also an objective moral good. I think we can still understand humanity as an unconditioned condition (in a certain sense) and still have it be subjectively valuable. The sense in which humanity is an unconditioned condition is that, for all of us, in making any sort of decision—whether it be about what to value or how to act—we make use of our humanity. We effortlessly make use of our rational reflective abilities quite frequently. The fact that we effortlessly do this makes it unconditioned in a sense—we value it to the extent that we use it so frequently and, perhaps, because we can't help but use it.

Much more could be said about and to this response—it is only a sketch. If Korsgaard's theory does suppose, after all, that humanity is objectively valuable, then the position might be in danger of collapsing into some kind of realism. But I think the sketch I've given here should give us reason to think that there are plausible ways out of such a charge. I will briefly come back to this topic in Chapter Three when I talk about the task of weighing the problems of one position

against the problems of another position. With Korsgaard's position on the table, let's now turn to Street's position so that we can see how they compare.

2.3 Street's Constructivism

Street calls her position the practical standpoint characterization of constructivism. I will often just refer to it as *Street's constructivism*. By the end of this section I will have argued that according to Street's position moral facts are subjective in the same way that Korsgaard's are. But before I explicitly make that argument I will have to do some setup.

Street begins with a stipulation about our recognition of the attitude of valuing (Street, 2010, p. 366). Street claims that no matter whether we understand what value is, we are able to understand the attitude of valuing. Why do we have this understanding? Street says it's because we live in a world full of valuing creatures and that we frequently are witness to their valuing—the attitude of valuing is observed so often that it would be a surprise if we did not recognize it. Street claims that when a creature values something he or she enters the practical point of view, which for Street is, “the point of view occupied by any creature who takes at least some things in the world to be good or bad, better or worse, required or optional, worthy or worthless, and so on” (2010, p. 366). While this is a slightly different formulation of the practical point of view than the one I gave in Chapter One, this shouldn't slow us down much. In Chapter One I said that the practical point of view was the position we occupy when we have to make a practical decision and need to weigh the reasons we have for acting in accordance with each relevant path of action. While this formulation makes explicit reference to decision making in the face of challenges and Street's characterization does not, I still take the two to be compatible. Consider that when we face problems and have to decide which actions to take, we necessarily take some things to be “good or bad, better or worse, required or optional, [and] worthy or worthless”

(Street, 2010, p. 366). The characterization of the practical point of view that I provided in Chapter One has an emphasis on weighing reasons for action in the face of problems that Street's does not share, but ultimately, this difference is merely skin-deep. Thus, we can move on to the actual statement of her position.

According to Street's general statement of constructivism, the truth of a moral claim "consists in that claim's being entailed from within the practical point of view, where the practical point of view is given a *formal* characterization" (Street, 2010, p. 369). I'll explain what a formal characterization is in the next section when I compare the positions of Korsgaard and Street, but for now we should get clear on what Street's position is. A slightly different (and clearer) formulation of Street's claim is as follows: in order for any moral claim to be true, it must survive the scrutiny of a given valuer's entire set of values once they are corrected for inconsistencies and non-moral mistakes (Street, 2006, pp. 110-111). Tropman provides the following helpful example of what this would look like using the example of Smith: "It is wrong for Smith to cause himself personal harm because we can suppose that the judgement, 'Causing myself personal harm is morally wrong', would withstand scrutiny from Smith's other informed judgements" (2014, p. 131).

It's worth a moment to consider the benefit of requiring that the truth of moral claims must survive the scrutiny of a given valuer's entire set of values *once they are corrected for inconsistencies and non-moral mistakes*. Consider Val. She has recently become a vegetarian. For her, the judgement 'Eating meat is wrong', is a judgement that follows from her other judgements. When she goes out to eat with friends at restaurants, she now tries to be more aware of what goes into her meals than she used to. Unfortunately, one Friday night at a particularly busy restaurant in town, Val orders a burger and forgets to substitute the veggie patty in for the

ground beef patty. While Val makes the judgement ‘I ordered a morally permissible meal’, she is actually mistaken about this due to a non-moral mistake. In fact, the judgement that follows from her other informed judgments *once corrected* for non-moral mistakes is that ‘It was wrong of me to order the meal I did’. While this example is an unfortunate one for Val, it serves to highlight a benefit of Street’s constructivism—that is, it acknowledges that people make errors in moral judgement.

From the above discussion we can gather that there are two things people must consider when they find themselves in a situation that calls for coming up with an answer to a moral question. The first thing that you’ll have to do is try to rid your values of inconsistencies and non-moral mistakes. After all, the rightness or wrongness of an act is determined by whether it would withstand the scrutiny of one’s entire set of values *once they are corrected for inconsistencies and non-moral mistakes*. The other thing you’ll have to do is check which values you hold as well as the relative strength of each value. If Val actually values not bothering people with her vegetarian food requests more than she values not eating meat (highly implausible, I admit), then her judgement ‘I ordered a morally permissible meal’ is actually correct. Even though this sounds strange to us, if her value of not bothering people with her food requests really is stronger than her value of not eating meat then this is what follows according to Street’s position.

2.4 Comparing and Contrasting Korsgaard and Street

The biggest difference between the constructivisms of Korsgaard and Street is that moral obligations, according to Korsgaard’s position, must not violate the value of humanity—Street’s position does not require this. To that extent, we may regard Street’s position as the less constrictive constructivism. However, this lack of constriction may not be as nice as it sounds.

Since moral facts do not have to be in line with the value of humanity according to Street's position, there will likely be cases where intuitively wrong behavior ends up being morally right. Before discussing this point, it will be helpful to compare the two views to see what other aspects the respective positions do and do not share. We begin with a discussion of the subjectivity of moral facts according to each position. In section 2.4.2 I will argue that Street's position has tools for handling the fact that intuitively wrong behavior appears to be occasionally entailed by her view—although, I don't think the tools can entirely save her position from the objection. In section 2.4.3 I'll show how constructivists avoid the epistemological problem facing realists.

2.4.1 Humean Constructivism vs Kantian Constructivism

Street's position is subjectivist insofar as it fixes moral truth to what each individual thinks. But it is not the same as the crude sort of subjectivism that we toured in Chapter One and in section 2.1 of the current chapter. That position does not allow for, among other things, moral error. Problems with this position have already been discussed, so I will skip it here. The constructivisms of Korsgaard and Street also share in the fact that they are not ideal observer theories. I discussed this in section 2.1 with respect to Korsgaard's theory, but I have yet to address the possibility that Street's position is an ideal observer theory. An ideal observer theory is one in which we are to imagine what an individual (or a group of people) with certain idealized characteristics would do in a given situations. For example we might imagine what a group of people with perfect non-moral information and a dedication to fairness would do if tasked with distributing wealth within a country. Street's view is not like this. Street is very explicit in laying out her theory that substantive values must ultimately "be supplied by the particular set of values with which one finds oneself alive as an agent" (Street, 2010, p. 370). Recall the example of Smith. We said that it was wrong for him to cause harm to himself

because the judgement ‘Causing myself harm is morally wrong’ withstands scrutiny from his other informed judgements (Tropman, 2014, p. 131). We only need to assume that Smith is an actual person—or we can even just substitute an actual person in for Smith—to see that Street’s constructivism is not an ideal observer theory. True, Street’s theory does require that one’s values must be corrected for inconsistencies and non-moral mistakes, but this is a far cry from having perfect idealized knowledge or any other idealized qualities.

So, at this point we can say that the constructivisms of Korsgaard and Street are the same insofar as they share the claim that the truth of moral facts is determined by actual human agents in the practical point of view. However, this does not mean that we will get the same moral obligations according to each respective position. One helpful way to highlight this point is to distinguish between Kantian and Humean versions of constructivism.²⁰ Korsgaard’s constructivism is Kantian, while Street’s is Humean.²¹ We can characterize Kantian constructivisms in the following way:

According to Kantian versions of metaethical constructivism ... we may start with a purely formal understanding of the attitude of valuing and demonstrate that recognizably moral values are entailed from within the standpoint of any valuer as such (Street, 2010, p. 369).

Recall that Korsgaard argued that moral obligations come from a central feature of human agency. We can rephrase in the verbiage of the definition just provided and say that Korsgaard’s constructivism holds that moral values are entailed from within the practical standpoint as such.

²⁰ Indeed, this is a distinction Street makes use of when distinguishing her position from Korsgaard’s (2010, pp. 369-370).

²¹ Kant’s emphasis on rationality is shared by Korsgaard, so that connection should be easy enough to see. Street’s view is Humean in that both Hume and Street hold that human attitudes are the ultimate creators of moral facts (Hume, 1975). In Hume’s case, he speaks of the passions creating morality while in Street’s case she speaks of the attitude of valuing being prior to values.

The value that is entailed according to Korsgaard's theory is the value of humanity, and, from that value, moral obligations follow.

Given a Humean constructivism such as Street's, there are no moral values or obligations that follow from within the practical point of view as such. In other words, this means that there will be no substantive values (including the value of humanity) that follow from the practical point of view given a formal characterization. One might be puzzled by me saying that on Street's Humean constructivism 'there will be no substantive values that follow from the practical point of view as such'. The key to sorting out this confusion is to understand that right now we are only talking about the practical point of view *as such*. Substantive values do follow from the practical point of view according to Street's view, it's just that Street thinks that there are no values that follow from the practical point of view when we understand it in a purely formal manner—substantive values must ultimately "be supplied by the particular set of values with which one finds oneself alive as an agent" (Street, 2010, p. 370). An example should help here. Recall Smith from section 2.3 and the fact that Street's constructivism requires that moral truth is fixed by an individual's moral judgements (once they are corrected for inconsistencies and non-moral mistakes). If Smith undergoes an extreme transformation and it becomes the case that the judgement 'Causing myself personal harm is morally permissible' would survive the scrutiny of his other judgements, then it would be morally permissible for Smith to harm himself. Smith would almost always be mistaken according to Korsgaard's Kantian constructivism because causing oneself harm is almost always a violation of the value of humanity. But because Street's Humean position does not require that Smith value humanity, he is not making a mistake when he harms himself (as long as this act passes the reflective scrutiny test).

2.4.2 Street's constructivism and intuitively wrong acts

While it's true that Korsgaard's theory treats moral facts as subjective (as does Street's), Korsgaard's theory is thought to capture some merits that usually belong to metaethical theories which treat moral facts as objective (such as moral realism). Street's position sometimes entails that acts that we generally think of as intuitively wrong acts are permissible or right—this is a knock against her theory. The ability of Korsgaard's constructivism to save itself from these cases is a point in favor of Korsgaard.

Korsgaard's way of keeping intuitively wrong acts from winding up morally permissible is to include a value that all people are necessarily committed to. According to her constructivism, the fact that we all necessarily value our humanity—and that we all share in this rational, reflective humanity—makes it so that there are many actions that will never be able to count as morally permissible; this is because many intuitively wrong actions violate the value of humanity. Why is this? Let's look at the murdering of innocent people as an example. This should always be wrong according to Korsgaard's position. Why? Because the act of murder is an act that takes away the rational reflective capabilities (that is, the humanity) of another person; and in valuing or contemplating the act of murder we are, necessarily, valuing our own rational, reflective abilities (that is, our own humanity). And since there is no relevant difference between the humanity of any two people, to value one person's humanity and not the other's is inconsistent. This analysis will likely hold true in all or nearly all cases of intuitively wrong acts.

In this section, I'd like to argue that Street's position may have a way of battling against cases where it seems like her theory entails that intuitively wrong actions will count as permissible or right. To be clear, my goal in this section is not to argue that Street's theory can

secure the claim that there will never be instances of intuitively wrong behavior that will count as permissible or right. Instead, my goal is to provide a strategy that might minimize such cases.

Let's consider the case of Caligula, the lascivious and violent Roman emperor.²² There is reason to think that Caligula genuinely valued torturing others for fun (even if this isn't true it is easy to imagine a person who does highly value torturing others for fun—we could easily replace Caligula with a person like this if we needed to). Korsgaard's theory is not going to allow Caligula's acts of torture to count as morally permissible. This is because torturing another person is likely always a violation of his or her humanity. Thus, according to Korsgaard's theory, Caligula's acts of torture would be wrong. Caligula would be making a mistake in his actions. But what about Street's constructivism? Things don't turn out the same way. According to Street's constructivism, an act's rightness consists in its surviving the scrutiny of a given valuer's entire set of values once they are corrected for inconsistencies and non-moral mistakes. So, if it is true that Caligula highly values torturing others for fun and that his acts of torture would survive the scrutiny of his entire set of values (once corrected for inconsistencies and non-moral errors), then these acts are right—Caligula's acts of torture are not morally wrong.

Moreover, this discrepancy between the two positions is bound to turn up in many other situations as well. Any time a person's values are composed in such a way that they highly value some sort of paradigmatically immoral behavior, the two theories are likely going to give a different answer to the morality of the act in question—and almost always Korsgaard's theory is going to be the one that aligns with our intuitions about the scenario. This is because many of the acts that Street's view will allow to count as morally right or permissible will be in violation of the value of humanity.

²² See Street (2010, p. 371) for a discussion of the different treatment Caligula gets according to Kantian and Humean constructivists.

These considerations are likely why Shafer-Landau classifies Korsgaard's theory as a kind of objectivist constructivism (2003, pp. 39-41). While I think this classification is ultimately wrong, it is worth seeing why he characterizes her constructivism this way—regardless of the classification, it will bring out some advantages that Korsgaard's view seems to have over Street's.²³ According to Shafer-Landau, objectivist constructivisms require some degree of idealization for the attitudes and responses that fix moral truth (2003, pp. 39-41). You'll recall from section 2.1 that it is the attitudes of actual individuals that fix moral truth according to Korsgaard's position, so we disagree with how to best understand her theory. However, regardless of the classificatory disagreement we have, Shafer-Landau rightly notes that Korsgaard's theory has the following three virtues: 1) it captures the impartial nature of morality, 2) it maintains the categorical nature of moral demands and 3) it provides a plausible view of the nature of moral error. This third virtue is something that I have already argued that Street's position can do, so we can forego any more discussion about moral error.

Let's see why Korsgaard's theory secures these virtues and why Street's allegedly does not. Consider the claim that Korsgaard's theory captures the impartial nature of morality. We might rightly wonder what we mean by impartiality here. Following Shafer-Landau, we can say that the kind of impartiality we are talking about here is, "that which forbids giving one's own parochial outlook any priority in shaping the contours and shapes of substantive moral theory" (2003, p. 40). As we saw above in our discussion of Caligula, Korsgaard's theory maintains this impartiality because in order to preserve the value of humanity, we will often not be able to consistently favor our own outlooks over those of others—even though Caligula highly values

²³ Full disclosure: Shafer-Landau ultimately finds all forms of metaethical constructivism problematic and instead argues for moral realism. However, this is not our concern. Regardless of whether he thinks constructivism is a plausible view or not, his comments about the relative strengths and weaknesses of the different kinds of constructivisms are what matters for my purposes.

torturing others for fun, he is still wrong to do so. However, according to Street's view, there is nothing wrong with Caligula's actions—since they survive the scrutiny of his entire set of his (corrected) values, his tortuous actions are permissible.

A similar analysis can be given to the claim that Korsgaard's theory maintains the categorical nature of moral demands. As Shafer-Landau points out, we tend to think that moral demands do not go away if they are not instrumental to satisfying our desires or self-interest (2003, p. 40). The moral demands that are entailed by Korsgaard's theory will be like this because our valuing of humanity will always entail that certain acts are wrong—in particular, intuitively wrong acts—because these acts will almost always violate the value of humanity. Even if I sustain a brain injury that changes my values in such a way that I now value hurting others above almost everything else, those actions will still be wrong according to Korsgaard's theory. The same cannot be said for Street's constructivism. When I sustain that brain injury and I begin to value harming others, as long as the acts of violence I want to commit survive the scrutiny of my entire set of corrected values, these acts will be permissible.

It is unclear how much of a knock against Street's theory this should count as. While Street acknowledges that her view has "implications some find deeply counterintuitive", she thinks that accepting her view "is what we are forced to by the untenability of realism plus the failure of Kantian versions of metaethical constructivism" (2010, p. 370). She says this at a point in an exposition of her position where she has just claimed that substantive values "must ultimately be supplied by the particular set of values with which one finds oneself alive as an agent" (Street, 2010, p. 370). I believe that the 'counterintuitive implications' she mentions (and which she does not specify) are almost certainly in reference to the sorts of cases we have been considering. She makes this concession at this point in her exposition because—I think—it is at

this point that the reader will be starting to see that there aren't really any constraints on "the particular set of values with which one finds oneself alive as an agent". By this concession, Street is acknowledging that there are going to be some values that people hold that will yield, to put it lightly, hard-to-accept results when put through the constructivist procedure (that is, reflective scrutiny).

That, of course, is just Street's own take on the issue. I think—without trying to quantify the seriousness of the problem—that it is at least problematic enough to warrant an attempt at finding a way out of. Indeed, Shafer-Landau approves of the fact that usually when constructivists try to characterize their respective constructive procedures they try to show that the initial conditions of construction and the constructed results match up with our deep ethical convictions—he says this is an "impeccable way to proceed" (2003, p. 42).²⁴ I presume that this is at least in part because he thinks that it would be problematic if a metaethical theory were to entail that the morality of certain acts were completely opposite of some of our deepest ethical convictions.

Luckily, I think there are ways Street can begin to bolster her view against these charges. Indeed, I think within Street's body of work there are already several places where a better response is in the making. However, there is no value that all people necessarily share according to Street's view; recall that her view does not entail any substantive values from the practical point of view as such. On the other hand, she does note that there are deep, widely shared contingent similarities in each individual's evaluative starting points that might yield similar constructed results for most people—indeed, these similarities in starting points may very well depend on "the existence of a shared human nature ... to the extent that there is such a thing, rather than on anything entailed by the practical point of view as such" (Street, 2010, p. 370).

²⁴ Full disclosure (again): Shafer-Landau ultimately rejects constructivism in favor of realism.

While Street does take the time to note this, she doesn't spend any time on pursuing this point in the 2010 paper in which she lays out her position. I think that this line of pursuit will make her constructivism more appealing and will bring it closer to securing the same advantages usually associated with moral realism.

This strategy will appeal to some values that all humans might plausibly share. At the outset I would like to say that I will not be attempting to argue that there has never been and never will be an intuitively wrong act that Street's position entails as permissible. Instead, I would like to argue that there are plausibly some other values that we might all share given our evolutionary history and, furthermore, that these might be enough to show that some of the counterintuitive implications of her view aren't as much of a worry as we might have thought.

Let's consider the following list of claims:²⁵

- (1) In general, I value acts that would promote my survival.
- (2) In general, I value acts that would promote the interests of my family members.
- (3) In general, I value helping my own children over helping complete strangers.
- (4) In general, I value reciprocity; for example, if someone treats me well, then I want to treat them similarly in return.
- (5) In general, if someone has done me or my kin deliberate harm, I value shunning that person or seeking his or her punishment.

These values are widely shared across time and cultures (Street, 2006, p. 115). It is plausible that these values were selected for by evolution. Recall Street's evolutionary story from Chapter One. She begins by noting that, evolutionarily speaking, there were some motivational tendencies or proto-judgments that were clearly better than others for the reproductive success of

²⁵ The original version of this list comes from Street (2006, p. 116). I have modified the language a bit, but in spirit it is much the same.

our ancestors (2006, p. 13-21). The creature that values putting itself in harm's way is clearly going to be less reproductively successful than the creature that avoids harm and instead values acquiring sufficient nutrition for itself and its kin ((1) and (2) from the list above). Having motivational tendencies to avoid harm, to seek out proper nutrition, and to help one's kin accomplish these goals would essentially be to enjoy a higher level of fitness than one's contemporaries. These motivational tendencies were heritable and over time became frequent in the population while the motivational tendencies of less fit creatures became decreasingly common. What is important about this is that these motivational tendencies, these proto-judgements, influenced the capacity of our more recent ancestors to form full-fledged cognitive judgements about morality. In fact, Street argues, the selection of certain motivational tendencies also affected the *content* of our current moral judgements. For example, Street would say that we now generally think that stealing from our kin is wrong because having this motivational tendency contributed to the reproductive success of our ancestors. Street's picture renders our judgements about moral facts and properties largely dependent on the evolutionary history of moral agents.²⁶

If we accept this evolutionary story then we will see that there are at least some values that everyone has a strong disposition towards accepting. Importantly, many of these values clash with the sorts of intuitively wrong acts that we've been considering in this chapter. Consider the act of torture. Most people are going to have no inclination towards committing such an act. However, there will be some who will genuinely consider doing it. Caligula is an example of such a person. The following is a way to handle this case. Anyone who wants to

²⁶ Street acknowledges that we can sometimes use rational reflection to change our moral beliefs and judgements. However, she thinks that even rational reflection has been influenced by evolutionary forces since rational reflection ultimately—she argues—requires the weighing of some evaluative or moral judgements against others (2006, pp. 123-124).

commit torture will realize that everyone else likely subscribes to (4) and (5) from above. When one realizes this, they will also see that committing torture will likely conflict with their own subscription to (1) because they will expect people to retaliate for such tortuous behavior. Thus, committing torture might not survive the scrutiny of Caligula's informed values because it may very well be inconsistent with his other values—he subscribes to (1) and committing torture will likely decrease his likelihood of survival, so torture may very well end up morally wrong for him.

Ultimately such a strategy may work in other cases of intuitively wrong behavior, but probably not all of them.²⁷ It is too likely that there will be some people whose values are so atypical and crude that what follows from them is intuitively wrong. However, in every case of intuitively wrong behavior that Street's theory seems to deem as right, this strategy is available. People are often bad at sorting through their own values, so any time an intuitively wrong act seems to be entailed as right by Street's position, it is open for discussion whether that act actually does follow from a given person's entire set of informed values; that person might just be overlooking an inconsistency in their values.

2.4.3 Constructivism and the epistemological challenge

In Chapter One I presented the epistemological problem for realism and I claimed that constructivism avoids that problem; but I did not say why that was the case. I have saved that discussion for the present section because it is a strength shared by the respective positions of Korsgaard and Street (and section 2.4 is a compare and contrast section!).

The epistemological problem for realism is that—even if we grant that objective moral facts exist—there is likely no way for our moral beliefs to count as knowledge (and realists want there to be moral knowledge). According to Street, we can't know the objective moral facts

²⁷ See Street (2009) for a more developed discussion of cases like Caligula.

because our moral judgement-making capacities are tainted by evolutionary forces in such a way that our judgements are not likely to be true (and since truth is a necessary condition for knowledge, we cannot say that what we come to believe via our moral judgements counts as knowledge). The reason our moral judgements are probably not true is that, according to Street, it seems like our current moral beliefs are not accurate enough to count as knowledge. They are not accurate enough to count as knowledge because it is more plausible to think that our ancestors acquired certain motivational tendencies due to the fact that those tendencies made it more likely that our ancestors would survive, rather than that our ancestors gained certain motivational tendencies because those tendencies allowed our ancestors to respond to objective moral facts. Even if our moral judgements do occasionally hit on the objective moral facts, it seems like this would be by sheer chance.

Take for example the prevalent moral judgement that we ought to help our kin survive. A realist might posit that there is an objective fact of the matter—that is, they might posit that it is an objective fact that we ought to help our kin survive (excluding atypical cases where our kin are, say, deranged or dangerous). If you accept Street’s objection and its evolutionary setup, realists cannot claim to have knowledge of this fact (even if we grant that it exists) because our judgement that we ought to help our kin survive is best explained by the fact that our evolutionary past has made us such that we favor making such a judgment—because of this, it is not likely that this judgement is true. To suppose it is true would be to suppose that we just happened to make the correct moral judgement by coincidence, knowing full well that there is no reason to suppose that evolutionary forces formed us so that we would respond to objective moral facts—such a coincidence would be “astoundingly convenient” for the realist to maintain (Street, 2006, p. 122).

Constructivists avoid this problem by claiming that because moral facts are constructed from the practical point of view, humans ultimately create morality. According to Street, the constructivist “understands the evolutionary causes to be prior, in the sense that these causes (along with many others) gave us our starting fund of evaluative attitudes, and evaluative truth is understood to be a function of those attitudes” (2006, p. 154).²⁸ Remember that moral values are evaluative values. Thus, Street is saying that moral facts are a function of our moral values. Constructivists have a better shot at gaining moral knowledge because, unlike realists, moral truths are set by what people value and believe to be moral. Our contingent values fix what the moral facts are and so they aren’t so far out of our grasp. For example, say we value altruism; according to constructivists, it will be a moral fact that altruism is (usually) good and that altruistic acts are (usually) morally right. According to constructivists, these moral facts are a function of our valuing of altruism. We create them, and because we create them, we can likely come to know them. If the moral facts are not a function of our moral values (as realism would have it), then it seems unlikely that we could come to know them.

2.5 Conclusion

While the constructivisms of Korsgaard and Street each have a difficulty or two, it’s perhaps more important to end this chapter by restating something good that the theories share: each theory sidesteps the epistemological problem that faces realists. This is a significant point in favor of constructivism because without side-stepping the problem, you’re looking at a moral theory which cannot claim that we have moral knowledge—given that realists and so many other people generally think that we have at least some moral knowledge, this implication is a problem.

²⁸ For Street, moral attitudes are a kind of evaluative attitude. Thus, when she says that evaluative truth is a function of our evaluative attitudes, we should also understand that moral truth is a function of our moral attitudes.

Of course there were still the problems for each constructivism. Korsgaard's theory faced the charge that it might endorse the existence of an objective value and Street's theory can't rule out that there will likely be some counter-intuitively wrong actions that end up morally right. In either case, I think that each of these issues is less severe than the epistemological problem facing realism.

In the next chapter we will turn to a challenge posed at constructivists that has received little attention in the literature.

CHAPTER THREE: CONSTRUCTIVISM AND DELIBERATION

In the last chapter, we saw how the positions of Street and Korsgaard compare. In this chapter I will suggest that there is a previously unnoticed advantage that Korsgaard's position has over Street's. The advantage lies in the fact that Korsgaard's position does a better job of maintaining a plausible phenomenology of deliberation. By 'phenomenology' I mean to indicate the experience of what it is like for someone to do something. In this case, we will look at the experience of what it is like to deliberate in moral situations and straightforward situations regarding typical facts. There's no reason to be convinced of my claim right now, but by the end of this chapter it's my goal to have made a persuasive case for why Korsgaard's position secures this advantage.

In section 3.1, I will present David Enoch's objection that global constructivism cannot make sense of deliberation.²⁹ In section 3.2, I'll present Enoch's case for why he thinks realism does a good job of maintaining what we tend to think happens in deliberation.³⁰ I think Enoch is right in suggesting that realism does a good job of capturing what it is like to deliberate so it will be advantageous for the respective positions of Street and Korsgaard to offer a similar characterization of moral deliberation. In section 3.3, I will tailor the objection raised against global constructivism to challenge Street's constructivism. I present this objection to Street's constructivism because it reveals something interesting about how moral deliberation must be characterized according to Street's position. In section 3.4 we will see a response to the objection on behalf of Street's constructivism. This response will ultimately reveal a problem that Street's position has with maintaining a plausible phenomenology of deliberation; this will be the focus

²⁹ Enoch (2009).

³⁰ Enoch (2011).

of section 3.5. In section 3.6, I will argue that Korsgaard's constructivism does a better job of maintaining a plausible phenomenology of deliberation. Section 3.7 is a conclusion. It will call back to earlier chapters and give perspective to the respective challenges faced by Korsgaard and Street.

Before I jump into Enoch's objection to global constructivism in the next section, I'd like to give a specific example of a situation in which someone faces a situation that calls for moral deliberation. Let's say you have a friend named Donna who is a well-intentioned though sometimes capricious person. Over drinks, Donna tells you how frustrated she is with the distribution of wealth in the U.S. and how she thinks that there should be more social safety nets in place so that no one has to worry about going hungry or affording housing and insurance. A tear glistens in her eye as she laments about the terrible effects poverty has on children—she is an elementary school teacher and you know how much she cares about kids. Perhaps the drinks have gone to Donna's head, because what she says next surprises you. She proposes that the two of you should rob a bank and clandestinely distribute the money to local organizations dedicated to helping poor families. You stare into her eyes and see that she is totally serious.

Whatever your immediate reaction to Donna's proposal, at least one thing seems true: you need to deliberate about whether robbing the bank with Donna is morally right or wrong. Perhaps you could dismiss the proposal as 'Donna just being Donna', but the bigger point should be taken that when it comes to moral decision making, we often deliberate about whether an act is morally right or wrong. Of course, there *are* situations in which one action is obviously better than others so deliberation isn't needed. (For example, say you are given two ways to save a human life. The first way to save the life is to pay a dollar; the other way is to attempt to balance a glass vase on your head for three hours, and if you fail the person dies. Here the first choice

seems so obvious that very little, if any, deliberation is needed.) But it seems like there are plenty of other situations in which moral decision making isn't so easy. Situations which give rise to the following sorts of questions seem to require deliberation: Is it wrong of me to drive to work each day given that I could easily walk? Given what I know about factory farming and given that I have the means to easily buy food that does not come from factory farms, what food should I buy? Would it be wrong of me to borrow my friendly neighbor's lawn mower without asking? Situations like these, in which there isn't any obvious choice of action, seem to require deliberation.

These are the kinds of questions we should have in mind when we talk about moral deliberation throughout this chapter. But, of course, there are other, non-moral situations in which deliberation seems called for. In the following section we will see a challenge from David Enoch posed at global constructivism, a theory which (for now) can be roughly characterized as a constructivist view about all normative facts, not just moral ones.³¹ This is a necessary detour because Enoch's objection to global constructivism, once I modify it, raises a different challenge for metaethical constructivists.

3.1 Enoch's objection: can global constructivism make sense of deliberation?

David Enoch suggests that deliberation is difficult for global constructivism to adequately capture (2009, pp. 333-335). Before presenting the objection, we should get clear on some terminology. The main difference between global and metaethical constructivism is that global constructivism claims that *all* normative facts are constructed, while metaethical constructivism does not. To bring out this difference, consider the following claim: "You ought to proportion your beliefs to the evidence you have for them". Such a claim is normative (in that it is supposed to be behavior-guiding), but it is not obviously moral. A global constructivist is committed to

³¹ Enoch (2009).

saying that this claim is constructed in the same way that moral facts are constructed according to metaethical constructivism. So, a global constructivist might say that you are doing something wrong if you don't proportion your beliefs to the evidence you have for them, but the 'wrong' here will not be the same as 'morally wrong'. We can think of metaethical constructivism as the narrower view of the two because it does not require its advocates to have any stance on the question of whether normative, non-moral facts are the product of constructive procedures.

Non-moral, normative facts enter into our lives perhaps even more often than moral facts do. We think about them any time we ask ourselves what we should do (in non-moral situations). Think of all the questions that arise throughout the day that involve trying to figure out what you should do: should you have oatmeal or something else for breakfast? Should you drive or bike to work? Should you go this or that way on your way to the store? These are trivial, but there are other more serious questions we ask ourselves: should I stay in this relationship? Should I take a new job at a company located across the country? Should I go back to school?

Enoch claims that global constructivists cannot make sense of the way we typically deliberate about the questions raised in the last paragraph. Enoch says that the answers to normative questions like those in the last paragraph—and, indeed, normative questions in general—do not seem to be created by a constructive procedure:

When engaging in this deliberation, when asking yourself these questions, you assume, so it seems to me, that they have answers. These answers may be very vague, allow for some indeterminacy and so on. But at the very least you assume that some possible answers to these questions are better than others. You try to find out what the (better) answers to these questions are ... You are not trying to create these answers (2011, p. 72).

The thought is that in order for us to successfully engage in the deliberation that is required for answering such questions, we are committed to there being answers to these questions which

exist independently of whatever we happen to decide at the end of deliberation. We can put Enoch's point as a rhetorical question: if you are deciding about whether to go back to school to get a degree, aren't you trying to find out, rather than construct, what the best course of action is?

Enoch thinks that this is a problem for the following reason: if it is the case that we are the ones who ultimately decide what the correct course of action is (and if it is the case that we know this), then we can't deliberate—instead, we would either get stuck or would have to resort to just picking a course of action. You get stuck if you are trying to find out what the correct (or better) answer is because you simultaneously know that whatever you construct as the correct (or better) answer is in fact the better answer. How are you supposed to deliberate if the answer you are looking for is yet to be constructed? Instead, we might just pick which answer seems right—but this is to abandon deliberation. Both getting stuck and just picking are not what we generally take to be involved in successful deliberation. I will expand on this idea of getting stuck in deliberation or having to resort to just picking in a moment.

Before I show how Enoch's objection to global constructivism can be modified in such a way as to reveal another issue for metaethical constructivists, it will be good to see why Enoch thinks that a realist stance is in better position to preserve what we tend to think happens in deliberation.

3.2 Enoch's case for why robust realism captures deliberation

Enoch calls his position *robust realism* and it bears the same relationship to metaethical realism as global constructivism bears to metaethical constructivism (Enoch, 2011). This is to say that while both robust realism and global constructivism are positions that make claims about the existence of moral facts, they are ultimately concerned with the entire normative realm, not just morality.

Enoch thinks that his position does a good job of capturing the experience we tend to think people go through when they deliberate about normative questions.³² Why is this? The first thing to note is that robust realists hold that normative facts are objective (while global constructivists hold that they are subjective). In order for a moral fact to be objective, its truth must not depend on any agent's or group of agents' attitudes regarding the object of moral evaluation. Thus, according to realists, the answer to your question regarding the morality of robbing the bank with Donna will obtain regardless of what your other values and attitudes are and regardless of whether you decide joining in the heist is right or wrong. Other normative non-moral cases will play out similarly.

Now imagine that you fancy yourself a robust realist and you have to deliberate about Donna's proposal. You know that the morality of the heist will not depend on any of your own attitudes about the act. Unlike global constructivism, the rightness or wrongness of robbing the bank will not depend on what you happen to decide. Instead, while deliberating, you will know that there is a fact of the matter that you can aim towards finding out. Thus, if you subscribe to realism, moral deliberation is possible.

Being able to capture moral deliberation is important too, according to Enoch. A theory that cannot capture deliberation fails to explain something which, according to Enoch, is rationally non-optional for us (by which he means that we will necessarily be irrational if we don't engage in deliberation) (Enoch, 2011, p. 70). Consider Enoch's remarks about why deliberation seems to be rationally non-optional for us:

Perhaps this is so partly because we are *essentially* deliberative creatures. Perhaps, in other words, we *cannot* avoid asking ourselves what to do, what to believe, how to reason, what to care about. We can, of course, stop deliberating about one thing or

³² See Enoch (2011) for a discussion of this point.

another ... It's opting out of the deliberative project as a whole that may not be an option for us (2011, pp. 70-71).

I agree with Enoch that deliberation is likely a project that we humans cannot stop doing. However, I don't need this claim in order to continue on with my discussion. If it turns out that deliberation isn't as important as Enoch thinks it is, it will only make constructivism's alleged problem that much smaller.

Now that we've seen why realism seems to be in a good position to capture what we tend to think happens in moral deliberation, we can see why it might appear as if metaethical constructivism cannot do the same.

3.3 The objection regarding deliberation as tailored to Street's constructivism

Broadly speaking, the charge against metaethical constructivism is that it does not easily preserve the way we typically think moral deliberation is experienced by deliberators. More specifically, the charge is this: if rightness and wrongness are constructed by some sort of deliberative process and if the deliberating agent knows that rightness and wrongness are constructed in such a way, it becomes unclear how the deliberator is supposed to genuinely deliberate about the issue. I will give a clarifying example in a moment, but first we should check if Street's position includes a deliberative process that fixes rightness and wrongness. Recall the following formulation of Street's constructivism: an act's rightness consists in its surviving the scrutiny of a given valuer's entire set of values once they are corrected for inconsistencies and non-moral mistakes. Thus, according to Street's constructivism, it seems that in order to find out whether an act is right or wrong, an agent has to deliberate about whether or not that act would follow from her entire set of values.

Enoch provides us with an analogy to help highlight the problem (2009, p. 33). To be clear, however, his objection was aimed at global constructivism. Here I am using his analogy

for my own purposes. He asks us to imagine that we are judges and have to decide the outcome of a case while knowing the following two pieces of information: 1) We have to decide the verdict of the case only according to what is the legally right decision and 2) the legally right decision is whatever decision we come to via deliberation. The second piece of information is what brings about the problem. The worry is that if I know that whatever decision I make is the legally right one, then I can't really aim at finding a correct decision to the case. Enoch says that, "Such a judge is now (sort-of) stuck – she cannot deliberate about what the legally right answer is, knowing that the conclusion of her deliberation is the only thing that will make the legally right answer right" (2009, p. 33).³³

To complete the analogy we need to see how this example maps on to a case of moral deliberation by a constructivist. Remember Donna? She has just asked you to rob a bank with her. You are now in the following situation: 1) You have to decide whether robbing the bank is morally right or wrong and 2) You subscribe to a metaethical position which holds that an act's rightness consists in its surviving the scrutiny of your entire set of values (once corrected for inconsistencies and non-moral mistakes). In this case—as in the judicial case—it seems that you won't be able to engage in deliberation. You will want to look for a correct answer to the question of whether robbing the bank is right or not, but at the same time you will know that there is no answer to this question until you decide what the answer is. Enoch says the following thought should occur to a constructivist who goes to deliberate about the rightness of an act:

I know that the conclusion of my deliberation will make my decision to do something the right decision, that there is no correct (right, rational, whatever) decision I can *aim* at, one

³³ To be sure, in reality judges do end up making decisions—this is because, presumably, that judges think there is a correct answer to legal questions and that the correct answers exist independently of whatever they (the judges) happen to decide. They will often (and hopefully always) think that they are making the correct decision, but they will likely realize that their decision is correct not because it is what they have decided but because, roughly, it follows from just laws.

that is independent of my very deliberation. And this makes deliberation impossible... (2009, p. 333).

Thus, goes the objection, constructivism cannot make sense of the common phenomena of moral deliberation.

There are a few important points to make regarding the possibility of moral deliberation with respect to Street's constructivism. The first point is that the purported problem constructivists have with deliberation only applies to agents who actually subscribe to constructivism. If you are an average person who does not read much academic philosophy, then even if constructivism is true you will not face the problem of being stuck when you go to engage in moral deliberation. Why is this? It is because you will likely not have the belief that the answer to the question you are trying to answer is determined by the conclusion you come to at the end of deliberation. What made deliberation impossible in the case of the judge, according to Enoch, "was not the mere fact that the conclusion of her deliberation will determine what the legally right decision is. Rather, it was that she knew (or at least believed) as much" (2009, p. 334). The same is true of constructivism. If you believe that the rightness or wrongness of acts is determined by what you decide at the end of deliberation, then you will be stuck with no answer towards which you can aim to find.

There are two initial replies that constructivists could try. Ultimately, I think each fails. First, constructivists could just say that when we face challenging moral situations like Donna's proposal we just arbitrarily decide whether the action is right or wrong. This would keep us from getting stuck in deliberation. So what would this look like? With respect to Donna's proposal, an example of an arbitrary decision would be the case in which you decide which course of action is morally right without considering any of the pros or cons of each path of action. This decision would be akin to the choice you make when you go to choose a wine glass from a set of

practically identical glasses. Enoch (rightly) rejects this constructivist response (Enoch, 2009, p. 333). He claims that a constructivist view that incorporates mere picking as its constructive procedure is ‘highly implausible’. Why should we think that moral facts are as arbitrary as our choice of which wine glass to use out of an identical set? In that case we just pick a glass and pour—there is no deliberation that happens in such a decision.

Now for the second initial response that one might want to offer on behalf of constructivists. You might want to just reply by shrugging your shoulders and saying, “So what? There are hardly any philosophers. Very few people would be affected by the implications of this argument.” After all, in order for the objection to work it requires that deliberators believe that it is the result of their deliberation that will fix moral truth. And since it seems like there aren’t very many people who have that belief, deliberators will not get stuck. I appreciate the sentiment of this objection, but the problem isn’t so simple. A plausible requirement of any moral theory is that public knowledge and popularity of it must not undermine it (call this the publicity requirement). Morality is a normative subject. This means that it is largely concerned with guiding human behavior. If a theory that is supposed to be action guiding can only successfully guide actions if it remains arcane, then the theory’s own popularity would undermine it as an action-guiding theory. I think this is a *prima facie* reason to reject such a reply.

Some scholars argue against the publicity requirement, however. For example, Sidgwick considers that utilitarians in government might not want to publicly endorse their theory (for various reasons) and instead use a common-sense morality to guide public policy (as cited in Smart, 1956, pp. 347-348). One might think that constructivists could adopt this move and say “There is no problem here. The view is true, it’s just the kind of view that works best when it is arcane”. This move fails, however. Utilitarianism—at least a basic understanding of it—is a view

that strives to drive human behavior. Its maxim is “Do the most good (which is identified as pleasure) and avoid as much badness (which is defined as pain) as possible!” The reason why some proponents of this view might want to keep public knowledge of the view at a minimum is that public knowledge of the view might undermine its goal. That is, public knowledge of utilitarianism might bring about less pleasure and more pain than if it were arcane. There is no analogous goal that proponents of Street’s position strive for. Utilitarianism is a consequentialist view; this means that is a view which is concerned with bringing about good consequences and avoiding bad consequences. Street’s position is in the business of explaining what moral facts are and how people relate to them; it should have no preference about which consequences are good and which are bad. The position itself is not a person and does not have values. Thus, Street’s position does not have some consequence that it needs to protect by keeping itself hidden.

In the next section I will discuss a more plausible reply to the objection.

3.4 A reply to the objection on behalf of Street’s constructivism

The objection to Street’s constructivism offered in the previous section is that it prevents at least some of its adherents from engaging in moral deliberation. Specifically, the objection asks how you could deliberate about the rightness or wrongness of an act knowing full-well that the rightness or wrongness of an act is determined merely by its surviving the scrutiny of your entire set of values (once they are corrected for inconsistencies and non-moral mistakes). Ultimately, I think that Street’s constructivism can avoid this criticism. However, in dodging the criticism I think that something interesting and problematic is revealed about her view. Namely, that the moral deliberation that her constructivism is able to preserve does not look like what we generally take deliberation to be; this will become clear by the end of this section.

The way to begin the response to the objection from the last section is to note that the objection only works if Street's view is understood in a particular way. The kind of constructivism the objection addresses, goes this line of response, is one in which moral deliberators have to *actually* go through the deliberative process in order for an act to become right or wrong. We see this in the analogy between constructivism and the judge who knows that her decision will determine the legally right verdict of the case she is considering. In that analogy we assumed that in each case it is the actual deliberation of a person that will create a fact (in the former case it is a moral fact that is created and in the latter case it is a legal fact).

The next step of this reply is to say that the assumption that actual deliberation is what fixes moral rightness and wrongness is not a claim made by Street. Instead, Street's view is more plausibly characterized as a position which holds that an act is morally right if it *would* survive the scrutiny of a given valuer's entire set of values once they are corrected for inconsistencies and non-moral mistakes. Call this the *hypothetical reading of Street*. This is to be distinguished from what I will call the *actual reading of Street*. This position holds that an act is morally right if it *actually* survives the scrutiny of a given valuer's entire set of values once they are corrected for inconsistencies and non-moral mistakes.

So which way should we understand Street? This one's easy: the characterization of Street's position that I presented in Chapter Two is the hypothetical reading of Street. Recall the case of Smith in order to see this. In Chapter Two we said that it was wrong for Smith to cause harm to himself because the judgement 'Causing myself harm is morally wrong' withstands scrutiny from his other informed judgements (Tropman, 2014, p. 131). Regardless of whether or not Smith actually deliberates about whether causing himself harm is wrong, the judgement 'Causing myself harm is morally wrong' may still be entailed by his entire set of values. Indeed,

Street claims that, “Quite apart from whether we think a given set of values is *correct* ... we can nevertheless think about and discuss what *follows*, as a purely logical and instrumental manner, from a given set of values in combination with the [non-moral] facts.” (2010, p. 367). The case of Smith is a fine example of this.

Now that we’ve distinguished the hypothetical and actual readings of Street, we can see that constructivists who engage in moral deliberation and subscribe to the hypothetical reading of Street will not get stuck. This is because they will not believe—contra those who subscribe to the actual reading of Street—that the rightness or wrongness of an act is determined by what they actually decide at the end of deliberation. Only people who subscribe to the actual reading of Street will have this belief. Instead, those who subscribe to the hypothetical reading will think that the rightness or wrongness of an act may already be determined before they go to deliberate. I say ‘may already be determined’ because it is possible that one may change their values in deliberation. But, as in the case of Smith, many acts will be entailed as right or wrong by your set of values before you go to deliberate. For example, there may already be a fact about whether or not it would be wrong of you to join Donna in her heist—it is determined by whether it would survive the scrutiny of your values. Thus, when you go to deliberate about the morality of the heist, you will believe that there is a fact of the matter towards which you can aim at finding. Thus, you will not be stuck when you go to deliberate; therefore, the hypothetical reading of Street avoids the objection that deliberators will get stuck. In having this answer towards which deliberators can aim to find, the hypothetical reading of Street also avoids the worry that deliberators might just end up having to pick what is right and wrong.

3.5 Why moral deliberation according to Street's position is flawed

This is where we find out something interesting and, ultimately, problematic about Street's view. A little setup is required. In establishing that Street's position is the hypothetical reading of Street (and not the actual reading), we have established that there is an answer towards which we can aim at finding in our moral deliberations. We also know that the best way to find out the morality of a given act will be to go through the very reflective process that is supposed to determine whether acts are right or wrong (for why would we want to use a different deliberative procedure than the procedure used to construct facts in the first place?)

This reflective process has two components to it, as we're familiar with by now. The first component is the checking of which values you hold as well as the relative strength of each value. In deliberating about Donna's proposal you will have to look into yourself to see whether you ultimately value justice more than helping poor children or vice versa (perhaps other values will enter into it, such as the value you place on helping friends, but those will be the primary two values you will have to examine). The second component is that you will have to try and rid your values of inconsistencies and non-moral mistakes. After all, the rightness or wrongness of an act is determined by whether it would withstand the scrutiny of one's entire set of values *once they are corrected for inconsistencies and non-moral mistakes*.

Some of Enoch's remarks about the phenomenology of deliberation help us see what is lacking from a deliberative process such as Street's. As we've seen, Enoch thinks that when we go to deliberate we have to assume that the questions we are occupied with (for example, "Is this act right?") have answers—answers that it doesn't seem like we construct (2011, p. 72). Even if the answers are vague, we still seem to need to assume they exist in order to deliberate, according to Enoch (2011, p. 72). In addition to this, he points out that this sort of deliberation is

strikingly similar to how we go about finding answers to straightforward factual questions, such as the question of whether Denver is farther north than San Francisco:

When trying to answer a straightforwardly factual question ... you try to get things right, to come up with an answer that is—independently of your settling on it—the right one. When deliberating, you also try to get things right, to decide as—independently of how you end up deciding—it makes most sense for you to decide (2011, p. 73).

This picture of what goes on in deliberation seems to me highly plausible. Even if I didn't know that Denver was north of San Francisco, I would assume that there was an answer to my question even if I couldn't deduce the answer from what I know about the geography of the American West or if I didn't have access to the internet (indeed, I checked the web and found out that Denver is the northernmost of the two). Just as you might yet be unsure of whether helping Donna commit a robbery would be right or wrong, you suspect that there is an answer to your question. This characterization of the phenomenology of deliberation will serve as a template to which we can compare the respective deliberative procedures of Street and Korsgaard. Some readers might want more justification for my use of this characterization of deliberation. I will provide such justification at the end of this section when I consider replies that defenders of Street's position might be tempted to voice. I hope for now you'll share my intuition that it is a plausible characterization of the phenomenology of deliberation.

Now that we have this plausible picture of what goes on in deliberation, we can check to see how closely Street's position resembles it. The second component of Street's deliberative procedure seems to capture some of what Enoch is getting at. In trying to find out the answer to a moral question, deliberators will have to make sure that they are avoiding inconsistencies and non-moral mistakes. This seems to capture some of what Enoch means when he says "you try to get things right" when you deliberate. You can't get things right if there are inconsistencies and non-moral errors in your beliefs.

Notice here that the thing we are trying to get right is what our own values are and how strongly we hold them. We are trying to get consistency and avoid errors and we are only doing so with respect to the things we already value. Thus ‘getting things right’ here only means being true (in some sense) to your own values. This is what we might expect from a theory which holds that moral facts are subjective. If the moral facts are determined by your values about the question at hand, then it would make sense that the deliberative process would involve a weighing of your relevant values. I think there are at least two issues for Street’s position here.

The first is that it seems like individuals would be able to change their values during deliberation. If you’re deliberating about Donna’s proposal, you might realize that you’ve been over- or under-valuing, say, the public’s confidence in bank security. Perhaps you think that robbing the bank with Donna might cause a lot of people to withdraw their money from banks, potentially leading to serious economic problems. But to do this—or any other value revising—seems to clash with the characterization of deliberation we got from Enoch. According to that characterization, we are trying to come up with answers that are independent of our settling on them. But if deliberation involves coming up with the answer as you deliberate, then we’ve strayed from Enoch’s characterization of deliberation. While this concern with the public’s confidence in banks does seem to be a relevant concern, according to Street’s position your concern with this factor ultimately goes into constructing whether or not helping Donna would be right or wrong. But this doesn’t seem to be what it feels like to deliberate. When we consider the public’s confidence in banks as a relevant factor in our deliberation, it seems to be in an effort to discover information that might be relevant to finding out the answer—discovering information in an effort to find out an answer is certainly not the same as *constructing* an answer.

There are surely some retorts to be had here on behalf of Street's position, but I am going to hold off on them for a moment.

The other issue with deliberation according to Street's view is that it seems like in deliberation we are trying to do more than just get consistency and error-avoidance in our own set of values. While there are certainly some situations in which deliberation might consist solely in weighing our own values (and correcting for inconsistencies and non-moral errors), there are other times when deliberation is more complicated. If you are deliberating about which ice cream flavor you should order, then you may very well be in a deliberative situation which calls for nothing more than some introspection about which flavor is likely to tickle your tongue the most. However, many moral deliberations don't seem to be like this. When we face tough decisions about moral matters, we want to know what's right and what's wrong and it doesn't feel like we have the answer hidden somewhere in our own set of values—instead, it's as if we were searching for an answer that can only be found outside of our own values, thoughts, etc. If it were the case that morality is fixed by whether or not an act passes a reflective scrutiny test—as Street's constructivism would have it—then we would expect that the reason moral deliberation is often tough is because we are having trouble weighing our values against one another or making them consistent. But again, I think the more plausible explanation of why many moral deliberations feel tough is that, insofar as the phenomenology is concerned, we feel as if we are searching for an answer outside of our own values, one that is to be found out or discovered. The question “What do I value most?” never agonizingly crosses our minds.

Objections are hot on my heels here, so it's time to stop and address them. First off, one might want to say that I've set up an unfair and maybe even impossible task by essentially asking a metaethical theory which holds that moral facts are subjective to try to make its deliberative

process resemble the straightforward answering of questions about objective facts (the fact that Denver is farther north than San Francisco is objective). There are two things to say here. First, I do not think the way I've set up the problem makes it impossible for subjectivist metaethical theories to answer. Here's a way we could test my claim: if another subjectivist theory can demonstrate that its deliberative process looks like Enoch's characterization of deliberation, then we will have shown that the task is not impossible. Indeed, this is what I intend to do in the next section with Korsgaard's position. Secondly, I do concede that I am putting a lot of confidence in the phenomenological account of deliberation that Enoch gives us. If this phenomenology turns out to be wrong or misguided then I think we should indeed return to this discussion and perhaps revise some of what I've said. I'll say more about the benefits of examining the phenomenology of deliberation momentarily in my answer to the next objection.

The next objection: one might want to dismiss Street's problem with deliberation as a minor flaw that shouldn't be given much weight when comparing the virtues and vices of different metaethical theories. I'll admit that it is sometimes hard to say exactly how significant certain criticisms are. However, I think that examining what deliberation would have to be like according to each theory brings to light some important aspects of the theories we've been examining. This is because it gives us insight into what the experience is like for deliberators without having to first make arguments about whether moral facts are subjective, intersubjective, or objective (or even extant at all). Investigating the phenomenology of moral deliberation may give us insight into the nature of moral facts (assuming there are some). If people deliberate about moral questions and the experience is one of focusing exclusively on their own values (and perhaps logical consistency), then we have at least some prima facie reason to suspect that moral facts are subjective. But if people experience moral deliberation in the same way that they

experience thinking about straightforward factual questions regarding objective facts, then we have some prima facie reason to suspect that moral facts are objective. Of course, there are ways in which this strategy can go awry. It might be the case that our experience as moral deliberators is not at all indicative of the nature of moral facts—our experiences of deliberation might not reveal anything about the nature of moral facts. We might think that we are searching for objective moral facts and it might feel as if this is what we are doing, but we nevertheless might be mistaken about this. Nothing that I've said rules this out. However, I submit that the starting position of our investigation into deliberation should not begin with the assumption that we are mistaken. It might turn out that we are mistaken, but I think the burden of proof is on the one who proposes this claim. Moreover, I do not mean to claim that investigations into deliberation are the best or only way to go about learning the nature of moral facts (assuming there are any)—but it's certainly one way to go about it.

I think I should also address the reader who might be wondering, "Isn't this a thesis in favor of constructivism? Why is the author making an argument that seems to favor the existence of objective moral facts when constructivists oppose this claim?" This is a very understandable thought to have. In response, I would like to say that when it comes to debates in metaethics (and indeed other areas of study), it is common practice and maybe even unavoidable, that you have to make concessions to your opponents—it is a further debate how much the concessions matter and whether one problem is more harmful to a position than a different problem is to a different position. At this point, proponents of Street's position can still contest that this one problem is practically nothing compared to the epistemological problem for realism. And, although I am going to argue in the next section that Korsgaard's theory fares much better with respect to the deliberative challenge we've been considering against Street, those in favor of Street's position

could certainly still argue that there are problems unique to Korsgaard's constructivism that should ultimately push us towards accepting her (Street's) view.

Defenders of Street's position might want to question my assumption that phenomenology matters at all. After all, couldn't we just be mistaken about our experiences of moral deliberation? I admit that we are not infallible when it comes to making sense of our experiences. It could be true that we are mistaken when it comes to our phenomenology regarding moral deliberation. However, moral theorizing has to have some starting points. If a metaethical theory which holds that moral facts exist is so revisionist that it rejects much of what we think is going on in moral situations, then that position risks losing much plausibility. It seems like a good idea to begin by tentatively assuming that many of the things we take to be the case are actually that way. We can always come back and revise this, but until a positive case is made for why we should distrust our experiences of moral deliberation is made, it seems like a fine way to go about investigating morality.

3.6 Moral deliberation: Why Korsgaard does it better than Street

In the last few sections I used Enoch's objection to global constructivism to show that if we apply a similar version of the argument to Street's constructivism it would reveal something interesting about Street's position.³⁴ This revelation was that her position doesn't do a great job of capturing a plausible phenomenology of moral deliberation. This section will be an attempt to see what deliberation is like for Korsgaardian deliberators. Ultimately, Korsgaard's position does a better job than Street's at capturing a plausible phenomenology of moral deliberation. In

³⁴ The difference between Enoch's objection to global constructivism and the one I've been aiming at Street's constructivism is as follows: Enoch used his objection to show how global constructivists could not even create the facts that they were supposed to be deliberating about. In the previous section I argued that Street's position avoids this problem. My objection was meant to show that even if there are some facts that people can aim at finding according to Street's position, the phenomenology of this deliberation would not be at all like what we generally take the phenomenology of moral deliberation to be.

making this case we will also see why Korsgaard's position avoids the problem of deliberators getting stuck in deliberation.

Let's consider your experience of deliberation with respect to Donna's proposal. Recall that according to Korsgaard, moral obligations come from our shared valuing of humanity. Furthermore, according to Korsgaard, you are necessarily committed to valuing humanity. When you deliberate about Donna's proposal, you must make sure that no one's humanity is violated (unless you're in a situation where it can't be avoided). The deliberation might get quite complicated as there are many people who might be affected by it—but we need not go into all these details. The point is that you will have to consider how other people will be affected by your actions.

This is why Korsgaard's position does not have the same problem as Street's. When we deliberate under a Korsgaardian framework, our valuing of humanity will always play a significant role in moral deliberation; it is a value that has a lot of weight when compared to others (it may even be a value which trumps all other values). Our valuing of humanity dictates that we must consider how other people will be affected by the actions that we are considering committing. This will get us much closer to the phenomenology of deliberation that I've been claiming is desirable. Again, that characterization said that deliberation feels as if we are trying to get the right answer and that the right answer feels as if it is independent of whatever we happen to decide. This is exactly what happens with deliberation in Korsgaard's position. Now, the answer you are looking for according to Korsgaard's position is not entirely independent of your moral attitudes since her position holds that moral facts are subjective. However, before you consider the humanity of each person that might be affected by the actions you are considering committing (and you must do this according to Korsgaard's position), you will (obviously) not

know how these people will be affected by the actions you are considering committing. And thus, while deliberating—while figuring out how others will be affected by your actions—you will be aiming to find an answer that you did not know before your deliberation. This was the main feature of the plausible phenomenology of deliberation that I've been endorsing. And because the value of humanity is so weighty when compared to other values, it needs to be considered in every moral deliberation.

By now it should be emerging why Korsgaard can easily dodge any objection that might be raised about deliberators getting stuck in deliberation. The answer that the deliberator is looking for won't be available to the deliberator until all the people who might be affected by the action are considered; specifically, until their humanity is considered. If the deliberator forgets that an action would violate the humanity of many people, it's not as if whatever she comes to at the end of deliberation is right—she will have gotten the answer wrong because it violated the humanity of other people. Because the answer the deliberator is looking for isn't available to her until she considers the humanity of other people, she will not get stuck in deliberation.

Now for a few objections one might have in response to what I've said in this section. Here's one: you might want to object here that there is still a big dissimilarity between the characterization of deliberation I've been endorsing and the Korsgaardian version of deliberation I've just characterized. The dissimilarity is that in order to find out whether we are obligated to perform or refrain from performing some act, we must always consider other people (in particular their humanity). According to the phenomenology of deliberation that I've been claiming is attractive, we need not necessarily consider other people. My response: this shouldn't concern us. The crucial part of the plausible phenomenology of deliberation is that it feels as if we are searching for an answer that is independent of what we happen to settle upon.

Of course, the answer one has to look for according to Korsgaard's position is not entirely independent of the moral attitudes of deliberators since Korsgaard's position holds that moral facts are subjective. However, before a deliberator considers the humanity of each person that might be affected by the actions s/he is considering committing, the deliberator will (obviously) not know how these people will be affected by the actions s/he is considering committing. And thus, while deliberating—while figuring out how others will be affected by her/his actions—s/he will be aiming to find an answer that s/he did not know before deliberation.

Here's another objection: defenders of Street's position might be tempted here to claim that the phenomenology of deliberation according to Street's position will often look like Korsgaard's phenomenology of deliberation. This is because so many people value the humanity of other people. If this is the case—and I'll concede that it really does seem to be the case—then when people deliberate under a Streetian framework, their phenomenology will be very similar to the phenomenology of deliberators under a Korsgaardian framework. However, this is only a contingent possibility according to Street's view. According to Korsgaard's position, if people do not properly consider the humanity of other people when they deliberate about moral issues they are making a mistake. This is not the case according to Street. I think Korsgaard's position still has the advantage here because her position *requires* that if deliberators want to get things right they *have* to go through a deliberative procedure of the kind Enoch characterized and I endorsed. At best, Street's position just gets to claim that many people will just happen to deliberate in this way.

3.7 Conclusion

The main focus of this chapter was to show that Street's position faces a problem maintaining a plausible picture of the phenomenology of moral deliberation. In the previous

section, I argued that Korsgaard's position avoids this problem. By now you may notice a common theme in much of my discussion of Korsgaard's position: her position has a lot of the same benefits that metaethical realism has. Her position avoids the deliberative problem because deliberators will feel as if they are searching for an answer that is independent of whatever they end up deciding. If you are a deliberator under a Korsgaardian framework, you will have to think about whether the humanity of other people will be violated (because you must value humanity). In having to consider how other people might be affected by your possible actions, you will feel as if the answer you are looking for is independent of whatever answer you end up deciding (since you can't possibly know all the implications your actions will have on people before you deliberate).

In Chapter Two, we saw a comparison of the metaethical positions of Korsgaard and Street. I pointed out that some people think that Korsgaard's assumption that rationality is unconditionally valuable is a realist move. I argued that there is reason to think that Korsgaard can give a decent response to this charge. In Chapter Two we also saw that Korsgaard has an advantage over Street: Street's position cannot guarantee that counter-intuitively wrong acts will always count as morally wrong. Korsgaard is able to protect her position from this problem because for any counter-intuitively wrong act against a person, that person's humanity will be violated; and since her position requires that we avoid doing this, all counter-intuitively wrong acts against people will be morally wrong (perhaps barring some extraordinary situations). I argued that there are some resources available to Street to make this problem a little less severe; but, ultimately, there is no way to secure the claim that all counter-intuitively wrong acts will always count as morally wrong according to Street's position.

In Chapter One, I argued that metaethical realists face an epistemological problem that prevents their theories from being able to claim any knowledge of moral facts. In Chapter Two I argued that both Korsgaard and Street avoid this problem. Overall, Korsgaard's position has an edge over Street's. Korsgaard's position has a plausible phenomenology of deliberation, it closely aligns with our deep ethical convictions regarding immoral behavior, it captures the impartial nature of morality, and it maintains the categorical nature of moral demands—Street's position does not secure these benefits. These benefits, along with the fact that the objection we saw against Korsgaard's position seems solvable, are enough to give her position an advantage over Street's. If the objection to Korsgaard seemed absolutely unsolvable or if Street had any superb advantages over Korsgaard then perhaps it would be toss-up between the two views. But this isn't the case. Korsgaard's position has an appreciable edge.

REFERENCES

- Audi, R. (1996). Intuitionism, pluralism, and the foundation of ethics. In W. Sinnott-Armstrong & M. Timmons (Eds.), *Moral knowledge? New readings in moral epistemology* (pp. 101-136). Oxford: Oxford university press.
- Ayer, A. J. (1952). *Language, truth, and logic*. New York: Dover Publications.
- Bagnoli, C. (2017) Constructivism in metaethics. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. (Fall 2017 ed.). Retrieved from <https://plato.stanford.edu/archives/fall2017/entries/constructivism-metaethics/>
- Blackburn, S. (1984). *Spreading the word: Groundings in the philosophy of language*. Oxford; New York; Clarendon Press; Oxford University Press.
- Brink, D. (1989). *Moral realism and the foundations of ethics* (Cambridge studies in philosophy). Cambridge; New York: Cambridge University Press.
- Enoch, D. (2009). Can there be a global, interesting, coherent constructivism about practical reason?. *Philosophical Explorations*, 12(3), 319-339.
- Enoch, D. (2011). *Taking morality seriously: A defense of robust realism*. Oxford: Oxford University Press.
- Glassen, P. (1959). The cognitivity of moral judgments. *Mind*, 68(269), 57-7.
- Huemer, M. (2005). *Ethical intuitionism*. Palgrave Macmillian.
- Huemer, M. (2008). Revisionary intuitionism. *Social Philosophy and Policy*, 25(01), 368-392.
- Hume, D., Norton, D., & Norton, Mary J. (2000). *A treatise of human nature* (Oxford philosophical texts). Oxford; New York: Oxford University Press.

- Korsgaard, C. (1996a). *Creating the kingdom of ends*. Cambridge; New York, NY: Cambridge University Press.
- Korsgaard, C. (1996b). *The sources of normativity*. Cambridge; New York: Cambridge University Press.
- Mackie, J. (1977). *Ethics: Inventing right and wrong*. New York: Penguin.
- Rawls, J. (1999). *A theory of justice* (Rev. ed.). Cambridge, Mass.: Belknap Press of Harvard University Press.
- Sayre-McCord, G. (1986). The many moral realisms. *The Southern Journal of Philosophy*, 24(S1), 1-22.
- Sayre-McCord, G. (2014) Metaethics. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. (Summer 2014 ed.). Retrieved from <http://plato.stanford.edu/archives/sum2014/entries/metaethics/>
- Shafer-Landau, R. (2003). *Moral realism: A defense*. Oxford University Press on Demand.
- Smart, J. (1956). Extreme and Restricted Utilitarianism. *The Philosophical Quarterly* (1950), 6(25), 344-354.
- Street, S. (2006). A Darwinian dilemma for realist theories of value. *Philosophical Studies*, 127(1), 109-166.
- Street, S. (2008). Constructivism about reasons. *Oxford Studies in Metaethics*, 3, 207-45.
- Street, S. (2009). In Defense of Future Tuesday Indifference: Ideally Coherent Eccentrics and the Contingency of What Matters. *Philosophical Issues*, 19, 273-298.
- Street, S. (2010). What is constructivism in ethics and metaethics?. *Philosophy Compass*, 5(5), 363-384.

- Tropman, E. (2014). Evolutionary debunking arguments: moral realism, constructivism, and explaining moral knowledge. *Philosophical Explorations*, 17(2), 126-140.
- Tubert, A. (2011). Korsgaard's Constitutive Arguments and the Principles of Practical Reason. *The Philosophical Quarterly*, 61(243), 343-362.
- Wallace, R. (2014). Practical Reason. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Summer 2014 ed.). Retrieved from <https://plato.stanford.edu/archives/sum2014/entries/practical-reason/>