

University of Richmond

UR Scholarship Repository

Honors Theses

Student Research

2020

So What: The Justification of Morality in Christine Korsgaard's The Sources Of Normativity

Julian Scott
University of Richmond

Follow this and additional works at: <https://scholarship.richmond.edu/honors-theses>



Part of the [Philosophy Commons](#)

Recommended Citation

Scott, Julian, "So What: The Justification of Morality in Christine Korsgaard's The Sources Of Normativity" (2020). *Honors Theses*. 1515.

<https://scholarship.richmond.edu/honors-theses/1515>

This Thesis is brought to you for free and open access by the Student Research at UR Scholarship Repository. It has been accepted for inclusion in Honors Theses by an authorized administrator of UR Scholarship Repository. For more information, please contact scholarshipprepository@richmond.edu.

So What: The Justification of Morality in Christine Korsgaard's *The Sources Of Normativity*

By

Julian Scott

Honors Thesis

Submitted to:

Philosophy Department
University of Richmond
Richmond, VA

May 1st, 2020

Advisor: Dr. Will Reckner

Part One: Introduction

Phillipa Foot once described the case of a Sudeten farm boy who, in 1944, had to choose between joining the SS and being executed (Foot 2). Perhaps it is clear that he was right to choose death rather than join such an evil organization. But if he disagreed, what are we supposed to say to him? Suppose he said that he knew the SS was evil and that joining would require him to do evil things, but when faced with the alternative of execution, why did he have to do the right thing?

This problem can be generalized: what are we to say to someone who is in between a rock and a hard place? Someone who has the choice between doing what is right and staying alive? It is in those moments when a particular problem begins to press. We can see what is right, but we find ourselves driven to ask why we must do what is right. That is, we ask: why is morality normative for us? Why does morality guide, command, and obligate us? What justifies the claims morality makes on us?

In her book, *The Sources of Normativity*, Christine Korsgaard takes up the problem of the normativity of morality. She argues that we can push a regress anytime we think we have a reason to do something and the only place that regress can end is with our nature as reflective beings (what she terms our humanity [see Korsgaard 121]), a nature we share with each other. Luckily, she finds, this nature serves as a firm foundation. If we are to have any reasons at all, we must regard our nature as a reflective being as a source of reasons (in Korsgaard's terminology, we must value it [see Korsgaard 121]). So long as our reflective nature is a source of reasons for us, we have nothing to fear. The moral law is a law for us.

But there's a snag. Maybe you've seen it already. A simple way to notice it is to ask what Korsgaard would have to say to the Sudeten farm boy. She might explain her argument down to the conclusion that we must regard reflective consciousness as a source of reasons if we are to have any reasons at all. The farm boy might agree with all of that, saying: "Annnnnnd?" Because saying that he must treat reflective consciousness as a source of reasons only addresses half of his problem. It remains

unclear why we are thus committed to choosing death rather than evil. So, Korsgaard must have something more to say to the Sudeten farm boy. She must offer some explanation of treating reflective consciousness as a source of reasons that explains our commitment to morality. If she cannot, she will have failed to justify morality in anything like the form we experience it. This would mean the failure of her project to justify morality.

I will first summarize Korsgaard's view as presented in *Sources*. Then, I will explain in more detail how the problem arises and what it is. Next, I will consider one alternative for how we might solve this problem. This first solution will revolve around arguing that our reflective consciousness gives us reasons to protect and preserve reflective consciousness. But this argument will fail because it is neither supple enough to accommodate actions that, while usually wrong, are permitted in certain circumstances nor strong enough to explain why we think many actions are wrong (notably, coercion). Last, I will present and defend my alternative theory.

I think Korsgaard's account has resources that we can extract to frame an answer to the Sudeten farm boy. In my view, when we are doing evil, we are ignoring the reasons that other people present to us and thus implicitly denying that there are any reasons at all. Thus, if we are to have any reasons at all, we must never do evil. Since we are the kind of being that needs reasons in order to do anything, we are required to avoid evil.

Part Two: Korsgaard's Argument

I will provide a foundation for my argument by summarizing Korsgaard's argument in *Sources* as I understand it. The summary will help clarify why Korsgaard has not gone far enough to show that commonsense morality is in fact normative for us. This is intended to be sketch of Korsgaard's argument. I will provide the main structure of the theory, not go detail on all points, and I will simply earmark objections. My goal is to get the main architecture of the theory on the table.

This architecture starts with the perspective of someone who has to make a decision. For example, let us say that a decision has to be made between spending the afternoon working and spending it playing. A person who is attempting to make this decision is aware of their desires as separate from themselves. They are aware that they desire to play because it would be fun but also that they desire to work (in some suitable sense) because it is important to do the work at some point. This is the kind of standpoint a being with reflective consciousness (as Korsgaard calls it) takes (Korsgaard 93).

We should make a couple of notes about this position. First, we can push a regress against any consideration that disposes us one way or the other. If I want to play, I can ask why. If I answer that it would be fun, I can ask why that matters for my decision (and so on and so on until I run out of answers) (Korsgaard 92-93). This is why, Korsgaard says, we act under the idea of freedom, because, when we are deciding, we cannot simply accept dispositions to act in a certain way. We must have reasons (Korsgaard 94). Second, when we deliberate, it is as if there is some being that is separate from and above our impulses that will pick between them and that being is us (Korsgaard 100). When I must decide between working and playing, it seems as if there is an *I* that will decide, even if to the outside observer it may seem that my actions always follow my strongest desire. This means that I must have a conception of myself because I have an awareness of this being which is me (Korsgaard 100). Because we think of the decision maker as our self in a particularly deep sense, our choices are expressive of ourselves (Korsgaard 100).

Korsgaard points out some further facts about what happens when we make a decision. The first is that we are making reasons for ourselves when we decide things. We doubt that we have reasons when we are frozen in reflection. Thus, if our impulses survive that reflection and we find we can act on them, they must be reasons (Korsgaard 93). Second, when we choose in this way, we must choose in a principled manner. To see this, consider an analogy between reasons and causality. Both

involve some idea of power and some idea of regularity. Why are power and regularity joined in causation (Korsgaard 225-226)? As Hume points out, we could never recognize causality without regularity (Korsgaard 226). Korsgaard adds that, in order to be agents, we must think of ourselves as the causes of actions. Thus, we must be able to distinguish between us being the cause of actions and something in us (like a compulsion) being the cause of actions (Korsgaard 227-228). Therefore, in order to be agents, we must think of ourselves as connected with actions in a regular way (Korsgaard 227). Thinking of ourselves as connected with actions in a regular way requires that we choose in a principled way (Korsgaard 229). Thus, Korsgaard concludes (though perhaps we need not), impulses that are suitable to be decided on (and thus to be reasons) must have the form of laws. That is, they must conform to the Universal Law Formulation of the Categorical Imperative (Korsgaard 97-98).^{1,2,3}

But, Korsgaard thinks, it is still reasonable to ask what domain our principles are to operate in (that is, what kind of being they apply to) (Korsgaard 98-99).⁴ What domain our principles are to operate in makes a practical identity (Korsgaard 100). This, she says, is a description under which we value ourselves and find our lives worth living and our actions worth doing (Korsgaard 101). It seems to me that the most important characterization of a practical identity is that it is the identity we adopt in deliberation. A particular practical identity is simply a conception of ourselves that we have as we deliberate (which Korsgaard has already pointed out we must have). Reasons are given to us based on the principles we adopt from our practical identities. Studying for exams could be a law for students, and so, as a student, I endorse my impulse to study. Obligations, crucially, spring from what our

¹ It is important to note that Korsgaard does not think (at least for the purpose of argument in *Sources*) that this gives us the entire content of morality (see 99).

² For another version of this argument see Korsgaard 94-98.

³ The Formula of Universal Law is: "act only according to that maxim through which you can at the same time will that it become a universal law." (Kant 34)

⁴ In this, she differs from Kant in a way that has sometimes attracted criticism.

identities forbid (Korsgaard 101-102). If my identity as a student forbids me from failing to turn in an assignment, I am obligated to turn in that assignment.

Thus, Korsgaard takes herself to have shown that we have reasons. The next step is for her to show that we are committed to seeing reflective consciousness as a source of reasons. I think this argument is supposed to go something like this: most of our practical identities are particular. I am a student but not everyone is. But whatever practical identity supports morality would have to be universal if morality is to be universal. So, we must also ask whether there are any universal practical identities (Korsgaard 113). Korsgaard thinks that a reflective being is a necessary practical identity and that it is the practical identity that grounds morality (Korsgaard 121). She offers what seems to be a set of closely related arguments for this position. I am not entirely sure how they are supposed to work and for this reason at least, not quite sure whether they do work. I think the simplest way of getting to the heart of her argument is to consider a situation in which we have a principle but are unsure if we should decide on this principle. That is, we are trying to decide whether to decide. Now, we must make a decision in order to act. If we cannot make a decision, then we are frozen, that is a fact about what it is like to be us. If we decide to make a decision, we are recognizing that we cannot act without making a decision and adopting this fact as a practical identity for us. This practical identity is none other than a being with a reflective consciousness for it is this identity that makes it such that we must decide in order to act. Thus, insofar as we decide in light of the fact that we must decide in order to act, we must affirm that reflective consciousness is a practical identity that we have. As Korsgaard puts it, we are reflective animals who need reasons “in order to act and to live.” (Korsgaard 121). That is, reflective consciousness is a source of reasons and obligations for all of us.⁵

But this presents a new problem: I must regard reflective consciousness in me as a source or reasons, but must I similarly regard reflective consciousness in you? Korsgaard thinks that if reasons are

⁵ For the arguments I attempt to capture here, see 118-125.

essentially shareable (that is, public, not private), then your reflective consciousness is a source of reasons for me as well. We could only think that reasons were private if we thought that reflective consciousness were private (Korsgaard 136).

So Korsgaard sets out to show that reflective consciousness is not private. If reflective consciousness (and thus our own individual consciousnesses) were private, then the meaning of words would be private as well. But meaning is not private (Korsgaard 138). If I say to you, the grass is green, you know what I mean. Moreover, it is hard for you to not know what I mean if you understand the language in which I address you. You cannot hear what I have said as mere noise. This means that your consciousness is not private. For I can intrude on your consciousness. In fact, it's quite easy to do so. All I have to do is address you in a language you understand (Korsgaard 139). Because consciousness is not private, reasons are not private either.⁶ As Korsgaard puts it: "We do not seem to need a reason to take the reasons of others into account." (Korsgaard 140) Therefore, your consciousness (i.e. you) are a source of reasons for me.

Part Three: The Problem

To see what is wrong with Korsgaard's argument, let's take a step back and consider why we thought such an argument was needed in the first place. Korsgaard introduces the problem she is trying to solve by saying:

"... the day will come, for most of us, when what morality commands, obliges, or recommends is hard: that we share decisions with people whose intelligence or integrity don't inspire our confidence; that we assume grave responsibilities to which we feel inadequate; that we sacrifice our lives, or voluntarily relinquish what makes them sweet. And then the question- why?- will press, and rightly so.

⁶ Not essentially private anyway. Korsgaard seems to want to leave it open (plausibly) that we might not be aware of all reasons at all times and this creates a way in which reasons are private. But they are only incidentally private when they are private on epistemic grounds like this.

Why should I be moral? ... We are asking what justifies the claims that morality makes on us. This is what I am calling 'the normative question'." (Korsgaard 9-10)

Would someone in this position be satisfied with the answer that Korsgaard has given? It seems unlikely. Even if we concede that everything Korsgaard says in *Sources* is right, all we have is that we must regard reflective consciousness (in ourselves and in others) as a source of reasons.

But it is difficult to see how this is supposed to establish any part of commonsense morality. Someone who is in the position of asking why they must do what morality demands may, it seems, be forgiven for asking why regarding reflective consciousness as a source of reasons commits them to doing this thing in particular. The Sudeten farm boy, who knows he will be killed if he does not join the SS, could reasonably ask why the thing that we have shown to be normative and named morality requires him to not join the SS. He needs to hear more about what regarding reflective consciousness as a source of reasons commits him to.

Thus, we are looking for an answer that will tell us why treating reflective consciousness as a source of reasons commits us to something like commonsense morality. This answer will take the form of an explication of treating reflective consciousness as a source of reasons; it will either tell us something about what it is to do so or what doing so commits us to. If there is not an answer, then Korsgaard will have failed to justify morality.

I want to articulate two ways we might find that Korsgaard's account fails to match our commonsense morality. These will be important criteria in looking for an effective explication of treating reflective consciousness as a source of reasons. The account might fail to be sufficiently strong or it might fail to be sufficiently supple.

When I say that it must be sufficiently strong, I mean that it must be able to account for a wide range of cases of immoral action. A guiding example here will be that of coercion. Say that I greatly desire that you turn in a particular essay. But you refuse to turn it in. I want you to turn in the essay so

badly that I take out a gun, point it at your head, and shout that I will shoot if you do not turn in the essay with the intention that this threat will result in you turning in the essay. It seems that me trying to coerce you with a threat of death is wrong, even if I have no intention of actually carrying out the threat.

But it is not obvious why being committed to regarding reflective consciousness as a source of reasons rules out coercion. Is coercion contrary to these reasons? If so, why? For us to be obligated not to coerce people, we would need to say that it was forbidden by our identity as reflective beings. Why is coercion forbidden by reflective consciousness? It does not seem like it would destroy reflective consciousness or even damage it. Perhaps all I have done is give you a reason to turn in the essay (a reason I expect to trump all others). A satisfactory answer to the question of why we are committed to morality will tell us why treating reflective consciousness as a source of reasons forbids coercion. That is, it must be strong enough.

Similarly, when I say that our theory must be sufficiently supple, I mean that it must account for cases in which otherwise impermissible actions turn out to be warranted. A guiding example here will be killing in self-defense. I take it that in commonsense morality, there is a strong presumption against killing. But, consider a case like this: one day you are walking down the street and you see me running towards you with a knife. Knowing that I have sworn to kill you and that I have been single-minded in pursuit of that goal, training night and day so that I will be able to kill you no matter what, you know that the only way to save your life is to kill me. So, you take out a weapon and kill me at the first opportunity. Now, there are three critical elements to this story. First, killing me is the only way to end a threat I willingly pose to your life. Second, you kill me. Third, your action in killing me is permissible.

But why is it permissible? It is still killing after all. Every killing, well, kills someone. They all seem to affect reflective consciousness in exactly the same way. A satisfactory explication of treating reflective consciousness as a source of reasons must be able to explain why some killings are

impermissible and others are permissible. That is, the account must be supple enough to conform to commonsense morality.

If we can arrive at an account of treating reflective consciousness as a source of reasons for us that can answer both these cases, then that account is probably strong and supple enough to accommodate a significant portion of our ordinary moral judgements. If a proposal can do that, the proposal is a promising one for how we are to understand seeing reflective consciousness as a source of reasons.

Part Four: Why Having Reasons Not To Harm Reflective Consciousness Does Not Get Us Morality

One way that we might try to answer this problem on Korsgaard's behalf is to say that we are obligated not to do immoral actions because they in some way harm or damage reflective consciousness. I shall consider two ways of cashing out this kind of account. The first holds that all immoral actions damage reflective consciousness, where reflectivity is a structure of consciousness, and that immoral actions are wrong because they damage reflective consciousness. The second holds that all immoral actions damage reflective consciousness, where reflection is a kind of characteristic activity, and that immoral actions are wrong because they disrupt this characteristic activity.

The answer based around a structure of consciousness might start with a thought like this: the fact that certain actions damage reflective consciousness is important because we need to have reasons in order to act and to live. This is the kind of being we are. But we can only get reasons from the exercise of reflective consciousness. So, we must have a reflective consciousness in order to get reasons. Because we need a reflective consciousness to arrive at reasons and we must live and act for reasons, we must protect our capacity for being able to arrive at reasons. Thus, we are obligated to protect reflective consciousness in us. Because the space of our consciousness is essentially public and our reasons essentially shareable, it seems to follow that we must protect reflective consciousness in others as well.

Now, this kind of answer can explain certain important moral judgements. For instance, it might explain why murder is wrong. Murder damages reflective consciousness because it destroys reflective consciousness in a particular person. Once someone is dead, the reflective structure of consciousness cannot exist in them. It has been eliminated. And we think that murder is wrong.

This explanation, as an interpretation of Korsgaard, also maps onto some interpretations of Kant. As Korsgaard is a latter-day Kantian, this seems to speak in favor of this explanation as an interpretation of Korsgaard. If we think the requirement in first formulation of Kant's categorical imperative that we be able to will our maxims as universal laws is in some important way about safeguarding the preconditions of our own human agency,⁷ then it seems that Korsgaard's supposedly parallel argument should also involve safeguarding the preconditions of human agency. In particular, the precondition of human agency is reflective consciousness. Reflective consciousness must function in order for us to continue as agents. If reflective consciousness in me was destroyed (for instance, if I was stuck in a car under a landslide for three days and never recovered the ability to think coherently or separate myself from immediate fears and desires), I would cease to exist as an agent, even if my physical body persisted. Thus, reflective consciousness is a precondition of human agency. If Korsgaard is concerned with protecting it on this account, this would be consistent with Kant.

Moreover, it seems that one interpretation of Kant's second formulation of the categorical imperative ("So act that you use humanity, in your own person as well as in the person of any other, always at the same time as an end, never merely as a means." [Kant 41]) is that it is about protecting the preconditions of valuing. That is, we must grant humanity some special value and this value requires never treating humanity merely as a means because humanity is a precondition for valuing anything at all. Korsgaard herself recommends this interpretation of Kant (see Korsgaard 122). Given that Korsgaard takes her argument to be parallel to Kant's (Korsgaard 122) and that Korsgaard herself endorses this

⁷ As Barbara Herman seems to think (see Herman 122, for instance).

view of the formula of humanity, it seems like Korsgaard's argument is intended to secure a similar kind of special status for reflective consciousness.⁸ This special status is a kind of protection. Because reflective consciousness is a precondition for having any reasons at all, we must protect reflective consciousness or lose the ability to have reasons. Thus, as we must continue to have reasons in order to live our lives, we must protect reflective consciousness.

Another way to understand reflective consciousness is to think of it as a kind of characteristic activity. Our consciousness is reflective insofar as we reflect on what we have reason to do and in light of the fact that we need reasons in order to act and to live. Merely having a mind sufficiently powerful or structured to make such reflection possible is not sufficient. I could be in a coma and still have the capacity or mental structure for reflection, but if I absolutely never actually reflect, in some important way, I do not have a reflective consciousness.⁹ Anything that prevents us from reflecting on reasons harms reflective consciousness in that it prevents us from carrying out the activity of reflective consciousness. Thus, moral obligations arise to protect reflection as an activity. So, we have a moral obligation not to do things that interfere with reflection.

This second explanation of the obligations that arise from our reflective identity is similar to the first version. As such, it is plausible for many of the same reasons. If we can accept that reflection is a characteristic activity and not merely a structure of consciousness, it would seem that we must be able to actually reflect if we are to have reasons. Furthermore, if we accept that reflection is a characteristic

⁸ Remember especially that she considers the terms humanity and reflective consciousness to be equivalent.

⁹ Now, this may look like a substantial revision of Korsgaard's view. She does after all write about reflective consciousness as a capacity that makes reflection possible (see Korsgaard 92-93). I do not intend to explore how important this revision would be. I want to suggest that this route, as an interpretation, would explain Korsgaard's focus on the necessity of reflection for a reflective consciousness (see Korsgaard 93) and her insistence on the importance of actual endorsement (see Korsgaard 254). I further want to suggest that if this is a viable way of vindicating morality, then that counts in favor of it as a route that Korsgaard could take. It seems to me that this route, despite differences, builds on Korsgaard's system.

activity, then, as reflection would then be a precondition of agency, in order to parallel Kant's argument, we must think that we are obligated not to interfere with reflection.

Most importantly, it seems that acts we ordinarily think of as morally wrong are those that interfere with reflection. Murder cuts short current reflections and prevents future reflection. And it is wrong to murder someone. Causing someone unbearable pain stops them from reflecting. And causing someone high degrees of pain is morally wrong (usually at least). It may also at first appear as though this conception can explain the wrong involved in coercion. For, we might be inclined to think that me pointing a gun at you and demanding that you turn in an essay or I'll shoot would interfere with your reflection. On this view it would therefore be wrong to do so.

As it happens, I think neither of these explanations of treating reflective consciousness as a source of reasons really goes far enough to vindicate morality. It seems unlikely that either of them is sufficiently supple or strong to explain important moral judgments. Because of this, I suggest that we abandon these accounts in favor of stronger and more supple accounts like the one I provide in the next section. But first, why do I think these accounts are likely to fail?

I will start with the strength of the account based on reflections as a characteristic activity because I said before that there is some reason to think that it might be able to explain the wrong in coercion. Despite an apparent solution to the problem of coercion, saying that reflection is a characteristic activity does not eliminate all difficulties with the immorality of coercion. The continued difficulty arises from the possibility that coercion is merely offering people reasons to do things they would not otherwise do.

To see why, consider this: If I put a gun to your head and tell you to turn in the essay or I'll shoot, I have given you a reason to turn in the essay. You certainly have reason to preserve your own life. Your life projects and goals are unfinished. Me pointing a gun at your head and saying that I will shoot you if you do not turn in the essay gives you a reason to believe that you will die if you do not turn

in the essay. These two things surely give you a reason to turn in the essay. If, when I threaten you like this, I am only giving you a reason to turn in the essay, then we should not say that I have interfered with your reflective consciousness. For, all I have done is give you an input into reflection. In this way, it is similar to me asking you to pass me the salad when we sit down to dinner. That is, in combination with certain background facts (like common American customs and that you hold me in a certain kind of regard) gives you a reason to pass the salad. While presumably the reasons you have to turn in the essay are better when I threaten you than your reasons to pass the salad when I ask, I cannot see why this fact makes it from a permissible way of acting into a wrong.

In order to say that coercion interferes with reflection I think we would need to say one of three things was true about the coercion scenario. We would need to say that when I coerce you 1) I intend to interfere with reflection (and this makes it impermissible), or 2) that me coercing you tends to interfere with reflection (and this fact makes it impermissible), or 3) that it is the fact that the action I threaten to do is wrong that converts what would otherwise be merely offering an input for reflection into interfering with reflection.

My intentions are not a good way to explain why coercion is wrong. For, while it is possible that I could intend to circumvent your reflection by threatening you, it is not necessary that I do. Perhaps I am a coercer of an unusually philosophical bent and very specifically intend to give you a remarkably strong reason to turn in the essay and that I expect you to respond to this reason by turning in the essay. In this case, I do not intend to circumvent your reflection. But even if I have this intention, it seems that coercing you to turn in the essay is wrong. If this is true, whether or not I intend to interfere with your reflection is irrelevant.

Nor is the tendency of my action a good way to solve this problem. If we think that a tendency or probability that an action will have a have the consequence of interfering with reflection is relevant, then it will be permissible to coerce someone anytime that consequence is impossible. So, if I know you

to be of steely will and mind and know that you are unafraid of death (but do not want to die), I can threaten you with impunity, confident that I am only giving you a reason to turn in the essay and that I shall in no way disrupt your reflection by threatening you. But this seems like an unwelcome consequence and so relying on the tendency of an action to have certain consequences does not seem to be a good way of explaining why coercion is wrong.¹⁰

Similarly, the moral valence of the action I threaten you with is irrelevant. To see this, consider the contrast between two cases. In the first, I threaten you with death. In the second, I threaten to break something that I own that you value greatly (perhaps a family heirloom I received from our common grandparents). It seems clear that both involve coercion and are both forbidden because of that. But the threatened action, while forbidden in the first case, is permissible in the second (I have a right to destroy my own property, it seems). Given that the permissibility of the threatened action does not impact whether the threat is objectionable as coercive, it seems unlikely that that the permissibility of the threatened action makes it such that the threat interferes with your reflection.¹¹ Thus, none of the three ways of cashing out harming reflective consciousness as interfering in reflection that we have reviewed seems able to explain why coercion is wrong.

Nor does it seem that the explanation of harming reflective consciousness where we take reflective consciousness as a structure of consciousness performs any better. The first account cannot explain why me putting a gun to your head and saying that I will shoot unless you turn in the essay is wrong. For coercing you like this is not deleterious to your reflective consciousness because it is not

¹⁰ And anyway, as Korsgaard is a prominent latter-day Kantian, it seems we should regard any interpretation that involves consequentialist style reasoning as distinctly undesirable.

¹¹ I put it this way because I find it plausible even before we consider examples that the permissibility of a threatened action does not affect whether the threat interferes with reflection. An alternative reading of this example is that the permissibility of the threatened action does impact whether the threat interferes with reflection but, because the threat is wrong in both cases but regardless of whether the threatened action is permissible, interference with reflection does not have the effect of making actions impermissible. This reading second seems to be a stronger claim, so I have opted for the other alternative.

deleterious to your capacities. Telling you that I will shoot does not change the structure of your consciousness or render you incapable of reflection. Even if I do threaten you, you will be able to reflect on your impulses, still need reasons to act, and still be capable of arriving at reasons (even if you do not do these things, having responded to my threat out of fear). So, this account will struggle to explain why coercion is wrong.

Moreover, neither way of cashing out harming reflective consciousness is sufficiently supple to explain the permissibility of killing when killing is required for self-defense. Both rule out wanton killing. Were I to kill you, for fun let's say, I would be destroying reflective consciousness in you and interfering with your reflections (by ending them). Because I would thus destroy reflective consciousness in you and interfere with your reflections, wanton killing goes against my obligation to protect reflective consciousness. Now, the question before us is what distinctions between self-defense and wanton killing these theories can make.

It does not seem that these theories can draw a distinction based on the fact that in self-defense, you are defending yourself from a moral wrong. To see why, consider what the basic prescriptions of each theory would be. For the first: protect the capacity for reflection. For the second: protect against interference in reflection. These prescriptions are, according to the explications of treating reflective consciousness as a source of reasons that I have put forward here, at the heart of all of morality. Neither of them conditions whether their prescription holds on any circumstance whatsoever (that is, they are categorical or unconditional). Under these theories, our obligation to protect reflective consciousness in someone else is not conditioned on whether that person is willing to do the same for us. So, neither of them seems to support a distinction between self-defense and aggression based on simply the defense against moral wrong.¹²

¹² Perhaps this is wrong and they can. But that would require a substantial elaboration of the theory and I suspect such an effort would ultimately prove unsuccessful. Given the elaboration required and the uncertainty of success, I am inclined to accept this *prima facie* failure.

Maybe though, we might want to adopt a line like this: the prescriptions of the two theories we are thinking about do not offer determinate guidance in cases where self-defense of this kind is appropriate. That is, when you must kill me in order to save yourself, there is no way that you will be able to protect all reflective consciousness. So, our moral rules do not offer definite guidance in this case. But you still have to choose between protecting reflective consciousness in yourself and protecting it in me. Therefore, because you have no definitive guidance and you have to pick, you can pick whichever. This explanation would capture the permissibility of self-defense and also the fact that we do not think (generally at least) that, when self-defense requires killing your attacker, you are required to defend yourself.

But this explanation fails to pick out quite the right issues. This is because, sometimes when we have to choose between dying ourselves and killing others, we should prefer to die than to do evil. This is the conclusion of the Sudeten farm boy that we are trying to do justice to, remember. But if morality offered no determinate advice in cases where we had to choose between dying and killing (as the above analysis suggests), then morality would offer no determinate advice to the Sudeten farm boy. And this is not what we think about that case. Nor is this an idiosyncratic fact about the unique evil that the farm boy would be participating in by joining the SS. In general, we think it is wrong to kill when we are coerced to do so. If I threaten to kill you unless you kill someone else and you do kill them, you have still done something wrong. You might escape punishment and reprobation, but that is because my threat excuses you for the wrong. But, in our case of self-defense, you have not committed a wrong in killing me. What you did was justified and permissible, not merely excusable. Thus, answering that we have no determinate guidance when we must choose between reflective consciousnesses to protect does not correspond to our moral judgements.

Our attempts to draw the line between self-defense and wanton killing do not seem to succeed. Therefore, we should return to our initial conclusion: that these two ways of explicating not harming

reflective consciousness are both unlikely to be able to distinguish wanton killing from self-defense. Therefore, they are unlikely to be supple enough to account for critical moral judgments.

Part Five: A Solution

Korsgaard has argued that we are committed to seeing reflective consciousness as a source of reasons. I have pointed out that it is difficult to see how this satisfies Korsgaard's project, which is to vindicate morality. Because, after all, what does seeing reflective consciousness as a source of reasons commit us to in terms of morality? In the last part, we considered two ways of thinking about an obligation to protect reflective consciousness. We rejected those views because it did not seem like those obligations would be strong enough or supple enough to explain key moral judgments. In this part, I will propose a radically different answer. The core idea of this solution is that immoral action involves disregarding the reasons of other people and is immoral on that account. I will argue that this account is both strong enough and supple enough to explain the wrongness of coercion and the permissibility of self-defense (our paragon examples).

Let's start by considering this case: I am hurting you (let us suppose for fun). You tell me that it hurts and tell me to stop. You are expressing reasons for me to stop. At first, they are just your reasons. It hurts after all and the fact that it hurts is a reason for you to avoid it. But I hear and understand your complaint; your meaning penetrates my consciousness. Thus, you share your reasons with me and they become my reasons as well. Because you have given me reasons to stop in this way, I am obligated to stop (Korsgaard 142-143).

I want to proceed by raising some questions about this case that I hope will draw out the contours of the solution. Let's start with this: why am I obligated rather than just given a reason to stop? As we have seen an obligation is different from a reason because an obligation arises from the rejection of an action inconsistent with identity whereas reasons arise from the positive recommendations of a practical identity. Korsgaard also thinks that the practical identity of a reflective

being is the moral identity (meaning the practical identity that grounds morality). So, the answer to our question must be that my action is inconsistent with my identity as a reflective being.

But why is hurting you for fun inconsistent with my identity? There are a number of steps to answering this question. First, we should note that it does not seem that my actions are necessarily deleterious to the reflective consciousness of either me or you. To see this, consider if the hurt comes from me shooting rubber bands at you. This could cause pain, annoyance and be something that I am obligated not to do. But shooting rubber bands at you does not seem to be the kind of activity that could destroy reflective consciousness in either of us. I will not be reduced to a vegetative state by shooting rubber bands and nor will you by being hit with them.

If my hurting you for fun does not necessarily destroy reflective consciousness, how are my actions inconsistent with regarding reflective consciousness as a source of reasons? One potential answer seems to be that my actions simply deny that your reflective consciousness can produce reasons. In calling upon me to stop, you have given me reasons to stop. These reasons to stop (we will suppose for now at least, I discuss how to determine the strength of reasons below) are better reasons than my reasons to go on hurting you. And they are no longer just your reasons. They are shared reasons. Because you have presented these reasons to me (by speaking, you have forced me to think about them), I cannot simply fail to notice them. I must actively reject their status as reasons. That is, I must deny that they are reasons.

When I deny that they are reasons, I deny that reflective consciousness in you can produce reasons. Can I think that in general your reflective consciousness is capable of producing reasons but that it failed this time? For instance, perhaps I think that you do not actually endorse the impulse you express. This raises epistemic issues about how we are to know whether other people are expressing reasons or acting on impulse. I do not have easy answers to these kinds of questions. However, I want to make two suggestions. First, it seems to me that we generally have a strong perceptual capacity for

when other people act with purpose, rather than impulsively. Second, it is possible we should give people the benefit of the doubt as regards whether they express reasons or merely impulses. This seems to underlie some of our pause when someone has agreed to something in advance but expresses doubts at when it comes time to do it (as people sometimes do for medical procedures like surgeries). We have to check, to make sure they have not changed their mind. With these, admittedly gestural, remarks, I hope to have explained why I should think, when I hurt you, that you have merely failed to produce a reason on this particular occasion. Thus, I should think in general that reflective consciousness in you does not produce reasons. I should think your reflective endorsement is meaningless.

But, *ex hypothesi*, because I am acting, I must regard reflective consciousness in me as a source of reasons. Because we can share our consciousness, it cannot be that we are unlike beings. As Korsgaard points out, by speaking to me, you force me to acknowledge you as someone (Korsgaard 143). Thus, I am inconsistent when I deny reflective consciousness in you can produce reasons while saying that reflective consciousness in me does produce reasons. This is why Korsgaard says her argument does use an appeal to consistency (Korsgaard 143).

Why is the denial that reflective consciousness in you can produce reasons inconsistent with the moral identity? It is because you are denying that reflective consciousness is a source of reasons. When I deny that reflective consciousness in you is a source of reasons, I am in a position where, to be consistent with myself, I must also deny that reflective consciousness in me is a source of reasons. But I cannot do that without losing grip on my own identity and all of my reasons. Thus, my identity is inconsistent with my denial that reflective consciousness in you is a source of reasons.

Is this account strong enough to match our commonsense morality? I think we can answer this question by asking why we have reasons that others not treat us in certain ways. We can start with asking why we have reason not to be hurt and build that answer up into an answer to why we have

reason not to be coerced. If we have reason not to be coerced, then others have reasons not to coerce us. If they ignore those reasons, they have failed in their obligations.

We have a reason not to be hurt if we have reason not to be in pain. Thus, if we have reason to not be in pain, we have reason to not be harmed. To that end, Korsgaard argues that the first instance of pain (like getting hit with a baseball bat or shot) is an unreflective rejection of a threat to our physical self. When we are in pain, we have reason to change our condition because the physical basis of our consciousness will be destroyed if we let whatever is causing this kind of pain go on. Damage to our physical body is a kind of threat to our self. This causes a strong impulse to change our condition. This impulse is part of what we feel as pain. The fact that pain is an impulse to change our condition is what causes us to suddenly draw away from sources of pain (as when we unexpectedly burn ourselves on a pan handle) and is the reason that pain is frequently less serious when we can simply accept that it hurts (Korsgaard 147-149. Moreover, insofar as we are animals, we have reason to perpetuate our own existence. Therefore, whenever we are in this kind of pain, we have reason to change our condition (Korsgaard 1491-151).

But not all pains comprise a mortal threat to our physical identity. Some are even required for our health (like getting vaccines). What are we to make of those pains that do not represent existential threats? What can we say about shooting rubber bands? In addition, what are we to think of emotional pains like grief, which are not linked to our physical identity at all? Korsgaard's answer is that pain, all pain, is the rejection of some aspect of our condition. We reject conditions when we have reason not to be in them (Korsgaard 153-154). I want to apply this to a couple different kinds of pain before turning away from pain and back towards coercion.

Let us turn to the case of pains that do not correspond to existential threats. Why do we reject having rubber bands shot at us? Simple: it is a distraction. If we made a game of it, it would not hurt so much.

How about grief? What reason do we have to reject the loss of a loved one? I think there are at least two fairly clear answers. First, when we love people, we care about them. This means that their death, insofar as it negatively affects them (for instance, by cutting short their life plans) is also a loss for us. This explains why deaths that are less bad (for instance, when we think that 'it was their time to go,' or if we believe in an afterlife) frequently inspire less grief (than, for instance, untimely deaths). These kinds of thoughts are at least offered as consolations to the bereaved. Second, and, I think, much less importantly, we derive pleasure from interacting with people we love. The withdrawal of this pleasure is also something that we reject.

Something important to note is that all of these reasons arise from particular practical identities. Physical pain is about my particular body. Grief is about loving a particular person. The reasons we have against being in certain conditions all seem to arise from our particular practical identities. This is why Korsgaard never says what reasons arise from being required to regard reflective consciousness as a source of reasons. Having to regard the impulses that have been endorsed by reflective beings as reasons is enough to get us crucial aspects of morality without any further "special" reasons to be deduced from our commitment to regarding reflective consciousness as a source of reasons.

Now, having finished our discussion of pain, let's turn back to coercion. If we have reason to reject being coerced (much as we have reason to reject being in pain), then we have reason not to be coerced by others (just as we have reason not to be hurt by others). Thus, if we have reasons not to be coerced, there are reasons for other people not to coerce us and ignoring those reasons is inconsistent with regarding reflective consciousness as a source of reasons. Therefore, I am obligated not to coerce you if you have reasons not to be coerced. To see that you do, take the case where I threaten you with a gun in order to get you to turn in an essay. You had reasons not to turn in the essay, those reasons explain why you did not do so already. Maybe you have not finished it to your satisfaction. Maybe you know that you will be gripped with crippling fear once you turn it in. Any reason you have to not turn in

the essay is a reason not to be forced to turn it in. I think this is an instance of a general principle: if you have reason not to do something, you have reason not to be forced to do it. If you are concerned that this principle may not hold, I think it helps to consider another instance of it. Say that someone is pleading with a hit man to spare their life. They might say that their family relies on them for income or that they have not had a chance to see their children grow up. These are the reasons they have to not die. But also, these are reasons for the hit man not to kill them (i.e. that they not be forced to die). Thus, if you have reason not to turn in the essay, you have reason not to be forced by me to turn in the essay.

More than that, we also have reason not to be coerced because coercion causes fear of the threatened consequence. You fear being shot because being shot would be a threat to the physical basis of your identity. But we also have reason not to be afraid. Much like being shot with rubber bands, it is a distraction. Moreover, I take it that fear is a kind of pain. Fear is a rejection of a certain condition that has yet to occur. But, just as with pain, it is the rejection of a condition that we have reason not to be in. Just as we then have reason not to be in pain, we similarly have reason not to be in fear. And similarly, just as we have reason not to be hurt, we have reason not to be frightened.

Thus, me coercing you into turning in the essay would involve me denying that the reasons not to coerce you are reasons at all. Denying that the reasons not to coerce you are reasons at all is inconsistent with regarding reflective consciousness as a source of reasons. Thus, I am obligated not to do it. Because this explanation of what regarding reflective consciousness as a source of reasons commits us to can explain what is wrong with coercion, it looks like it is strong enough to match commonsense morality.

Now, you have probably noticed that everything I have said about having to ignore your reasons comes crashing down if I have good reasons for what I do (better reasons than the ones you provide against, anyway). For instance, sometimes, when you hurt me, you have good reasons for doing so.

Maybe your stabbing me with a knife will prevent me from killing you. Sometimes those reasons are such good reasons that they overcome my reasons to not be stabbed. When they do, you need not deny the reasons not to stab me in order to stab me. When you do not need to deny the reasons against stabbing me to stab me, then you have not done anything inconsistent with reflective consciousness when you do stab me. Thus, you were not acting against your obligations. I want to try and say what it takes to have better reasons. I think, once we are clear about this, it will also be apparent why my account is supple enough.¹³

What makes some reasons better is a difficult question to answer. We cannot simply say that whatever reasons seem to me to be better are actually the better reasons. Then I could never be wrong when I act on my choices and there would be no normativity.^{14, 15} So, we need some standard apart from how an agent perceives reasons. However, reference to any standard outside the reflecting agent would simply reintroduce substantive realism which Korsgaard has rejected already (this is part of what the second half of lecture one is about, see 47-48 for Korsgaard's conclusion). Thus, our answer must somehow be internal to the reflecting agent.

Korsgaard seems to suggest a heuristic for determining when the reasons against something are better than the reasons for it. She writes, "I invite you to consider how you would like it if someone did that to you. You realize that you would not merely dislike it, you would resent it. You would think that the other has a reason to stop, more, that he has an obligation to stop." (Korsgaard 143)¹⁶ But this is

¹³ Another way to answer this question would be to talk about reason defeaters. Korsgaard may suggest something like this answer when she says that we need a reason not to take the reasons of others into account (see 140-141). I will not follow up trying to answer this way because 1) it seems my account would need to have an explanation of what makes reasons better or worse just as a matter of theoretical neatness, 2) I think the picture I offer based on the strength of reasons can solve the suppleness problem, and 3) it is very unclear how reason defeaters could be integrated in the framework (especially given that Korsgaard does not explicitly mention anything about them).

¹⁴ Korsgaard insists that normativity requires the possibility of error. See 137.

¹⁵ If you are worried here about akrasia, I must say that I suspect that Korsgaard would say that akrasia is about acting on impulse and not on our choices.

¹⁶ It possible that I am wrong when I assert that Korsgaard offers this as a heuristic for determining when the reasons against something are greater than the reasons for. However, it seems likely that this is the argument

merely a mechanism to help us approximate the weight of reasons. After all, it still relies entirely on our perception of reasons and simply asks us to broaden our perspective by imagining that we are on the receiving end of the actions we are contemplating undertaking. If we are looking for an answer that allows us to determine the strength of reasons without relying on our perceptual capacity, we obviously cannot stop here.

I will consider two potential answers to this question. First, I will consider an answer that holds that the better reasons come from more basic or fundamental practical identities. Second, I will consider an answer that holds that reasons to do something overcome reasons against doing it when they are shared in the right kind of way. I will endorse this second way as more promising.

It sometimes seems that Korsgaard thinks that there is a kind of basing relation that exists between practical identities. The most obvious instance of this is the relation in which particular practical identities stand to our identity as a reflective being. For instance, Korsgaard says, "...our identity as moral beings- as people who value themselves as human beings- stands behind our more particular practical identities."¹⁷ Perhaps we can extend this idea to other practical identities as well. For example, it seems that my identity as an American stands behind my identity as a Washingtonian. If my identity as an American ceased to be normative for me, so would my identity as a Washingtonian. But if my identity as a Washingtonian ceased to be normative, it is not necessary that my identity as an American would cease to be normative for me as well. Maybe I became a Virginian.

Korsgaard thinks that the reasons that arise from the moral identity are more powerful reasons than reasons that arise from our contingent practical identities. This is presumably why we can say that an assassin who valued themselves as a fearless killer should not murder people, despite the importance

Korsgaard means when she says, "There is an appeal to consistency in this argument; it is meant to remind you of what the value of humanity requires. But it is not what makes you take my reasons into account..." (143) In addition, if this is to be our rule for judgment in all circumstances, it is gravely lacking (what if I am a masochist after all).

¹⁷ Korsgaard 121. There are many references to this kind of idea on 121 and 122.

of murder to their practical identity as an assassin. Perhaps similarly, all reasons originating from a more basic practical identity are better than reasons originating from less basic ones. So, on this hypothesis, the reasons I have as an American are better than the reasons that I have as a Washingtonian. Presumably, this hierarchy could be extended in both directions. At bottom, the hierarchy must arrive at the moral identity, upon which all of our particular practical identities are based.

There are a couple of problems with this as an answer to what reasons are better in the context of moral issues. First, hierarchies of particular practical identities need not match between people. This is because not everyone has the same practical identities. For instance, you may consider yourself a stamp collector while I do not consider myself a stamp collector. When our practical identities diverge like this, there does not seem to be a determinate answer about what reasons are better. Moreover, the contingent practical identities that are in fact shared need not stand in the same place in the hierarchy for both people. Perhaps you consider yourself a Virginian and a Richmonder (in that order) but I instead consider myself a Virginian merely in light of the fact that I am a Richmonder. If these hierarchies of practical identity do not match, then it seems again that there is no determinate answer about what reasons are better than others. After all, the reasons I have from being a Richmonder are hierarchically on a par with the reasons you generate from being a Virginian. So, which is the better reason? It would seem that we have no way of determining if place in the hierarchy is solely what matters. But being able to say that some actions are wrong requires being able to come to a settled judgment about which reasons are better. Thus, reasons being of indeterminate strength is a critical weakness for justifying morality.

Even if the first problem could be surmounted, we are faced with a problem: what about practical identities that are themselves wrong or if people have improper orderings of practical identities? If we think there are practical identities or hierarchies that fall into these categories, then we

cannot appeal to the place of a practical identity in the hierarchy to explain the force of reasons. Consider if I had a very fundamental practical identity of being a sadist. Say that this practical identity rests only on my practical identity as a reflective being and that numerous other practical identities depend on it. If the quality of reasons is determined by their place within the hierarchy, then the reasons that emerge from my being a sadist are quite good reasons. But this is highly counterintuitive. It might imply that my identity as a sadist would provide reasons strong enough to overcome pretty much any reason you could throw at it. This would mean that I could do pretty much anything to you in the name of my sadism. But this obviously does not conform with our ideas about morality. We think that, despite the importance to me of causing suffering, I am not justified in inflicting suffering on you. Thus, I reject attempting to found the strength of reasons on the position of their associated practical identity within a hierarchy.¹⁸

The second answer to the question of what makes a better reason better is that the reasons for doing something are better than those against it when they are endorsed by the agent and the object at a suitable level of generality. To start to see how this answer is plausible, consider the importance of consent. I take it that we think that orthodontists, when they act with their patient's consent, act permissibly. But the very same actions, done without consent, would be wrong. A mad orthodontist who ran around applying braces to people's teeth would be a strange villain, but a villain, nonetheless.

What is it about consent that makes it morally relevant? When we consent, we endorse the action. To carry on with our example, when you consent to have me put braces on your teeth, you endorse me putting braces on your teeth. Why do you consent to having braces put on your teeth? Because you appreciate the reasons to do so. You agree that your condition will be improved in thus and such a way by the operation.

¹⁸ At times I am tempted to think that there is another answer, one that revolves around the degree to which particular practical identities are essential to the person. This answer, however, would be vulnerable to the same objections and, as it less obviously relates to Korsgaard's position in *Sources*, I feel no need to explore it.

But it is not just improving your own condition that leads you to consent to things. You might, consent, for instance, to donating some of your blood, some marrow, or perhaps a kidney, to someone who needed a transfusion or transplant. You would do this if you could be brought to see how the other person's condition would be improved if you consented to the operation. In this kind of case as well, you consent because you appreciate the reasons for the action to be done, even though you will suffer from the action.

But what if you did not endorse a particular action, even though you endorse that kind of action in general? One way to start thinking about this is to ask ourselves what we would think of a person who told us over and over again that people should be legally required to give blood when a disaster requires additional supplies yet objected strenuously when, after a disaster, they were so required to donate blood. Well, we would think them a hypocrite.

It seems to me that, if we disagree and think that people should not be legally required to give blood even in extraordinary circumstances, then we would merely think that the hypocrite had shown their beliefs to be hollow and themselves to be lacking integrity. But if we agreed that people should be required to give blood, we would think that they had failed, and perhaps that they were cowards.

What this seems to show is that we require people to stick to their principles in adjudicating cases that involve them. When their principles license a certain kind of action (when they endorse a certain kind of action at an appropriate level of abstraction), we expect them to follow those principles even when they themselves are on the receiving end.^{19, 20}

¹⁹ I think this is also an upshot of Korsgaard's argument that we must endorse in a principled way or risk losing a grip on our self as an agent.

²⁰ Perhaps also, this is what I proposed was a heuristic was supposed to capture. The "how would you like it if I did it to you," could be interpreted as trying to bring a malefactor's attention to the ways in which their actions differ from their principles. But note that this way of interpreting the passage is strikingly different from the interpretation I gave it before when I introduced it as a heuristic. There, this called on the malefactor to broaden the perspective about reasons, but to continue to rely on their perception of the quality of reasons.

With this analysis in tow, let us turn back to the case of self-defense. Say that you see me and I see you and you know that I will kill you unless you kill me first. Here, you propose to act on me in a way that presumably does not appeal to me. But say that I am an ordinarily self-interested being. If I am self-interested to an ordinary degree, then I endorse self-defense as a legitimate reason for action. If someone were to attack me as I am about to attack you, I would probably think it an appropriate time to defend myself. That is, in a generalized or archetypical case, I endorse self-defense. From here, there are two possibilities. Either I endorse this particular instance of (relevantly similar) self-defense or I do not. If I do endorse this specific instance of self-defense, then I have no objection to make. You are perfectly free to kill me in self-defense. Note how odd it would be for me, the assailant, to berate you for breaking the rules when you struck me down. Curse you, yes, say you acted wrongly in killing me, no. If I do not endorse this particular instance of self-defense, then I am being inconsistent. As we saw above, we think that someone should stick to their principles even then they are going to be negatively affected by the actions the principles license. So, I should still stick to my principle and can have no objection to you killing me in self-defense.

I hope it is clear how this answer deals with the problems of suppleness that beset our previous attempts at solutions to the problem of why immoral actions are wrong. We can say that reasons are stronger when they are endorsed by both parties. If you hypocritically refuse to endorse an action, that reason is weaker. This means that even actions that are generally forbidden can be licensed when the person who is subjected to the action endorses the action. While this may not be enough to fully determine the strength of reasons, it seems enough for our purposes (which is to show the kind of suppleness that our moral judgments do).

Part Six: Conclusion

So, what are we to say to the farm boy who must choose between doing evil and death? With Korsgaard, we say to him that, if he is to have any reasons at all, if he is to be able to act at all, he must

not act in ways that are incompatible with seeing reflective consciousness as a source of reasons. If he were to do evil, he would have to deny that the reasons other people give him to stop are reasons at all. And if he were to do that, he would deny that their reflective consciousness could produce reasons. In doing so, he would deny that reflective consciousness is a source of reasons. This he must not do. Thus, he must not do evil. Morality is normative for him.

Bibliography

- Cholbi, Michael J. "On Hazing." *Public Affairs Quarterly* (2009): 143-160.
- Foot, Philippa. "Rationality and Goodness." O'Hear, ed. Anthony. *Modern Moral Philosophy*. Cambridge: Cambridge University Press, 2004. 1-13.
- Fuller, Lon. "The Case of the Speluncean Explorer." *Harvard Law Review* (1949): 616-645.
- Herman, Barbara. "Murder and Mayhem." Herman, Barbara. *The Practice of Moral Judgment*. Cambridge: The President and Fellows of Harvard College, 1993.
- Hieronymi, Pamela. "The Wrong Kind of Reason." *The Journal of Philosophy* (2005): 437-457.
- Kant, Immanuel, eds. and trans. Mary Gregor and Jens Timmerman. *Groundwork of the Metaphysics of Morals: Revised Edition*. Cambridge: Cambridge University Press, 2012. Print.
- Korsgaard, Christine. "The Normativity of Instrumental Reason." Korsgaard, Christine. *The Constitution of Agency*. Oxford: Oxford University Press, 2008. 27-68.
- Korsgaard, Christine, et al. *The Sources of Normativity*. Cambridge: Cambridge University Press, 1996. Print.
- Luban, David. "Natural Law as Professional Ethics: A Reading of Fuller." *Social Philosophy and Policy* (2000): 176-205.
- Rawls, John. "Lecture VIII." Rawls, John. *Political Liberalism*. Chichester: Columbia University Press, 2005. 289-371.