

A Defense of Common-Sense Deontology

Jacob Gordon

Northwestern University

BIOGRAPHY

I am a third-year undergraduate at Northwestern University, where I study Philosophy, Critical Theory and Political Science. I am currently pursuing an honors thesis on Nietzsche's Critique of Morality. After graduation, I plan on continuing my interest in moral and political philosophy through either a PhD program or a JD program.

ACKNOWLEDGMENTS

I would like to thank Stephen White, an Assistant Professor of Philosophy at Northwestern, for advising this project. I would also like to thank Northwestern's Weinberg College of Arts and Sciences, whose Summer Research Grant funded my research.

ABSTRACT

In this essay, I attempt to defend common-sense deontological ethics from within a value maximizing framework that is often associated with consequentialism. I begin by exploring and defending what I take to be the two components of what is often called consequentialism's "compelling idea": first, a maximizing conception of rationality, and second, a commitment to the priority of value. After defending these two aspects of the "compelling idea", I argue that the proper conception of "value" is plural and agent-centered. Then, I use the value of "respectedness" to show that deontology can be a rational response to value, if value is conceived in a sufficiently complex way. My ultimate goal in this essay is to show that we can retain both the intuitions of deontology and the sensibility of the compelling idea, by understanding deontology as a rational guide to practical reasoning that arises from our basic commitments.

KEYWORDS

Consequentialism, Deontology, Value

INTRODUCTION

Deontological considerations have a substantial role in common-sense moral reasoning. Imagine that a woman named Jill watches a broken train barrel towards five people who are stuck on the track. Jill can push Tom, a stranger, in front of the train, and doing so will stop the train before it can reach the five. Jill will probably at least hesitate to push Tom. Even if she decides to push him, she will almost certainly be reluctant to do so. Jill, like most of us, thinks her actions should be restrained in some way, dependent not only on the simple calculation of maximizing life, pleasure, or any other apparent metric.

While many of us share Jill's deontological intuitions, these intuitions are notoriously difficult to justify. Critics argue that, if Jill's refusal to push Tom comes from an aversion to death, such an aversion should actually compel Jill to push Tom, because doing so will minimize the number of deaths Jill is able to produce. These critics contrast deontological commitments to a basic conception of rationality at the heart of consequentialism - a part of consequentialism's "compelling idea" - which states that, in Samuel Scheffler's words, "if one accepts the desirability of a certain goal being achieved, and if one has a choice between two options, one of which is certain to accomplish the goal better than the other, then it is, *ceteris paribus*, rational to choose the former over the latter" (Scheffler 1982, 414). If Jill accepts the goal of preventing death, she should simply produce fewer deaths, all else equal. Contrasting with this compelling idea, Jill's intuitive deontology irrationally "present[s] as desirable [a] non-relative goal whose maximum accomplishment it then prohibits" (Scheffler 1982, 416).

The first section of this essay will explore and defend the two parts of consequentialism's compelling idea: first, its maximizing conception of rationality, and second, its commitment to the priority of value. After defending the compelling idea, I will argue that the "value" we adopt should be plural and agent-centered. Then, I will show why deontology is a rational response to such plural and agent-centered value. Ultimately, this argument will show that we can and should retain both the intuitions of deontology and the sensibility of the compelling idea, by understanding deontology as a rational guide to practical reasoning that arises from our basic commitments. Finally, I will reflect briefly on how my argument relates to larger questions in moral philosophy.

1 - THE COMPELLING IDEA

The first principle at the heart of consequentialism's "Compelling Idea" is a maximizing conception of rationality. As above, Scheffler's conception of the compelling idea holds that "if one accepts the desirability of a certain goal being achieved, and if one has a choice between two options, one of which is certain to accomplish the goal better than the other, then it is, *ceteris paribus*, rational to choose the former over the latter" (Scheffler 1982, 414). For a moral theory to accord with this basic rationality, it must recognize that the better a certain action fulfills its stated goal, the more preferable it is.

This maximizing conception of rationality is incredibly intuitive. Making this point, Douglas Portmore has presented the obvious implausibility of a "deontological egoism", which forbids an agent from performing a single self-sacrifice even if doing so will minimize her total self-sacrifice. Insofar as an agent's goal is to minimize her self-sacrifice, such a deontological egoism, by forbidding the fulfillment of its stated goal, "seems paradoxical" (Portmore 2006, 14). Indeed, the maximizing conception of rationality merely states that an action A is preferable to an action B, from the perspective of a stated goal X, insofar as it better fulfills the stated goal X. A theory denies this only if it presents a goal X, admits that an action A will better fulfill X than action B, and still insists that B is preferable to A, without referring to any motivation besides X. The inconsistency of such a theory is undeniable.

Still, Scheffler's "*ceteris paribus*" clause opens the door to criticism from deontologists. "What if things are not equal, and there is a rule that simply forbids murder?", one could ask. There is nothing irrational, in Scheffler's sense, about such a rule. If taken as primary to our other goals, this rule could simply sway the prescribed action of our moral theory. In Jill's case, a primary and unconditional prohibition on murder would adequately forbid pushing Tom. Yet such a rule conflicts with the second prong of consequentialism's compelling idea: that a rule may only be justified if it advances a primary 'value'. While Scheffler's formulation of the compelling idea does not make this commitment clear, Jennie Louise's does. In her words, "Consequentialism, broadly construed, says only that agents should produce as much value as possible" (Louise 2014, 520). Consequentialism insists not only that one must be consistent in prescribing actions as to fulfill basic goals (as in Scheffler), but also that every rule or goal should be measured by its maximization of value.

compos mentis

This second prong of the compelling idea is, like the first, hard to deny. How might a deontologist who denies the primacy of value defend the rules she endorses? She cannot, given her denial of the primacy of value, argue that her rules are superior to others because they better achieve any given thing, because that something which her theory achieves would then represent a primary value. Further, she cannot even argue that there is value in her rules themselves being upheld, because doing so would admit that there is some value that morality aims to bring about, just that the following of rules constitutes that value. Admitting even this would lapse the deontologist into accepting the compelling idea. A theory's 'value' provides the reason for its prescriptions, so denying the primacy of value seems to entail denying that prescriptions have a reason. A theory admits the primacy of value insofar as its rules are rooted in their ability to further whatever that theory regards as basically important. It is hard to imagine a tenable normative theory that fails to root its prescriptions in value, construed as such.

To this point, I have defended the compelling idea's maximizing conception of rationality, as well as its insistence that value underlies moral rules. It might seem that this compelling idea does not dispel some classically "non-consequentialist" theories. Indeed, the literature examining which theories can and cannot be 'consequentialized', or made consistent with the compelling idea, is growing.¹ My interest here, rather than to directly contribute to these higher-level distinctions, is to illustrate that certain rules - our common-sense deontological ones specifically - are not only consistent with the compelling idea, but are its reasonable offspring, if value is properly conceived.

2 - VALUE

2.1 - Potential Plurality and Agent-Centeredness

The two principles of the compelling idea, while illustrating that the moral status of an action should be a function of that action's promotion of value, say little about the nature of value. First, the compelling idea does not exclude the possibility of value plurality. As written, the compelling idea's rationality condition

1. See Doug Portmore's "Consequentializing Moral Theories", Jennie Louise's "Relativity of Value and the Consequentialist Umbrella", James Dreier's "Structures of Normative Theories", Mark Schroeder's "Teleology, Agent-Relative Value, and 'Good'", and Campbell Brown's "Consequentialize This", among others.

insists that it is irrational to hold that “an action A is preferable to an action B, though B better fulfills our only value, X, than does A.” However, it is perfectly rational to hold that, “though B better fulfills X than A, A is preferable to B, because there is a value besides X, namely Y, which A better fulfills than B, and which takes priority over X.” In other words, there being multiple intrinsic sources of value does not contradict the compelling idea, which only insists that we be consistent within our preference ordering, given our value. Consequentialism, as Louise writes, “does not say anything about what is to be regarded as valuable” (Louise 2014, 520).

Second, the compelling idea does not exclude the possibility of agent-centered value. A moral theory need not demand that all individuals maximize the same things. A simple example of a moral theory that utilizes agent-centered value, while abiding by the compelling idea’s conditions, is ethical egoism. Egoism demands that every individual maximize only her own “good”, and therefore formulates value in an entirely agent-centered way. Despite that, it still accords with both the rationality, and value-first reasoning, demanded by the compelling idea (Scheffler 1982, 416).

The fact that value could be plural and agent-centered does not mean that it actually is, but it does mean that the notorious spell-binding nature of the compelling idea does not eliminate theories of value that are plural and agent-centered. This point is significant because unity and agent-neutrality have often been sold as necessary accessories to the compelling idea. Scheffler begins his *Consequentialism and Its Critics* by selling this bundle. While presenting the compelling idea, Scheffler calls consequentialism “impersonal” (Scheffler 1988, 1), and insists that consequentialist theories “all share” the insistence that we “ought to...minimize evil and maximize good” (Scheffler 1988, 1). As shown, though, the conditions of agent-neutrality and of a unified “good” are not entailed by consequentialism’s compelling idea.² While I have not yet argued that agent-centered pluralist theories are preferable to their opponents, I have shown that they cannot be dismissed as non-consequentialist insofar as that label implies an inconsistency with the compelling idea. Answering yes to “Do I maximize value?” does not require one to answer a certain way when asked “Where does value lie?”

2. I have no problem with defining “consequentialism” to include only moral theories that offer an agent-neutral and unified theory of value, so long as we differentiate this condition from consequentialism’s “compelling idea”, which makes no such demand.

2.2 - Agent-Neutrality, Unity, and Alienation

The strengths of agent-centered and pluralist theories of value are best drawn by the weaknesses of agent-neutral or unified theories of value. Consider Peter Railton's "Alienation, Consequentialism, and the Demands of Morality", which begins by describing a husband, John, who treats his wife Anne as a perfect husband would, but who explains his behavior as follows: "I've always thought that people should help each other when they're in a specially good position to do so. I know Anne better than anyone else does, so I know better what she wants and needs...Just think how awful marriage would be, or life itself, if people didn't take special care of the ones they love" (Railton 1984, 135).

John's reasoning seems twisted and inhuman. As Railton puts it, "that he devotes himself to her because of the characteristically good consequences of doing so seems to leave her, and their relationship as such, too far out of the picture" (Railton 1984, 135). While John commits himself to a kind of rule-utilitarian practical reasoning, which aims at acting as will tend to make life better, his problem would remain if he were committed to a more directly utilitarian reasoning, through which he argued that "I should care for Anne, because doing so improves the world."³ John strays more basically, in actually valuing the wrong things. Instead of valuing Anne, on the most primary level, John values the entire world. John arrives at the conclusion to care for Anne second, only after calculating that he may care for the world best by caring for his wife. John is in this way "alienated" from those he should love. This kind of alienation is inevitable if we locate value in any agent-neutral place. If the thing I aim to bring about in moral action would always bind another agent as well as it binds me, then my actual relationships cannot serve as my primary motivations. This is a fundamentally inhuman way to value others.

Theories of unified value alienate similarly to those of agent-neutral value. Consider how Railton's John might respond when asked why he cares for his daughter. To satisfy us, his response would need to be about his daughter specifically. Alternatively, if John sees his wife and his daughter as different sources of the same value (as is necessary on a unified theory of value), he will be

3. Railton's essay ultimately argues that the alienation problem does not doom utilitarianism. I think his argument fails, but exploring that is beyond the scope of this essay, which is not capable of responding to every response to the alienation problem. If readers wish to explore defenses of utilitarianism against the charges I present here, they can read the second half of Railton's essay. I should note that I do not know why some feel so compelled to save utilitarianism in the first place.

alienated from each, only caring for each as a factor of her being instrumental for that value. To genuinely care for his wife and his daughter, John should love each irreducibly. Insofar as we are unsatisfied with any theory that fails to see John's wife and daughter as independent and irreducible sources of value for John, we will be unsatisfied with any theory that places value in a single, or agent-neutral, place.

2.3 - A More Faithful Alternative - Agent-Centered Plurality

Fortunately, we may view value in a different and more elegant way, as arising from our individual relationships. If we believe that John should care for his wife simply because he loves his wife, and care for a stranger simply because he values that stranger, then we believe value ought to be agent-centered and plural in nature. This admission allows us to easily avoid some of consequentialism's famous problems. If I need to choose between saving my child and saving two other children, I can simply save my child. Doing so is justified because the value of my child, to me, is irreducible to my caring about the world, and is simply greater (to me) than the values of the two other children. We retain this ability under an agent-centered and pluralist theory of value without having trouble explaining general benevolence. We ought to care for those who suffer simply because we value those persons.

This picture is more faithful to our genuine moral reasoning than any agent-neutral or unified theory of value can be. We are more inclined to answer, "why do you care about the world?" with "because I care about the individuals in it", than to answer, "why do you care about individuals?" with "because I care about the world." If this is correct, the foundation of value is in our individual relationships, and is therefore agent-centered, arising from our feelings of commitment, and plural, located in each individual, rather than in some mass of aggregate "moral stuff".

The plurality of value goes further even than that it is located in separate individuals; it seems located in different feelings and attributes of different individuals. A good friend wants to sponsor both her friend's autonomy and her joy. These two values are distinct, but each independently draws the good friend. It would be wrong to reduce autonomy to joy, or joy to autonomy.⁴ What we

4. The plurality of value, as I have described it, delivers a response to Mark Schroeder's allegation that "Agent-Relative Teleology" relies on the incoherent concept of "good-relative-to", which

are left with on this picture is appealing and simple - we ought to value what we actually value.

2.4 - Defending a Faithful Conception of Value

This conclusion might seem unsettling for a few reasons. One could ask, "Isn't the purpose of moral theory to tell us what we "should" value, rather than what we "do" value?" But how would we be convinced that we "should" value something new? Such an argument needs to be based on a set of shared premises, which we already do accept. An argument calling on us to change our understood sources of value is really an argument that we are already committed to accepting different sources of value. The goal of moral theory is to show us, at root, what we already are committed to, and to direct our moral beliefs and actions towards those commitments.

I have intended to show, largely through Railton's example of John the husband, that our primary sources of value are agent-centered and plural, meaning that no agent-neutral or unified theory of value can connect with our most basic commitments. We are individual relationships "all the way down." This is the foundational point of agent-centered and pluralist theories of value. These theories, aside from their more intuitive prescriptions which I will outline below, undoubtably accord better with our humanity.

It could still be argued that the picture of value I am drawing is too permissive. Is it true that all morality requires is for an agent act in accordance with her own value? If this is the case, does nobody ever act wrongly? I have two responses to this critique. First, my argument is not committed to accepting that morality is subjective in the alleged way. Even if we do accept a kind of value-objectivity, we should accept that value lies where it seems to given our most basic commitments, which is, as shown, in agent-centered and plural places.

Second, even if we do accept value-subjectivity, the concern about over-permission applies to far fewer cases than it may seem to, because an agent cannot simply choose whether an action accords with her deepest values. A son

is allegedly less clear than the concepts of "the good" or "good for" (Schroeder 2007). Even if Schroeder's critique is sound (I am inclined to think that all three of those concepts are incoherent), it presents no problem for the agent-relative pluralist theory of value that I have described. Perhaps we cannot understand what it means for an outcome to be "good-relative-to" a given agent, but we can understand what it means for an outcome to include more of an agent's friend's pleasure, or an agent's son's autonomy.

can tell his father that he was right not to bring a jacket to the park because he does not feel cold, but his insistence will have no effect on whether he really feels cold. Subjective value would act in much the same way. When we act against our values, nothing we say will make it otherwise. What this means is that the subjective agent-centering of value is not as terrifying a prospect as it might seem. "Coldness" is subjective, but we understand, for the most part, what causes others to feel cold. As we can reliably conclude when another person "should have brought a jacket", we can conclude then they should have acted differently. For example, if Jesse is rude to a stranger, even the subjectivist can criticize her rudeness, not because Jesse's rudeness fails to align with the critic's values, but because Jesse's rudeness probably fails to align with Jesse's own values. If Jesse looked into the eyes of the stranger, and heard about his life, she would surely regret being rude. To be clear, my theory is still entirely consistent with an objectivist view, but in case the agent-centering of value seems a slippery slope towards the 'subjectivizing' of value, I am not sure the terrain of that slope is so rough.

Lastly, it is worth noting that an agent-centered and pluralist theory of value is much better at responding to general moral skepticism than is an agent-neutral or unified one. A person with strong moral intuitions can always feel uncompeled to a utilitarianism, which calls her to accept a foreign theory of value. If she asks that pesky question of "why?", an agent-neutral unified theory often has no clear response. She cannot similarly ask "why?" to a theory of value whose structure is determined by the entailments of her most primary commitments.

In this section, I have attempted to illustrate that a plural and agent-centered theory of value is consistent with the compelling idea (in 2.1), that such a theory of value is faithful to our actual values (in 2.2 and 2.3), and that we are right to adopt, in moral reasoning, the theory of value that is faithful to our actual values (in 2.4). From this point, I will explain how an agent-centered and plural theory of value can reasonably lead to the deontological conclusions we intuitively hold.

3 - DEONTOLOGY

3.1 - Setting the Stage: Why Not Deontology?

I will begin this section by illustrating the alleged irrationality of deontology. Then, I will describe how the value of subjects' being respected, while being a

reasonable value for a pluralist agent-centered theory to adopt, also leads to deontological conclusions.⁵

Deontology appears irrational chiefly because of the difficulty it faces in responding to the “lesser evil” problem. If we accept that a certain value ought to be promoted, it seems irrational to establish rules relating to that value that neglect the net effect agents can have on that value’s promotion. If we, for example, agree that people have a right not to be slaves, it seems that an agent’s primary drive should be to minimize the total number of slaves, rather than to avoid using slaves herself. And if one is forced to use a slave to minimize the total number of slaves, it seems that she should use that slave. Deontology, by insisting that an agent never use a slave, even to minimize the number of slaves, allegedly endorses a worse outcome and forbids a better one. It therefore “present[s] as desirable [a] non-relative goal whose maximum accomplishment it then prohibits” (Scheffler 1982, 416), and fails to abide by the compelling idea, as above.

The proper analysis of value, however, shows that this critique over-simplifies the situations that draw on our deontological intuitions. Consider Jill, from the beginning of this essay. She must choose whether or not to push Tom in front of the runaway train. Contrary to the picture illustrated by the critic above, pushing Tom does not simply promote more of the value of “life” or “pleasure” than not pushing Tom, because, as established, each individual Jill is in a position to affect is a separate source of value, irreducible to the others she is able to affect. It would be different if, for example, Jill saw an unsuspecting Tom stuck in the middle of a road, with a bus coming toward him. There, since Tom’s pleasure, or life, is a discrete and single source of value that can be maximized, Jill would undoubtedly be right in acting as the “utilitarian” would towards him. She should push him out of the way of the bus, perhaps hurting him in the process, because doing so would maximize his net pleasure, or life. But, again, the case Jill faces to begin this essay is quite different, because Jill must choose what to do, given the plurality of values to which she is drawn. This point alone does not justify deontology, but it is significant to note, in order to make our depiction of “the lesser evil” problem more genuine. Even in cases that do not call on our deontology, such as one in which we must choose whether to save one drowning person or two other

5. This method is preferable to one that lays out a more complete list of the values that a proper theory would recognize for two reasons. First, I cannot claim to know what this complete list would look like. Second, drawing up such a list would likely spark a substantial amount of disagreement that is irrelevant for the purpose of justifying deontology.

drowning people, we always deal with discrete and independent persons, who should be recognized as such.⁶

3.2 - Deontology Through "Respectedness"

The more significant way in which critiques of deontology fail, though, is by under-recognizing subtle sources of value, particularly the value of subjects' being respected. Our intuitive deontology, only appearing to endorse worse outcomes, actually endorses better ones, given its acknowledgment of the value that lies in subjects' being respected, or in their "respectedness" (for lack of a better term). If respectedness is a credible source of value, and if its maximization leads to intuitive deontological conclusions, then those intuitive deontological conclusions, far from being irrational, are particularly compelling.

We ought to value respectedness because humans have real personal bonds to each other, which are inspired by mutual compassion for and understanding of others' experiences. Disrespect represents the harm done by the very breaching of one of these bonds, which is a harm that exists in addition to whatever harm is otherwise done to lead to that breach. Insofar as we value our very compassion for one another, we ought to accept that when human bonds are broken, real harm follows. If a man's arm is broken in a natural disaster, he may experience intense physical pain, but if that man's arm is broken willingly by another person, he will also face an accompanying emotional pain. Such a personal attack will, in a real way, "add insult to injury." Because its knowing intent causes an additional harm to the pain otherwise caused, human-driven suffering is generally worse than naturally-driven suffering.

Importantly, this value of respectedness magnifies harm more than it can magnifies 'help.' While a person helping another person might strengthen the human bond between them, and thereby represent an added benefit over and above the benefit otherwise given, that added benefit does not impact the subject

6. While some argue that we are not so clearly compelled to save, for example, two strangers rather than one (See John Taurek's "Should the Numbers Count?"), I will assume here that the solution to this moral problem is as straightforward as most suppose it is. It might seem that this de-commits me from the true plurality of value - that I have chosen in this assumption to treat three strangers' lives as simply countable. I reject this characterization, because I do not deny the real moral remainder in failing to save the one, even in saving the two. I acknowledge that hard questions sometimes have answers, but that does not put me in the same boat as those who deny that there are hard questions.

nearly as much as would that bond's breaking. Since the human experience deeply requires independence, we care more about not being unnecessarily hurt by others than we do about being unnecessarily helped by them. That our mutual bonds are composed chiefly of our not harming each other is exemplified by the fact that we are angrier when pushed to the pavement than we are relieved when helped off of it. Since each of us largely wants to be left alone by strangers, our being respected depends on our not being harmed by others much more than it depends on our being helped by them. Agents therefore respect others chiefly by not harming them.

What does maximizing respectedness look like in Jill's case? If we assume that Jill cares equally for all six people she may affect (between Tom and the five on the track), and if we accept the compelling idea, then whether pushing Tom is right will depend on whether Jill will bring about more value by pushing him or by not pushing him. The value in these six individuals, though, lies partially in their being respected. Jill can only maximize Tom's respectedness, as above defined, by refusing to push him in front of the train, since his death on her hands would break her human bond with him. Interestingly, though, Jill need not push Tom in front of the train to respect the five on the track. Since these five have found themselves in the path of a runaway train by accident, their deaths will not involve a lack of respect, because no person will have put them in a position to die, which would involve a unique harm over and above that of dying. Further, it is not as though Jill can greatly increase the five's respectedness by saving them, since their being respected is dependent, as above, chiefly on their not being harmed by another person, rather than on their being helped by another person. If Jill refuses to push Tom, she will maximize the net respectedness of the subjects she can impact. Thus, respectedness provides Jill with a perfectly justifiable value whose maximization is only achievable if she refuses to push Tom. Through respectedness, Jill can retain her intuitive deontological commitments, without betraying the compelling idea's maximizing conception of rationality.

Some might be unwilling to accept that Jill should worry about maximizing respectedness in a decision whose consequences involve multiple deaths. First, it should be noted that Tom's respectedness may not outweigh the lives of the five; perhaps Jill would be right to push Tom. My analysis of respectedness simply shows that our deontological intuition is right to recognize that something muddies the waters of Jill's choice. We are right that Jill's case is different than one in which she

must either save one drowning person or five drowning people. Second, I think respectedness becomes a more powerful value the more deeply it is considered. We do frequently judge situations that involve disrespect as worse than similar ones that lack disrespect. News of a person's being murdered is worse than news of their being killed by a fallen tree (even if the fallen tree may cause more physical pain than the murder), precisely because there is disvalue in someone's being disrespected. As so much of our lives are built on mutual trust and respect, the value of maintaining that respect should not surprise us. It may seem an odd value to systematize, as I have, simply because it is subtler than something like pleasure or pain. But the subtle importance of being respected may still rival the more obvious importance of avoiding physical pain. I write this to suggest that Jill's refusal to push Tom, on the grounds of maximizing his respectedness, really can be a perfectly reasonable choice. She could say to the five on the track "I am deeply sorry that nature has given you this fate, but the damage done when a person ends another person's life is so great that it should not be done."

Furthermore, the non-parallel relationship Jill maintains with Tom's respectedness relative to the respectedness of the five on the track would be partially retained if the runaway train were directed by a murderous conductor, rather than by a natural accident. Such a murderous conductor's actions, while more intentional than those of nature, are still likely less personal than Jill's. As this conductor is driven to murder in the first place, Jill may be right to assume that he lacks some basic moral apparatus that she has, and that he therefore never maintained the moral bonds whose breaching constitutes disrespect. Being killed by this murderous conductor, almost like being harmed by a child, lacks the degree of true and sophisticated disrespect that being killed by Jill would involve. We would rather be harmed by the school bully (the conductor), who lacks the basic capacity for compassion and understanding, than by another victim of the bully (Jill), who understands the harm caused by the bully, and who we thought we could trust. We can imagine that, in the path of the mad conductor, the individuals on the track might ask "why has fate befallen us this way?" as they would if struck by a natural disaster, rather than "why have you, another person, done this? Don't you care about me?", as Tom would ask Jill if she were to kill him. This means that refusing to murder to prevent more murders, in addition to preventing more deaths, can be justifiable on the grounds of respectedness.

It could be argued that this concept of “respectedness” is circular. Specifically, I rely on the fact that Tom would be disrespected if killed by Jill to argue that Jill would be wrong to kill Tom, even though Tom would perhaps only be disrespected in being killed by Jill if Jill would be wrong to kill him in the first place. I do not think this critique is correct, because Jill’s personal bond with Tom exists independently of her bonds with the other five. As such, if Jill kills Tom, a real bond will be broken, and he will be disrespected, regardless of whether Jill saves the five. The fact that Tom will be disrespected if killed comes first, as an input to be weighed against the lives of the five, rather than second, as dependent on how compelled Jill is to save the five.

For the reasons above, it seems that maximizing respectedness can ground many of our deontological intuitions. I will not argue that our deontological intuitions can always be justified in this way, but it seems that they often can be. While Jill’s case deals with the question of murder, it could just as well deal with questions of torture, stealing, or lying. Our common-sense deontology, I think, never actually forbids the lesser evil. It only prescribes rational action based on its recognition of value plurality, and of respectedness’ importance.⁷

3.3 - Evaluating Agent-Centered Justifications of Deontology

I have attempted to justify deontological rules by citing plural values, rather than the agent-centeredness of decision making. Agent-centered justifications of deontology are possible, and frequently given, but they are often either circular or selfish. The circular argument for deontology supposes that we have a good reason to remain “morally pure” (meaning that we should, for example, refrain from killing one to save five), without first showing that refusing to kill is the “purer” choice. Certainly, a utilitarian would argue that “moral purity” is maintained by killing the one, because doing so is the right thing to do, and because moral purity should entail doing the right thing. The deontologist who relies on “moral purity” often presupposes that the rule she is defending is right in her argument that it is right.

7. It is possible that there is some value that overlaps with or exists separately from respectedness, which can do the same work for deontology that respectedness does. While I think respectedness alone can justify deontology, I do not think it must be alone in doing this. As noted in footnote 5, I welcome the possibility of new values, and do not claim to have a grasp on all the subtle values that would belong in a comprehensive moral theory.

Deontologists may make a subtler argument for moral purity, which defends refraining from causing harm, rather than refraining from doing 'wrong'. While it is possible that killing one to save five is not wrong, it certainly causes harm. And so, if refraining from causing harm is justified, a blanket prohibition on killing might be too. The problem with this position is that any argument for refraining from causing harm, that does not devolve into an argument for refraining from doing wrong, is selfish. Rather than being primarily about others, any such argument is primarily about ourselves, and the actions we would like to take. If we cared primarily about others, we would commit ourselves simply to do the right thing, even if the right thing caused some harm. For example, if I were forced to amputate someone's finger to save her arm, and I refused to amputate the finger because "I wished not to cause harm", my refusal would clearly be sourced in my own squeamishness, rather than in my care for that person. Perhaps these personal aversions can maintain marginal weight in our decision-making (that one might receive serious psychological trauma by killing someone is significant), but if these selfish considerations were the foundation of our deontology, that deontology would be indefensible.

This point detracts neither from my argument that value ought to be agent-centered, nor my argument that human-driven harm is generally worse than natural similar natural harm. Those two considerations affect our evaluation of consequences and impact our consideration of which actions are right. Meanwhile, justifications of deontology that cite the agent-centeredness of decision-making go one step further. They argue that, even once we have figured out what is right, we can refuse to act rightly, simply because of the actions we wish to take.

Alternatively, the justification of deontology that I have provided remains consistently about others. Besides being generally more defensible, this position tracks better with our intuitive deontological conclusions than does a system of deontological rules, like "do not kill", that address the actions themselves that an actor is allowed take. To show this, imagine that Tom's entire family is stuck to the train track, and that Tom asks Jill to push him in front of the train, which will still kill him, so that his family can live. Tom, here, can only get in the train's way with Jill's push, because he needs help getting past a guard-rail. Since Tom asks to be killed here, and since he has a good reason for his request, this case intuitively differs significantly from the case that began this essay. However, any deontology concerned only with the actions Jill can take will be unable to explain the clear

difference between pushing Tom here and pushing Tom in the original case, since the difference between the two scenarios lies in the subjects of Jill's acts, rather than in the acts themselves. Meanwhile, since my theory sources Jill's deontology in Tom's being respected, and since respecting Tom requires considering his requests, my theory is able to explain the difference between this case and the initial one. By focusing on subjects rather than actions themselves, my theory tracks with our intuitive deontology in a way agent-centered explanations of deontology do not.

3.4 - Addressing "Gimmicky-ness"

My goal in this essay has been to justify our intuitive deontology, by showing that it arises from a reasonable understanding of value. Some critics argue that justifications of deontology like mine, which attempt to demonstrate that deontology is consistent with the compelling idea, fail to do justice to the deontological commitments they claim to defend. Robert Nozick, for example, once accused these justifications of being "gimmicky" (Nozick 1968, 4). The core of Nozick's charge, I think, is that these justifications of deontology use reasoning that differs substantially from the reasoning that genuinely drives our intuitive deontology. This criticism could be substantial. Insofar as my goal is to justify common-sense, it would be deceptive for me to reach the same conclusions as common-sense, but from completely different grounds. If supposed justifications of deontology do this, they do not really defend our intuitions. Luckily, I think the reasoning I have used tracks well with our intuitive reasoning. By locating the importance of deontology in the importance of independent individuals being respected, I seem to have utilized the same sort of reasoning that Robert Nozick himself used in his defense of "side constraints": "Side constraints upon action reflect the underlying Kantian principle that individuals...may not be sacrificed or used for the achieving of other ends without their consent"(Nozick 1988, 138). Using an individual 'for the achieving of other ends without their consent' seems to be a particular way of disrespecting them, such that the image of value I have drawn captures, in a systematized way, the Kantian intuition expressed by Nozick. Furthermore, my theory's sincerity to our intuitive reasoning is reflected by the fact that its analyses of situations track better with our intuitions than do the analyses of common agent-centered justifications of deontology, as shown above. In sum,

I understand and respect the risk of “gimmicky-ness”, but do not think that the particular picture of deontology that I have drawn is “gimmicky”.

CONCLUSION

In this essay, I have sought to show that our intuitive deontological commitments rationally arise from a commitment to agent-centered and plural value, because the subtle value of “respectedness” compels us to act as our deontological intuitions also drive us to act. While my arguments address a number of topics, such as the viability of deontology and the nature of value, my chief motivation in writing this essay has been to explore whether our moral intuitions are defensible. I have argued that that they are.

If I am right, we can understand our common-sense morality as sincerely founded on a genuine interest people have in others’ well-being. Among other things, this might indicate that moral theorists typically underestimate the complexity of the average person’s moral reasoning (this essay might be a 7,000+ word explanation of a reasoning process we all make without even thinking). If I am wrong, we will be forced to accept that our moral common-sense is fundamentally inconsistent with the things we claim to most basically value. This would be a heavy charge to accept. It would force us either to abandon our deep moral intuitions, or, if we wish to keep these intuitions, to admit that our morality is fundamentally irrational, not actually arising from our care for others, but instead from something like selfish calculation or predetermined evolutionary psychology.

Our answer to the question of whether moral common-sense is defensible, therefore, informs and reflects our answers to moral and social philosophy’s deepest questions. Regardless of whether my arguments have been compelling, this is a question that deserves to be asked.

REFERENCES

- Brown, Campbell. 2011. “Consequentialize This”. *Ethics* 121 (July 2011). 749–771.
- Dreier, James. 1993. “Structures of Normative Theories”. *The Monist* 76 (1993). 22–40.
- Louise, Jennie. 2014. “Relativity of Value and the Consequentialist Umbrella”. *The Philosophical Quarterly* 54 (October 2014). 518–536.

- Nozick, Robert. 1968. "Moral Complications and Moral Structures". *Natural Law Forum*. 1–50.
- Nozick, Robert. 1988. "Side Constraints". Printed in *Consequentialism and its Critics*, edited by Scheffler, Samuel. Oxford: Oxford University Press. 134–142.
- Portmore, Douglas. Draft from April 2006. "Consequentializing Moral Theories".
- Railton, Peter. 1984. "Alienation, Consequentialism, and the Demands of Morality". *Philosophy and Public Affairs* 13 (Spring 1984). 134–171.
- Scheffler, Samuel. 1982. *The Rejection of Consequentialism*. Oxford: Clarendon Press. revised edition.
- Scheffler, Samuel. 1988. "Introduction". *Consequentialism and its Critics*, edited by Scheffler, Samuel. Oxford: Oxford University Press. 1–14.
- Schroeder, Mark. 2007. "Teleology, Agent-Relative Value, and 'Good'". *Ethics* 117 (January 2007). 265–295.
- Taurek, John. 1977. "Should the Numbers Count". *Philosophy and Public Affairs* 6 (Summer, 1977). 293–316.