12-2020

Relativized rankings

Matthew HAMMERTON
Singapore Management University, mhammerton@smu.edu.sg

Relativized Rankings[1]

Matthew Hammerton

mhammerton@smu.edu.sg

All consequentialist theories contain an axiology—a theory of the good that ranks various outcomes as better or worse. Although consequentialism has historically been associated with certain axiologies, there are in principle no restrictions on the axiological structures a consequentialist theory can contain. This chapter is concerned with one important distinction in axiological structure—the distinction between position-neutral and position-relative axiologies. When a state of affairs is evaluated, there are various connections that may hold between the position of the evaluator and the state evaluated. For example, the evaluator may be an agent who acts in the state she is evaluating. Or, she may have the same temporal location as the state she is evaluating. In a *position-neutral axiology* these kinds of connections are irrelevant to the correct ranking of states. In a *position-relative axiology* these kinds of connections are sometimes relevant to the correct ranking. In other words, the rankings are indexed to certain positional facts such that a certain ranking may be correct relative to one position but not relative to another. Different kinds of position-relative rankings are available depending on what positional facts are relevant. *Agent-relative* rankings make certain connections between the numerical identity of the evaluator and the state evaluated relevant. *Time-relative* rankings make certain connections between the temporal location of the evaluation and the temporal location of the evaluated state relevant. Further positional facts related to spatial location, possible world location, and patient identity may also result in relativized rankings and are discussed in this chapter.

It is notable that position-relative axiologies have been neglected for much of the history of consequentialism. The classical utilitarians did not recognize the possibility of position-relative value and assumed by default that all value is position-neutral. Subsequent generations of consequentialists followed them in making this assumption and it was not until Amartya Sen published "Rights and Agency" in 1982 that a position-

[1] This is the penultimate draft of a chapter that appears in *The Oxford Handbook of Consequentialism*. New York, USA: Oxford University Press. pp. 46-66 (2020). Please cite the published version.

relative form of consequentialism was first proposed.[2] By contrast, several other significant axiological assumptions made by the classical utilitarians were challenged much earlier in the history of consequentialism. For example, the classical utilitarians adopted an axiology that is monadic (there is only one intrinsic good), welfarist (all goods are good for someone), and atomistic (the good of a whole is a sum of the good of its parts). Yet by 1903 a form of consequentialism that abandons each of these axiological assumptions had already been put forward by G.E. Moore in his *Principia Ethica*. The delayed recognition of position-relative consequentialism may be explained by the fact that, more so than any other axiological structure, position-relative rankings appear to lack a precedent in everyday moral talk and thought. Some have used this to argue that position-relative values are implausible, while others claim that they are unprecedented. These issues are discussed in §3.1.

There are many reasons why consequentialists are interested in position-relative rankings. However, one reason stands out above all others. It has long been held that consequentialism, despite its axiological openness, is unable to capture commonsense moral verdicts involving deontic constraints and special duties. On these grounds many have rejected consequentialism despite finding its theoretical framework attractive. If position-relative consequentialism is viable then it appears to be a game-changer in this debate. For position-relative consequentialism can produce deontic constraints and special duties. Thus it opens the way for an attractive form of consequentialism that preserves key aspects of the consequentialist framework yet captures the intuitive verdicts of commonsense morality.

One final thing to note. This article will follow the dominant contemporary practice of extending the term "consequentialism" to all theories that tie the deontic status of acts (or rules or motives) to the promotion of outcomes ordered by some evaluative ranking. Some

[2] When I talk of "recognizing relative values" I am talking about recognizing the possibility of *moral* value being relative. It has long been recognized that *prudential* value ("good for") is relative. However, this is insufficient for position-relative consequentialism (see §1.2), even though some regard ethical egoism, which requires each agent to maximize what is *good for her*, as a kind of proto agent-relative consequentialism. The reader should also note that, although I suggest that Sen (1982) is the first consequentialist to advance a position-relative axiology, Hurka (2003: 609-612) has argued for an unorthodox interpretation of Henry Sidgwick's *Methods of Ethics* in which Sidgwick accepts a kind of agent-relative moral good.

have insisted on restricting the term "consequentialism" to theories with neutral evaluative rankings because of the strong historical associations between consequentialism and agent-neutrality.[3] However, we need not get caught up in a debate about how to use the word "consequentialism." What ultimately matters is not the terms we use but rather that we clearly see the logical space of theoretical options. Those who prefer the restricted use of "consequentialism" should still admit that there is an important family of theories in normative ethics that borrow key structural features from traditional consequentialism yet employ a position-relative axiology. This family deserves careful examination in a volume like this.

1. Agent-Relativity

When Sen (1982) first proposed a position-relative form of consequentialism he was proposing an agent-relative form. Although other categories of relative value have since been explored, the agent-relative variety is by far the most significant and most widely discussed. Therefore, this section will focus on agent-relative value, with other varieties of relative value being explored in the next section.

1.1. Agent-Relative Rules and Theories

Before examining agent-relative value we should first consider what it is for a moral rule, reason, or theory to be agent-relative. I will use Derek Parfit's (1984) intuitive account of the distinction, although several formal accounts are also available.[4] Parfit takes all moral rules to give agents certain substantive aims. For example, a rule that prohibits lying gives each agent the aim that she does not lie. Parfit defines an agent-neutral rule as one that gives all agents the same aim and an agent-relative rule as one that gives different aims to different agents. So, for example, consider the following rules:

(1) Each agent must not tell lies.
(2) Each agent must ensure, to the best of her ability, that her family is honest.
(3) Each agent must ensure, to the best of her ability, that everyone is honest.
(4) Each agent must minimize general violations of (1).

[3] For example, see Scheffler (1982: 1), Pettit (1997: 129), and Brown (2011).

[4] For a formal account see the account I defend in Hammerton (2019), which is based on the work of McNaughton and Rawling (1991).

Rules (1) and (2) are generally regarded as agent-relative rules whereas (3) and (4) are generally regarded as agent-neutral. Parfit's account explains this as follows. Rules (1) and (2) are agent-relative because they give different aims to different agents. For example, (1) gives Akira the aim that *Akira does not tell lies* and Mia the different aim that *Mia does not tell lies*. Likewise, (2) gives me the aim that *my* relatives are honest and you the different aim that *your* relatives are honest. By contrast, (3) and (4) give the same aims to all agents. For example, (3) gives you and me (and everyone else) the common aim that everyone is honest. Likewise, (4) gives everyone the common aim that no one violates (1), and that if there are any violations of (1) then they are minimal.[5]

Parfit extends his account to moral theories. He says that an *agent-neutral theory* gives all agents the same aims, whereas an *agent-relative theory* sometimes gives different aims to different agents. It follows that any theory containing at least one agent-relative rule will itself be an agent-relative theory, whereas only theories with exclusively agent-neutral rules are agent-neutral.

Understanding the agent-relative/neutral distinction helps us to clarify two important features of commonsense morality—*deontic constraints* and *special duties*. A deontic constraint prohibits agents performing acts of a certain type even if doing so is the only way to prevent more acts of that type being performed by others. A special duty requires agents to give priority to people that they have some kind of special relationship with (e.g., parent and child, siblings, friends, compatriots etc.), even if not giving this priority would result in more prioritization of special relationships by others. Let's consider examples of each. First, consider the following case:

[5] A couple of clarifications are helpful here. First, when we talk of "aims" we are talking about *ultimate aims* and not *derived aims*. Given our different circumstances, I may best minimize lying by aiming to punish liars and you may best minimize it by aiming to reward honesty. However, despite our different derived aims our ultimate aim remains the same. Second, one could interpret (3) and (4) as giving different aims to different agents. For example, (3) might give me the aim that "*I* ensure that everyone is honest" and you the aim that "*you* ensure that everyone is honest." All agent-neutral rules can be reinterpreted as giving different aims in this way. However, no agent-relative rule can be reinterpreted as giving everyone the same aim (for example there is no interpretation of (1) which gives a common aim to everyone). Thus, strictly speaking, we should say that agent-relative rules are rules that cannot be interpreted as giving the same aim to all agents, whereas agent-neutral rules are rules that can be given this interpretation.

> MURDER: Belle is in a murderous rage and is about to kill two innocent people. The only way to stop her is for you to kill a random innocent person (who is not one of Belle's potential targets).

Many people judge that it would be wrong for you to kill the innocent person in this case even though it will result in fewer innocent people being killed overall. In making this judgement, they are endorsing a deontic constraint on killing innocent people. Second, consider the following case:

> ANTIVENOM: Kimbo's daughter and two other children he has no ties to have all been bitten by a brown snake. Kimbo has one dosage of antivenom available. To save his daughter's life he will need to give her a full dosage. The other two children can each be saved with a half dosage. Kimbo could allow the parents of these children to administer a half dosage to each child, in which case their children live and his daughter dies. Or he could administer the full dosage to his daughter, in which case she lives and the other children die.

Many people judge in this case that Kimbo ought to administer the antivenom to his own child even though offering it to the other parents would result in more children being saved by their parents. In making this judgement, they are endorsing a special duty requiring parents to favor their own children.

Deontic constraints and special duties are both agent-relative rules as they give different aims to different agents. For example, a deontic constraint on killing gives me the aim that *I do not kill* and you the aim that *you do not kill*. Likewise, a special duty for parents to favor their children gives me the aim that *I favor my child* and you the aim that *you favor yours*. Some have attempted an agent-neutral interpretation of deontic constraints and special duties by suggesting that a common aim could account for them. For example, Dougherty (2013) and Setiya (2018) have argued that killing in order to prevent more killings is a different act from killing for some other end such as revenge. A rule requiring us to prioritize minimizing the first kind of killing over minimizing the second kind gives us all a common aim and yet appears to produce a deontic constraint on killing. However, in Hammerton (2017) I show that this strategy is inadequate. Suppose that I can either ensure that I do not vengefully kill someone or ensure that you do not vengefully kill someone. The deontic constraint on killing requires me to ensure first that I do not kill, prioritizing my own non-killing over yours. Yet this cannot be captured by the common

aim that *no one kills to prevent more killings*. Only an agent-relative rule giving each agent the aim that *she does not kill* can capture this deontic verdict.[6]

Deontic constraints and special duties are core features of our commonsense morality (as our intuitive reactions to MURDER and ANTIVENOM demonstrate). As such, they show that commonsense morality is necessarily agent-relative. This presents a challenge to all agent-neutral theories. Such theories are incompatible with core features of commonsense morality and need to be defended against the charge that this incompatibility makes them implausible. Thus, the distinction between agent-neutral and agent-relative moral theories is thought to mark an important distinction between theories that can capture certain core features of commonsense morality and those that cannot.

When it was assumed that a consequentialist axiology is agent-neutral, the aforementioned distinction was thought to divide consequentialism from its main rivals, which all appeared to be agent-relative theories. However, the possibility of agent-relative rankings opens the way for agent-relative versions of consequentialism that accommodate deontic constraints and special duties. Thus, the possibility of agent-relative rankings is especially relevant to the question of whether consequentialism can accommodate these core features of commonsense morality.

1.2. Agent-Relative Rankings

The key idea behind agent-relative rankings is that an ordering of possible outcomes as better or worse could be relativised to agents. This relativisation allows for the correct ranking of states to depend on the agent who is evaluating those states such that, relative to one agent, state S_1 might be correctly ranked *above* state S_2, whereas, relative to another agent, S_1 might be correctly ranked *below* S_2.

[6] What if everyone shares the common aim that *every agent ensures that she does not kill*? In the example we are considering this common aim might explain why I must ensure that I do not kill rather than ensuring that you do not kill (for only the former contributes to "every agent ensuring her own non-killing"). However, this is still inadequate as it does not give the correct verdict in a case in which I can either ensure that *I do not kill* or ensure that *you ensure that you do not kill*. So, an agent-relative interpretation is still necessary. See Nair (2014) for a discussion of further technical issues related to agent-neutral theories capturing deontic constraints.

If such relativized rankings are adopted, then the value that they are supposed to capture must itself be relative. Things are not just morally good or bad, they are morally good or bad relative to particular agents. It follows that moral goodness is not a property possessed by states of affairs but a two-place relation between a state of affairs and an agent.[7]

Admitting agent-relative value into the axiology of a consequentialist theory allows it to capture deontic constraints and special duties. To see this let's consider the deontic constraint on killing. According to this constraint, each agent must not kill an innocent person even if it is the only way to prevent more killings by others. We saw earlier that an agent-neutral theory cannot capture this constraint. Yet, a consequentialist theory with an agent-relative axiology can. It does this by holding that, for each agent, that *she* kills an innocent person is *worse-relative-to her* than any number of innocent people being killed by others.[8] The result is that when each agent maximizes what is *good-relative-to her*, she refrains from killing innocent people even in cases like MURDER where it can prevent more killings by others. A similar move can capture special duties. The consequentialist can hold that, for each agent, that she neglects her child is *worse-relative-to her* than several other parents neglecting their children. Thus, maximizing what is *good-relative-to her* requires each parent to not neglect her child even when doing so can prevent more child neglect by others.[9]

Focusing on deontic constraints and special duties suggests that agent-relative value is based on a relation of identity between the evaluator of a state and an agent acting in that state.[10] For example, if I am comparing my killing an innocent person with two strangers killing an innocent person, I rank the former state as worse because I am identical to the agent who kills in that state. However, further examples suggest that other kinds of connections between the evaluator and the evaluated state could justify agent-relative

[7] For an alternative, non-standard picture of agent-relative value that takes properties rather than states of affairs to be the bearers of value, see the appendix in Dreier (2011).

[8] This captures an *absolute* constraint. To capture a *threshold* constraint the consequentialist must hold that, for each agent, her killing one innocent person is *worse-relative-to her* than other agents killing *n* or fewer innocent people (where *n* is the number at which the threshold ends).

[9] Examples of agent-relative consequentialists endorsing these kinds of axiological claims include: Sen (1982, 1983), Broome (1991), Dreier (1993, 2011), Portmore (2001, 2011), Smith (2003, 2009), and Louise (2004).

[10] In previous work (Hammerton 2016) I characterized agent-relativity in terms of numerical identity.

rankings. For example, compare a state in which my mother drowns to one in which an equally worthy stranger drowns. Now suppose that you have no personal connections to either my mother or the stranger. How ought we to rank these states? Given my personal connection to my mother, it appears that I ought to rank her drowning as worse than a stranger drowning. Given your lack of personal connection, it seems that you ought to rank these states as equally bad. Neither of us acts in the states we are assessing. However, my special relationship with a moral patient whose welfare is at stake in these states (and your lack of such a relationship) requires us to rank these states differently.[11]

To better understand agent-relative value, there are several further points that we need to clarify. First, we must not confuse what is *good-relative-to* an agent with what is *good for* an agent.[12] Both are categories of value that connect an agent with a state of affairs. However, the former is a kind of *moral* value whereas the latter is a kind of *prudential* value. Although some moral philosophers hold substantive views that connect moral value and prudential value (e.g., derive one from the other), they nonetheless remain conceptually distinct. This is why, in a situation where several people are drowning and either Azra will be rescued or five others will be rescued, it is not incoherent to say that "It is better for Azra that she is rescued, yet, from a moral perspective, it would be better if the five others were rescued instead." An agent-relative consequentialist might want to make similar distinctions. For example, in a situation where a murderer will kill Azra and another innocent person unless Azra herself kills a random innocent person, the consequentialist might want to say that Azra killing the innocent person is *better for* her, yet, from a moral perspective, it is *worse-relative-to* her. Even a consequentialist who does not want to say this because she accepts a tight connection between prudential goods and moral goods should at least admit that it is a possible view that some might want to hold because *good for* and *good-relative-to* are distinct concepts.

Second, that something is *good-relative-to* an agent does not entail that its value is constituted by, or dependent on, the mental states of that agent. Some philosophers who

[11] An upshot of this is that it may be a mistake to describe the kind of relativity we are concerned with as an "agent" kind. Mightn't a being who does not meet the full conditions for moral agency (e.g. a three-year-old human child) nonetheless correctly evaluate states differently depending on whether their content is appropriately connected to herself? If so then we should speak of "evaluator-relativity" rather than "agent-relativity."

[12] On this point see: Regan (1983: 96), Broome (1991: 8), and Schroeder (2007: 272).

endorse agent-relative value hold that it is mind-dependent and some hold that it is mind-independent. Just as is the case with agent-neutral-value, both positions are possible. However, although agent-relative value does not entail mind-dependence, there are certain mind-dependent accounts of value that entail that value is agent-relative. We will look at this in more detail later when we discuss an important example given by Michael Smith.

Third, on the standard understating of agent-relative value, each agent is supposed to recognize the correctness of each set of rankings relative to the agent it is indexed to, rather than treating her own relativized rankings as the universal yardstick by which everyone is assessed. Thus, if you act in a way that maximizes that which is *good-relative-to* you but does not maximize that which is *good-relative-to* me, then I must recognize you as acting correctly even though, from my evaluative perspective, things would have gone better if you had acted differently. A consequence of this is that it is only assessments of value that are relative and not assessments of an agent's actions or attitudes. Although different agents may correctly rank the same state of affairs differently, they must accept that there is only one correct deontic status for each specific act based on what is *good-relative-to* the agent who performs the act.

Fourth, some have found it illuminating to explain agent-relative value in terms of the fittingness analysis of value.[13] Independently of debates about agent-relative value, some have found the following analysis of value appealing:

FIT: A state of affairs is valuable if, and only if, it is fitting to desire its realization.

If this analysis is accepted then it provides an attractive account of why value comes in agent-relative and agent-neutral varieties. For it is plausible to hold that which states it is fitting to prefer at least sometimes depends on who you are and what relations hold between you and the people or things in the states you are assessing. For example, we might think that it is fitting for a parent to prefer that her child has a pleasant experience over another child having an equally pleasant experience, yet fitting for a neutral bystander to be indifferent about which child has the experience. If our evaluative ranking of states is derived from what it is fitting to prefer, and what it is fitting to prefer can vary from agent to agent, then the correct evaluative ranking can also vary from agent to agent.

[13] For example, see: Garcia (1986), Smith (2003, 2009), Portmore (2007), and Cullity (2015).

Several other analyses of value that are structurally similar to FIT also appear to explain agent-relative value. For example, analyzing value in terms of which states we have most reason to prefer, or which states it is appropriate to love for their own sake, or which states an idealized version of ourselves would desire all make room for relative values.[14] By contrast, the Moorean account of value—often regarded as the main rival of the fittingness analysis—is not as accommodating of agent-relative value. According to the Moorean account, value is unanalyzable—it cannot be explained in terms of other, more fundamental, normative properties. If the fact that a state is valuable is just a brute normative fact then it is difficult to see why that value would be relativized to agents. Nonetheless, a Moorean might offer independent reasons for postulating relative values, such as their theoretical utility. These issues will be discussed further in §3.1.

Finally, earlier we made the simplifying assumption that if value is relativized to agents then different evaluative rankings will be correct for different agents. However, this assumption is false (see Smith 2003, 2009) and this has important implications for how we characterize agent-relative value. To see that it is false consider a version of the dispositional theory of value holding that to judge that a state S_1 is better than S_2 is to judge that you would prefer S_1 to S_2 if you had a set of preferences that were maximally informed, coherent, and unified. This analysis of value entails that claims about value are indexed to agents because they are analysed in terms of what the person making the value judgement would ideally prefer. Yet, it does not entail that different rankings are correct relative to different agents. For it could be the case that the idealized preferences of all possible agents converge, and if they did coverage then a single evaluative ranking would be correct for all agents. Thus, holding the dispositional theory of value, and endorsing convergence among all possible agents, results in a theory in which the "better-than relation" is relativized to agents and yet there is a single correct evaluative ranking for all agents.

This example shows us that indexing evaluations to agents is not sufficient for there to be different correct evaluative rankings for different agents. However, the former is at least necessary for the latter because, if an evaluative ranking is correct for one agent but not correct for another, then that ranking must be indexed to the first agent but not the second.

[14] See Hurka (2001: Ch.7) and Smith (2009: 268).

Given this result, we need to clarify what we are trying to capture when we contrast agent-relative axiologies with agent-neutral axiologies. Many important things thought to follow from agent-relative rankings (such as the possibility of consequentializing deontic constraints and special duties) actually require that the relativized rankings are different for different agents. Thus, I suggest that we make a distinction between axiologies that are agent-relative in *only* a weak sense and those that are agent-relative in a strong sense. An axiology is agent-relative in the weak sense if the evaluative rankings it contains are relativized to agents. It is agent-relative in the strong sense if the evaluative rankings are relativized to agents and different rankings are correct relative to different agents.

With this distinction in place, we can note that the strong sense of agent-relativity is the more important one when it comes to the contemporary interest in agent-relative axiologies. For example, advocates of agent-relative consequentialism have all had the strong sense in mind when developing their theories. Similar points will apply to other kinds of position-relative axiologies. Therefore, whenever I refer to any kind of position-relative axiology in this article, it is the strong sense of relativity that I have in mind.

2. Other Kinds of Relativized Rankings

2.1. Time-Relativity

Earlier we saw that deontic constraints and special duties are agent-relative rules that can only be captured by an agent-relative moral theory. However, some moral philosophers have argued that agent-relativity alone does not fully capture the deontic verdicts that follow from these rules. They have been motivated by cases like the following:

> RASKOLNIKOV: Raskolnikov's mind is gradually being overcome by homicidal thoughts. He judges correctly that: (1) in the next few days he will give in to these urges and commit multiple homicides, and (2) the only option available to him that will prevent this killing spree is to commit a single homicide now (perhaps this will lead to his immediate imprisonment, or satiate his homicidal urges, eradicating them from his psychology). Thus, Raskolnikov must choose between killing one innocent person now or several innocent people later.

If you accept that there is a deontic constraint on killing, what verdict should you reach in this case? Whatever Raskolnikov does, he will end up violating the constraint at some point. Nonetheless, what does the constraint require him to do right now? According to

one view he ought to kill the one now. After all, the deontic constraint gives him the aim that he does not kill and thus, in a situation in which all available options involve him killing, he does best to at least minimize the killing that he does.

According to another view he must *not* kill the one now. After all, if preventing another agent from murdering several innocent people cannot justify killing an innocent person then why should preventing your future self from murdering several innocent people justify it? To permit the sacrifice only when it prevents your killings appears to focus excessively on your own clean hands, while neglecting the rights of the victim you sacrifice.[15]

Those who endorse the second view must interpret the constraint on killing as not just an agent-relative rule, but also a time-relative rule. In Parfit's terms, it doesn't just give different aims to different agents, it also gives agents different aims at different times. Thus, my theory-given aim is not just that I do not kill innocent people, but that I do not kill innocent people at the current moment. By contrast, those who favor the first view are interpreting the constraint as agent-relative yet time-neutral.

To incorporate the time-relative interpretation of deontic constraints into consequentialism the consequentialist must relativize evaluative rankings not only to agents but also to times. Thus, there are not only different evaluative rankings for different agents but also different evaluative rankings for the same agent at different times. It can then be held that, for each agent, *her* killing an innocent person at a time *t* is *worse-relative-to her at t* than any number of innocent people being killed by others at any time, or by her at times other than *t*.[16] It follows that at each moment, when an agent maximizes what is *good-relative-to her* at that moment, she refrains from killing even in cases like RASKOLNIKOV where she can minimize her total killings by killing now.

Although time-relative rankings look most plausible when combined with agent-relative rankings, accepting time-relativity without agent-relativity is also possible. Broome (1991: 8) gives the example of a rule requiring us to expend all available resources saving any miners who are currently trapped in mines even if doing so will, by diminishing our

[15] The first view is defended in Lopez et al (2009) and Portmore (2011: 103-108). The second view is defended in Kamm (1989), Brook (1991), Kamm (1996, ch.9), Louise (2004), and Johnson (2019).

[16] This captures an absolute constraint. To capture a threshold constraint, we replace talk of "any number" with talk of "fewer than *n*," where *n* is the number at which the threshold ends.

resources, prevent us from saving more trapped miners in the future. Such a rule is agent-neutral (all agents are given the same aim of saving trapped miners), yet time-relative (our aims are different at different times). A time-relative, yet agent-neutral ranking could capture this rule. It would hold that at any time *t*, saving any trapped miners at *t* is better-relative-to *t* than saving more trapped miners after *t*.

In the examples discussed so far the relevant connection between the evaluator's position and the evaluated state has been *identical* temporal locations. However, other types of connections could also produce time-relative rankings. For example, a painful experience in my past appears to be better than a painful experience in my future. Indeed, at any particular moment, many would prefer that they had a very painful experience yesterday over being about to have a moderately painful experience.[17] This indicates that we are ranking states differently at different times. When the very painful experience is behind me and the moderately painful one is ahead of me, then I rank the latter as worse. However, if both are behind me or both are ahead of me, then I rank the former as worse. Notably, what makes the difference here is not identity in times but a time being before or after the temporal location of the evaluator. Although these examples are instructive, they appear to only apply to prudential evaluations we make about our own lives. For example, if I am told that my loved one either had a very painful operation yesterday or is about to have a moderately painful operation today I prefer that she was spared yesterday's suffering and has the lessor suffering today.

2.2. Location-Relativity and World-Relativity

Recently, two further types of position-relativity have been explored by Dreier (2018). Position-relativity involves a kind of indexical reference built into rules or values. Indexicals may be indexed to persons, times, places, or worlds (e.g., "I", "now", "here", and "actually"). This suggests that in addition to agents and times, there could be relativity based on physical location (location-relativity) or possible world location (world-relativity). Are there any reasons for adopting an axiology that is position-relative in either of these ways?

Dreier (2018: 34) suggests that location-relativity could be supported by the view that the physical distance between an agent and a person needing aid is a morally significant factor

[17] See Parfit (1984, Ch.8).

that determines the strength of the obligation to aid. Because Dreier rejects this view, he rejects location-relativity. However, his discussion suggests that those like Kamm (2000, and 2007: chap. 11), who endorse the view, have grounds for endorsing location-relativity. I will now argue that this is incorrect. Even if physical distance is morally significant in the duty to aid, this is best accounted for without postulating location-relativity.

The key idea behind position-relativity is that evaluators in relevantly different positions will give different rankings of states that are appropriately related to their position. This needs to be distinguished from a certain positional fact being morally significant, which can sometimes be accounted for in position-neutral terms. For example, suppose I think that it morally matters that parents help their own children. This makes agential positions morally significant but can be accounted for in agent-neutral terms (everyone shares the common aim that parents care for their children). We recognize parental duties as agent-relative because we make the further judgment that a parent must not neglect her child even if it will result in less child-neglect by other parents. Thus, each parent must assess *her neglecting her child as worse* than *other parents neglecting their children.*

Applying this to the duty to aid, note that the supposed moral significance of physical distance can be accounted for in location-neutral terms. Irrespective of our location we can all recognize the value of an agent prioritizing those physically closer to her when deciding who to aid. Location-relativity is required only if we make the further judgement that an evaluator at location L_1 should assess states where an agent at L_1 prioritizes those closer to L_1 as better than states where an agent at L_2 prioritizes those closer to L_2. But this is intuitively implausible and the arguments offered by Kamm and others do not suppose anything of this sort. Given that there are no other rationales for location-relativity, we should conclude that it is implausible.

Let's now consider world-relativity. Suppose that we must choose whether or not to bring a large number of extra lives into existence. If we bring these extra lives into existence, then our lives will be slightly worse as there will be more people to share resources with. However, the new people who exist will have very good lives that outweigh the slight loss of welfare in our lives. One might think that this example results in a kind of paradoxical situation where, whatever option we choose, its outcome is better than the other option. If we bring into existence the extra lives we have done the best thing

because, although our lives are slightly worse, this is outweighed by the good lives lived by the extra people. If we don't bring these lives into existence, then we have done the best thing because we have avoided making our lives worse and the extra welfare we might have created is irrelevant because it belongs to people who do not exist and the potential welfare of non-existent people has no moral significance.

Dreier (2018) suggests that this verdict is best interpreted as endorsing a form of world-relativity. In this example, we are essentially choosing which of two possible worlds to bring about. If we bring about the first world (with the extra people) then, from our position as evaluators located in that world, we must rank it as better than the second world. If we bring about the second world (with no extra people) then, from our position as evaluators located in that world, we must rank it as better than the first world. Thus different rankings are correct relative to different worlds.

One interesting feature of world-relativity is that, unlike agent and time-relativity, it appears to play no role in practical deliberation. Our choices determine which possible world is actual, so knowing that different things are best or worst relative to different worlds we can bring about cannot make one world more choiceworthy than another. Nonetheless, Dreier suggests that it is still a useful concept for the following three reasons. First, it can explain why we should feel a certain way about a choice we have made. For example, in the case we just considered it explains why, regardless of which option we choose, we can feel good about our choice. Second, it can make sense of a strategy used by Portmore (2011) to consequentialize prohibition dilemmas. Third, it can explain Temkin's (2012) notion of "essentially comparative value" without appealing to the counterintuitive idea that "better-than" is intransitive. These are substantial claims that I cannot explore further here. However, their significance suggests that world-relativity deserves further exploration in the future.

2.3. Patient-Relativity

In Hammerton (2016) I used the following analogy to argue that "patient-relativity" is possible. In agent-relative and time-relative rules numerical identity between *agents* or *times* makes a difference to the deontic status of acts. For example, killing one to prevent the killing of two may be permissible when the killer of the one and the killer of the two are identical, and impermissible when they are distinct. Numerical identity between *moral patients* can also make a difference to the deontic status of acts. For example, perhaps it

is wrong to break a promise to one promisee in order to keep comparable promises to several other promisees. However, if all the promises at stake are owed to the same promisee then it seems permissible to break one promise to that promisee in order to keep several comparable promises to her.

I now recognize that it was a mistake to describe this patient-related phenomenon as a kind of "relativity." In agent-relative and time-relative rules, identity between agents or times might make a difference to an act's deontic status. However, this is not what makes them relative rules. They are relative because they give different substantive aims to agents at different positions (positions picked out by the agent or the time). Yet, there is no "patient" position that can result in different aims being given to different patient identities.[18]

We can also make this point in terms of relativised rankings. If the patient in one state is identical to the patient in another state this may make a difference to how we ought to rank those states. However, the identity between patients is a neutral fact accessible from all evaluative positions, so it does not result in multiple correct rankings of states. By contrast, if the fact that I act in a state I am evaluating makes a difference to how I ought to rank that state, but not how others ought to rank it, then the correct rankings of states will vary between agents.

A lesson we can draw from this is that we must be careful not to confuse *identity difference-making* (identity between two things making a difference to the deontic status of acts) with *position-relativity*. Standard cases of agent and time-relativity may involve identity difference-making, but this is not what makes them position-relative. A further source of confusion is that, in the cases where patient identity makes a difference to an act's deontic status, it also turns on or off certain kinds of position-relativity. For example, in the promise-keeping case, when the patients are the same the constraint on promise-breaking appears to be agent-relative yet time-neutral. And when the patients are different the constraint appears to be both agent and time-relative.

[18] Of course, the evaluator assessing states might be a moral patient. Furthermore, her connection to moral patients in the states she assesses might be the basis for a kind of relativity. However, this is best categorized as "evaluator relativity" (see footnote 10 above), and, in any case, is not the phenomenon that I previously described as "patient-relativity."

2.4. Combining Relativities

All four kinds of position-relativity that we have reviewed so far are logically independent of one another. This means that a moral theory can, as a matter of logic, contain any combination of different kinds of position-relativity/neutrality. Nonetheless, Derek Parfit (1984) has argued that only theories that are neutral in every domain or relative in every domain are plausible. He calls such theories respectively "fully neutral" and "fully relative" and contrasts them with theories that are "incompletely neutral/relative." Parfit says:

> This claim can appeal to the analogy between oneself and the present, or what is referred to by the words 'I' and 'now'. This analogy holds only at a formal level. Particular times do not resemble particular people. But the word 'I' refers to a particular person *in the same way in which* the word 'now' refers to a particular time. And when each of us is deciding what to do, he is asking, 'What should *I* do *now?*' Given the analogy between 'I' and 'now', a theory ought to give to both the same treatment.[19]

Parfit applies his analogy only to agents and times, presumably because he was unaware of the possibility of location-relativity and world-relativity. However, the analogy he draws seems to extend to places and worlds (e.g., "What should *I* do *now*, *here*, in the *actual* world?"). Therefore, if Parfit's argument works what it shows is that moral theories must be either neutral with respect to agents, times, locations, and worlds, or relative with respect to them.

What should we make of Parfit's argument? I suspect that many will find the analogy unconvincing. Why should the fact that persons, times, places, and worlds are all elements of indices that are relevant to practical decision making mean that they must all be given identical treatment? If we have independent reasons for thinking that one of these domains is relative and another is neutral, then why not accept incomplete relativity?

Some philosophers have appealed to full relativity to justify giving agent-relative deontic constraints a time-relative interpretation.[20] However, this use of Parfit's argument seems problematic. Deontic constraints are generally defended by appeal to commonsense

[19] Parfit (1984: 140).

[20] See: Louise (2004: 535) and Johnson (2019).

morality. Yet a fully relative moral theory is in conflict with fundamental features of commonsense morality. I note some of these conflicts in Hammerton (2016). Our discussion of location-relativity in §2.2 suggests a further conflict. A fully relative moral theory must be location-relative. Yet location-relative deontic verdicts are deeply counterintuitive and there appear to be no theoretical reasons (other than the appeal to full relativity) to endorse them. On these grounds, we might conclude that if the appeal to full relativity succeeds then it pushes us toward position-neutral theories, which, although also in conflict in commonsense morality, can at least be supported by independent theoretical considerations. Those wishing to defend agent-relative moral theories probably ought to reject the appeal to full relativity.

3. The Debate on Relativizing

In this chapter, we have noted several reasons in favor of adopting a position-relative axiology. First, it allows certain intuitive verdicts of commonsense morality to be captured from within a consequentialist framework.[21] Second, it is entailed by the combination of a fittingness analysis of value and our intuitions about what is fitting. And third, it allows our moral evaluations to reflect our position as particular agents, acting at a particular time in the world we inhabit.[22] In this final section we will consider several arguments against position-relative axiologies and assess how successful they are at undermining position-relative consequentialism.

3.1. The Incoherence Argument

One argument against position-relativity questions the coherence of relativized values. Because our ordinary commonsense notion of moral value is firmly position-neutral, it is alleged that there is no conceptual space for position-relative moral values.[23] Here is Broome (1991: 8) expressing this concern:

> [The agent-relative consequentialist] claims that breaking my promise is worse *for me* than keeping it, and one might doubt that such a claim makes good sense.

[21] Portmore (2011:103-109) and Hammerton (2020) go a step further on this point and argue that agent-relative consequentialism does *better* than non-consequentialist alternatives at accounting for our commonsense intuitions around deontic constraints.

[22] On this last point see especially Sen (1982: 29-30).

[23] See, especially Regan (1983) and Schroeder (2007).

> Certainly, there is one clear sense it could have. It could mean that breaking my promise is against my interest: that it would make me less well off than I would otherwise have been. But that is not what it is supposed to mean in the argument. The good in question is not my private wellbeing, and the argument is not appealing to my own self-interest. The good is general good, not mine; but general good evaluated from my own special position as an agent. And one might doubt that good can really be agent-relative in this way. Agent-relative teleology appeals to agent-relative good, and one might be dubious of such a concept.

And here is Schroeder (2007: 291) expressing it more stridently:

> …. I don't understand what "good-relative-to" talk is all about, I don't understand how it could be appealing to think that you shouldn't do something that will be worse-relative-to you. I don't even understand what that means! Until the [agent-relative consequentialists] give me some reason to think that the *good-relative-to* relation is somehow very much like the *good* property and the *good for* relation, I don't see why I should remotely find such an idea deeply compelling.

On one version of this argument, the critic claims that position-relative values are incompatible with our ordinary conception of the moral good and concludes that position-relative axiologies are not viable. However, Schroeder (2007) rightly notes that this is too quick. For, even if it lacks a commonsense precedent, the logical structure of position-relative value is clear enough to allow it to be defined by stipulation as a kind of theoretical posit. Schroeder usefully draws an analogy to "electron." We have no pre-theoretic concept of an electron. Nonetheless, we are justified in employing this concept in scientific theory because it serves a useful role, allowing us to give a unified and elegant explanation of several empirical observations. If a useful theoretical role can also be found for position-relative value, then that could justify its postulation. Schroeder then goes on to argue that no useful theoretical role has been provided by those who defend relative values. For example, many agent-relative consequentialists appeal to the so-called "compelling idea" that it is always permissible to do what will lead to the best outcome. Schroeder argues that this can provide no independent support for postulating agent-relative values because the intuitive appeal of the compelling idea is based on our ordinary, agent-neutral understanding of "good" and "best."

Two different lines of response have been offered to Schroeder's argument. One line attempts to show that, contrary to Schroeder's claim, a kind of position-relative value is already present in our ordinary moral thought and talk. For example, suppose that during a flood a parent decides to rescue two random children instead of rescuing her own child. Someone attempting to explain why it was morally wrong for the parent to act this way might say: "She should have recognized that, from her position as a parent, her child drowning is worse than two random strangers drowning." This talk appears to commit the speaker to agent-relative value as "worse" is tied to the agent's position and yet is being used in a *moral* rather than *prudential* sense (the context is one of moral blame). Therefore, insofar as people sometimes talk this way, there is room in ordinary moral discourse for position-relative values. More sophisticated versions of this line of response can be found in Dreier (2011) and Cullity (2015).

A second line of response appeals to the fittingness analysis of value. Recall from §1.2 that this analysis of value, and our intuitions about what is fitting, imply that there are position-relative values. This means that, even if we have no pre-theoretical concept of position-relative value, we can nonetheless justify it as a theoretical posit by pointing to independent reasons that support the fittingness analysis of value.[24]

Schroeder anticipates this line of response and offers a reply. He points out that, if the agent-relative consequentialist is to explain the paradoxical nature of deontic constraints, then she needs to explain why an agent must not violate a constraint even when doing so is agent-neutrally better. She can do this by supposing that, for each agent, the relative value of her not violating the constraint outweighs the neutral value of there being fewer constraint violations by others. However, Schroeder argues that if she endorses the fittingness analysis then there is a problem. Talk of "what it is fitting to desire" requires an agent-place ("what it is fitting for *x* to desire") and so the fittingness analysis requires *all* value claims to be relativized to agents. Thus, the only way for it to accommodate neutral values is to say that a state is agent-neutrally good when it is *good-relative-to* every agent. But this makes the earlier claim about the agent-relative value of a state sometimes outweighing its agent-neutral badness incoherent. This claim is interpreted as follows: There is a state that is *better-relative-to me* but *worse-relative-to everyone*

[24] See Portmore (2007) and Smith (2009).

(including me). Schroeder concludes from this that basing agent-relative consequentialism on the fittingness analysis is not viable.

Two different counter-arguments have been developed against Schroeder's reply. Suikkanen (2009) argues that Schroeder overlooks an alternative way of accommodating agent-neutral values into the fittingness analysis. Perhaps a state is agent-neutrally good when it is *good-relative-to* the impartial spectator. Under the fittingness analysis this translates as it being fitting for the impartial spectator to desire that state. This allows for clashes between neutral and relative values to be explained as cases where a state's *value-relative-to the impartial spectator* conflicts with its *value-relative-to me.*

A second counterargument comes from Smith (2009). Suppose that a child narrowly escapes being crushed by a falling tree. Everyone can recognize that this state of affairs is good. What makes it good is that someone has escaped serious injury. Now suppose that the child is my child. I have an additional reason for recognizing this state as good that others lack. What makes it good is that *my* child has escaped serious injury. Smith suggests that there is an important distinction behind these two different value-making features. The former is "neutral" because it makes no indexical references to anyone and thus can be appreciated by all agents. The latter is "relative" because it contains an essential indexical reference and thus can only be appreciated by the persons picked out by this reference. This distinction between neutral and relative value-making features applies even if all value terms are indexed to agents. We saw in §2.2 that indexing to agents is not sufficient for what I called "strong agent-relativity" because it doesn't guarantee that different rankings will be correct for different agents. However, having relative value-making features is sufficient for strong agent-relativity as it entails that there will be different rankings for different agents. Thus, we can use Smith's value-maker distinction to characterize strong agent-relativity. This provides a reply to Schroeder. In the fittingness analysis of value, value-making features can be understood as the features that make certain responses fitting. Some of these features may have the indexical reference that Smith refers to and others may lack it. Thus, we can interpret the claim that the agent-relative value of a state sometimes outweighs its agent-neutral badness as follows. Sometimes neutral value-making features count against a state yet stronger relative value-making features count in favor of it.

The strong replies offered against Schroeder's argument suggest that the incoherence argument against position-relative axiologies is not decisive.

3.2. Theoretical Arguments

Another kind of argument against position-relative axiologies points to theoretical considerations that favor neutrality. These arguments are usually offered by agent-neutral consequentialists and, if successful, undermine not only position-relative forms of consequentialism but all non-consequentialist theories that are relative. I will consider two such arguments here.

First, Derek Parfit (1984: Ch.4) argues that all agent-relative moral theories face a problem he calls "direct, collective self-defeat."[25] A theory is *directly collectively self-defeating* if and only if, when all of us successfully follow that theory, our theory-given aims are worse achieved than they would have been if none of us had successfully followed that theory. Parfit demonstrates this problem with various Prisoner Dilemma-style cases in which agents following agent-relative requirements do worse at bringing about their theory-given aims than they would have done if they had instead followed an alternative agent-neutral rule. For example, in a case he calls the "parent's dilemma," parents following an agent-relative requirement to minimize harm to their *own* child end up in a situation where each of their children suffers more harm than he would have suffered if the parents had instead followed the agent-neutral rule of minimizing harm to *all* children.

This is a serious problem for agent-relative theories. However, Parfit shows that it can be resolved if agent-relative theories are revised so that they require agents to follow agent-neutral rules in cases of direct collective self-defeat. This solution can be adapted to agent-relative consequentialism.[26] It can escape direct collective self-defeat by fixing its axiology such that, in circumstances where direct collective self-defeat would otherwise occur, outcomes where agents cooperate to bring about that which is best relative to all

[25] Regan (1983: 109), and Forcehimes and Semrau (2019) also endorse versions of this argument against agent-relative consequentialism.

[26] I show how this can be done in Hammerton (forthcoming). Parfit doesn't consider this possibility because he does not recognize the possibility of agent-relative consequentialism.

are ranked higher on each agent's relative rankings than all outcomes where they do not cooperate.[27]

Although, this solution is viable, it does have a bit of an ad hoc feel to it. Furthermore, it brings agent-relative theories closer to agent-neutral theories by having them adopt certain agent-neutral principles to avoid direct collective self-defeat. This leads to the obvious thought that a simpler and more elegant way to deal with this problem is to adopt a completely agent-neutral theory—a thought that Parfit has some sympathy with.

A second argument against position-relative axiologies, that has similarities with Parfit's argument, can be found in the work of Philip Pettit. Pettit argues that agent-relative moral theories result in a kind of "moral civil war" in which:

> … two or more agents are required, and required in a self-relativized manner, each to try to advance a cause that falls particularly to them. The cause may be that of looking after certain dependents, doing one's duty in some regard, respecting the rights of those with whom one happens to deal directly, or just trying to maintain a clean record in regard to certain neutral values. (1997: 149)

The result is that:

> … universalizing no longer means recognizing a commonly espoused goodness or rightness. All that it means is recognizing a common structure in the essentially different properties of rightness by which we are each required to steer and orientate. (1997: 149)

Pettit is right to emphasize the attractive "collective harmony" of agent-neural theories. However, the inter-agential conflicts that follow from agent-relative theories may be more limited than he suggests. Recall the point made in §2.2 that I must recognize you as acting correctly when you maximize that which is *good-relative-to you* even if your act is counterproductive to my relative good. We can add to this point the observation that commonsense morality prohibits forcing, manipulating, influencing, or aiding another to do the wrong thing even in cases where doing this maximizes your relative good. Thus suppose that you can either save your child or save mine. From my evaluative perspective it would be better if you saved my child. Nonetheless, I must recognize that it is right that

[27] This only covers cases where everyone will do what they ideally ought to do. To cover other kinds of cases we will need further axiological claims related to Parfit's principles R2 and R3.

you save your child (as it will maximize your relative good) and must not attempt to induce you to save mine instead. A consequentialist can capture this with the axiological claim that my getting you to save my child in these unscrupulous ways is *worse-relative-to me* than you not saving my child.[28] This avoids the kind of "moral civil war" in which agents regularly interfere with each other's attempts to do the right thing. Our mutual respect for each other's need to pursue her own relative good insulates us against such conflicts. Thus, although Pettit has a fair point to make about the attractive "collective harmony" of agent-neutrality, it is not as strong as it initially appears.

Parfit and Pettit remind us that there are theoretical considerations that push us toward position-neutrality. What we have is a clash between lower-order intuitions telling us, on a case by case basis, that morality should respect the partiality of our situated lives and higher-order intuitions telling us that morality should move beyond our partial interests to some impartial perspective. Whether a consequentialist should endorse a position-relative axiology may ultimately depend on how she resolves this clash.[29]

References

Brook, R. (1991). "Agency and Morality." *Journal of Philosophy* 88: 190-212.

Broome, J. (1991). *Weighing Goods*. New York: Blackwell.

Brown, C. (2011). "Consequentialize This." *Ethics* 121: 749-771.

Cullity, G. (2015) "Neutral and Relative Value." In *The Oxford Handbook of Value Theory*, edited by I. Hirose and J. Olson, 96-116. Oxford: Oxford University Press.

Dougherty, T. (2013). "Agent-Neutral Deontology." *Philosophical Studies* 163: 527-537.

Dreier, J. (1993). "Structures of Normative Theories." *The Monist* 76: 22-40.

Dreier, J. (2011). "In Defense of Consequentializing." *Oxford Studies in Normative Ethics*, vol. 1, edited by M.Timmons, 97-118. Oxford: Oxford University Press.

Dreier, J. (2018). "World-Centered Value." In *Consequentialism, New Directions, New Problems*, edited by C.Seidel, 31-50 New York: Oxford University Press.

Forcehimes, A. and Semrau, L. (2019). "Non-Compliance Shouldn't Be Better." *Australasian Journal of Philosophy* 97: 46-56.

Garcia, J. (1986). "Evaluator Relativity and the Theory of Value." *Mind* 95: 242-245.

[28] Sen (1983) appears to endorse something like this in response to an objection in Regan (1983).

[29] For valuable comments on an earlier draft of this chapter, I would like to thank Ryan Cox and Daniel Nolan.

Hammerton, M. (2016). “Patient-Relativity in Morality.” *Ethics* 127: 6-26.

Hammerton, M. (2017). “Is Agent-Neutral Deontology Possible?” *Journal of Ethics and Social Philosophy* 12: 319-324.

Hammerton, M. (2019). “Distinguishing Agent-Relativity from Agent-Neutrality.” *Australasian Journal of Philosophy* 97: 239-250.

Hammerton, M. (2020) “Deontic Constraints are Maximizing Rules.” *Journal of Value Inquiry* 54: 571-588.

Hammerton, M. (2020) “Agent-Relative Consequentialism and Collective Self-Defeat.” *Utilitas*. doi: 10.1017/S0953820820000096

Hurka T. (2001). *Virtue, Vice, and Value*. New York: Oxford University Press.

Hurka, T. (2003). “Moore in the Middle.” *Ethics* 113: 599-628.

Johnson, C. (2019). “The Intrapersonal Paradox of Deontology.” *Journal of Moral Philosophy* 16: 279-301.

Kamm, F.M. (1989). “Harming Some to Save Others.” *Philosophical Studies* 57: 227-260.

Kamm, F.M. (1996). *Morality, Mortality: Volume II: Rights, Duties, and Status*. New York University Press.

Kamm, F.M. (2000). “Does Distance Morally Matter to the Duty to Rescue.” *Law and Philosophy* 19: 655-681.

Kamm, F.M. (2007). *Intricate Ethics: Rights, Responsibilities, and Permissible Harms*. Oxford: Oxford University Press.

Lopez, T., Zamzow, J., Gill, M., and Nichols, S., (2009). “Side Constraints and the Structure of Commonsense Ethics.” *Philosophical Perspectives* 23: 305-319.

Louise, J. (2004). “Relativity of Value and the Consequentialist Umbrella.” *Philosophical Quarterly* 54: 518-36

McNaughton, D. & Rawling, P. (1991). “Agent-Relativity and the Doing-Happening Distinction.” *Philosophical Studies* 63: 167-185.

Moore, G.E. (1903). *Principia Ethica*. Cambridge: Cambridge University Press.

Nair, S. (2014). “A Fault Line in Ethical Theory.” *Philosophical Perspectives* 28: 173-200.

Parfit, D. (1984) *Reasons and Persons*. Oxford: Clarendon Press.

Pettit, P. (1997) “The Consequentialist Perspective.” In *Three Methods in Ethics*, edited by M.W. Baron, P Pettit, and M. Slote, 92-174, Oxford: Blackwell.

Portmore, D. (2001). “Can an Act-Consequentialist Theory be Agent Relative?” *American Philosophical Quarterly* 38: 363-377.

Portmore, D. (2007). “Consequentializing Moral Theories.” *Pacific Philosophical Quarterly* 88: 39-73.

Portmore, D. (2011). *Commonsense Consequentialism: Wherein Morality Meets Rationality*. Oxford University Press.

Regan, D. (1983). "Against Evaluator Relativity: A Response to Sen." *Philosophy & Public Affairs* 12: 93-112.

Scheffler, S. (1982). *The Rejection of Consequentialism*, Oxford: Clarendon Press.

Schroeder, M. (2007) "Teleology, Agent-Relative Value, and 'Good'." *Ethics* 116: 265-295.

Sen, A. (1982). "Rights and Agency." *Philosophy & Public Affairs* 11: 3-39.

Sen, A. (1983) "Evaluator Relativity and Consequential Evaluation." *Philosophy & Public Affairs* 12: 113-132.

Setiya, K. (2018). Must Consequentialists Kill? *The Journal of Philosophy* 115: 92-105.

Sidgwick, H. (1874). *The Methods of Ethics*. London: Macmillan.

Smith, M. (2003). "Neutral and Relative Value after Moore." *Ethics* 113: 576-598.

Smith, M. (2009). "Two Types of Consequentialism." *Philosophical Issues* 19: 257-272.

Suikkanen, J. (2009). "Consequentialism, Constraints and the Good-Relative-to: A Reply to Mark Schroeder." *Journal of Ethics and Social Philosophy* 1-9.

Temkin, L. (2012). *Rethinking the Good. Moral Ideas and the Nature of Practical Reasoning*. Oxford: Oxford University Press.