Against Consequentialist Interpretations of Deontological Principles

AMIR SAEMI

Abstract
The distinction between doing and allowing (DDA); intentional and non-intentional agency (DDE); and agent-relative versus agent-neutral values have been invoked to support the deontological nature of moral thought. Consequentialists claim that, since these distinctions can be incorporated into value theory, all plausible deontological theories ultimately reduce to consequentialism. I argue: first, that while the DDA and DDE can be framed in agent-neutral terms, they must be interpreted in agent-relative terms; second, that even when interpreted in agent-relative terms, the DDA and DDE compel non-consequentialist understandings; and third, that these three distinctions function as non-consequentialist constraints on deliberation about action, thus resisting attempts to 'consequentialize' them. The conclusion is that, far from being a mere variant of consequentialism, deontology is a distinct moral theory—one that offers a principled rejection of the consequentialist 'compelling idea' that it is always permissible to bring about the outcome with the best possible consequences.

A longstanding debate persists between consequentialists and deontologists over whether the moral status of actions depends solely on the value of their outcomes or whether moral principles govern the will independently of 'what it brings about or accomplishes' (Kant, 2019 [1785], p. 10). According to act-consequentialism, the action yielding the optimal outcome is morally permissible. However, various moral principles and distinctions have been claimed to demonstrate the deontological nature of our moral thought. Notable among them are the distinctions between *doing* and *allowing, intentional agency* and *non-intentional agency*, and *agent-relative* and *agent-neutral values*.[1] Using these distinctions, deontologists

[1] There is a large literature on each of those distinctions. The early modern formulation of the doctrine of doing and allowing can be traced back to Foot (1967) and Quinn (1989a). For a review of the literature on the doctrine see Woollard (2012a and b). An early formulation of the doctrine of double effect can be found in Quinn (1989b). For a literature review on the doctrine, see Fitzpatrick (2012). Many philosophers have held that these distinctions are essentially non-consequentialist. Quinn (1989a, b) is

doi:10.1017/S0031819125000075 © The Author(s), 2025. Published by Cambridge University Press on behalf of The Royal Institute of Philosophy

https://doi.org/10.1017/S0031819125000075 Published online by Cambridge University Press

present cases where an action resulting in the best outcome is impermissible.

On the other hand, consequentialists, also known as 'consequentializers,' argue that these distinctions can be accommodated within value theory and that consequentialism does not need to be tied to a particular value theory. Accordingly, consequentialists contend that all plausible deontological theories are a version of consequentialism, thereby establishing consequentialism as the clear victor in the debate.

In the following three sections of this paper, I will present arguments for the following claims:

(1) While it is conceivable to construe the doctrine of doing allowing (DDA) and the doctrine of double effect (DDE) in agent-neutral terms, we are driven to their agent-relative interpretations.
(2) Despite the possibility of agent-relative consequentialist understandings of the DDA and the DDE, we are driven to non-consequentialist interpretations of them.
(3) Drawing upon a process-oriented understanding of actions, I will argue for conceptualizing moral norms such as DDE and DDA as norms governing the process of deliberation about action. These norms do not pertain to producing the best outcomes but to the validity of our moral reasoning.

The first two theses show that consequentialists' efforts to integrate the DDE and DDA into either agent-neutral or agent-relative consequentialism are bound to fail. The third thesis explains why: properly understood, the DDE and DDA are inherently non-consequentialist constraints on the process of deliberation about action and resist attempts to be 'consequentialized'. Deontology thus emerges as not just a variant of consequentialism but a genuinely distinct moral theory, one that offers a principled rejection of the consequentialist's 'compelling idea' that it is always permissible to bring about the best possible outcome.

---

among the first philosophers who make this point with respect to the distinctions between doing and allowing, and intentional and non-intentional agency. Scheffler (1982), Darwall (1986), and Kagan (1989) are among the first philosophers who seek to make a distinction between deontological ethics and consequentialism by appealing to agent-relative values.

https://doi.org/10.1017/S0031819125000075 Published online by Cambridge University Press

1. Agent-Relativity of the DDA and DDE

According to the doctrine of doing and allowing (DDA), it is harder to justify doing harm than merely allowing harm. Quinn (1989a) characterizes doing harm as *positive harmful agency* and allowing harm as *negative harmful agency*. Using this terminology, the DDA says that positive harmful agency is harder to justify than negative harmful agency. It is tricky to figure out the necessary and sufficient conditions for positive and negative agency, and the literature offers various proposals. Considering that I will only focus on the core cases, we do not need to evaluate these different proposals. To demonstrate the distinction between positive and negative agency, consider the following case:

> *Rescue inside the Trolley*. I am a passenger on a runaway trolley, trying to save two lives inside as it heads toward a person. The driver is gone; I can go to the control panel to stop the trolley before it hits that person, but doing so prevents me from saving the lives of the two persons at the back.

Intuitively, I am not required to stop the trolley, when doing so prevents me from saving the lives of the two persons at the back.[2] My agency toward the one person in this case is negative harmful agency. But now contrast this case with a classic trolley case in which a runaway trolley is heading towards *two* persons. The trolley will kill those two persons unless it is diverted to an alternative track where it will kill another person.[3] I am in the position to pull a lever which

[2] One might think that an occupant of the runaway trolley has a special responsibility for what the trolley does. This thought might cloud our intuition about the case. For the sake of argument, let's suppose that this is false. If that thought affects your intuition, we can modify the case so that the dying passengers, I, and the control panel are all outside the runaway trolley.

[3] Quinn raises concern about the tension between the intuition to pull the lever in the Classic Trolley scenario (described below) and the DDA. In Quinn (1989a), it's argued that not pulling the lever also constitutes positive harmful agency. However, philosophers like Woollard (2012a), counter Quinn's position, suggesting that while refraining from pulling the lever isn't positive agency, Classic Trolley does not undermine the DDA. It suggests, at most, that *one* positive killing can outweigh allowing *five* persons to die. To sidestep controversy, my DDA examples involve killing one vs. letting *two* persons die. Modifying Classic Trolley to have only two people on the main track weakens the intuition for the permissibility of pulling the lever.

https://doi.org/10.1017/S0031819125000075 Published online by Cambridge University Press

will divert the trolley to save two persons from being killed at the cost of one person being killed. If I don't pull the lever, I exhibit negative harmful agency toward the persons being harmed by the trolley (*i.e.*, I will let them be harmed by the trolley), whereas if I pull the lever, I exhibit (non-intentional) positive harmful agency toward the person on the alternative track. Compared to Rescue inside the Trolley, it seems less justifiable to save *two* persons in this case, as doing so would require my positive harmful agency toward the one person on the track. However, as we will discuss below, if more people were on the main track, say 5 persons, my non-intentional positive harmful agency toward one could be justified by my attempt to save five. One might think that I could still pull the lever when there are two people on the main track. If so, we can construct scenarios where not diverting the trolley would result in only one death and one amputation. It is clearly permissible to continue treating the victims in the Rescue inside the Trolley whereas it is unclear whether it is permissible to turn the trolley when you have one versus two, or one versus one plus an amputation etc. That is enough to show a moral difference between negative harmful agency and positive harmful agency.

According to the doctrine of double-effect (DDE), it is harder to justify intentionally harming than non-intentionally harming. I am not concerned with the exact formulation of the doctrine. It is adequate to note that there is an intuition that there should be a stronger prohibition against intentional or direct harm than against non-intentional or indirect harm.[4] The following cases can serve as illustrations of the doctrine.

> *Classic Trolley*. There is a runaway trolley heading toward five persons. The trolley will kill the five persons unless it is diverted to an alternative track where it will kill a person. I am in a position to pull a lever which will divert the trolley to save those five persons
>
> *Footbridge*. I am standing on a bridge under which a runaway trolley is passing. I know that it will kill five people unless I

[4] In this paper, I will use 'direct agency' and 'intentional agency' interchangeably. Quinn (1989b) has famously argued that the DDE should be formulated in terms of direct agency, where 'harm comes to some victims, at least in part, from the agent's deliberately involving them in something to further his purpose precisely by way of their being so involved,' (p. 323) and not intentional agency, where harm is in fact intended. For an argument supporting the understanding of the DDE in terms of direct agency rather than intentional agency, see Tadros (2015).

https://doi.org/10.1017/S0031819125000075 Published online by Cambridge University Press

> push a very large man next to me over the bridge and onto the track. While pushing the large man causes him to die, I know that I can save those five lives by pushing him onto the track.

There is a strong intuition that I may pull the lever in Classic Trolley, whereas I may not push the large man in Footbridge. While some philosophers attempt to explain (or explain away) this intuition using notions other than intention, the DDE explains the intuition based on the difference between the cases in terms of intentional agency. In Classic Trolley, I have no intention to harm the person on the alternative track (I only intend to divert the trolley), but, in Footbridge, I intend to save those five people by pushing the large man. My intention includes harming the large man. I *assume* the DDE advocates are correct in their belief that the prohibition against pushing the large man is grounded in my intentional harmful agency toward him.

The DDA and the DDE are principles that assert a moral distinction between forms of harmful agency: positive and negative agency in the former, and non-intentional and intentional agency in the latter. However, the moral distinction between those forms of agency can be understood in either *agent-neutral* or *agent-relative* terms. In contrast to the agent-relative construal, which involves an essential reference to the agent, the agent-neutral construal of the DDA and the DDE does not make a moral distinction between the different forms of *my* agency and those of *any other agent*. To clarify the notion of agent-neutrality, we can adopt the following conceptualization:

> *Agent-Neutrality*: My ranking of different classes of actions or outcomes is fixed by descriptions of them that make no indexical reference to me.[5]

By Agent-Neutrality, we can interpret the DDA and the DDE as doctrines that advocate preferring negative harmful agency (of a given agent) over positive agency (of a given agent) and

[5] This formulation is borrowed from Setiya (2018). I interpret agent-neutral consequentialism as a form of consequentialism aligned with Agent-Neutrality. The distinction between agent-relative and agent-neutral values lies in whether an outcome has *value* relative to an agent. Any consequentialism that satisfies Agent-Neutrality adopts agent-neutral values. However, Howard (2022) contends that, since some *outcomes* are essentially individuated indexically, a consequentialism embracing agent-neutral values need not fulfill Agent-Neutrality (see also note 15).

https://doi.org/10.1017/S0031819125000075 Published online by Cambridge University Press

non-intentional harmful agency (of a given agent) over intentional harmful agency (of a given agent) when the harms are equal. It is easy to see that the DDA and the DDE, when construed agent-neutrally, can be incorporated into agent-neutral consequentialism.

While we might initially think that the DDE or the DDA should not naturally be understood agent-neutrally, a number of philosophers have invoked, either implicitly or explicitly, agent-neutral construals to support their arguments. For example, Bazargan-Forward (2014) argues that 'the intention/foresight distinction can have third-personal relevance: by committing φ, [agent 1] brings to fruition [agent 2]'s wrongful intentions, which is why it should receive as much weight in [agent 1]'s deliberations as an intentionally committed harm'. In the same vein, addressing a classic trolley case in which a trolley is heading toward five innocent people unless it is diverted to another track to which Y wrongly tied four people, Tadros (2016, p. 117) writes, 'If [the agent] is not permitted to turn the trolley, the five will be victims of Y's wrongdoing. If X is permitted to turn it, the four will be victims of Y's wrongdoing. But if [the agent] is not permitted to turn the trolley, and does not do so, Y will have done something morally equivalent to preventing the five from being saved. In contrast, if [the agent] is permitted to turn it, and does so, Y will have done something equivalent to killing the four. It is plausibly worse to kill a person than it is to prevent the person from being saved. If so, the reason against [the agent] being permitted to turn the trolley is more powerful than the countervailing reason for his being permitted'. (Implicit appeals to agent-neutral construals can be found in the works of Frowe *e.g.*, see Frowe 2014, p. 136).

The agent-neutral interpretations of the DDA and the DDE are not plausible, though. Let's first discuss DDE. Consider the following scenario.

> *Sabotaged Trolley*. Everything is like Classic Trolley, except that Killer who is the enemy of the person in the alternative track has sabotaged the trolley to run away toward five people, in anticipation of me pulling the lever and diverting the trolley to the alternative track. All along, Killer's plan was to cause a trolley accident to kill his enemy.

In Sabotaged Trolley, pulling the lever would result in carrying out Killer's plan, and thus, Killer's intentionally killing the man on the alternative track. Therefore, the consequences of pulling the lever include saving five people but also the intentional killing of one person by Killer. Compare this scenario with Footbridge, where pushing the large man involves one intentional killing by me

https://doi.org/10.1017/S0031819125000075 Published online by Cambridge University Press

and saving five people. If the DDE is construed agent-neutrally, in accordance with Agent-Neutrality, where outcomes are ranked without indexical reference to me, there is no moral difference between my intentional killing in Footbridge and Killer's intentional killing in Sabotaged Trolley. Thus, if the DDE is construed in an agent-neutral manner, Sabotaged Trolley would be morally akin to Footbridge, making pulling the lever in Sabotaged Trolley equally impermissible as pushing the large man off the bridge.

I find it very counterintuitive, however, to say that pulling the lever in Sabotaged Trolley is impermissible. It seems very strange to say that I may pull the lever in Classic Trolley, since the trolley is out of control spontaneously, but not in Sabotaged Trolley, since the trolley is sabotaged. The man on the alternative track in Sabotaged Trolley does not have a stronger claim not to be harmed than the man in Classic Trolley simply because the cause of his misfortune is his enemy rather than just pure ill fortune.

However, one might say that we have stronger reasons to prevent wrongs from *being done* than reasons to prevent bad things from *happening*, as humans perform evil deeds; nature does not. Deterring human transgressions can be beneficial for both the wrongdoers and society at large. But let's set the social benefits aside, and focus on the claim of the victim to be saved. Once an innocent non-threatening individual, through no action of their own, finds themselves ensnared in a wretched predicament, the account of how they got there is a piece of irrelevant history in relation to the strength of their claims to be saved on others. To hold it otherwise would be unfair to the victim. For example, consider two babies: one born with a fatal illness resulting from a random mutation, and the other afflicted with the same disease due to the mother's consumption of alcohol. If there were a pill that could only save one of them, would it not be unfair to assume that the second baby has a stronger claim to be saved? I shall refer to this intuition as the *irrelevance of historical origin*. The irrelevance of historical origin posits that the historical origin of the predicament of an innocent non-threatening individual is generally irrelevant when assessing the strength of their claim to be saved or protected from harm by those who bear no responsibility for creating said predicament. It would be unfair to victims if we allow our decision to be guided by the historical origin of the predicament.[6]

[6] The agent-neutral consequentialist might argue that the same intuition pulls us toward a view in which whether the agent would kill or simply

https://doi.org/10.1017/S0031819125000075 Published online by Cambridge University Press

The agent-neutral interpretation of the DDA can be similarly challenged in light of the irrelevance of historical origin. Consider Rescue inside the Trolley again. Intuitively, I'm not required to stop the trolley. The same intuition holds even if the movement of the trolley and the resulting death are caused by a technical mistake made by the operator. The cause of the movement of the trolley in this case is morally irrelevant to the permissibility of my action. I am still not obligated to abandon the rescue of the two lives to prevent a single killing. However, the outcome of not abandoning the rescue of two lives includes a harmful positive agency (by the operator), and thus the agent-neutral DDA goes against our intuition by implying that I ought to abandon the rescue of the two lives at the back to stop the trolley.

The arguments presented so far against the agent-neutral interpretations of the DDE and the DDA are arguments from harm. However, by appealing to the irrelevance of historical origin, we can also make arguments *from rescue* to reject those principles. To illustrate this, let us consider a scenario involving two runaway trolleys, one heading toward two innocent individuals, and the other toward one innocent individual. If the trolleys continue on their current path, they will result in the deaths of those innocent persons. Suppose an unpredictable accident causes the trolley aimed at two persons to go off course, while a technical error by an operator results in the deviation of the other trolley towards one person. Now, imagine that I have the choice to save either the two persons or the one person. Opting not to save the two persons would lead to an outcome where a killing by the operator occurs, while saving the one person would involve allowing two persons to die. According to the agent-neutral construal of the DDA, the preference would be to save the one over the two, as one killing is worse than two acts of letting die. However, the irrelevance of historical origin undermines the plausibility of this conclusion.

A similar argument from rescue can be given against the agent-neutral construal of the DDE. Suppose that instead of a technical error, the operator makes an intentional mistake with the aim of

---

refrain from saving, and whether the agent would strictly intend or merely foresee their death, is a piece of irrelevant background. However, I believe there can be a plausible deontic theory that doesn't fully succumb to this pull. In such a theory, while the ethical insignificance of the historical origin of a *victim*'s predicament is acknowledged, the character of *the agent's* response to the situation has ethical importance. Thanks to the anonymous reviewer of Philosophy for raising the concern.

https://doi.org/10.1017/S0031819125000075 Published online by Cambridge University Press

killing the one person. To make it more similar to Footbridge, let's imagine that instead of two individuals, there are now five innocent people on the second track. Similar to the Footbridge case, if I choose to rescue the five, it would result in one intentional killing, whereas if I help the one person, it would lead to five deaths. I find it highly plausible to believe that I should prioritize saving the five innocent persons, whereas the agent-neutral DDE suggests a different outcome.

The irrelevance of historical origin compels us to adopt an agent-relative interpretation of the DDA and the DDE. The agent-relative construal of the DDA asserts that in my evaluation of actions, *my* harming is harder to justify than my allowing harm, while the agent-relative interpretation of the DDE posits that *my* intentional harmful agency is harder to justify than my non-intentional ones. However, an agent-relative consequentialist might argue that these interpretations can be accommodated within consequentialism by using agent-relative rankings of outcomes.[7]

2. Against Agent-Relativity

This section presents arguments against a consequentialist agent-relative interpretation of the DDA and the DDE. Let me first clarify how I understand agent-relative consequentialism.

Any consequentialist theory acknowledges that the action that produces the best outcome is permissible. However, it is helpful

[7] Setiya (2018) argues that consequentialists can resist the pressure – such as that raised by scenarios like *Footbridge* – to adopt agent-relative consequentialism. His key claim is that, from an agent-neutral perspective, killing one person to save five is worse than five random killings. Assuming that *Sabotaged Trolley* involves a random killing by the villain, Setiya might contend that it differs from *Footbridge* because the latter, unlike the former, involves an intentional killing to save five lives. However, Setiya's argument does not undermine my rescue-based objections to the agent-neutral DDA and DDE. Furthermore, *Footbridge* can be slightly modified to avoid his response: suppose I push the large man, not to save five, but out of vengeance. In this version, both pushing the large man and diverting the trolley in *Sabotaged Trolley* involve non-instrumental, intentional killings. Yet, it still seems intuitive that diverting the trolley is permissible, while pushing the man out of vengeance is not. Moreover, Setiya's claim – that killing one to save five is worse than five killings – remains deeply counterintuitive (see Howard 2021 for a detailed critique).

https://doi.org/10.1017/S0031819125000075 Published online by Cambridge University Press

to differentiate between two types of consequentialism. *Substantive* consequentialism not only deems the action with the best outcome as permissible, but also asserts that the evaluative rankings of outcomes take explanatory priority over the deontic statuses of the actions that produce them (Portmore 2022). On the other hand, *notational* consequentialism does not make a claim about explanatory priority; rather, its objective is to demonstrate that any reasonable moral theory can be formulated within a consequentialist framework (Dreier 2011, 1993, Louise 2004). While there may be pragmatic reasons for formulating plausible moral theories as consequentialist (Schroeder 2017), I will specifically refer to substantive consequentialism throughout this paper when using the term 'consequentialism'.

How should we understand the priority claim? According to the priority claim, figuring out what our ends are takes precedence over figuring out what we should do. In other words, the theory tells us, first, what our ends should be and, then, derives the deontic status of the action by determining how well it can suitably achieve those ends. Portmore calls this 'the teleological approach' to moral constraints (forthcoming). According to teleologists, 'the permissibility of an agent's φ-ing depends on what they could achieve by φ-ing and how that compares to what they could achieve by instead refraining from φ-ing.' Given this broad understanding of the teleological approach, substantive consequentialism is teleological consequentialism.

Furthermore, Portmore observes that consequentialists are not required to claim that all agents must share the same set of ends. It is possible, for example, that my end, but not yours, must be to minimize the harm caused by my actions. If all agents must adopt an end, it is agent-neutral, while if only I must adopt an end, it is agent-relative. Since my preferences determine the ranking of my ends, agent-relative (teleological) consequentialism can be expressed in terms of the ends that I ought to prefer. Portmore (forthcoming) formulates agent-relative consequentialism as follows:

> Agent-Relative Consequentialism: For any subject S and any act available to them φ, S's φ-ing is morally permissible if and only if there is no available alternative act whose prospect they ought to prefer to that of their φ-ing.

Underlying the teleological approach is the '*action as production*' understanding of actions, according to which all there is to action, from a moral perspective, is its outcome (see Muñoz 2021

https://doi.org/10.1017/S0031819125000075 Published online by Cambridge University Press

for demonstrating that action as production is central to consequentialism, the original idea goes back to Korsgaard 2008, ch. 7). As Portmore puts it, 'actions are attempts to affect how the world goes' (Portmore 2022, forthcoming, 2011, p. 56). Moreover, consequentialism has a broad understanding of an act's outcome, which consists of the possible world that would be actual if the act were to be performed. Call this the *broad outcome*. It does not only include the causal consequences of the action but also the very fact that the action is performed.[8]

The understanding of action as production can be contrasted with '*action as processes*', where actions are viewed as processes unfolding over time. An action-process begins when the agent forms the intention to φ. The process continues when the agent is φ'ing, and the process is complete when the agent has φ'ed. The agent engages in various types of instrumental reasonings during the action-process to achieve her goal.[9] I will come back to the moral significance of these two contrasting understandings later.

Scenarios like Sabotaged Trolley pose challenges for the consequentialism that adheres to the agent-neutral DDE. On one hand, the broad outcomes of both options that the agent faces involve intentional agency. On the other hand, it seems that the intentional harmful agency of the third party in these scenarios is irrelevant. To disregard the intentional agency of the third party, consequentialists are compelled to argue that only *my* own intentional harmful agency holds moral significance in this case. While a number of philosophers have criticized appealing to agent-relative values as being excessively self-indulgent (Ridge 2001, 2009; Mack 1998; Persson 2013; Howard 2021), Portmore holds that there is nothing pejoratively self-centered about appealing to agent-relative values, as 'agents have a greater responsibility for what *they* do than for what *others* do.' (Portmore 2023).

However, this consequentialist solution is problematic not merely because of self-indulgence, but rather due to the possibility of scenarios where the third party involved – the party whose actions are distinct from the agent's action-process – could be the agent themself. That is, the misfortune of the people whom I am trying to save by intentionally killing someone else might actually be a result of

[8] The term broad outcome is borrowed from Muñoz 2021, with the same definition.

[9] For this conception of action, see Thompson (2008). Stout (2018) compiled papers addressing diverse facets of this conceptualization (see especially contributions by David Charles and Anton Ford).

https://doi.org/10.1017/S0031819125000075 Published online by Cambridge University Press

my own agency; nevertheless, this fact does not justify me killing an innocent person as a means to save others. Let me explain the gist of this objection by appealing to a bit of sci-fi fantasy.

> *Time-Traveler*: In the year 2000, as a time-traveler from the future (let's say 2030), I discover a disturbing situation. There is a Killer who is about to take the lives of five innocent people. The only way to prevent this is by sacrificing an innocent bystander, whom I'll refer to as A. If I choose to kill A, the Killer will be prevented from harming the five individuals. Conversely, if I refrain from killing A, the Killer will proceed with their plan, resulting in the deaths of those five people. And then I come to the chilling realization that the Killer is actually me in 2000.

I suppose that in Time-Traveler, I may not kill A. It does not matter that my 2000 self is about to kill 5, I can't use A to save those 5. To simplify the matter, I assume an *endurance theory*, wherein the person traveling from 2030 and the person living in 2000 are the *same person*. But the main point stands, *mutatis mutandis*, even if we assume a perdurance theory.[10] I am well aware that many moral philosophers hesitate to accept an argument based on intuitions related to time-travel scenarios. To address this concern, I will soon present an alternative argument that does not rely on such scenarios. However, the inclusion of Time-Traveler may effectively illustrate my points.

Consider my practical deliberation – from the point of view of my 2030 self – when deciding whether or not to kill one person in order to save five from being killed by my 2000 self. It may be the case that the fact that I am the cause of harm to those five individuals generates a special reason for me to protect them (so the cause of their misfortune is not completely irrelevant here). However, the special reasons I have to protect those five people are not strong enough to justify using one person as a means.

From *this* deliberative standpoint, it is not within my power to change the fact that my 2000 self is killing those five individuals. What lies within my control is the choice not to treat this innocent person as a mere means to save the others. Yet, agent-relative consequentialism fails to accommodate this conclusion. For I should prefer the broad outcome where I have killed one person rather than

[10] According to the perdurance theory, these two persons are different time-slices of a four-dimensional object.

https://doi.org/10.1017/S0031819125000075 Published online by Cambridge University Press

the broad outcome where I have killed five. It is preferable by me and others (except the one victim) that I kill fewer people. The special disvalue that prevents killing one to save five in other interpersonal cases can't do the work here because there aren't multiple agents to whom value can be indexed. Therefore, agent-relative consequentialism would deem it permissible for me to kill one to save five in this scenario.

The Time-Traveler scenario demonstrates two important points: *First*, the ranking of outcomes from my perspective does not necessarily align with the moral rules that determine which means I am allowed to choose when I practically deliberate about how to proceed to achieve my goals. While in my outcome ranking, the broad outcome that I killed 1 may appear preferable to the broad outcome that I killed 5, I do not have moral permission to choose this possible world to actualize. *Second*, what is morally significant is not that *I* exhibit harmful agency, but rather the realization that within *this* action-process, the use of the one victim as an instrument is not permissible.

To argue for the aforementioned two points without recourse to a time-travel scenario, we may avail ourselves of temporal variations of the scenarios discussed in the paper. Consider Setiya's *Murderous Footbridge* – a variation of the 'guilty agent' case originally introduced by Kamm (1996, p. 242). This scenario is like *Footbridge*, except that a villain has directed the trolley toward five people with the intent to kill them. If I choose not to push the large man, the villain will succeed in killing all five. However, if I push the large man, only one person will be killed – by me. Now consider this temporal version of *Murderous Footbridge*.

> *My Earlier Footbridge*: Everything is like the murderous footbridge, except that it was I who had directed the trolley at the five, intending to kill them. I now regret my action and recognize the possibility that I can stop the trolley by pushing a large man off the bridge onto the track in front of the trolley.

Morally speaking, My Earlier Footbridge is no different than Footbridge as both scenarios evoke the intuition that I should not kill the large man. He has a valid claim not to be used in a very harmful way against his will to save others. To kill an innocent person to make up for my past mistakes would be entirely self-indulgent. This scenario, however, has been addressed by agent-relative consequentialists, who have presented two categories of responses to yield the result that I should not kill the large man.

https://doi.org/10.1017/S0031819125000075 Published online by Cambridge University Press

(1) *Temporal-relativity*. Louise (2004) argues that agent-relativity and temporal-relativity often go together. That is, the agent-relative theorist can say that values are both agent-relative and temporal-relative, that is, killing one person *by me now* is worse than killing five persons by me at some other time. But this notion appears peculiar, as Persson (2013, p. 109) queries, 'why should what we are doing at the present time have this special importance?'. Portmore observes that there are cases where the agent's preference should lean towards refraining from infringing upon a constraint multiple times in the future rather than refraining from infringing upon a constraint in the present. Portmore illustrates this point through an example where I have made three promises to individuals A, B, and C. Upon realizing that keeping my promise to A now would necessitate breaking my promises to B and C later (possibly due to time constraints), it becomes intuitive that I may break my promise to A now in order to honor my commitments to B and C later. Examples such as this cannot be adequately accommodated within an agent-temporal relative framework. Discussing cases involving allowing oneself to cause harm, Persson (2013, pp. 103–107) also presents compelling arguments similar to Portmore's against the temporal relativity move. So, I assume the use of the agent-temporal relativity framework is misguided.

(2) *The end of not treating others as a mere means*. Portmore argues that consequentialism can yield the result that I should not kill the large man in *My Earlier Footbridge* without appealing to time-relativity (see Côté 2021 for a similar point). A consequentialist can say that an important goal of the agent is to never treat others as mere means. In *My Earlier Footbridge*, 'with respect to the five, [I don't] face the choice of whether to treat them as mere means, for [I have] already done that and nothing [I] can do now can change that. By contrast, with respect to [the large man], [I do] face the choice of whether to treat [him] as a mere means to preventing [myself] from being the murderer of the five.' (Portmore, forthcoming). For the sake of argument, let's suppose that Portmore's has provided an adequate consequentialist response to *My Earlier Footbridge*. The problem is that the same kind of response is not available to the future-oriented variant of Murderous Footbridge. Consider this scenario:

> *My Later Footbridge*: I am aware of my multiple personality disorder, which leads to episodes during which my alternate personality makes decisions contrary to my values. Let us suppose that my unwelcome personality maintains sufficient

https://doi.org/10.1017/S0031819125000075 Published online by Cambridge University Press

> psychological connectedness to my present self, ensuring continuity even within the neo-Lockean account of personal identity. As I am on the verge of an episode, I possess ample evidence to justifiably believe that I will intentionally direct the trolley towards five individuals, intending to kill them. However, if I were to push the large man off the bridge now, the track would be blocked, preventing me from carrying out the act of killing five later.

Morally speaking, there is no difference between Footbridge, My Earlier Footbridge, and My Later Footbridge. I should not kill the one man to prevent five killings. Yet, contrary to what Portmore said regarding My Earlier Footbridge, here I have not yet treated the five as mere means and by killing one now I can avoid treating them as a mere means later.

Addressing scenarios involving future-killing (offered by Sturgeon 1996), Portmore says 'it seems even more difficult – perhaps, impossible – to imagine a case in which the only way that I can ensure that I do not *voluntarily* commit two murders in the future is to commit one murder now. If these two future murders are indeed going to be murders that I voluntarily commit, then it must be that I can refrain from committing them—that is part of what it means for them to be voluntary. And if I can refrain from committing these two murders in the future even after refraining from committing murder now, then I should refrain from murder both now and in the future'. But in My Later Footbridge, the killing of five by me later is voluntary, it's just that I have no control over my decisions during my episodes. Consequently, it is not within my control to do anything to actualize the possible world in which I don't kill one now, and I don't kill five later.

Not only do I retain the same body during my episodes, but I also maintain sufficient psychological connectedness to my personality throughout. This ensures that, according to any account of personal identity, I remain the same person during my episodes. One might argue that although I maintain personal identity, during my episodes, my deliberative standpoint undergoes such a significant shift that it resembles that of a different moral agent. And this is exactly the point of the example. While in agent-relative consequentialism values are tied to the agent's personal identity, what ultimately determines the permissibility of my actions is the deliberative standpoint from which this action-process unfolds. Time-travel or temporal shift scenarios such as Time-Traveler, My Earlier or Later Footbridge, depict cases where the agent's deliberative

https://doi.org/10.1017/S0031819125000075 Published online by Cambridge University Press

standpoints vary while their personal identity does not. In such instances, only the former holds moral significance for determining the permissibility of actions, whereas agent-relative consequentialism mistakenly assigns importance to the latter. Moreover, it may well be the case that in my later footbridge, I prefer the broad outcome in which I kill fewer people, yet, within my deliberative framework, I have no moral permission to choose the possible world I ought to prefer the most to actualize.

To challenge the agent-neutral interpretations of the DDA and the DDE, I presented two structurally similar *rescue-based arguments* in the previous section, involving scenarios where there were two runaway trollies, one approaching an innocent individual on one track and the other heading towards multiple innocent people on the other track. I had to make a choice between rescuing one person on the first track or several individuals on the second track. By constructing similar rescue-based arguments, we can further reinforce the rejection of the agent-relative interpretations of the DDA and the DDE.

For the sake of brevity, I will focus on the rescue argument specifically related to the DDE (the argument against the agent-relative DDA would be similar). In this argument, an unpredictable accident causes one trolley to deviate towards five innocent people, while the second trolley is directed towards one person intentionally by the operator—who happens to be me. Now, the question arises: which individuals should I rescue, the one person or the group of five?

Considering that I would bear direct responsibility for the death of the one person, one might argue that it could be permissible for me to allow the five individuals to perish and attempt to save only the one person on the other track. However, it is also plausible to believe that it would be equally permissible for me to save the group of five on the track. In other words, I have moral permission to rescue one, but I have also moral permission to opt for saving the five. But the agent-relative DDE cannot account for the latter permission. If I save the five, I have intentionally killed one, but the agent-relative DDE implies it is very hard to justify my intentional killing, even when it is required to save five. One might think that in light of my direct responsibility for the death of one person, it is not obvious that it is permissible for me to rescue the five. However, agent-relative consequentialism should leave it clear that the agent must save the one. Nonetheless, it appears quite plausible that the obligation to save the one (if there is any) is not obvious at all. The lack of clear intuition by itself indicates that the agent-relative DDE

https://doi.org/10.1017/S0031819125000075 Published online by Cambridge University Press

mistakenly assumes that this is, morally speaking, a case morally akin to Footbridge.

Neither of the two strategies previously adopted by agent-relative consequentialism, as discussed in the context of My Earlier Footbridge, adequately accounts for the intuition that saving the five is permissible in this case. When considering time-relativity, choosing to save the five would involve my intentionally killing one person *now*.[11] Furthermore, the strategy of avoiding treating others as mere means does not apply here. Regardless of my choice, I would not be treating anyone as a mere means anymore (and I can't change the fact that I treated one as a mere means before). The options I face now are either the intentional killing of one person by me or the allowing of five to die; my intentional killing would be worse according to the agent-relative DDE. Therefore, the agent-relative DDE appears to conflict with the permissibility of saving 5 in this case.

However, it is possible to reconcile the DDE with the permissibility of saving 5 in this case by adopting a non-consequentialist understanding of it. According to this interpretation, the DDE serves as a governing norm for practical deliberation during an action-process, requiring a higher level of justification for actions that harm others instrumentally to achieve our ends. In the context of our case, neither rescue option involves treating individuals as

[11] One might think that the decision I face *now* is whether to let myself kill and that it is not clear where my act of killing should be temporally located. It might be thought that I performed the action in the past even if it is only now that *that action* becomes a killing (thanks to the anonymous reviewer for this objection). However, since I, *ex hypothesi*, have *control* over not killing the one, it seems that saving the five would entail my killing the one person now. To illustrate, imagine that last night I placed a letter into my private mailbox to be collected by the mailman this morning. Intuitively I mailed the letter this morning - and not last night. While the physical act occurred last night, it is only when combined with the omission of not retrieving the letter during the night that it becomes an act of mailing in the morning. Similarly, as Persson (p. 106) remarks about a similar case, '[my] letting him be killed (by [myself]) secures the *completion* of the process that is [my] killing of him which started with [my directing the trolley towards him] since this act would not be a killing without [my] letting it to be a killing'. The earlier act compounds into an act of killing only when it is combined with the omission not intervening (to save the five). So, it is not implausible to say that the killing is occurring now – at the time of omission. However, if one argues that the time of killing is uncertain, temporal relativity may be invoked to justify the permissibility of saving the five. Nevertheless, as argued previously, temporal relativity lacks plausibility.

https://doi.org/10.1017/S0031819125000075 Published online by Cambridge University Press

mere means, and thus the DDE does not prohibit either choice. In the subsequent section, I will present arguments elaborating on how this interpretation conflicts with teleological consequentialism.

3. Process-Oriented Principles

Deontological norms like the DDA or DDE, I argue, should be viewed as norms governing action-processes. In the previous section, I distinguished between the 'action as production' and 'action as process' conceptions of action. Teleological consequentialism perceives actions as broad outcomes, while deontology views actions as processes. The clearest difference between these two conceptions arises when considering an agent's *inaction*. Let's take the example of Murderous Footbridge. Suppose the agent decides *not* to push the large man. The broad outcome of this action is five killings. Identifying the act of not pushing the man with the broad outcome amounts to identifying that action as the act of producing five killings. However, when the agent chooses not to push the large man, they do nothing. There is no action-process in which the agent is involved. The DDA or DDE, understood as norms governing action-processes, typically do not apply to an agent's inaction.

The DDE may apply to inactions in cases where one's inaction becomes part of a larger action-process. For instance, it is impermissible to refrain from giving someone an antidote to a disease in order to observe the progress of the disease and devise a remedy for many others (Quinn 1989b). Similarly, withholding food from a starving person with the intention of allowing their death for organ donation is also impermissible.[12] In these instances, inaction is the conclusion of a practical deliberation regarding the larger action and its goal, and the DDE applies to this inaction as it serves as a means to the intended outcome. However, generally, inactions, such as not pushing the large man, are not the product of practical inference in a larger action-process.

The DDA and the DDE are not about identifying the best broad outcomes, rather they govern practical inference within an action-process, pertaining not only to outcomes but also to means I may choose to achieve my ends. Herman (2014) observes that 'the framework for choice on a Kantian view isn't directly about [possible

[12] Thanks to the anonymous reviewer for drawing my attention to these cases.

https://doi.org/10.1017/S0031819125000075 Published online by Cambridge University Press

outcomes] but about the duties and obligations that bear. These duties and obligations place demands on agents' reasoning, laying out connections between moral premises and possible actions that follow'. To see how the DDE should be understood in accordance with Herman's Kantian perspective, let's examine the agent's reasoning in the scenario of *Murderous Trolley*. Their practical deliberation might proceed as follows: '(1) I desire to save those five lives (or I desire to prevent five killings). (2) The only means of accomplishing this is to push the large man off the bridge. (3) Hence, let's push him off the bridge.' This reasoning holds instrumental validity but lacks moral validity. The DDE requires a stringent justification for intentionally and opportunistically killing an individual as a means, which is not met by the end of saving five lives. Deontological principles like the DDA and the DDE determine what types of practical inferences can be morally valid.[13]

However, consequentialists have countered this line of thinking by asserting that such rules reflecting duties and obligations are themselves bare appeals to intuition (Smith, 2003) unless they can be provided with plausible rationales, and have argued that the need for such a plausible rationale leads ineluctably to an outcome-centered rationale. But it is important to note that a value-based rationale for deontological principles does not need to collapse into a *broad* outcome-centered rationale. For instance, we can consider the DDA and the DDE as different manifestations of the Kantian idea of dignity, guiding our treatment of others during action-processes. They specify the extent to which we should regard others as, to use Kamm's term (1996), inviolable in our practical inference. *Within our practical deliberation,* individuals must be regarded as possessing a certain degree of inviolability against being used as a means, and a lesser degree of inviolability against non-intentional harm.

But can't we consequentialize this view of deontological principles? Of course, we can. To accommodate agent-relative values, the consequentialist allows for rankings to be relative to agents, *i.e.*, agents can have their own rankings of outcomes. To accommodate process governed norms, on the other hand, the consequentializer needs to use process-centered rankings. Action-processes are goal-directed events that happen in a certain time and place. They differ from each other based on their goals, their times, and their places. By

[13] In the same vein, according to Herman (2022, Ch. 1), the formula of universal law determines whether a practical inference is morally valid.

https://doi.org/10.1017/S0031819125000075 Published online by Cambridge University Press

allowing rankings to be relativized with respect to action-processes (individuated by their agent's intention, time and locations), we would have a more fine-grained view than simply relativizing to agents or time. Such a view can yield correct results in examples discussed in the previous section.[14] But the problem is that while we might have pragmatic or notational incentives to consequentialize process-centered deontological norms, the resulting theory would not be teleological consequentialism.

Teleological consequentialism holds that broad outcomes should be first ranked based on their *preferability to the agent*, and the action that produces the best outcome is always permissible. Nevertheless, it is conceivable, as demonstrated by my examples where there are multiple action-processes, but there are not multiple agents to index value to, that the ranking of broad outcomes, based on what the agent ought to prefer, deviates from the deontic ranking of actions. If such divergences can occur, it is possible to conceive a plausible deontological view that is not teleological, in the sense that the right action is not determined by the preferability of the broad outcome to the agent. While teleological consequentialism is all about identifying the best broad outcome with respect to an agent, deontological principles are also about whether the means taken, within an action-process, are morally problematic.[15]

[14] There is a worry, however, that yielding the correct results is not enough to consequentialize a deontological view. What is distinctive about Kant's view is that 'an action from duty has its moral worth not in the purpose to be attained by it but in the maxim in accordance with which it is decided upon' (Kant, 4:400. As Hurley correctly notes that 'If the agent refrains but not for the right reasons, she does not act as duty requires on the Kantian account, she only performs an action of the same type—merely in accordance with duty. This is a fundamental form of *deontic* failure for Kant, but the [consequentialist] counterpart fails to capture either such fundamental deontic evaluations of actions or the features relevant to them in its ranking of outcomes, hence it fails to provide deontic equivalents to such deontic failures' (Hurley 2022).

[15] While my objections primarily target views like Portmore's, Nathan Howard (2022) offers a consequentialist view that, if successful, could get past my objections. If Howard's approach succeeds, consequentialists can incorporate the DDA and the DDE within an agent-neutral framework. He argues that this can be achieved by recognizing that some outcomes are essentially first-personal (*de se*) – private and inaccessible from a third-person perspective. Consider Time-Traveler: I must choose between killing one person or allowing my past self, Killer, to kill five. According to Howard, the outcome of my killing one is first-personal, inaccessible to

https://doi.org/10.1017/S0031819125000075 Published online by Cambridge University Press

The core idea, often referred to as *the compelling idea*, of consequentialism is that 'it is always morally permissible for an agent to perform the option that would bring about the outcome that they ought (or that it is fitting for them) to prefer to every available alternative' (Portmore 2022). But examples of the last section demonstrated the possibility that the ranking of the broad outcomes from my perspective does not necessarily correspond to moral rules that determine which means I may choose when I practically deliberate about how to proceed to achieve my goals.

When we view actions as just productions, the core idea of consequentialism looks almost irrefutable. If there is nothing to actions other than the production of broad outcomes, how can producing the best outcome be impermissible? However, when we individuate actions as a process which incorporates certain instrumental reasoning, the core idea of consequentialism does not appear as compelling anymore.[16] Consider the case of my Earlier Footbridge, where I deliberate whether to push the man. The broad outcome of my inaction includes the fact that I kill five people. However, my action of killing five people, which was started earlier, is distinct from that act I presently deliberating about, which is subject to deontological principles. While my inaction does not violate the DDE, my pushing the large man would. If deontological principles govern actions, understood as processes, it looks much more questionable to think

---

Killer. From my standpoint, killing one results in one first-personal kill, while letting Killer kill five results in five third-personal kills. If a first-personal kill is worse than a third-personal one, I should let Killer kill five, even though Killer is still me. However, Howard's view faces a challenge: While Killer's *de se* thought – expressed by 'I killed five' – is private in the sense that I cannot entertain it, I am *aware that* Killer has such a thought. Much like how *evidence of private evidence* grants me indirect epistemic access to that private evidence, my awareness that there is a private outcome renders that outcome indirectly accessible to me. For Howard's view to work, it must assume that in my ranking of outcomes, directly known private outcomes take precedence over those indirectly known. This raises a challenge: How can Howard consistently discount the de se thoughts of other agents? In teleological consequentialism, the broad outcome – prioritized for its explanatory role – encompasses not only the state of affairs involving my own de se thoughts, but also those of other agents, though my knowledge of those states is acquired through different methods: directly in my case and indirectly in theirs.

[16] In a similar vein, Hurley (2017) argues that although 'the compelling idea' may initially appear compelling, it might lose its force when we distinguish between the best outcome and the best option.

https://doi.org/10.1017/S0031819125000075 Published online by Cambridge University Press

that the broad outcome of my inaction irrefutably determine the deontic status of my action.

The fundamental idea of deontology is that, although I may have a preference, or ought to have a preference, to bring about a certain possible world, I am not at liberty, morally speaking, to choose that possible world to actualize. Deontological principles, such as the DDA or the DDE serve as constraints on whether we can actualize the best possible outcome, whether it be broad or causal.

References

Saba Bazargan-Forward, 'Moral coercion', *Philosophers' Imprint*, *14*.11 (2014), 1-18.

Nicolas Côté, 'A Diachronic Consistency Argument for Minimizing One's Own Rights Violations', *Ethical Theory and Moral Practice* 24.5 (2021):1109-1121.

Stephen Darwall, 'Agent-centered restrictions from the inside out', *Philosophical Studies*, *50*.3 (1986), 291-319

James Dreier, 'Structures of Normative Theories', *The Monist* 76 (1993): 22–40.

James Dreier, 'In Defense of Consequentializing.' In *Oxford Studies in Normative Ethics*, vol. 1, edited by Mark Timmons, 97–119. (Oxford: Oxford University Press, 2011).

Richard Feldman, 'Reasonable Religious Disagreements,' in *Philosophers Without God: Meditations on Atheism and the Secular Life*, (Oxford University Press, 2007), p.208.

William FitzPatrick, 'The doctrine of double effect: Intention and permissibility.' *Philosophy Compass* 7.3 (2012): 183-196.

Philippa Foot, 'The Problem of Abortion and the Doctrine of the Double Effect', *Oxford Review* 5 (1967), 5–15.

Helen Frowe, *Defensive killing*. (Oxford University Press, 2014).

Barbara Herman, 'Impermissibility and Wrongness', *Norton Introduction to Philosophy*, ed. Gideon Rosen et al., (W. W. Norton & Company, 2014)

Barbara Herman, *Kantian Commitments: Essays on Moral Theory and Practice*. (Oxford: Oxford University Press, 2022)

Christopher Howard, 'Consequentialists Must Kill'. *Ethics* 131.4 (2021), 727-753.

Nathan Robert Howard, 'Consequentialism and the Agent's Point of View', *Ethics*, *132*.4 (2022), 787-816.

Paul Hurley, 'Why Consequentialism's "Compelling Idea" Is Not'. *Social Theory and Practice* 43.1 (2017), 29-54.

Paul Hurley, 'The Consequentializing Argument Against... Consequentializing?' *Oxford Studies in Normative Ethics* 12 (2022), 253-275.

https://doi.org/10.1017/S0031819125000075 Published online by Cambridge University Press

Shelly Kagan, *The Limits of Morality* (Oxford: Oxford University Press, 1989).
Immanuel Kant, I. *Groundwork of the metaphysics of morals*, (Oxford University Press, 1785/2019).
Francis Myrna Kamm, *Morality, Mortality. Vol.2: A, Rights, Duties, and Status*, (New York: Oxford University Press, 1996).
Christine M Korsgaard, *The Constitution of Agency: Essays on Practical Reason And Moral Psychology*. (New York: Oxford University Press, 2008).
Jennie Louise 'Relativity of Value and the Consequentialist Umbrella', *Philosophical Quarterly* 54 (2004): 518–536.
Eric Mack, 'Deontic Restrictions are Not Agent-Relative Restrictions.' *Social Philosophy and Policy*. 15 (1998), 61–83.
Daniel Muñoz, 'The Rejection of Consequentializing', *Journal of Philosophy* 118.2 (2021), 79-96.
Ingmar Persson, *From* Morality to the End of Reason*: An Essay on Rights, Reasons, and Responsibility* (Oxford: Oxford University Press, 2013)
Douglas W Portmore, 'Consequentializing Moral Theories', *Pacific Philosophical Quarterly* 88 (2007), 39–73.
Douglas W Portmore, *Commonsense Consequentialism: Wherein Morality Meets Rationality*. (Oxford: Oxford University Press, 2011).
Douglas W Portmore, 'Consequentializing', *The Stanford Encyclopedia of Philosophy* (Fall 2022 Edition), Edward N. Zalta & Uri Nodelman (eds.), URL = https://plato.stanford.edu/archives/fall2022/entries/consequentializing/
Douglas W Portmore, 'Consequentializing agent-centered restrictions: A Kantsequentialist approach', *Analytic Philosophy*, 64.4 (2023), 443-467.
Warren S Quinn, 'Actions, Intentions, and Consequences: The Doctrine of Doing and Allowing' *The Philosophical Review*, 98.3 (1989a), 287-312
Warren Quinn, 'Actions, Intentions, and Consequences: The Doctrine of Double Effect', *Philosophy and Public Affairs 18*.4 (1989b), 334-351
Michael Ridge, 'Agent-Neutral Consequentialism From the Inside-Out: Concern for Integrity Without Self-Indulgence', *Utilitas*, 13.2 (2001), 236–254.
Michael Ridge, 'Consequentialist Kantianism', *Philosophical Perspectives*, 23 (2009), 421-438
Samuel Scheffler, *The Rejection of Consequentialism* (Oxford: Oxford University Press, 1982).
S Andrew Schroeder, 'Consequentializing and Its Consequences', *Philosophical Studies*, 174.6 (2017), 1475–1497.
Kieran Setiya, 'Must Consequentialists Kill?' *Journal of Philosophy* 115.2 (2018), 92-105.
Nicholas L Sturgeon, 'Anderson on Reason and Value', *Ethics* 106 (1996), 509–524.
Victor Tadros 'Wrongful Intentions Without Closeness', *Philosophy and Public Affairs*, 43 (2015), 52–74.

https://doi.org/10.1017/S0031819125000075 Published online by Cambridge University Press

Victor Tadros, 'Permissibility in a World of Wrongdoing', *Philosophy and Public Affairs*, *44*.2 (2016), 101-132.
Fiona Woollard, 'The Doctrine of Doing and Allowing I: Analysis of the Doing/Allowing Distinction', *Philosophy Compass*, 7.7 (2012a), 448-458.
Fiona Woollard, 'The Doctrine of Doing and Allowing II: The Moral Relevance of the Doing/Allowing Distinction', *Philosophy Compass*, 7.7 (2012b), 459-469.

AMIR SAEMI (saemi.amir@gmail.com) *is Professor of Institute for Research in Fundamental Sciences (IPM).*

https://doi.org/10.1017/S0031819125000075 Published online by Cambridge University Press