

# A Robust Defence of the Doctrine of Doing and Allowing

XIAOFEI LIU

*University of Missouri*

Philosophers debate over the truth of the Doctrine of Doing and Allowing, the thesis that there is a morally significant difference between doing harm and merely allowing harm to happen. Deontologists tend to accept this doctrine, whereas consequentialists tend to reject it. A robust defence of this doctrine would require a conceptual distinction between doing and allowing that both matches our ordinary use of the concepts in a wide range of cases and enables a justification for the alleged moral difference. In this article, I argue not only that a robust defence of this doctrine is available, but also that it is available within a consequentialist framework.

## I. INTRODUCTION

The Doctrine of Doing and Allowing (DDA) states that there is a significant moral difference between doing harm and merely allowing harm to happen. The alleged moral difference is widely understood as the ‘asymmetrical constraints’: it requires achievement of a much greater moral good for a harm-doing to be permissible than for a harm-allowing to be permissible.<sup>1</sup> Deontologists tend to accept DDA, whereas consequentialists tend to reject it.<sup>2</sup> In this article, I argue not only that DDA is defensible but also that it can be defended in a consequentialist approach.

Those who reject DDA believe that there is no conceptual distinction between doing and allowing that both matches our ordinary use of the concepts in a wide range of cases and enables a justification for the alleged moral difference.<sup>3</sup> Thus, any robust defence of DDA has to meet

<sup>1</sup> David McCarthy, ‘Harming and Allowing Harm’, *Ethics* 110 (2000), pp. 749–79, at 749. There are also other interpretations of DDA. One interpretation is that doing harm makes an agent more reprehensible, which seems to be the interpretation assumed in James Rachels, ‘Active and Passive Euthanasia’, *New England Journal of Medicine* 292 (1975), pp. 78–80, especially in his discussion of the famous Smith–Jones case. Another interpretation is that doing harm constitutes a greater moral evil than allowing harm. This interpretation is assumed in Judith Lichtenberg, ‘The Moral Equivalence of Action and Omission’, *Canadian Journal of Philosophy* 8 (1982), pp. 19–36; and also in Winston Nesbitt, ‘Is Killing No Worse than Letting Die’, *Journal of Applied Philosophy* 12 (1995), pp. 101–5. The ‘Asymmetrical Constraints’ interpretation, as I will show, is actually equivalent to the ‘greater moral evil’ interpretation.

<sup>2</sup> For challenges to DDA, see Rachels, ‘Active’; Michael Tooley, *Abortion and Infanticide* (Oxford, 1983); and Lichtenberg, ‘Moral Equivalence’.

<sup>3</sup> For a helpful survey of various defences of DDA and their problems, see Alastair Norcross’s introduction to *Killing and Letting Die*, 2nd edn., ed. B. Steinbock and A. Norcross (New York, 1994), pp. 1–23; and Frances Howard-Snyder, ‘Doing vs. Allowing

these two challenges: first, to provide a conceptual distinction between doing and allowing that matches our ordinary use of the concepts in a wide range of cases; second, to show that this conceptual distinction also enables a justification for the alleged moral difference between doing and allowing.<sup>4</sup> The second challenge is especially compelling, because the lack of a completely satisfactory conceptual distinction by no means suggests that a crude distinction between doing and allowing does not exist – after all, many of our deeply held concepts resist precise analysis (e.g. knowledge), yet the lack of a satisfactory justification for the alleged moral difference would be fatal to DDA. In this article, I propose a value-based analysis of DDA that successfully meets both challenges.

Although I intend a robust defence, the version of DDA that I defend here is a moderate one – other things being equal, an act<sup>5</sup> of doing a certain degree of harm to a certain kind of well-being requires a greater moral good to justify it than an act that allows the same degree of harm to that same kind of well-being. For example, the version of DDA that I defend would say that killing one innocent human being requires more to justify it than allowing one innocent human being to die. However, I leave it open whether killing one innocent human being requires more to justify it than allowing, say, five innocent human beings to die.

## II. THE RIGHTS-BASED ANALYSIS

Various attempts have been made to justify DDA.<sup>6</sup> The one that holds the most promise to meet both challenges is the rights-based analysis.<sup>7</sup> In what follows, I will examine two rights-based analyses, one by Philippa Foot and one by Kai Draper. I will show why neither analysis is

Harm', *Stanford Encyclopedia of Philosophy*, <<http://plato.stanford.edu/entries/doing-allowing/>> (2007).

<sup>4</sup> The doctrine can also be defended in a less robust way: even if there is no sound conceptual distinction that matches our ordinary use of the concepts in a wide range of cases, the doctrine can still be defended by showing that there is at least a justification for the alleged moral difference between cases that are typically regarded as harm-doing and cases that are typically regarded as harm-allowing.

<sup>5</sup> By 'act', I mean any intentional behaviour. I understand 'act' in a very loose sense here: an intentional non-performance can be called an *act* of allowing harm.

<sup>6</sup> For proposals of the conceptual distinction, see Jonathan Bennett, 'Negation and Abstinence: Two Theories of Allowing', *Ethics* 104 (1993), pp. 75–96; Foot, 'Killing and Letting Die', *Killing and Letting Die*, pp. 280–9; Alan Donagan, *The Theory of Morality* (Chicago, 1977); Warren S. Quinn, 'Actions, Intentions, and Consequences: The Doctrine of Doing and Allowing', *Killing and Letting Die*, pp. 355–82; and Jeff McMahan, 'Killing, Letting Die and Withdrawing Aid', *Killing and Letting Die*, pp. 383–420. For justifications of the alleged moral difference, see Philippa Foot, 'Euthanasia', *Philosophy and Public Affairs* 6 (1977), pp. 85–112; Nesbitt, 'Letting Die'; McCarthy, 'Harming'; Samuel Scheffler, 'Doing and Allowing', *Ethics* 114 (2004), pp. 215–39; and F. M. Kamm, *Intricate Ethics: Rights, Responsibilities, and Permissible Harm* (Oxford, 2006).

<sup>7</sup> See, for example, Foot, 'Euthanasia', 'Killing'; and Quinn, 'Actions'.

satisfactory. Drawing lessons from these analyses, I propose a different analysis in the next section.

Let us start with Foot's analysis. First, to explain the moral difference between doing and allowing, Foot appeals to the distinction between negative rights, which she calls *rights to non-interference*, and positive rights, *rights to goods or services*.<sup>8</sup> A violation of negative rights involves initiating a harmful sequence or sustaining an existing one. Setting a building on fire, for example, is initiating a sequence harmful to its residents, and adding gasoline to a fire that is about to go out is sustaining a harmful sequence. A violation of positive rights, on the other hand, involves merely allowing a harmful sequence to complete. Negative rights, according to Foot, are morally more important than positive rights. Furthermore, on Foot's view, the conceptual distinction between doing and allowing perfectly corresponds to the distinction between these two types of rights violation. Doing harm initiates or sustains a harmful sequence and constitutes a violation of negative rights; whereas allowing harm merely allows the completion of a harmful sequence and constitutes a violation of positive rights.<sup>9</sup> Since negative rights are more important than positive rights, doing harm is harder to justify than allowing harm – this is the justification of the moral difference between doing and allowing according to Foot's rights-based analysis.

Leaving aside whether Foot's analysis of the conceptual distinction is satisfactory, her attempt to justify the moral difference is not. On Foot's view, an interference is harder to justify than a withholding of goods or services, but this is not because the former is a greater evil than the latter. Foot admits that Smith's killing his cousin might be just as evil as Jones's watching his cousin die.<sup>10</sup> The difference is simply because of the nature of the two types of rights: negative rights are simply morally more important than positive rights. But there is a legitimate question: why are negative rights necessarily more important than positive rights? Foot does not give a theoretical explanation; she simply appeals to our moral intuition as illustrated by a few cases.<sup>11</sup> However, this appeal to our intuition in a few cases seems unsatisfactory. It is possible, as Alastair Norcross contends, that our intuition that my right, say, not to be poisoned is stronger than my right to be given food to survive may well derive from the intuition that killing is worse than letting die.<sup>12</sup> Moreover, one may ask why violating my positive right to

<sup>8</sup> Foot, 'Killing', p. 284.

<sup>9</sup> Foot, 'Morality, Action and Outcome', *Morality and Objectivity: A Tribute to J. L. Mackie*, ed. T. Honderich (London, 1985), pp. 22–38, at 24.

<sup>10</sup> Foot, 'Killing', p. 286.

<sup>11</sup> One example Foot uses is the wounded soldier case. See Foot, 'Euthanasia', p. 100.

<sup>12</sup> Norcross, *Killing and Letting Die*, p. 16.

adequate nutrition by destroying a piece of food that I need for survival is not as morally offensive as violating my negative property right to a comparable piece of food that I need for survival. It is not obvious that the negative right here is necessarily more important than its corresponding positive right. Without an explanation of precisely what is so unique about negative rights and why it makes them morally more significant, the alleged moral difference between doing and allowing remains a mystery. Foot's analysis is thus *explanatorily incomplete*.

Kai Draper has recently developed a different analysis of why negative rights are in general stronger than positive rights.<sup>13</sup> He takes it that there are two basic moral values: the value of well-being and the value of autonomy. The value of well-being constitutes the foundation of positive rights, while the value of autonomy, together with the value of well-being, constitutes the foundation of negative rights. According to Draper, it is the value of autonomy that explains why negative rights, such as rights of self-ownership, are usually stronger than positive rights: positive rights are

typically more opposed to the value of autonomy than the typical negative right; for in requiring the performance of a specific sort of action, positive rights close off all alternative actions, whereas the typical negative right merely requires the nonperformance of a specific sort of action and so leaves open some number of alternatives.<sup>14</sup>

Draper understands autonomy as consisting in the self-mastery of one's own person and the freedom from control by others. All rights, thus, are opposed to autonomy or self-mastery of other people, because these rights give the right-holder authority over what is otherwise within the domain of others' self-mastery and close off their alternative actions. The moral difference between negative rights and positive rights, according to Draper, consists in the different constraints that they impose on autonomy: positive rights typically impose a greater constraint on autonomy because they demand a specific use of the property or body of another person, whereas negative rights do not typically impose such constraint on others' self-mastery. Since we value autonomy, the greater the constraint upon autonomy, the less

<sup>13</sup> Kai Draper, 'Rights and the Doctrine of Doing and Allowing', *Philosophy and Public Affairs* 33 (2005), pp. 253–80. Although Draper agrees that the rights-based theory provides an explanation of the alleged moral difference between cases that are typically regarded as harm-doing and cases that are typically regarded as harm-allowing, unlike Foot, Draper does not think the rights-based theory helps to establish a conceptual distinction between doing and allowing that matches our ordinary use of those concepts in a wider range of cases. So, Draper urges us to jettison DDA for a pure rights-based approach.

<sup>14</sup> Draper, 'Rights', p. 277.

morally desirable the rights – this explains why negative rights are more important.

However, Draper's explanation seems unsatisfactory too. While a positive right imposes a greater constraint on one's autonomy by closing off her alternative actions, it also provides additional protection for the right-holder's well-being by requiring others' assistance. So, it is unclear why the decrease in one's autonomy cannot be balanced out by the additional protection of another's well-being. Moreover, assuming that every person gets the same set of rights, while others' positive rights decrease one's autonomy, one's own positive rights provide additional protection of one's own well-being. So, why cannot the decrease in autonomy be balanced out by the additional protection of one's own well-being?<sup>15</sup> One answer might be that autonomy value is more important than well-being value. But, the same kind of problem facing Foot's analysis re-emerges – exactly why autonomy is more important than well-being. Why, for instance, is the burden of having to throw a life-preserver when another person is drowning not balanced out by the benefit of having someone throw you a life-preserver when you are drowning? A further explanation is definitely required. Thus, Draper's analysis is also explanatorily incomplete.

### III. A VALUE-BASED ANALYSIS: THE FIRST APPROXIMATION

However, Draper's discussion of the relation between rights and values sheds light on a different answer to the alleged moral difference. Like Draper, I think there are two basic types of moral value relevant to rights: autonomy and well-being. We have an intuitive understanding of what autonomy is: my autonomy consists in my exclusive authority over my life, body, property and whatever else that naturally falls in the domain of self-mastery. Autonomy is valuable because we value the self-mastery of our own person – the freedom to do whatever is within that domain of self-mastery and the protection against any unjust interference with that domain. So, any act that unjustly interferes with our authority of self-mastery negatively affects autonomy, and thus produces a moral evil. An act *unjustly* interferes with our authority over part of our domain of self-mastery just in case it interferes with

<sup>15</sup> This point mirrors Fiona Woollard's criticism of F. M. Kamm's argument for constraints against harming based on inviolability, in which Woollard argues that it is not clear, according to Kamm's account, why 'protection against being harmed is a better reflection of high moral worth than protection against being allowed to suffer avoidable harm'. Woollard, 'Intricate Ethics and Inviolability: Frances Kamm's Nonconsequentialism', *Ratio* 21 (2008), pp. 231–8, at 238.

that authority without our forfeiting or willingly transferring it.<sup>16</sup> By contrast, well-being consists in the class of things that fall naturally in the domain of self-mastery, such as life, body and property. An act that impairs those things negatively affects the well-being value, and also produces some moral evil.

Before I propose my own analysis, it is important to pay attention to one distinction. One might think that the only moral evil relevant to the autonomy value is the loss of authority. This is a mistake. We need to distinguish two distinct kinds of moral evil with respect to autonomy. One is the mere loss of authority over part of the domain of self-mastery. A moral evil of this kind can be caused by either a natural event or a deliberate action. Both a strike of lightning and an intentional shooting can deprive me of my authority over my life.

The other kind of moral evil relevant to autonomy is the unjust interfering itself. An unjust interfering with one's authority of self-mastery is a distinct moral evil from the mere loss of that authority. An act that does not actually deprive me of my authority over myself may nevertheless unjustly interfere with my autonomy and thus be morally objectionable. Suppose, for example, that a neuroscientist designed a radio-like device and used it to try to control my mind. However, through some miscalculation, the device failed to have any effect on my brain. The neuroscientist's act does not deprive me of my relevant authority – it does not affect my freedom of will. But it still constitutes an unjust interfering with my authority – it unjustly infringes my domain of self-mastery, which is certainly morally objectionable.

An unjust interfering also differs from the mere loss of authority in the following two respects. First, unlike the mere loss of authority, an unjust interfering can only be caused by intentional action. Second, an unjust interfering is also a distinct moral evil from any impairment of well-being – an unjust interfering can be morally objectionable even if it causes no impairment to any well-being; in contrast, a loss of authority is not always a distinct moral evil from an impairment of well-being – e.g. the impairment of my financial well-being (my loss of money) consists in nothing other than the loss of my authority over the money. Thus, there are two distinct kinds of moral evil – the unjust interfering with one's authority of self-mastery on the one hand and the loss of that authority on the other – that are relevant to autonomy.

<sup>16</sup> When a person interferes with others' autonomy with a positive justification, which overrides the moral need to protect autonomy, the interference is still unjust (though it is justified). For example, when you use my car for an urgent hospital visit without my permission, your act unjustly interferes with my autonomy (you wrong me) even if you may be justified in doing so. Here, I am following Jeff McMahan's distinction between 'just' and 'justified' in his *Killing in War* (Oxford, 2009).

With these notions at hand, I now propose the following analysis of the moral difference between a violation of negative rights and a violation of positive rights (a refinement will be added in the next section):

- (a) A violation of negative rights constitutes an unjust interfering with the victim's autonomy, and it usually results in a loss of her well-being and her corresponding authority;
- (b) A violation of positive rights usually results in merely a loss of the victim's well-being and her corresponding authority.

For instance, killing a person not only causes the loss of a life and the corresponding authority over that life, it also unjustly interferes with the victim's autonomy – in other words, the act of killing produces three distinct moral evils. By contrast, allowing a person to die only causes the loss of a life and the corresponding authority – in other words, it produces only two moral evils. The unjust interfering plus the loss of a life and its corresponding authority is definitely morally weightier than just the loss of another presumably equal life and its corresponding authority – this is why a violation of negative rights requires more to justify it than a violation of positive rights.

There is no need to stipulate that an unjust interfering with autonomy is necessarily a greater moral evil than the loss of some well-being and the corresponding authority. We could, in theory, remain neutral about the comparative weight of the relevant autonomy value and well-being value. Thus, we avoid the problematic comparison that Draper's proposal seems to require.

One worry might be that the moral significance of unjust interfering with autonomy seems to reside precisely in the moral significance of negative rights, and thus my analysis simply presupposes the moral significance of negative rights, instead of explaining it. However, this does not mean that my analysis of the alleged moral difference is circular. Recall that the problem facing the rights-based analysis is that it simply presupposes that a violation of negative rights is more significant than a violation of positive rights. My analysis, on the other hand, shows exactly why a violation of negative rights is more significant: a violation of negative rights results in one more moral evil than a violation of positive rights. Thus my analysis offers the crucial piece of explanation that is missing in the rights-based analysis. The problem of explanatory incompleteness facing both Foot's analysis and Draper's analysis is thus solved.

Therefore, since, as the rights-based theory correctly points out, a typical case of doing harm (e.g. stealing somebody else's food) violates a negative right whereas a typical case of allowing harm (e.g. refusing to share your food with a starving beggar) violates a positive right, we

now have a justification for the alleged moral difference between typical cases of harm-doing and typical cases of harm-allowing. According to the value-based analysis I just proposed, when we compare the values involved in doing harm and allowing harm, we realize that doing harm results in one more moral evil, namely, an unjust interfering with autonomy, and thus would require a greater moral good to justify it.

Of course, my analysis does not show that an act of doing any harm is always harder to justify than an act of allowing any harm. But it provides us with a *prima facie* reason why an act of doing a certain degree of harm to a certain well-being requires more to justify it than an act of allowing the same degree of harm to the same kind of well-being.

#### IV. OBJECTIONS AND THE FULL ANALYSIS

Let me now consider some objections to this value-based analysis. One objection is that some cases of allowing harm are just as hard to justify as certain cases of doing harm. For example, watching a baby freeze to death seems as hard to justify as killing a greedy businessman. I do not think cases like these are genuine counterexamples. One possible explanation is that the well-being values involved in these two incidents are different: between an innocent baby's life and a life of a morally corrupt adult, we usually place a greater moral value on the former, which can affect our moral assessment here.

A second objection is this.<sup>17</sup> When I violate someone's negative rights, I interfere with her authority of self-mastery, which is a moral evil. When I violate someone's positive rights – for example, Bob is drowning and I could throw him a life preserver but I choose not to, there seems to be a parallel authority that I interfere with, namely, Bob's authority\* over me that I throw him a life-preserver. So, one may argue that, in a violation of positive rights, there is a parallel moral evil, namely, the interfering\* with the right-holder's authority\* over the agent for his assistance. If this is right, then there is no difference between violations of negative rights and violations of positive rights after all. One could not simply respond that unjust interfering with one's authority over oneself is a relevant moral evil, but interfering\* with one's authority\* over other people is not, for this would be equal to presupposing the moral difference between doing and allowing.

However, there is good reason to believe that unjust interfering with authority is a distinct moral evil from the loss of well-being and the

<sup>17</sup> Thanks to the reviewer of this journal for raising this objection. It helped to shape an important distinction, which, as we shall see, is needed to solve a troubling issue in the next objection.



corresponding authority, whereas interfering\* with authority\* is not. Here is why. Unlike the authority over one's own self, which we value precisely because we value the control over ourselves, in the case of the authority\* over others for their assistance, we value this authority\* not because we value the control we have over them (which would be morally perverse). Rather we value this authority\* precisely because we value our well-being and the corresponding authority over ourselves – the authority\* is nothing but a protection of our well-being and the corresponding authority. Thus, the moral evil in an interfering\* with authority\* is exhausted by the loss of the relevant well-being and its corresponding authority.<sup>18</sup> By contrast, the moral evil of an unjust interference is not exhausted by the loss of well-being and the corresponding authority. The neuroscientist's attempt to control my mind is morally objectionable, even if it results in no loss of my well-being or authority at all. The unjust intrusion upon my domain of self-mastery is itself a moral evil. Therefore, unjust interfering with authority is a distinct moral evil from the loss of well-being and its corresponding authority, whereas interfering\* with authority\* is not. So, there is indeed an important asymmetry.

Let us consider a third objection, which will show that a further complication needs to be added to the proposed analysis. If you fatally shoot a person, her autonomy is unjustly interfered with and her life is lost. If you allow a person to be fatally shot by another person, it is also the case that her autonomy is unjustly interfered with and her life is lost. Thus, it seems that, on my analysis, there is no moral difference between you harming a person and you allowing a person to be harmed by another.

Two responses are available. The easier response is this. When you shoot a person yourself, the unjust interfering is a consequence of your harm-doing. By contrast, if you simply stand by when you could push the victim out of the course of the bullet, the unjust interfering is not a consequence of your act – the act of shooting has already taken place and you, by no means, causally contribute to its occurrence. Thus, there is indeed an important difference between these two cases – your harm-doing results in one more moral evil, i.e. an unjust interfering.

However, this easy response does not work for more sophisticated cases. Compare a case where you shoot someone yourself with a case

<sup>18</sup> One may say that there is indeed a distinct moral evil in the violation of positive rights, i.e. the agent's disregard to others' well-being. However, bad characters or attitudes belong to agent evaluation, and my primary focus here is act evaluation. Furthermore, even if we take agent evaluation into consideration, we can still say that in a violation of negative rights, the agent manifests not only a disregard to others' well-being, but also a disregard for others' autonomy, in particular their protection against unjust interfering. Thus, we still have an asymmetry here.

where you could prevent John from shooting someone else but you simply allow him to do it. It seems that my theory would say these two cases are morally equivalent because your act results in an unjust interfering in both cases.

At this point, one might start seriously to question this value-based analysis. I try to locate the unique moral weight of doing harm in *unjust interfering with autonomy*. But the occurrence of an unjust interfering does not seem to be agent-relative. In other words, it does not seem to matter, for my account, whether you are the source of the interfering or somebody else is. By contrast, there seems to be something agent-relative about the alleged moral difference between doing and allowing – it matters whether you do the harm or somebody else does it. And the aforementioned case just illustrated this problem. So, one may conclude that my focus on unjust interfering with autonomy is completely misplaced.

But this conclusion is too quick. The value-based analysis has the resources to accommodate this agent-relativity. Notice that in the case of your allowing John to shoot another, if we require you not to allow harm, we will deprive you of something morally valuable, namely, your freedom to choose what to do in that situation, by putting this constraint on your choice. So, despite the moral evil your allowing brings about (the unjust interfering with the victim's autonomy plus the loss of her well-being and authority), you, in doing so, enjoy a moral good, namely, the freedom to choose what to do (which is, of course, part of your authority over yourself). By contrast, no such moral good is preserved in doing harm. It is true that if we require you to refrain from shooting somebody, we would also constrain your freedom (i.e. the freedom to shoot somebody). However, such freedom – the freedom to interfere unjustly with others' autonomy – is certainly not part of the legitimate authority of self-mastery that any reasonable theory of autonomy would recognize. Otherwise, the idea of autonomy would be self-defeating: taking the self-mastery of one's own person as a serious universal moral value entails accepting a certain boundary of the freedom that one can legitimately have. Thus, freedom to interfere unjustly with others' autonomy is not part of the value of autonomy, and, for a similar reason, it would be equally absurd to take it as a well-being.<sup>19</sup> So, it is not a moral good at all in any value theory that recognizes the significance of autonomy.

<sup>19</sup> Here I focus on the intrinsic moral value of the mere freedom itself. I leave aside all the contingent moral goods that may result from the enjoyment of such freedom, and I also leave aside all the well-being-related functions that may be required to enjoy such freedom (such as the physical capacity to hold a gun). The freedom to kill someone may bring me his money, but the well-being value of the money is certainly not intrinsic to the freedom itself. The capacity to hold a gun has some well-being value, but it does not mean that the mere freedom to shoot someone, which requires that capacity, has that

But what about freedom to allow harm? Notice that positive rights give others authority\* over us for our assistance. So, would not their authority\* (over us) deprive us of the corresponding part of our authority (over ourselves) such that the freedom to allow harm is not part of the value of autonomy either? If so, there is no asymmetry after all.

The moral significance of the authority\* over others for their assistance, as I have argued, derives from the need to protect the victim's well-being and the corresponding authority. The control over others is not, in and of itself, a moral good. So, the authority\* does not work in a way that it completely *annihilates* the moral value of the authority over oneself; rather it simply *overrides* the latter. This is implicit in the widely shared view that the duty to assist arises only when the good to be preserved is considerably greater than the cost,<sup>20</sup> which clearly suggests a comparison between two sets of values. So, there is good reason to think that freedom to allow harm is indeed a moral good, which simply gets overridden by the need to protect other people's well-being in certain circumstances; whereas freedom to interfere unjustly with others' autonomy is not.

Therefore, my response to the third objection is this. Even if shooting a person oneself and allowing someone to shoot another both result in an unjust interfering with the victim's autonomy, there is still an important difference: the freedom to allow harm is itself a moral good, yet the freedom to do harm is not.<sup>21</sup>

The apparent lack of agent-relativity in my focus on unjust interfering with autonomy is also solved by adding this further complication into our assessment: the moral value of the freedom that the agent enjoys in performing the act in question. In committing an unjust interfering oneself, the agent enjoys a freedom of no positive moral value; whereas in allowing an unjust interfering by another person, the agent enjoys a freedom that has a positive moral value. The greater the freedom that is at issue, the easier the justification for allowing harm, and thus the greater the difference between doing and allowing.

Finally, we have the full analysis of the moral difference between doing and allowing:

well-being value itself. I think to view the mere freedom to shoot someone as a well-being, something that is entitled to respect, is to slap the idea of autonomy in the face.

<sup>20</sup> See, for example, Peter Singer, 'Famine, Affluence, and Morality', *Philosophy & Public Affairs* 3 (1972), pp. 229–43.

<sup>21</sup> There is no need to include further the loss of the victim's authority\* over the agent in our assessment, since the moral significance of the loss of that authority\*, as I have argued, is exhausted by the loss of the victim's well-being and her corresponding authority, which has already been included in the assessment.

- (a) A harm-doing constitutes an unjust interfering with the victim's autonomy, and it usually results in a loss of her well-being and the corresponding authority;
- (b) A harm-allowing usually results in merely a loss of the victim's well-being and the corresponding authority, and in allowing harm, the agent enjoys a moral good, namely, part of his authority over himself (or his freedom to allow harm).

As we can see, the key difference is that a harm-doing always constitutes an unjust interfering with autonomy, and though a harm-allowing may sometimes also result in an unjust interfering, the shift of agency, from one's own interfering to somebody else's interfering, changes the moral value of the freedom involved.<sup>22</sup>

Let me consider a final objection to the proposed analysis. One may object that I have, so far, only considered the harm and the benefit that is *necessarily* attached to harm-doing or harm-allowing, but a faithful consequentialist approach must also take into consideration the benefits that follow *contingently* from those acts. Killing one person, for instance, may help to save five other lives. When we move from simple cases where we only consider the harm and the benefit that is necessarily involved to cases where those values are also accompanied by some contingent benefits, the assessment becomes more complex than what I have considered.

I think this is certainly right. When we consider those more complex cases, especially cases where the benefits resulting from a harm-doing differ from the benefits resulting from a harm-allowing (e.g. an act of killing one person saves five other lives, but an act of allowing one person to die saves just two), my analysis does not show that the harm-doing is also harder to justify than the harm-allowing there. However, it is not my intention to defend such a strong version of DDA. My analysis, I believe, at least shows that there is a *prima facie* reason for believing that doing a certain degree of harm to a certain well-being requires a greater moral good to justify it than allowing the same degree of harm to the same kind of well-being. Moreover, I think the value-based analysis provides a framework for assessing even those more complex cases – after all, what is involved in those cases is nothing but a few extra instances of autonomy value and well-being value.

It is worth mentioning that it is this moderate version of DDA that most defenders of DDA have in mind. For example, some philosophers have tried to defend DDA within a consequentialist framework by

<sup>22</sup> Again, there is no need to know the comparative weight of the negative value and the positive value involved. All we need to know about the moral difference is that harm-allowing involves one less evil and one more good.

showing that they are compatible. According to them, to prove that DDA is compatible with consequentialism is to justify the following principle: 'An agent's appeal to cost has more weight in cases of granting permissions for allowing harm than in cases of granting permissions for doing harm.'<sup>23</sup> The idea is that if this principle is justified, then it might be permissible to allow harm in order to avoid paying \$1,000 dollars, yet impermissible to do harm in order to avoid the same amount of cost. In other words, doing harm, other things being equal, requires more to justify it than allowing harm. This strategy to defend DDA also aims at providing just a prima facie reason for the moral difference between doing and allowing. Moreover, I believe my value-based analysis provides a better justification for the aforementioned principle than those that have been offered. According to one consequentialist justification proposed by Bashshar Haydar, a person who actively engages in decreasing the overall good becomes 'especially responsible' for this decrease in the overall good, while this is not the case in merely allowing the decrease to happen – this is why appeal to cost has less weight in cases of granting permissions for doing harm.<sup>24</sup> However, this justification suffers the same problem as the rights-based analysis – it remains unclear exactly why actively engaging in decreasing the good makes one more responsible for the decrease, since in either case one seems to meet both the epistemic condition and the control condition for moral responsibility.<sup>25</sup> By contrast, my analysis offers a neat explanation – doing harm produces more evil and thus decreases the overall good to a greater degree.

## V. THE CONCEPTUAL DISTINCTION

We now have a reasonable justification for the alleged moral difference between cases that we usually regard as harm-doing and cases that we usually regard as harm-allowing. But in order to have a robust defence of DDA, we also need a conceptual distinction that matches our ordinary use of the concepts in a wider range of cases. In what follows, I want to suggest that the value-based analysis also provides such a conceptual distinction.

In the simple and paradigm cases, we can say that, following Foot's suggestion, an act that initiates a harmful sequence is doing harm; an

<sup>23</sup> See, for example, Scheffler, 'Prerogatives without Restrictions', *Philosophical Perspectives* 6 (1992), pp. 377–97; and Bashshar Haydar, 'Consequentialism and the Doing–Allowing Distinction', *Utilitas* 14 (2002), pp. 96–107.

<sup>24</sup> Haydar, 'Consequentialism', p. 103.

<sup>25</sup> For a discussion of these two conditions for moral responsibility, see, for example, J. M. Fischer and M. Ravizza, *Responsibility and Control: A Theory of Moral Responsibility* (New York, 1998), p. 13.

act that merely allows a harmful sequence to complete or to continue is allowing harm.<sup>26</sup> Initiating a sequence that is harmful to others usually constitutes an unjust interfering with their autonomy. Allowing a harmful sequence to continue, on the other hand, does not constitute such an interfering. It might result in someone else's interfering, but meanwhile it also preserves a moral good, which sustains the asymmetry.<sup>27</sup>

This initial account, however, fails to match our ordinary use of the concepts in a wider range of cases. In particular, it fails to account for a group of cases that involve *removal or withdrawal of aid or protection*. Consider the following two cases.

*Firefighter*: Jake trapped atop a high building that is on fire leaps off. Seeing this, a firefighter stations a self-standing net underneath and then goes on to assist others. Jake's enemy, Smith, happens to be nearby and, seeing his opportunity, swiftly removes the net. As a result, Jake hits the ground and dies.<sup>28</sup>

*Sealer*: An earthquake cracks a pipe at a factory, releasing poisonous chemicals into the water supply. Before a dangerous amount is released, a worker, Smith, seals the pipe. But a year later Smith returns and removes the seal. As a result, numerous people die from drinking contaminated water.<sup>29</sup>

Smith is, arguably, not the one who initiates the harmful sequence in either case. But our intuition is that his acts are harm-doing rather than harm-allowing. Thus, the initial account seems unable to accommodate complex cases like these.

Jeff McMahan has provided perhaps the most exhaustive treatment of the conceptual distinction between doing and allowing for cases involving removal or withdrawal of aid or protection. According to McMahan, three factors are relevant to whether a termination of aid

<sup>26</sup> I think 'allowing a harmful sequence to *complete*' is too strong a condition for allowing harm. Merely allowing a harmful sequence to *continue* should count as harm-allowing. Otherwise, cases in which the harm is allowed but fails eventually to occur owing to some other reasons would not count as harm-allowing, which seems incorrect.

<sup>27</sup> I think, in the simple cases, it is quite intuitive that, as Foot suggests, initiating a harmful sequence constitutes an unjust interfering with autonomy but merely allowing a harmful sequence to continue does not. One may object that my analysis is also incomplete because I simply appeal to our intuition to draw the conceptual distinction here. However, this appeal to intuition is quite different from the appeal to intuition that we saw in the rights-based analysis. It is problematic to explain an important moral distinction by simply appealing to some obscure intuitions. But it is much less problematic to explain a conceptual distinction by appealing to a seemingly unproblematic conceptual intuition. So, even if my conceptual analysis here is, in this sense, incomplete (as any analysis has to stop somewhere), I do not think my analysis of DDA as a whole involves any circularity or the kind of incompleteness that the rights-based analysis faces.

<sup>28</sup> This is a case originally discussed in McMahan, 'Killing, Letting'. A few details are changed.

<sup>29</sup> McMahan, 'Killing, Letting', p. 389.

or protection counts as doing harm. The first is ‘whether the person who terminates the aid or protection is the person who provides it’.<sup>30</sup> In the *Firefighter* case, if the firefighter, instead of Smith, removes the net, our intuition is that his act does not count as doing harm, and this is due to the fact that he is the person who provided that protection in the first place. The second factor is ‘whether the aid or protection is self-sustaining or requires more from the agent’.<sup>31</sup> In the *Sealer* case, though it is Smith who sealed the pipe in the first place, since the sealer is self-sustaining, Smith’s removal still counts as harm-doing. By contrast, if the sealer requires constant maintenance from him and he simply stops working on it one day, he merely allows harm. The last factor is ‘whether the aid or protection is operative or as yet inoperative’.<sup>32</sup> On McMahan’s view, if the aid or protection is already in operation (and it is self-sustaining), then terminating it counts as a harm-doing; if not, then harm-allowing. Based on these factors, McMahan then proposes a complex conceptual distinction between doing and allowing.<sup>33</sup>

I think McMahan is right about the relevance of most of these factors. However, his account seems disturbingly *ad hoc* in nature.<sup>34</sup> No principle unifies all the different factors into an integral conceptual schema. His account would be much more attractive if he could provide a unifying principle that shows how those different factors can be related. Is there such a unifying principle? I think the value-based analysis provides one.

Here is how those factors can be unified – they are all relevant to whether the act in question constitutes an unjust interfering with autonomy. First, whether the aid or protection is provided by the person who terminates it is relevant, because if the aid or protection is provided by others, then one’s removing it usually unjustly interferes with what is already at the victim’s disposal. One may argue that since the victim is already in harm’s way, there is no autonomy to be interfered with. This is mistaken. Suppose a patient’s life relies on her receiving an

<sup>30</sup> McMahan, ‘Killing, Letting’, p. 396.

<sup>31</sup> McMahan, ‘Killing, Letting’, p. 396.

<sup>32</sup> McMahan, ‘Killing, Letting’, p. 396.

<sup>33</sup> According to McMahan, ‘[W]hen an agent withdraws aid or protection from a lethal threat that he has not himself provided, or when he withdraws aid or protection that he has provided but which was complete and self-sustaining, his action counts as killing; but when an agent withdraws aid or protection that he himself has provided but which requires further contributions from him to be effective, then his action counts as letting the victim die.’ Furthermore, ‘[I]f a person requires or is dependent for survival on further aid from or protection by an agent, and if the person dies because the agent fails to provide further aid or withdraws his own aid either while it is in progress or before it becomes operative, and if the agent is not causally responsible for the person’s need for aid or protection, then the agent lets the person die.’

<sup>34</sup> Howard-Snyder, ‘Doing vs. Allowing’, sect. 8.

antidote in time, and I steal that antidote before she receives it. It would be absurd to say that my stealing is not an unjust interfering with the victim's autonomy simply because there already exists a life-threatening sequence.

Second, whether the aid or protection is self-sustaining or requires a continuous effort on the actor's part is also relevant. If the aid or protection is already self-sustaining, then the situation is like the one in which the aid or protection is provided by others – the act of terminating the aid would usually unjustly interfere with what is already at the victim's disposal. This is true even if it is the agent who provided the aid in the first place, because in most cases where a self-sustaining aid is willingly provided, *the relevant authority over the aid has been successfully transferred from the provider to the recipient*. Thus, if the aid or protection is self-sustaining, removing it would unjustly interfere with the victim's autonomy. By contrast, if the aid or protection requires further effort on the actor's part, then terminating the effort would not constitute an unjust interfering. Whether I want to continue my effort to help someone to whom I have no existing obligation is completely within my autonomy. It may be uncharitable to terminate the aid, or it might be immoral all things considered. But doing so certainly does not unjustly interfere with the victim's autonomy.

However, I disagree with McMahan on the third factor. Whether an aid or protection is already operative is not a relevant factor. Consider the following case.

*Coat:* My friend and I are stuck in the mountain because of a huge storm. I lend my fur coat to my friend to fend off the cold for a while. But eventually, the temperature drops so low that if I do not put on the coat, I will be frozen to death. But I also know that if I take the coat back the same thing will happen to my friend. I decide to take the coat back. Consequently, my friend dies instead of me.

In this case, the protection my friend had – the protection provided by the coat – is surely self-sustaining and operative. But our intuition is that I merely let my friend die by taking back the coat (provided that I did not intentionally put my friend in that danger when I actually foresaw the storm coming). Again, I think the relevant factor in this case is rather *whether the agent has proper authority to remove or withdraw the property needed for the aid or protection*.<sup>35</sup> If one has the authority over the property, then like the situation in which the aid or protection requires the agent's further effort, terminating the aid or protection is usually something within the agent's autonomy. Thus, in

<sup>35</sup> Draper makes a similar point in his criticism to McMahan. See Draper, 'Rights', p. 266.



*Coat*, my taking the coat back is an act of letting my friend die because the coat is still my property. If, by contrast, it is my friend's enemy who takes the coat away, it is a case of killing because he lacks the relevant authority to do so.

So, there is indeed a unifying principle behind all these factors – they are relevant to whether an act unjustly interferes with autonomy. Thus, based on McMahan's account and Foot's account, I propose the following conceptual distinction between doing and allowing.

An act is a doing harm if and only if

- (1) it initiates a harmful sequence; or
- (2) when there already exists a harmful sequence that is not initiated by that act, (a) the act removes an existing obstacle that will prevent the harmful sequence,<sup>36</sup> or prevents the creation of such an obstacle, and (b) sustaining or creating the obstacle does not require the actor's further effort, and (c) the part of the obstacle that is to be removed or prevented from being created by the act is not something that the actor has authority to remove or withdraw.

When these conditions are met, the act not only impairs certain well-being and the corresponding authority, but also unjustly interferes with the victim's autonomy.

By contrast, an act is an allowing-harm if and only if

- (1) it merely allows a harmful sequence to continue; or
- (2) it removes, or prevents the creation of, an obstacle that will prevent the harmful sequence, but (a) sustaining or creating the obstacle requires the actor's further effort, or (b) the part of the obstacle that is to be removed or prevented from being created by the act is something that the actor has authority to remove or withdraw.

When these conditions are met, the act does not constitute an unjust interfering with the victim's autonomy.

Let me close by considering one objection to this proposed conceptual distinction. Imagine that a lab assistant produces a poisonous reagent and leaves it on the table for future experiment. A child sneaks in, drinks the reagent and dies. One might say that the act of producing the reagent and the act of leaving it on the table initiate a harmful sequence, but intuitively they are not acts of killing the child. So, we

<sup>36</sup> I understand an obstacle that prevents a harmful sequence in a very broad sense. A car that happens to stand between a running trolley car and me is an obstacle to prevent the trolley car from harming me. A samaritan's effort to save a drowning child can also be regarded as an obstacle that prevents a harmful sequence, namely, the child's drowning.

need to add that only an act that is *normally* harmful can be a harm-doing. Yet there is a further complication. Suppose the lab assistant knows that the child will drink the reagent, so he purposely produces that poisonous reagent and leaves it on the table. In this case, even if producing a poisonous reagent in a lab is normally harmless, it seems that the lab assistant kills the child. Thus, the actor's intention also seems to matter for whether an act is a harm-doing. Now, one may argue that whether an act interferes with autonomy should not depend on what course that type of action normally takes or on what intention the actor has. So, these complications involved in the notion of doing harm cannot be accommodated by the notion of unjust interfering with autonomy: it is a mistake to analyse the concept of doing harm in terms of unjust interfering with autonomy.

I disagree. These complications are indeed implicated in the vagueness of the concept of unjust interfering with autonomy. Two different kinds of unjust interfering with autonomy can usually occur. I may unjustly interfere with another's autonomy by directly acting on what is within her domain of self-mastery, e.g. punching her in the face; or I may unjustly interfere by acting on something in the environment, which further intrudes on her autonomy, e.g. diverting a trolley car to a side track where she is working and thus causing her death. However, there is a considerable amount of vagueness about the interference of the second type. Not all cases where a change in the environment leads to an intrusion on autonomy are cases of unjust interfering with autonomy. For example, if I open a window and let a butterfly out of my room, which causes a storm in a certain area, which, in turn, sets off a trolley car and kills someone, it is absurd to say that I unjustly interfere with that person's autonomy. The difference between this case and the previous trolley-car case seems to be that the act of opening a window is normally harmless, but diverting a trolley car in that situation is not. Likewise, intention may also matter. If, for instance, I know for certain that my opening the window will lead to that harmful result, it might be said that my act is indeed an unjust interfering with the victim's autonomy. So, the complications we have in the concept of doing harm have their counterparts in the concept of unjust interfering with autonomy.

## VI. CONCLUSION

Drawing lessons from Foot's account and McMahan's account, I proposed an account of the conceptual distinction between doing and allowing that matches our ordinary use of the concepts in a wide range of cases. I also showed that there is a unifying principle underlying this account – harm-doing constitutes an unjust interfering with autonomy,

but harm-allowing does not. This distinction enables a justification for the alleged moral difference: harm-allowing usually involves one less evil – the unjust interfering with autonomy – and one more good – the freedom to allow harm. This is my robust defence of a moderate version of DDA. This value-based analysis provides a deeper and more satisfactory explanation of the alleged moral difference than the rights-based alternative. What is a little surprising perhaps is that consequentialism turns out to have better resources to justify DDA.<sup>37</sup>

xl8q5@mail.missouri.edu

<sup>37</sup> I am grateful to Peter Vallentyne, Peter Markie, Kai Draper, Robert Johnson, Crystal Allen, Daniel Marshall, Ashton Sperry and Wenwen Fan for constructive comments. I am especially indebted to an anonymous reviewer for this journal for insightful comments.