

Doing, Allowing, and Enabling Harm: An Empirical Investigation ^{*}

Christian Barry (Australian National University)

Matthew Lindauer (Yale University)

Gerhard Øverland (University of Oslo)

Forthcoming in *Oxford Studies in Experimental Philosophy*

(Joshua Knobe, Tania Lombrozo, and Shaun Nichols eds.)

Abstract

Traditionally, moral philosophers have distinguished between doing and allowing harm, and have normally proceeded as if this bipartite distinction can exhaustively characterize all cases of human conduct involving harm. By contrast, cognitive scientists and psychologists studying causal judgment have investigated the concept ‘enable’ as distinct from the concept ‘cause’ and other causal terms. Empirical work on ‘enable’ and its employment has generally not focused on cases where human agents enable harm. In this paper, we present new empirical evidence to support the claim that some important cases in the moral philosophical literature are best viewed as instances of enabling harm rather than doing or allowing harm. We also present evidence that enabling harm is regarded as normatively distinct from doing and allowing harm when it comes to assigning compensatory responsibility. Moral philosophers should be exploring the tripartite distinction between doing harm, allowing harm, and enabling harm, rather than simply the traditional bipartite distinction. Cognitive scientists and psychologists

^{*} We have benefitted from discussions of earlier versions of this paper with members of the Experimental Philosophy Lab at Yale, the Yale Global Justice Program, and participants in the Workshop on Enabling Harm at the University of Oslo in June 2012. The authors are grateful for comments on this paper that we received from Steve Guglielmo, Mark Sheskin, Brent Strickland, and three anonymous reviewers, and especially to Joshua Knobe and Tania Lombrozo for their detailed comments on successive written drafts. We would also like to thank Kevin Callender and Elizabeth Roberto for their assistance with the statistical analyses used in the paper. Barry, Lindauer, and Øverland are full joint authors of the paper. The theoretical aspects of the doing-allowing-enabling harm distinction are based on Barry and Øverland’s previous and forthcoming work. Lindauer had the primary role in conducting the experiments and positioning the paper in relation to cognitive scientific and psychological research.

studying moral judgment, who have so far largely adopted the bipartite distinction in this area of research, should likewise investigate the tripartite distinction.

1. Introduction

Moral philosophers examining the nature and significance of human conduct involving harm have tended to employ a bipartite distinction between doing harm and allowing harm. This distinction has generally been treated as exhaustive—all cases of human conduct involving harm have typically been thought to be descriptively and normatively subsumable under one or the other of these two categories. Enabling harm, in particular, has generally not been treated as an independent category of human conduct. By contrast, cognitive scientists and psychologists have studied the concept of ‘enabling’ as distinct from ‘cause’ and other causal terms. Yet ‘enabling’ and ‘enabling harm’ have not yet been the focus of empirical research on judgments about cases of human conduct involving harm. Empirical studies pertaining to these judgments have generally employed the traditional bipartite doing-allowing harm distinction.

In this paper, we marshal new empirical evidence from two studies to argue that well-known cases in the moral philosophical literature are cases of enabling harm, rather than doing or allowing harm. This suggests that enabling harm is judged to be a distinct category of human conduct, not descriptively subsumable under the categories of doing or allowing harm. A third study provides evidence that it is judged to be normatively distinct too. The results of this third study suggest that, in assigning compensatory responsibility, enabling harm is regarded as more like doing harm than allowing harm when an agent is culpable but more like allowing harm than doing harm when an agent is innocent. Given this evidence, we argue that moral philosophers and moral psychologists interested in human conduct involving harm should employ a tripartite doing-allowing-enabling harm distinction, rather than the traditional bipartite doing-allowing harm distinction.

1.1. *The Moral Philosophical Literature*

A vast literature in moral philosophy is devoted to examining the meaning and moral significance of the distinction between doing harm and allowing harm. While theorists have employed different terms in characterizing these notions (e.g., Bennett 1995; Donagan 1977; Foot 1994; Quinn 1989/1994), the contributors to this literature have generally taken themselves to be giving competing accounts of the same underlying conceptual distinction. In some instances, agents actively contribute to, do, or bring about harm, whereas in other instances, they merely allow or fail to prevent the occurrence of harm. A great many philosophers have held that all instances of human conduct involving harm are accurately described either as instances of doing harm or as instances of allowing harm. The notion of enabling harm has been given very little attention, and where it has been discussed, it has generally been thought to be descriptively subsumable under one or the other of these categories.¹ Whether a given case of enabling harm is viewed as a case of doing harm or allowing harm has generally been related to specifics of the philosopher's account of the doing-allowing distinction.

The interest in understanding this distinction between doing harm and allowing harm is due to the fact that it is generally thought to be normatively significant.² All other things being equal, reasons against doing harm are commonly thought to be more stringent than reasons against allowing harm. They are more stringent in the sense that they *constrain* agents: prospective doers of harm cannot so easily justify their conduct by appealing to the costs to themselves of refraining from doing harm, or by appealing to the overall good that their conduct will bring about. For example, we generally do not think that we can break someone's arm to prevent one of our own arms from being broken, or to prevent the breaking of two other persons' arms, all else being equal. But it seems more acceptable, at the very least, to allow someone's arm to be broken if it will prevent

¹ See, for instance, Foot (1994).

² In this paper we do not discuss the related act-omission distinction, which is also generally taken to be normatively significant. We remain neutral on the question of whether these distinctions perfectly map onto one another for the purposes of this paper. If they indeed do, it would seem that the act-omission distinction is unable to distinguish acts that are instances of doing harm from acts that are instances of enabling harm. Since we do not focus on the differences in moral judgments that are made about cases of doing harm, enabling harm, and allowing harm or acts and omissions, we also do not discuss in this paper the empirical literature on the "omission bias," describable as the preference for harms caused by omissions over equal or lesser harms caused by acts.

the same harm from coming to ourselves or to two other people. Reasons against doing harm are also thought to be more stringent than reasons against allowing harm in the sense that they *demand* more of agents after the harm has eventuated. If, for instance, one has broken another person's arm, we typically think that one must take on significant cost to compensate this person, whereas if one merely allowed the person's arm to break, one is far less or perhaps not at all responsible for compensating them. In other words, one has more demanding compensatory responsibilities if one has done harm than if one has allowed the same harm.

The main topic of this paper and the first two experiments that we will discuss is the descriptive classification of well-known cases of doing harm, allowing harm, and (what we will argue are cases of) enabling harm in the moral philosophical literature. While our aim here is not primarily to document nor elaborate upon purported normative differences between doing and allowing harm, the existence of a third category of human conduct could turn out to be important for normative assessment. It would be important, for instance, if enabling harm is normatively distinct from doing and allowing harm. We cannot fully argue in this paper that enabling harm is indeed normatively distinct from doing harm and allowing harm, but there is reason to believe that it may well be. For example, one normative characteristic that reasons against enabling harm appear to have in common with reasons against allowing harm is that they constrain conduct less with regard to shifting cost from one person or group of people to some other people than reasons against doing harm (Hanser 1999; Rickless 2011). At the same time, however, reasons against enabling harm appear to constrain conduct more with regard to shifting cost from ourselves to others than do reasons against allowing harm. If reasons against enabling harm are indeed normatively distinct from reasons against doing harm and allowing harm, then failing to focus on this category could invite quite serious moral misunderstanding. Barry and Øverland (2012) have argued, for instance, that by failing to recognize the category of enabling harm, participants in debates on global justice often tended to talk past each other, and to make implausible claims about the nature of obligations to address global poverty. The results of the third experiment discussed in this paper constitute suggestive evidence that enabling harm is regarded as normatively distinct from doing harm and allowing harm when it comes to compensatory

responsibility. We hope in further research to investigate how aspects of the normative status of enabling harm are generally understood.

1.2 Cognitive Scientific and Psychological Research

Cognitive scientists and psychologists have studied the representation of different kinds of causal relations. Wolff (2007) finds evidence for a “dynamics model,” based on the force-dynamics theory originally proposed by Talmy (1988), according to which causal relations are represented in terms of configurations of forces. In particular, ‘cause’ and other causal concepts are applied to causal situations on the basis of what kinds of interactions are represented as taking place between two main types of entities: affectors and patients that affectors act on. According to Wolff, the particular causal concept that is applied to a given representation of causal interaction between an affector and a patient depends on three factors: 1) whether the patient is represented as having a force-based tendency towards an endstate, 2) whether the forces represented as exerted by the patient and affector are concordant, and 3) whether the patient is represented as making progress towards the endstate.³ For instance, when a patient doesn’t have a tendency towards an endstate, the affector exerts a force on the patient towards the endstate (the forces of the patient and affector are not concordant), and the patient makes progress towards the endstate, the model predicts that the concept ‘cause’ will be applied to the representation. By contrast, when a patient does have a tendency towards the endstate, an affector exerts a concordant force on the patient towards that endstate, and the patient makes progress towards the endstate, the model predicts that the concept ‘enable’ will be applied.

Wolff (2007) attempted to show that the dynamics model has better extensional adequacy than competing models, more accurately picking out causal situations from non-causal situations and capturing distinctions among causal situations that the others cannot. Participants were asked to view computer simulations of interactions between affectors and patients—for instance, a boat and a bank of fans in a shallow pool, where a cone served as an endstate—and classify the scenario using various causal terms. The predictions of the model were borne out with respect to the application of the terms

³ Wolff (2007), pp. 87-89.

‘cause,’ ‘enable,’ ‘prevent,’ and ‘despite’ in a number of experiments. An additional experiment was conducted to disambiguate the dynamics model from a kinematics model that would make predictions based solely on the visible properties of the interaction events. A further experiment applied the model to social causation by looking at how participants classified simulations involving causal intent and communication in a street intersection, suggesting that ‘social forces’ are treated like physical forces.

Other views of causal representation have also been put forward.⁴ Among the most prominent alternatives to the dynamics model is the “causal model theory” developed by Sloman and colleagues. According to this theory, causal terms express beliefs about the structure of causal models (Sloman et al. 2009). To assert “A causes B” expresses the belief, roughly, that there is an asymmetric dependence relationship between A and B such that a change in A will bring about a change in B, but a change in B will not bring about a change in A. Asserting “A enables B” additionally implies that A is necessary for B, and also that B has additional causes.⁵ A third prominent approach is the “mental model theory,” developed by Goldvarg and Johnson-Laird (2001). The mental model theory holds that causal terms and relations are represented as logical relationships between possibilities or “models.” For instance, ‘A causes B’ corresponds to three possibilities: A and B are both present, A is absent and B is present, or both A and B are absent, with the temporal constraint that B doesn’t precede A. ‘A enables B’ has either a weak or strong sense, according to which either all four possibilities might obtain (both A and B present; A present and B absent; A absent, B present; both A and B absent), or, on the stronger interpretation, A is needed for B (and so the possibility that A is absent and B is present is removed).⁶

Despite the differences between each of these views, they each recognize a psychologically real distinction between the concepts ‘cause’ and ‘enable.’ Nonetheless,

⁴ Our discussion of these views has been aided by Lombrozo (2010). Lombrozo finds evidence for causal-explanatory pluralism, according to which theories of causal representation that emphasize physical connections and those that emphasize counterfactual or logical dependence are both psychologically real, but judgments in line with each approach are promoted more and less by different modes of causal explanation. Lombrozo argues for a more unified account of causation which she calls an “exportable dependence” theory. See Lombrozo (2010).

⁵ Sloman et al. (2009), p. 21.

⁶ Goldvarg and Johnson-Laird (2001) focus predominantly on enabling *conditions*.

they leave many questions about ‘enable’ and enabling in general open. For instance, proponents of these views have suggested that ‘enable’ has both a ‘help’ sense and an ‘allow’ sense roughly akin to ‘make possible,’ but it is not clear how they distinguish enabling from allowing in the sense of ‘not preventing.’ Further, the concept ‘enable’ has largely figured into studies in which authors, implicitly or explicitly, sought to control for the effects of moral evaluation on causal judgment. The distinction between ‘enable’ and other causal terms has yet to be explored with respect to cases of human conduct involving harm. Moral psychological research on the relationship between normative, including moral, judgments and causal judgments has not singled out the concept of enabling harm for investigation (e.g., Hitchcock and Knobe 2009; Knobe and Fraser 2008; Lagnado and Channon 2008). Cushman, Knobe and Sinnott-Armstrong (2008) demonstrate that moral appraisals can affect doing/allowing judgments—morally bad behavior is more likely to be construed as doing harm than allowing harm. Work in cognitive science and moral psychology directly investigating judgments about doing harm and allowing harm has thus far only employed the traditional bipartite distinction.

1.3 Doing, Allowing, and Enabling Harm

In this paper, we argue that new experimental data supports the view that some instances of human conduct involving harm are not descriptively subsumable under the traditional doing-allowing harm distinction. They are instances of enabling harm, which should be viewed as a distinct category. The cases that we have investigated are variations on well-known cases in the moral philosophical literature on doing and allowing harm that were first employed by Bennett (1995).⁷

Throughout the paper we will refer to Push and Stayback cases, which take the following general form:

Push: A cart stands at the top of a hill. John pushes it. The cart rolls down the hill and injures Tom, who is sitting at the bottom of the hill.

⁷ We have made changes to the cases so as to make the kind and amount of harm that results in each of them equal.

Stayback: A cart is rolling down a hill. John could put a rock in the way of the cart that would stop it, but does not. The cart rolls down the hill and injures Tom, who is sitting at the bottom of the hill.

Push and Stayback are straightforward in the sense that their classifications under the traditional doing-allowing distinction have been uncontroversial among philosophers. Philosophers generally agree that Push and Push-like cases are cases of doing harm, and Stayback and Stayback-like cases are cases of allowing harm.

By contrast, there are what we will refer to as Interpose and Remove cases. They take the following form:

Interpose: A cart accidentally starts rolling down a hill. Tom, who is sitting at the bottom of the hill, won't be injured by the cart if he can get out of the way of its path. John puts a rock on the ground. The rock stops Tom, and he is injured by the cart.

Remove: A cart is rolling downhill towards a point where there is a rock that would bring it to a stop. John removes the rock. The cart rolls down the hill and injures Tom, who is sitting there.

Interpose and Remove cases and cases like them—cases in which agents either create an obstacle that blocks a causal process that would prevent a harm or remove an obstacle that would block a causal process leading to a harm—are the subject of considerable descriptive disagreement among philosophers who employ the traditional doing-allowing distinction.⁸ We hold that theorists on either side of these disagreements are mistakenly employing a bipartite distinction where an additional, independent empirical category should instead be employed to accurately classify these cases.

Our argument is thus largely in agreement with the general tendency in cognitive scientific and psychological research on causal representation to treat enabling as an

⁸ A number of philosophers who emphasize an agent's movements, such as Jonathan Bennett, or actions, such as Warren Quinn, Kadre Vihvelin, and Terrance Tomkow, would say that John does harm to Tom in Interpose and Remove. See Bennett (1995), Quinn (1989), and Vihvelin and Tomkow (2005). Philippa Foot, by contrast, reduces the doing-allowing distinction to that between *initiating or sustaining* a harmful causal sequence and *allowing* a harmful sequence to run its course. Consequently, Foot would claim that John merely allows Tom to be harmed in cases such as Remove. It is perhaps less obvious what her account would say about Interpose, but she could say that by interposing a rock that prevents Tom from escaping the cart's path, John allows that harmful causal sequence to run its course. See Foot (1994).

independent category worthy of study. Why do Interpose and Remove cases belong to such a category? It is hard to provide an answer to this without giving a detailed account of the tripartite distinction, which is beyond the scope of this paper. In prior work, Barry and Øverland (2012) advance a particular view of what separates enabling harm from doing harm and allowing harm. According to this view, the presence or absence of two factors—relevant action and initiation of a complete causal process—determines the appropriate classification of cases of human conduct involving harm under these categories. Relevant action, the first factor, obtains if the question of how an agent is relevant to some harm refers to some action of theirs.⁹ In Push, for example, the correct answer to the question of how John is relevant to the harms suffered by Tom will refer to his pushing of the cart. The second factor obtains if there is a complete, intact causal process initiated by the agent’s action that links this action to the harm.¹⁰ Barry and Øverland argue that when an agent does harm, both of these factors obtain—we can explain how some action that agent has taken is relevant to the harm that occurs and there is a complete, intact causal process beginning with the agent’s relevant action that links this action to the harm. When an agent enables harm, they argue, the agent is connected to the harm that occurs by relevant action but without its being the case that there is a complete causal process between the agent’s action and the harm. By contrast with cases of doing harm and enabling harm, neither of these factors obtains in cases of allowing harm—it is the agent’s omission rather than an action that is relevant to the harmful outcome in these cases, and there isn’t a complete causal process between an omission and an outcome.¹¹

In this paper, we will not defend this view or discuss how it might relate to the theories of causal representation discussed in the previous section. We mention it here, instead, to emphasize that we do believe that some account can be given of the principles that underlie the differences between cases of doing, allowing, and enabling harm. Our contribution in this paper involves the use of experimental methods to argue that the very cases that moral philosophers employing the bipartite distinction have disagreed deeply

⁹ The basic idea of relevant action is due to Judith Jarvis Thomson (1996).

¹⁰ Hall (2002, 2004) refers to this kind of causal connection as exhibiting ‘locality’.

¹¹ This is not to say that there cannot be causation by omission.

about are more naturally regarded as cases of enabling harm. In making this argument, we go beyond the existing empirical research on enabling by studying what we take to be cases of enabling harm by human agents, and contrasting judgments about these cases with judgments about more straightforward cases. Of course, we believe that there are many different ways of doing, allowing, and enabling harm that are not found in the cases that we study. The value of focusing on the Interpose and Remove cases lies in the fact that they are structurally similar to what have been regarded by philosophers as clear-cut cases of doing harm and allowing harm: Push and Stayback cases respectively. This allows us to compare experimental participants' classifying judgments about conduct in these cases and conduct in Push and Stayback cases while holding other aspects of the cases fixed.

We will now present three key experiments whose results support the claim that moral philosophers, cognitive scientists, and moral psychologists should be investigating a tripartite distinction between doing, allowing, and enabling harm. The first experiment's results suggest that the traditional bipartite doing-allowing harm distinction cannot easily capture Interpose and Remove cases. The second experiment's results suggest that Interpose and Remove cases are more naturally classified as cases of enabling harm. On the basis of these results, we hold that the tripartite distinction between doing, allowing, and enabling harm better encompasses participants' intuitive classifications of these cases from the moral philosophical literature than the traditional bipartite distinction. Moral philosophers and researchers studying moral judgment should thus explore the nature and significance of the tripartite distinction rather than focusing only on the traditional bipartite distinction. The third experiment's results suggest that the tripartite distinction is also judged to be normatively significant.

2. Experiment 1: Doing-Allowing Harm Judgments Regarding Push, Stayback, Interpose, and Remove Cases

On our view, certain cases in the moral philosophical literature—Interpose and Remove cases—cannot easily be captured by the traditional bipartite doing-allowing

harm distinction. They are neither clear-cut cases of doing harm nor of allowing harm. In this experiment, we tried to show that this view is supported by the way participants using the traditional doing-allowing harm distinction classify Interpose and Remove cases in contrast to the way they classify Push and Stayback cases.

Participants were asked to read Push, Stayback, Interpose, and Remove cases and classify them as cases of either doing harm or allowing harm. We expected that Push cases would be more likely than the other types of cases to be classified as cases of doing harm, and Stayback cases more likely than the others to be classified as cases of allowing harm. We hypothesized, however, that participants' judgments about Interpose and Remove cases would be intermediate between their judgments about Push and Stayback cases, in the following sense. Interpose and Remove cases would be more likely to be classified as cases of allowing harm than Push cases, but less likely to be classified as cases of allowing harm than Stayback cases.¹² If this hypothesis turns out to be correct, we will have evidence that Interpose and Remove cases differ from Push and Stayback cases across people in the ease with which they can be subsumed under the traditional doing-allowing harm distinction. This evidence would lend support to the view that Interpose and Remove cases are neither cases of doing harm nor cases of allowing harm.

We employed two versions of each case. In one version an agent was culpable with respect to the harm suffered by another agent, while in the other version he was innocent. This manipulation of the first agent's moral status was introduced for two main reasons. First, we wanted to control for potential differences in the moral status of the first agent between the cases that might be inferred by participants if they were to read versions of the cases in which no information about the agent's moral status was provided. Second, prior work by Knobe and Fraser (2008), Hitchcock and Knobe (2009), and Cushman, Knobe and Sinnott-Armstrong (2008) has demonstrated that causal judgments and judgments about doing and allowing harm can be affected by moral and

¹² Because the traditional doing-allowing harm distinction is bipartite, the dependent variable in this experiment—classification of the case as one of doing harm or allowing harm—is dichotomous. Our hypothesis can thus also be expressed in terms of the comparative likelihood of judging cases as ones of doing harm. We predicted that participants would be less likely to classify Interpose or Remove cases as cases of doing harm than Push cases, but more likely to classify them as cases of doing harm than Stayback cases.

other normative judgments. We wanted to follow up on this work, even though our main interest in this paper is in how people classify Interpose and Remove cases as compared to Push and Stayback cases.¹³

2.1. Methods

2.1.1. Participants

The participants were 240 MTurk users. We limited our participant pool to MTurk users from the United States with a 98% or better approval rate for their work.

2.1.2. Materials

The experiment employed a 4 (Push, Stayback, Interpose, Remove) x 2 (culpable vs. innocent) factorial design. In the culpable version of each case, John knows that Tom will be harmed and it will not be costly for John to refrain from performing the action that will contribute to Tom being harmed (Push, Interpose, Remove) or to perform the action that would prevent Tom from being harmed (Stayback). In the innocent version of each case, John doesn't know and couldn't be expected to know that Tom will be harmed. All cases used in Experiments 1, 2 and 3 are listed in Appendix A. Below are two sample cases:

Push, Culpable: A cart stands at the top of a hill. John pushes it. John knows that Tom is sitting at the bottom of the hill, and refraining from pushing the cart would not be costly for John. The cart rolls down the hill and injures Tom.

Push, Innocent: A cart stands at the top of a hill. John pushes it. John doesn't know and couldn't be expected to know that Tom is sitting at the bottom of the hill. The cart rolls down the hill and injures Tom.

The names given to the cases did not appear in the study.

Participants were asked to read the case they were assigned and to respond to the prompt "This is most appropriately described as a case of:" by selecting "doing harm" or "allowing harm." The order of the choices was counterbalanced.

¹³ Foreseeability of harm and cost to self of refraining from taking a harmful action (Push, Interpose, Remove) or engaging in harm-preventing action (Stayback) are used as proxies for culpability. It is conceivable, however, given prior work by Lagnado and Channon (2008) that the effect of foreseeability on causal judgments is driving the effects that we observe and suggest are attributable to moral judgments. We are grateful to an anonymous reviewer for this point.

2.1.3. Procedure

The experiment employed a between-participant design. Participants were asked to read and respond to only one of the eight cases, to which they were randomly assigned. Responses to each case were capped at $n = 30$.

2.2. Results

Figure 1 represents the proportion of participants who judged each case to be one of allowing harm rather than doing harm (see Table 1 in Appendix B for the proportions of each response in Experiment 1). We used logistic regressions to measure the influence of the independent variables, case type (Push, Stayback, Interpose, Remove) and moral status (coded as 0 = culpable, 1 = innocent), on the dependent variable, classification of the case as one of doing harm or allowing harm. For the first logistic regression, responses to Push were used as the reference group, and for the second logistic regression, responses to Stayback were used as the reference group. We chose to run two regressions with these reference groups so that we could compare judgments about Interpose and Remove cases to judgments about Push cases and judgments about Stayback cases, which would be sufficient to test our main hypothesis. Tables 1 and 2 show the results of these regressions. Table 1 presents the main effect of case type and interaction model results, and Table 2 presents the main effect of moral status.

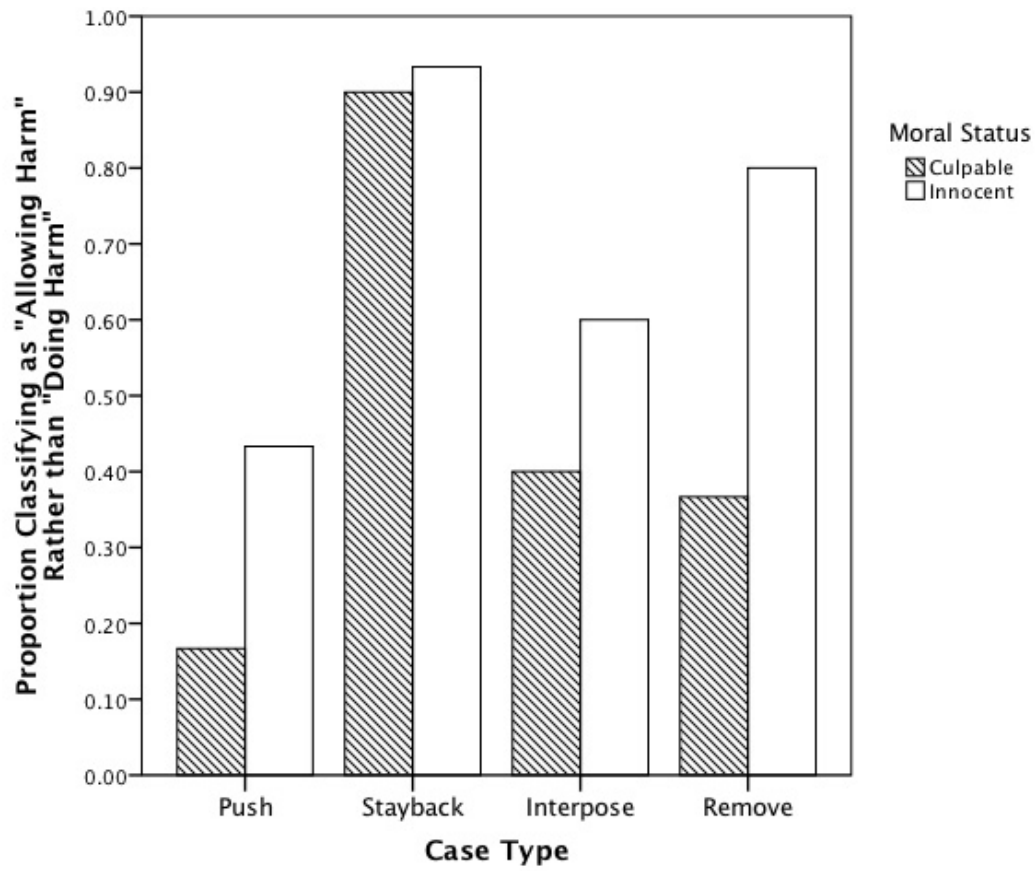


Figure 1. Proportion of Participants Who Classified Each Case as One of Allowing Harm Rather than Doing Harm.

Table 1

DV: Classified as Allowing Instead of Doing	Odds Ratio	p-value	95% C.I. for Odds Ratio	
			Lower	Upper
<u>Main Effects Model</u>				
Innocent vs. Culpable	3.49	<.001	1.89	6.44
Stayback vs. Push	32.52	<.001	10.68	99.04
Interpose vs. Push	2.53	.021	1.15	5.54
Innocent vs. Stayback ^a	.078	<.001	.03	.23
Remove vs. Push	3.66	.001	1.65	8.09
Remove vs. Stayback ^a	.11	<.001	.04	.33
<u>Interaction Model</u>				
Innocent X Stayback	.41	.427	.04	3.74
Innocent X Interpose	.59	.512	.12	2.87
Innocent X Remove	1.81	.488	.34	9.62
Innocent X Interpose ^a	1.45	.734	.17	12.20
Innocent X Remove ^a	4.44	.184	.49	40.02

^a indicates the use of Stayback as the reference group (rather than Push)

Table 2

Main Effect of Moral Status

DV: Classified as Allowing Instead of Doing	Odds Ratio	p-value	95% C.I. for Odds Ratio	
			Lower	Upper
Innocent vs. Culpable in Push Condition	3.82	.029	1.15	12.71
Innocent vs. Culpable in Stayback Condition	1.56	.643	.24	10.05
Innocent vs. Culpable in Interpose Condition	2.25	.124	.80	6.32
Innocent vs. Culpable in Remove Condition	6.91	.001	2.16	22.10

Unsurprisingly, participants responded to Push and Stayback cases as we expected. Controlling for moral status, participants were significantly more likely to classify Push cases than Stayback, Interpose, or Remove cases as cases of doing harm, and significantly more likely to classify Stayback cases than Push, Interpose, or Remove cases as cases of allowing harm.

Our main hypothesis was confirmed. Controlling for the effect of moral status, judgments about Interpose and Remove cases were intermediate between judgments about Push and Stayback cases. Participants responding to an Interpose case were significantly more likely to classify the case as one of allowing harm than participants who received a Push case (OR = 2.53 [1.15, 5.54], $p = .02$), and significantly less likely to classify an Interpose case as one of allowing harm than participants who received a Stayback case (OR = .08 [.03, .23], $p < .001$). Participants responding to a Remove case were significantly more likely to classify the case as one of allowing harm than participants who received a Push case (OR = 3.66 [1.65, 8.09], $p = .001$), and significantly less likely to classify the case as one of allowing harm than participants responding to a Stayback case (OR = .11 [.04, .33], $p < .001$).

There was an overall main effect of moral status whereby participants assigned to an innocent version of a case were significantly more likely to classify the case as one of allowing harm than participants assigned to a culpable version of a case (OR = 3.49 [1.89, 6.44], $p < .001$). Participants assigned to the innocent version of Push were significantly more likely to classify the case as one of allowing harm than participants assigned to the culpable version of Push (OR = 3.82 [1.15, 12.71], $p = .03$). Similarly, participants assigned to the innocent version of Remove were significantly more likely to classify the case as one of allowing harm than participants assigned to the culpable version of Remove (OR = 6.91 [2.16, 22.10], $p = .001$). There was no main effect of moral status on responses to the Stayback cases or the Interpose cases. No significant case type x moral status interaction effects were found in the analyses conducted using either the Push or Stayback reference groups.

2.3. Discussion

The results of Experiment 1 confirmed our main hypothesis about the classification of Interpose and Remove cases as compared with the classification of Push and Stayback cases under the traditional doing-allowing harm distinction. Whereas Push and Stayback cases were more likely than any other of the kinds of cases to be given the classification that philosophers have mostly regarded as uncontroversial for them, Interpose and Remove cases were treated differently. An Interpose or Remove case was significantly more likely to be classified as a case of allowing harm than a Push case, but significantly less likely to be classified as a case of allowing harm than a Stayback case. Judgments about these cases, in other words, varied more than judgments about what have been regarded as clear-cut cases of doing harm, on the one hand, and clear-cut cases of allowing harm, on the other. This evidence suggests that the traditional doing-allowing harm distinction is not able to provide easy classification of Interpose and Remove cases, or at least that individual variation with respect to doing and allowing judgments is more sensitive to features of Interpose and Remove cases than to features of Push and Stayback.

It is worth noting that there was a significant main effect of moral status on judgments about the Push cases. While the majority of participants in each version of Push classified the case they read as one of doing harm, participants assigned to the innocent version of the case were significantly more likely to classify their version as one of allowing harm than participants assigned to the culpable version. This fits with the effect of moral judgments on doing-allowing judgments observed in Cushman et al. (2008), and in Knobe's work on the influence of normative judgments on what are thought to be non-normative judgments more generally.¹⁴ The analogous main effect of moral status obtained for Remove cases. No similar effect of moral status was found for Stayback or Interpose cases.

The results of Experiment 1 suggest that the traditional bipartite doing-allowing harm distinction cannot capture Interpose and Remove cases. In the next experiment, we

¹⁴ Knobe first reported the effect of people's moral judgments on their judgments about intentional action in Knobe (2003). For an overview of this research program, see Knobe (2010).

investigated whether the tripartite doing-allowing-enabling harm distinction is better able to do so.

3. Experiment 2: Doing-Allowing-Enabling Harm Judgments Regarding Push, Stayback, Interpose, and Remove Cases

In Experiment 1, we found strong evidence to support the claim that the doing-allowing harm distinction cannot accurately classify Interpose and Remove cases. With Experiment 2, we sought to show that our proposed doing-allowing-enabling harm distinction does a better job of capturing these cases. Experiment 2 differs from Experiment 1 only in virtue of adding the category of enabling harm to the options that participants can choose to classify the cases. We hypothesized that participants would generally classify Interpose and Remove cases as instances of enabling harm and, controlling for moral status, would be significantly more likely to do so than participants classifying Push and Stayback cases. If this hypothesis holds, we will have good evidence to support our view that the tripartite doing-allowing-enabling harm distinction is better able to capture these cases than the traditional bipartite doing-allowing harm distinction.

3.1. Method

3.1.1. Participants

The participants were 240 MTurk users. We limited our participant pool to MTurk users from the United States with a 98% or better approval rate for their work. Participants from Experiment 1 were not allowed to participate in Experiment 2.

3.1.2. Materials and Procedure

The experiment employed a 4 (Push, Stayback, Interpose, Remove) x 2 (culpable vs. innocent) between-participant factorial design. Participants read only one of the eight cases, to which they were randomly assigned, and were asked to respond to the prompt “This is most appropriately described as a case of:” by selecting “doing harm,” “allowing

harm,” or “enabling harm.” The order of the choices was counterbalanced. Responses to each case were capped at $n = 30$.

3.2. Results

Figure 2 shows the proportion of enabling harm vs. doing harm or allowing harm responses to the four classes of cases (see Table 2 in Appendix B for the proportions of each response in Experiment 2). Because we were especially interested in judgments about enabling harm, we grouped doing harm and allowing harm together in our analyses. We used logistic regressions to measure the influence of the independent variables, case type (Push, Stayback, Interpose, Remove) and moral status (0 = culpable, 1 = innocent), on the dependent variable, classification of the case as one of enabling harm vs. either doing harm or allowing harm. In these analyses, we also used Push as the reference group for the first logistic regression and Stayback as the reference group for the second logistic regression. We chose to run two regressions with these reference groups so that we could compare judgments about Interpose and Remove cases to judgments about Push and Stayback cases, which would be sufficient to test our hypothesis. Tables 3 and 4 show the results of these regressions. Table 3 presents the main effect of case type and interaction model results, and Table 4 presents the main effect of moral status.




Figure 2. Proportion of Participants Who Classified Each Case as One of Enabling Harm Rather than Doing or Allowing Harm.

Table 3

DV: Classified as Enabling Instead of Other Options	Odds Ratio	p-value	95% C.I. for Odds Ratio	
			Lower	Upper
<u>Main Effects Model</u>				
Innocent vs. Culpable	2.87	<.001	1.61	5.10
Stayback vs. Push	.43	.057	.18	1.03
Interpose vs. Push	4.06	<.001	1.85	8.89
Innocent vs. Stayback ^a	9.42	<.001	3.93	22.55
Remove vs. Push	2.10	.057	.98	4.50
Remove vs. Stayback ^a	4.86	<.001	2.07	11.41
<u>Interaction Model</u>				
Innocent X Stayback	.11	.017	.02	.67
Innocent X Interpose	.39	.271	.07	2.10
Innocent X Remove	.40	.279	.08	2.09
Innocent X Interpose ^a	3.59	.145	.65	20.04
Innocent X Remove ^a	3.73	.125	.70	20.03

^a indicates the use of Stayback as the reference group (rather than Push)

Table 4

Main Effect of Moral Status

DV: Classified as Enabling Instead of Other Options	Odds Ratio	p-value	95% C.I. for Odds Ratio	
			Lower	Upper
Innocent vs. Culpable in Push Condition	7.43	.002	2.08	26.55
Innocent vs. Culpable in Stayback Condition	.80	.739	.22	2.97
Innocent vs. Culpable in Interpose Condition	2.88	.062	.95	8.72
Innocent vs. Culpable in Remove Condition	2.98	.041	1.04	8.53

Enabling harm was the modal response to the Interpose cases (Culpable Version: 53.33%, Innocent Version: 76.67%). This was also true for the Remove cases, although in the culpable version of this case this was only by a margin of one response over doing harm, (Culpable Version: 36.67%, Innocent Version: 63.33%).

Controlling for the effect of moral status, participants responding to an Interpose case were significantly more likely to classify the case as one of enabling harm than participants responding to a Push case (OR = 4.06 [1.85, 8.89], $p < .001$) or a Stayback case (OR = 9.42 [3.93, 22.55], $p < .001$). Participants responding to a Remove case were also significantly more likely to classify their case as one of enabling harm than participants responding to a Stayback case (OR = 4.86 [2.07, 11.41], $p < .001$). The comparison between enabling harm judgments about Remove cases and Push cases was marginally significant in the direction predicted by the hypothesis (OR = 2.1 [.98, 4.5], $p = .06$).

There was an overall main effect of moral status whereby participants assigned to an innocent version of a case were significantly more likely to classify the case as one of enabling harm than participants assigned to a culpable version of a case (OR = 2.87 [1.61, 5.1], $p < .001$). A significant main effect of moral status on participants' classifications was observed for the Push cases and Remove cases. Participants assigned to the innocent version of Push were significantly more likely to classify the case as one of enabling harm than participants assigned to the culpable version of Push (OR = 7.43 [2.08, 26.55], $p = .002$). Participants assigned to the innocent version of Remove were also significantly more likely to classify the case as one of enabling harm than participants assigned to the culpable version of Remove (OR = 2.98 [1.04, 8.53], $p = .04$). No similar main effect of moral status was found for responses to Stayback or Interpose cases.

The only significant interaction effect between John's moral status and the case that participants received was observed for the Push/Stayback comparison coding ($p = .017$, shown in Table 3 as Innocent x Stayback with Push as the reference group). Follow-up analyses showed that participants were significantly more likely to classify the case they were assigned to as one of enabling harm if they were in the Push condition and

assigned to the innocent version rather than the culpable version (OR = 3.45, [1.87, 6.37], $p < .001$).

3.3. Discussion

The results of Experiment 2 confirmed our hypothesis regarding how participants would classify the Interpose and Remove cases and apply the doing-allowing-enabling harm distinction to these cases as compared to the Push and Stayback cases. More participants regarded Interpose and Remove cases as instances of enabling harm than doing harm or allowing harm. Statistical analyses showed that participants were significantly more likely to classify Interpose cases as instances of enabling harm than Push and Stayback cases, cases which have been treated by philosophers as uncontroversial cases of doing and allowing harm respectively. Participants were also significantly more likely to classify Remove cases as instances of enabling harm than Stayback cases, and marginally more likely to classify Remove cases as instances of enabling harm than Push cases. These results suggest that participants generally regard Interpose and Remove cases as instances enabling harm, and that they apply the doing-allowing-enabling harm distinction discriminately in the way that we expected.

Interestingly, there was a significant main effect of moral status on judgments about the Push cases. Participants assigned to the innocent version of Push were significantly more likely to regard it as an instance of enabling harm than participants assigned to the culpable version. In fact, enabling harm was the modal response to the innocent version of Push (53.33%). There was also a significant difference in enabling harm judgments between the innocent and culpable versions of Remove. These results suggests that the effect of moral judgments on doing-allowing harm judgments that Cushman, Knobe and Sinnott-Armstrong have provided evidence for can also be observed with respect to doing-allowing-enabling harm judgments. While our main interest in this paper is not in the effect of moral judgments on classifying judgments, we suggest that this effect is worth studying in future research on the doing-allowing-enabling harm distinction.

4. Experiment 3: The Doing-Allowing-Enabling Harm Distinction and Compensatory Responsibility

We have been arguing that a tripartite doing-allowing-enabling harm distinction does a better job of classifying judgments about well-known cases found in the moral philosophical literature than the traditional bipartite doing-allowing harm distinction. A further question arises as to whether enabling harm is generally regarded as normatively like doing harm or allowing harm, or as distinct from both of these categories. While the main focus of this paper is on the descriptive classification of cases of human conduct involving harm, we wanted to explore evidence that bears on this question. If enabling harm is regarded as normatively distinct from both doing harm and allowing harm along the lines of some normative factor, this would lend further support to our view that research in moral philosophy and moral psychology should employ the tripartite distinction. We hypothesized that one factor in terms of which people regard enabling harm as normatively distinct from doing harm and allowing harm is the assignment of compensatory responsibility—responsibility to compensate for harm (Goodin 1989). It seemed to us that, like agents who culpably do harm, agents who culpably enable harm have a more stringent responsibility to compensate the victims of harm than agents who culpably allow harm. Yet it also seemed to us that, like agents who innocently allow harm, agents who innocently enable harm have less stringency compensatory responsibility than agents who culpably do harm. Using the eight cases from the previous two studies, our third study explored this possibility.

4.1. Methods

4.1.1. Participants

The participants were 240 MTurk users. We limited our participant pool to MTurk users from the United States with a 98% or better approval rate for their work. Participants from Experiments 1 and 2 were not allowed to participate in Experiment 3.

4.1.2. Materials and Procedure

The experiment employed a 4 (Push, Stayback, Interpose, Remove) x 2 (culpable vs. innocent) between-participant factorial design. Participants were asked to read only one of the eight cases, to which they were randomly assigned, and were asked to rate their level of agreement with the statement “John should compensate Tom for the injury” on a 7-point Likert Scale (1 = Not At All, 7 = Fully Agree). Responses to each case were capped at $n = 30$.

4.2. Results

Figure 3 represents the mean levels of agreement with the statement regarding John’s compensatory responsibility in each of the eight cases (see Table 3 in Appendix B for means with standard deviations in response to each case). We used a two-way ANOVA and post hoc tests to determine whether our hypothesis was borne out by the responses. All analyses for Experiment 3 were conducted using the Bonferonni correction for multiple comparisons. As our main hypothesis pertains to differences among culpable cases, on the one hand, and differences among innocent cases, on the other, we focused on the pairwise comparisons in post hoc tests on the interaction of Case Type and Moral Status.




Figure 3. Mean Levels of Agreement with the Statement Regarding Compensatory Responsibility in Response to Each Case.

The two-way ANOVA revealed that there was a significant main effect of Case Type on responses to the statement regarding compensatory responsibility, $F(3, 232) = 17.9, p < .001$. Controlling for Moral Status, means levels of agreement with the statement were significantly higher in response to Push ($M = 5.52, SD = 1.36$) than in response to Stayback ($M = 3.46, SD = 1.96$), Interpose ($M = 4.30, SD = 2.44$), or Remove ($M = 4.18, SD = 2.22$), $ps < .001$. There was also a significant main effect of Moral Status, $F(1, 232) = 143.45, p < .001$. Controlling for Case Type, mean levels of agreement with the statement in response to culpable versions of the cases ($M = 5.58, SD = 1.69$) were significantly higher than in response to innocent versions of the cases ($M = 3.15, SD = 1.88$), $p < .001$. These main effects were qualified by a significant interaction effect between Case Type and Moral Status, $F(3, 232) = 7.58, p < .001$.

Our hypothesis was confirmed by pairwise comparisons among the culpable and innocent cases, which also help to explain the interaction effect between Case Type and Moral Status. Means levels of agreement with the statement regarding compensatory responsibility were significantly higher in response to the culpable version of Push ($M = 6.17$, $SD = .91$), Interpose ($M = 6.00$, $SD = 1.70$), or Remove ($M = 5.87$, $SD = 1.43$) than in response to the culpable version of Stayback ($M = 4.27$, $SD = 1.87$), $ps < .01$. There were no significant differences between the mean levels of agreement with the statement in response to the culpable versions of Push, Interpose, or Remove. Mean levels of agreement with the statement were significantly higher in response to the innocent version of Push ($M = 4.87$, $SD = 1.43$) than in response to the innocent versions of Stayback ($M = 2.63$, $SD = 1.71$), Interpose ($M = 2.60$, $SD = 1.81$), or Remove ($M = 2.5$, $SD = 1.46$), $ps < .001$. There were no significant differences between the mean levels of agreement with the statement in response to the innocent versions of Stayback, Interpose, and Remove.

4.3. Discussion

The results of Experiment 3 confirmed our hypothesis that enabling harm is generally regarded as normatively distinct from both doing harm and allowing harm in terms of assigning compensatory responsibility. Among the culpable cases, on one hand, mean levels of agreement with the statement were significantly higher in response to the Push, Interpose, and Remove cases than in response to the Stayback case and not significantly different from one another. Among the innocent cases, on the other hand, mean levels of agreement with the statement were significantly lower in response to the Interpose, Remove, and Stayback cases than in response to the Push case and not significantly different from one another. Interpose and Remove cases, which we hold are cases of enabling harm, are regarded as normatively distinct from Push and Stayback cases, which have been treated by philosophers as straightforward cases of doing and allowing harm.

5. General Discussion

In this paper, we argued that the results of three experimental studies support the view that moral philosophers and moral psychologists should employ a tripartite distinction between doing, allowing, and enabling harm. In previous work, theorists and empirical researchers studying moral judgment have typically employed the traditional bipartite doing-allowing harm distinction. We have argued that this traditional distinction cannot capture well-known cases in the moral philosophical literature—Interpose and Remove cases—whereas the tripartite doing-allowing-enabling harm distinction is better able to do so. We also provided evidence suggesting that enabling harm is judged to be normatively distinct from doing harm and allowing harm, at least when it comes to assigning compensatory responsibility.

In his influential treatment, Jonathan Bennett described three possible outcomes of an investigation of the doing-allowing harm distinction:

“(1) The ordinary uses of these concepts reflect an underlying jumble, a big mess that resists analysis of any kind; (2) Ordinary uses express a simple and clear distinction; (3) People are guided in the use of these terms by simple and clear distinctions, but sometimes mischaracterize conduct due to distortions in judgment of various kinds.”¹⁵

The results of our studies suggest that none of the three possibilities that Bennett envisions obtains, because there is no one binary, exhaustive distinction between doing and allowing harm. There is, instead, a tripartite distinction that participants draw between doing, allowing, and enabling harm. Judgments about doing harm and allowing harm, to be sure, are not simply a mess. Participants are more likely to regard Push cases as cases of doing harm and Stayback cases as cases of allowing harm than the comparison class of cases that we have looked at in this paper. However, things do indeed get messy when one attempts to classify cases like Interpose and Remove as instances of doing harm or allowing harm. The results of the first two studies presented in this paper suggest that they belong to the distinct category of enabling harm. The doing-allowing-enabling harm distinction, a distinction that gives this category its proper

¹⁵ Bennett (1995), pp. 99-100.

standing, should replace the doing-allowing harm distinction in future research on the nature and significance of human conduct involving harm.

As noted above, the deep interest of philosophers in the doing-allowing harm distinction is due to the fact that it is often taken to be normatively significant. What then, is the normative status of enabling harm? Does it possess normative characteristics that are distinct from doing harm on the one hand, and allowing harm on the other? In previous work, Barry and Øverland (2012) argue that it is indeed a normatively distinct category. The results of experiment 3 suggest that experimental participants, too, seem to regard it as normatively distinct, at least when it comes to judgments about whether an agent bears compensatory responsibility for harm. Our primary focus has been on the descriptive classification of what we hold to be cases of doing, allowing, and enabling harm in the moral philosophical literature. We also showed that people judge the cases that we have argued are cases of enabling harm as normatively distinct from cases of doing harm and allowing harm in terms of this normative factor. It is an open question whether people tend to judge these cases as normatively distinct in other respects. We thus conclude by signaling directions for future research in moral philosophy and moral psychology.

Appendix A

All Cases Used in Experiments 1-3.

Push, Culpable: A cart stands at the top of a hill. John pushes it. John knows that Tom is sitting at the bottom of the hill, and refraining from pushing the cart would not be costly for John. The cart rolls down the hill and injures Tom.

Push, Innocent: A cart stands at the top of a hill. John pushes it. John doesn't know and couldn't be expected to know that Tom is sitting at the bottom of the hill. The cart rolls down the hill and injures Tom.

Stayback, Culpable: A cart is rolling down a hill. John could put a rock in the way of the cart that would stop it, but does not. John knows that Tom is sitting at the bottom of the hill, and it would not be costly for John to put the rock in the cart's way. The cart rolls down the hill and injures Tom.

Stayback, Innocent: A cart is rolling down a hill. John could put a rock in the way of the cart that would stop it, but does not. John doesn't know and couldn't be expected to know that Tom is sitting at the bottom of the hill. The cart rolls down the hill and injures Tom.

Interpose, Culpable: A cart accidentally starts rolling down a hill. Tom, who is sitting at the bottom of the hill, won't be injured by the cart if he can get out of the way of its path. John knows about the cart. He also knows that if he puts a rock on the ground, it will stop Tom from getting out of the cart's path and Tom will be injured. Not putting a rock on the ground would not be costly for John. John puts a rock on the ground. The rock stops Tom, and he is injured by the cart.

Interpose, Innocent: A cart accidentally starts rolling down a hill. Tom, who is sitting at the bottom of the hill, won't be injured by the cart if he can get out of the way of its path. John doesn't know about the cart and couldn't be expected to know that if he puts a rock on the ground, it will stop Tom from getting out of the cart's path and Tom will be injured. John puts a rock on the ground. The rock stops Tom, and he is injured by the cart.

Remove, Culpable: A cart is rolling downhill towards a point where there is a rock that would bring it to a stop. John knows about the cart. He also knows that if he removes the rock, Tom, who is sitting at the bottom of the hill, will be injured. Not removing the rock would not be costly for John. John removes the rock. The cart rolls down the hill and injures Tom.

Remove, Innocent: A cart is rolling downhill towards a point where there is a rock that would bring it to a stop. John doesn't know about the cart and couldn't be expected to know that if he removes the rock, Tom, who is sitting by the bottom of the hill, will be injured. John removes the rock. The cart rolls down the hill and injures Tom.

Appendix B

Table 1

Proportions of Participants in Experiment 1 Classifying Each Case as an Instance of “Doing Harm” and “Allowing Harm”

Case	Doing Harm	Allowing Harm
Push, Culpable	.83	.17
Push, Innocent	.57	.43
Stayback, Culpable	.10	.90
Stayback, Innocent	.07	.93
Interpose, Culpable	.60	.40
Interpose, Innocent	.40	.60
Remove, Culpable	.63	.37
Remove, Innocent	.20	.80

Table 2

Proportions of Participants in Experiment 2 Classifying Each Case as an Instance of “Doing Harm,” “Allowing Harm,” and “Enabling Harm”

Case	Doing Harm	Allowing Harm	Enabling Harm
Push, Culpable	.80	.07	.13
Push, Innocent	.33	.13	.53
Stayback, Culpable	.03	.77	.20
Stayback, Innocent	.03	.80	.17
Interpose, Culpable	.33	.13	.53
Interpose, Innocent	.07	.17	.77
Remove, Culpable	.33	.30	.37
Remove, Innocent	.13	.23	.63

Table 3

Means Levels of Agreement in Response to the Statement Regarding Compensatory Responsibility and Standard Deviations for Each Case

Case	<i>Mean</i>	<i>SD</i>
Push, Culpable	6.17	.91
Push, Innocent	4.87	1.43
Stayback, Culpable	4.27	1.87
Stayback, Innocent	2.63	1.71
Interpose, Culpable	6.0	1.70
Interpose, Innocent	2.6	1.81
Remove, Culpable	5.87	1.43
Remove, Innocent	2.5	1.46

References

- Barry, C. and Øverland, G. "The Feasible Alternatives Thesis: Kicking away the livelihoods of the global poor." *Politics, Philosophy and Economics* 11(1) (2012): 97-119.
- Bennett, J. *The Act Itself*. New York: Oxford University Press, 1995.
- Cushman, F.A., Knobe, J. & Sinnott-Armstrong, W. "Moral appraisals affect doing/allowing judgments." *Cognition* 108 (2008): 281-289.
- Donagan, A. *The Theory of Morality*. Chicago: Chicago University Press, 1977.
- Foot, P. "Killing and Letting Die." In *Killing and Letting Die*, eds. A. Norcross and B. Steinbock. New York: Fordham University Press, 1994.
- Godfrey-Smith, P. "Causal pluralism." In *The Oxford Handbook of Causation*, eds. H. Beebe, C. Hitchcock, and P. Menzie, Oxford: Oxford University Press, 2010.
- Goldvarg, E. & Johnson-Laird, P.N. "Naive causality: a mental model theory of causal meaning and reasoning." *Cognitive Science* 25 (2001): 565-610.
- Goodin, R.E. "Theories of Compensation," *Oxford Journal of Legal Studies* 9 (1989): 56-75
- Hall, N. "Non-Locality on the Cheap? A New Problem for Counterfactual Analyses of Causation." *Noûs* 36(2) (2002): 276-94.
- . "Two Concepts of Causation." In *Causation and Counterfactuals*, eds. J. Collins, N. Hall and L.A. Paul. Cambridge, Massachusetts: MIT Press, 2004. Pp. 225-76.
- Hanser, M. "Killing, Letting Die and Preventing People from Being Saved." *Utilitas* 11 (1999): 277-95.
- Hitchcock, C. & Knobe, J. "Cause and norm." *Journal of Philosophy* 106(11) (2009): 587-612.
- Howard-Snyder, F. "Doing and Allowing Harm." In *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, 2008.
<plato.stanford.edu/entries/doing-allowing/>
- Knobe, J. "Intentional Action and Side-Effects in Ordinary Language." *Analysis* 63 (2003): 190-193.
- Knobe, J. and Fraser, B. "Causal judgment and moral judgment: Two experiments." In *Moral psychology*, ed. W. Sinnott-Armstrong. Cambridge, MA: MIT Press, 2008. Pp. 441-448.
- Knobe, J. "Person as scientist, person as moralist." *Behavioral and Brain Sciences* 33 (2010): 315-329.

- Lagnado, D.A. and Channon, S. "Judgments of cause and blame: The effects of intentionality and foreseeability." *Cognition* 108 (2008): 754-770.
- Lombrozo, T. "Causal-explanatory pluralism: How intentions, functions, and mechanisms influence causal ascriptions." *Cognitive Psychology* 61 (2010): 303-332.
- Mackie, J.L. *The Cement of the Universe*. Oxford: Oxford University Press, 1980.
- McMahan, J. "A Challenge to Common Sense Morality." *Ethics* 108(2) (1998): 394-418.
- . "Killing, Letting Die and Withdrawing Aid." *Ethics* 103(2) (1993): 250-79.
- Quinn, W. "Actions, Intentions, and Consequences: The Doctrine of Doing and Allowing." (Originally published in *Philosophical Review* 98(3) (1989).) In *Killing and Letting Die*, eds. A. Norcross. and B. Steinbock. New York: Fordham University Press, 1994.
- Rickless, S.C. "The Moral Status of Enabling Harm." *Pacific Philosophical Quarterly* 92(4) (2011): 66-86.
- Slooman, S.A., Barbey, A.K. and Hotaling, J. "A causal model of the meaning of 'cause,' 'enable,' and 'prevent.'" *Cognitive Science* 33 (2009): 21-50.
- Talmy, L. "Force dynamics in language and cognition." *Cognitive Science* 12 (1988): 49-100.
- Thomson, J. "Critical Study of Jonathan Bennett's *The Act Itself*." *Nous* 30(4) (1996): 545-57.
- Vihvelin, K., and Tomkowiak, T. "The Dif." *Journal of Philosophy* 102(4) (2005): 183-205.
- Wolff, P. "Representing causation." *Journal of Experimental Psychology: General*, 136 (2007): 82-111.