

Is there really an omission effect?

Pascale Willemsen^{1*}, Kevin Reuter^{2*}

¹ Ruhr University Bochum

² University of Bern

* Both author contributed equally to this paper.

Corresponding Author:

Pascale Willemsen

Email: Pascale.Willemsen@RUB.de

Accepted for publication in *Philosophical Psychology*. Please do not cite without permission

Abstract

The omission effect, first described by Spranca et al. (1991), has been extensively studied in the proceeding decades and repeatedly confirmed (Cushman et al., 2012). Accordingly, all else being equal, most people judge it to be morally worse to actively bring about a negative effect than to passively allow that effect to happen. In this paper, we provide new experimental data challenging previous studies on the omission effect, both methodologically and philosophically. We argue that those studies have failed to control for the equivalence of rules that get violated by actions and omissions. Once equivalent norms are introduced, our results show that the omission effect is eliminated even if the negative outcome of the behavior is foreseen and intended by the agent. We will show that the omission effect does not constitute a basic moral disposition but occurs rather exclusively in more complex moral situations. Building on these empirical results, we put into doubt two influential explanations of the omission effect – the Causal Relevance Hypothesis and the Overgeneralization Hypothesis – and provide a novel explanation. Furthermore, we discuss various ramifications for the interplay between our understanding of omissions and legal systems.

1 Introduction

1.1 *Introducing the Omission Effect*

Is it morally worse to poison someone compared to not warning someone who is about to eat poisoned food? If you think it is, you are in good company. Many philosophers (Foot, 1985; Moore, 2010; Quinn, 1989) have even argued for the more general claim that it is always morally worse to actively harm someone than to passively allow such harm to occur, a claim known as the *Doctrine of Doing and Allowing*. This doctrine seems to reflect the intuitions of many laypeople. In an empirical study, Spranca et al. (1991) directly tested people's responses to scenarios that described morally condemnable actions and omissions: In order to win the final, tennis player John West intends to poison Ivan Lendl during dinner the night before the match, by making sure Ivan eats the salad dressing he is allergic to. The alternative endings of the story vary only in regards to whether John actively recommends the poisonous dressing or allows Ivan to eat the poisonous dressing (in the case in which Ivan himself chooses the dressing he is allergic to). Spranca et al. (1991) found that even though in both scenarios John's behavior is condemned, actively recommending the house dressing is considered morally worse than failing to warn Ivan. This asymmetry between actions and omissions has since been called *omission bias* (Ritov and Baron, 1999; Spranca et al., 1991) or *omission effect* (Cushman et al., 2012).¹ Accordingly, many people judge omissions less harshly than actions, and even prefer omissions with a bad outcome over actions with a lower amount of the same bad outcome (Cushman et al., 2006; Ritov and Baron, 1999; Spranca et al., 1991).

¹ In this paper we refer to it as 'omission *effect*' rather than 'omission *bias*'. We prefer to stay neutral on whether the omission effect is in fact a bias or reflects a reasonable principle that helps distinguish between different kinds of behaviour.

Not only do people show such an omission effect in controlled experiments, most legal systems issue more severe punishments for committing rather than omitting an action even if all other factors seem to be the same.² The current debate on euthanasia provides an interesting case at hand. Whereas many countries allow for passive euthanasia, active euthanasia is prohibited in most countries – notable exceptions include Belgium, Luxembourg, the Netherlands, and the State of Oregon. The US Supreme Court stated that the “distinction between letting a patient die and making that patient die is important, logical, rational, and well established: It comports with fundamental legal principles of causation, and intent; and has been widely recognized and endorsed in the medical profession, the state courts, and the overwhelming majority of state legislatures, which, like New York's, have permitted the former while prohibiting the latter.” (Vacco v. Quill, 1997). Thus, the *Doctrine of Doing and Allowing* is a philosophical distinction that seems to accurately map folk intuitions as well as some fundamental principles in codes of law.

However, several philosophers have raised serious concerns about the moral relevance of this distinction. Jonathan Bennett for instance claims that such a distinction does not draw any morally significant line, but rather creates a “complex, shallow mess” (Bennett, 2011, 68). According to Bennett, what really matters is the change in likelihood that a certain event occurs given the agent’s behavior. Raising the probability of a bad outcome is morally impermissible, irrespective of whether the probability was raised by committing or omitting an action.

Peter Singer (2011) also challenges the Doctrine of Doing and Allowing by adopting a utilitarian point of view. He argues that under certain conditions it is even worse to let

² This difference is particularly striking in American law, in which omissions are punishable in a very limited set of cases.

someone die than to kill that person, whereby the debate about euthanasia provides one of the most striking examples. Consider the case of Baby Doe, an infant born with Down's syndrome and additional severe disabilities, which gained much attention in the 1980s. Baby Doe was born with serious disfigurement of his gullet that made it impossible for him to eat. After long deliberation his parents decided against surgery. After five days, Baby Doe died of dehydration and starvation. While this form of passive euthanasia was a legal course for his parents to take, administering a lethal injection was not. Singer claims: "The death resulting from the failure to operate in these circumstances is neither swift nor painless. [...] To 'allow nature to take its course', withholding treatment but refusing to kill, would obviously be wrong." (2011, 186).³

1.2 Explaining the Omission Effect

In this paper, we shall not address the normative question of whether the Doctrine of Doing and Allowing is a desirable or even justifiable moral principle. Instead, we will investigate if and when people believe that actions are indeed worse than omissions when the same outcome occurs. Although several explanations have been offered to account for the omission effect stated above, we will focus in this essay on two of the most interesting and influential theories.

First, Spranca et al. (1991) found that many people explain differences in their moral judgments by a difference in the causal status of actions and omissions – call this the *Causal Relevance Hypothesis*. Several participants in Spranca et al.'s study reported that the agent's action was causally more relevant to the outcome than the omission, for not

³ While Singer makes a compelling argument in this specific case, it should be noted, however, that Singer compares two situations in which one outcome is radically different from the other: death with little pain vs. death with a lot of pain. Many people will not be surprised that certain legal principles lose some of their force if the consequences vary as strongly between two options as in the case of Baby Doe.

intervening simply means to let nature take its course. Cushman & Young also propose “that the action/omission distinction affects moral judgment principally via causal attribution” (Cushman & Young, 2011, 1069), whereas intentions and the outcome’s severity play a subordinate role. The results of both studies seem to provide some evidence that the agent’s causal relevance for the outcome triggers different moral evaluations. However, the direction of influence between causal and blame attributions is far from clear: Although their accounts differ in important respects, both Knobe and Fraser (2008) as well as Alicke (2008) argue that causal judgements are strongly influenced by moral judgements. Thus, when a subject conceives of another person’s behavior as morally blameworthy, then this is likely to also increase its perceived causal relevance. It is therefore at least controversial to identify causal attributions as the possible source of the omission effect.

A second explanation has been suggested by Ritov & Baron (1999) among others. Ritov and Baron argue that it can hardly be denied that actions are usually morally worse than omissions: Omissions often correlate with lesser amounts of malicious intent, the desired outcome of omissions is less certain, etc. According to the *Overgeneralization Hypothesis*, when subjects judge the two endings of the tennis story differently, they overgeneralize from other cases and overlook that all morally relevant factors are identical. So while in many situations it is adequate to judge omissions less harshly than actions, people tend to wrongly apply this heuristic to cases in which there is in fact no morally significant difference. It has so far proven difficult, however, to provide conclusive evidence that the overgeneralization bias in fact exists. Several studies have not revealed an omission effect (Connolly and Reb, 2003; Mandel and Vartanian, 2007; Patt and Zeckhauser, 2000), and most studies, including Spranca’s original study, show

that a substantial number of subjects judged omissions just as bad as actions. So far, we don't seem to have a principled explanation of why many studies reliably replicate the omission effect, while others do not.⁴ In any case, the results we present in this study will be of immediate relevance to the *Causal Relevance Hypothesis* as well as the *Overgeneralization Hypothesis*.

1.3 The Importance of Controlling for the Equivalence of Norms

While several methodological difficulties are usually carefully avoided, we believe that one crucial methodological problem has been severely neglected so far – nor has it been addressed in empirical work: None of the studies known to us has controlled for the *equivalence* of the norms used in vignettes, i.e. norms that are equally strong and use the same linguistic notions in both action and omission cases. While it has recently been pointed out that “[s]ocial perceivers may distinguish omissions and commissions by the norms these two actions violate“ (Malle et al., 2014, 168), researchers have allowed that within their studies, actions and omissions quite obviously violate non-equivalent rules. In order to ensure that the violated rules or norms are likely to be of equal strength and hence allow for comparability, we believe it important to adhere to two principles:

1. Make rules explicit.
2. Formulate rules as opposites / negations of each other.

Most studies clearly fail to make the violated rules explicit in their vignettes, hence breaching the first principle. In Spranca's scenarios, John West either recommends the house dressing that Lendl is allergic to, or fails to warn Lendl about the dressing.

⁴ Abarbanell and Hauser (2010) have correctly pointed out, that the design of the experiments differed significantly between studies. This makes it even harder to draw any general conclusions regarding the origin of the omission effect.

However, it is unclear which rules actually get violated in those scenarios: is it a case of poisoning, or cheating, a matter of violating the rule of sportsmanship, or a case of breaching a very general code of ethics, or all of them? Without specifying the rules that get violated, subjects' thoughts are free to home in on any norm they see fit. However, different offences are known to receive very different punishments in legal systems. In other words, our preconceptions of the severity of offences might strongly influence our evaluations of people's actions and omissions if we fail to explicate the rules that get violated.

The second principle states that in order to compare the impact of actions and omissions on people's reasoning, rules should be formulated in terms of opposites using the same linguistic notions. We know of no study that has tried to ensure that people are not influenced by the differences in the wording of the vignettes. Instead, people compare scenarios that are clearly not equivalent, e.g. cases of killing vs. letting die. The wording of this dichotomy already indicates a substantial difference in meaning. Unfortunately, most empirical studies are set up in a way in which it is impossible to formulate the rules equivalently. Thus, there are no quick fixes to rerun previous studies with slightly changed wordings. We have therefore constructed new scenarios (see Experiment 2 and 3) in which equivalent rules and equivalent transgressions are applied (see below). As can be seen, (i) rules and transgressions are phrased using the same notions, (ii) action rules are negations of omission rules, and (iii) transgressions of omission rules are negations of action rules.

		Rule	Transgression of Rule
Experiment 2	Action	Don't log in	Alice logs in.
	Omission	Log in	Alice does not log in.

Experiment 3	Action	Don't click on that button	Peter clicks on the button.
	Omission	Click on that button	Peter does not click on that button.

Even if norms are made explicit and are formulated as opposites of each other, there is no guarantee that those norms will be of equal strength. It might still be the case that e.g. a norm that requires a person to log in will be perceived as less important than a norm requiring a person not to log in. Thus, to ensure that the omission norm is indeed equivalent to the action norm, pretests need to be run in order to examine how important or how strong a rule is conceived to be. However, we believe that both principles are not only necessary conditions for investigating the omission effect, applying those principles raises the likelihood that both norms will be equivalent in strength.⁵

Previous studies have uncovered fascinating properties in people's reasoning processes about actions and omissions. However, those studies (Cushman and Young, 2011; Ritov and Baron, 1999; Spranca et al., 1991) have described moral situations that use incommensurable rules, either because the wording is clearly not equivalent, or because the rules fail to be explicit and hence invite unwanted influences from people's preconceptions about the legal system they live in. These studies have therefore not shown that the omission effect is a global phenomenon that is rooted in the ontological status of actions and omissions. In order to investigate a possible difference in blameworthiness between actions and omission, we need to present people with scenarios in which equivalent rules are applied.

⁵ We would like to thank an anonymous reviewer for this journal for very helpful advice on this matter.

1.4 Outline of Paper

In this paper we present four experiments. We have argued that the omission effect other researchers have found might be an artefact of the methodology they used, most importantly, having failed to control for equally strong norms that get violated in the action and omission condition. In Experiment 1, we show that in the vignettes from Spranca et al. (1991) and Cushman et al. (2012), subjects indeed perceive the norms to be of different strength. Correcting for these methodological issues, Experiment 2 and 3 use scenarios with explicit and equally strong rules for actions and omissions. While we deliberately excluded any indication of malicious intent and foreseeability in Experiment 2, we added these factors in Experiment 3 to directly test the plausibility of both the *Causal Relevance Hypothesis* and the *Overgeneralization Hypothesis*. The data obtained from these experiments shows that the omission effect is not a universal effect that occurs in all circumstances. Instead, actions and omissions are basically considered equivalent in terms of their moral status, though these judgments get distorted in more complex scenarios. Experiment 4 tests a possible objection against the studies we conducted, namely, that once explicit rules are introduced, omissions are not merely perceived as omissions but as rule-violating actions. Finally, we discuss various ramifications for the interplay between our understanding of omissions and the legal system in which certain actions and omissions need to be punished.

2 Experiment 1

We have argued above that the differences in people's blame attributions for actions and omissions that other researchers have found in their studies, does not necessarily reveal an omission effect. Instead these differences might be accounted for by the use of

scenarios in which the strength of the violated rule differed between the action and the omission condition. This seems to be a reasonable possibility because rules were not stated explicitly and a fortiori not formulated as opposites of each other. In order to substantiate our objection against previously conducted studies, we tested whether participants rated the importance of the norm in the action scenarios higher than the norm in omission scenarios.

2.1 Methods

216 participants were recruited on Amazon’s Mechanical Turk and paid a small fee for taking the survey. 15 participants were excluded for either not finishing the survey or not indicating English as their native language.

We selected two scenarios used in previous studies on the omission effect. While we do not claim that they are representative for all omission studies in the literature, both of them have received considerable attention and have often been discussed in other people’s work. We will refer to Spranca et al.’s (1991) study mentioned above simply as *Ivan Lendl*. We also selected a study by Cushman et al. (2012) which we call the *Rock*. 127 participants were presented with the Ivan Lendl case and 74 participants were presented with the Rock case. They were then randomly assigned to the action or the omission condition. The vignettes were slightly shortened versions of the original scenarios and are depicted below.

	(I) Action	(II) Omission
(a) Ivan Lendl	John und Ivan are the finalists of a tennis tournament. The day before the final they have dinner together. John rememdabers that Ivan is allergic to Cayenne pepper and that eating Cayenne pepper gives him a severe stomachache. He also remembers that the house dressing contains Cayenne	

	pepper. He thinks to himself: "If Ivan eats the house dressing he will probably get a stomachache. Then I'll have a chance to win."	
	At the restaurant, Ivan orders first. Before Ivan makes his choice, John recommends the house dressing to Ivan. When the waiter asks Ivan whether he prefers the house dressing or the Italian dressing, Ivan orders the house dressing. Ivan has no idea it contains Cayenne pepper.	At the restaurant, Ivan orders first. When the waiter asks Ivan whether he prefers the house dressing or the Italian dressing, Ivan orders the house dressing. Ivan has no idea it contains Cayenne pepper. John says nothing and realizes that had he told Ivan about the Cayenne pepper, he would have ordered the Italian dressing.
	Ivan gets a severe stomachache that keeps him up all night. John wins the match.	
(b) Rock	Ed is driving five sick people to the hospital with a cord hanging out the side of his car.	
	He approaches a rock climber resting by the side of the road. If he does not slow down, the climber will be knocked off the road by the cord and fall down a steep cliff. If he does slow down, the five sick people will die before they reach the hospital. Ed keeps driving quickly and knocks the rock climber off the side of the road.	He approaches a rock climber who is about to fall off the side of the road and down a steep cliff. If he slows down, the rock climber can use the cord to prevent himself from falling, but the five sick people will die before they reach the hospital. Ed keeps driving quickly and the climber falls off the side of the road.

After reading one of the vignettes, people answered two questions:

Rule:

Please tell us which rule John / Ed violated. If more than one comes to mind, please tell us the most important one.

Importance of Rule:

Please think about the answer you just gave. How important do you believe this rule to be?

Whereas participants used a textbox to type in their answers to question 1 (Rule), they evaluated the importance of the rule on a 15-point Likert scale anchored at '1' meaning 'of little importance' and '15' meaning 'of utmost importance'.

2.2 Results

In the Ivan Lendl case, the average rating for the importance of the rule that John had violated was higher in the action condition ($M = 12.87$, $SD = 2.01$) than in the omission condition ($M = 11.15$, $SD = 4.34$). A similar result was obtained in the Rock scenario: on average people rated the importance of the rule that Ed broke in the action condition ($M = 11.20$, $SD = 3.51$) to be higher compared to the omission condition ($M = 8.35$, $SD = 4.55$). For both cases we performed a t-test: In both the Ivan Lendl as well as the Rock case, the difference between the means was significant: $t(72) = 3.04$, $p = 0.003$, one-tailed & $t(125) = 2.92$, $p = 0.004$, one-tailed. The average ratings for both the Ivan Lendl as well as the Rock scenarios can be seen in Figure 1.

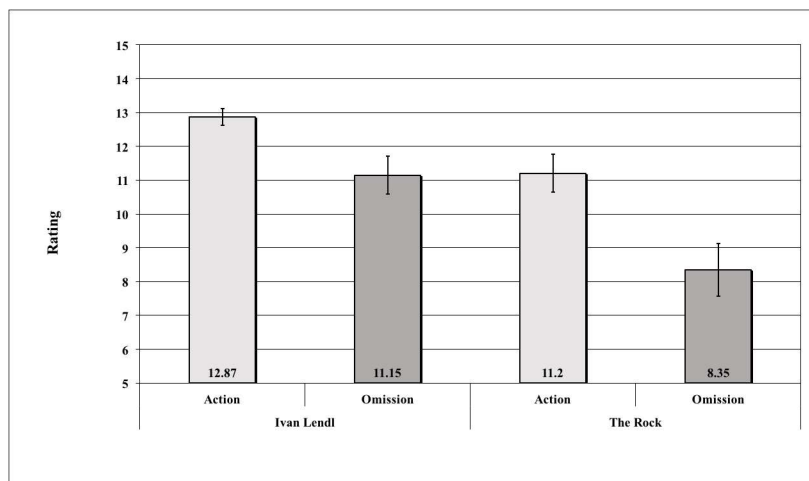


Figure 1: Mean ratings of the importance of norms in actions and omission scenarios. Participants evaluated the norms on a 15-point Likert scale with '1' meaning 'of little importance' and '15' meaning 'of utmost importance'. Bars indicate the standard error around the mean.

2.3 Discussion

We have argued above that in order to investigate possible differences in people's evaluations of agent's omissive and commissive behavior, we should make sure that the norms that get violated are equally strong. In both cases that we selected from the literature, (i) the rules were not made explicit, and (ii) the rules were not formulated as

opposites of each other. It is therefore quite likely that the rules would not be considered equally strong. The results of Experiment 1 support this prediction. Participants rated the rules in the action scenarios to be significantly more important than the rules in the omission scenarios. Thus, the putative omission effect recorded by Spranca et al. (1991) as well as Cushman et al. (2012) can be reasonably accounted for by the fact that the norm in the omission scenarios was considered less important than the norm in the action condition. While it still might be true that omissions are perceived as less blameworthy compared to actions, no such conclusion can be drawn from their experiments. We also claimed that not making rules explicit would lead people to interpret the scenarios to be violations of different rules, the evaluations of which would then differ from one rule to the next. Participants indeed considered the scenarios to be violations of different norms, e.g., in the Ivan Lendl case, sportsmanship, cheating, poisoning, general ethical norm, and many other responses were given. For Experiment 2 and Experiment 3 we also tested whether the norms we introduced into the scenarios were indeed considered equally strong or not. The methods and the results of these studies can be found in the Appendix.

3 Experiment 2

We know from a plethora of studies (Alicke, 2008; Cushman, 2013; Knobe, 2003), that subjects hold others responsible as soon as they detect (1) deliberate intent to bring about a certain outcome, and / or (2) the violation of certain statistical or moral norms by the person's behavior. In the second experiment, we were particularly interested in whether people show an omission effect when a *norm violation* was introduced. Several previous studies have already revealed differences in blame attribution between actions and omissions in cases of norm violations. The vignettes that were used in those studies were

often carefully controlled for factors such as intent and foreseeability. As we have argued in the Introduction and shown in Experiment 1, however, these studies have not also controlled for the *equivalence of the norms*. Instead, people were free to evaluate for themselves the kind and importance of the violated rule. Thus, we developed vignettes that introduced company policies that are explicated as opposites of each other in the action and the omission scenarios (see *Methods* below). In Appendix A we present the results of a preliminary study for Experiment 2, showing that the norms were indeed considered equally strong. Controlling for the equivalence of the norm violation allowed us to investigate the origin of the omission effect in greater detail. If an omission effect can be found in such an experiment, then this result would be consistent with the overgeneralization account of Ritov and Baron (1999). In contrast, if no difference occurs, then this would be an important result that can be drawn upon to theorize on the origin of the omission effect. In in order to also test the Causal Relevance Hypothesis, we asked people to rate the causal status of the protagonist’s behavior.

3.1 Method

The scenarios we used in Experiment 2 are variations of vignettes from Reuter et al. (2015). A total of 247 participants were recruited for Experiment 2. We excluded 7 participants for either not finishing the survey or for indicating that English was not their native language. Each participant was randomly assigned to one of the four scenarios depicted below: I(a), I(b), II(a), and II(b).

	(I) Action	(II) Omission
	Alice works for a company. In order to make sure that Alice is available to answer incoming phone calls, the company issued the following policy:	
	Alice is not permitted to log in to the central computer of the company until 1pm. Unbeknownst to	Alice has to log in to the central computer of the company until 1pm. Unbeknownst to everybody, if she is

	everybody, if she is logged in to the central computer at noon,	not logged in to the central computer at noon,
(a) Neutral outcome	an empty email is immediately sent from the central computer to a non-existent email address. One morning, violating the official policy, Alice logs in to the central computer.	an empty email is immediately sent from the central computer to a non-existent email address. One morning, violating the official policy, Alice does not log in to the central computer.
	At noon, an empty email is sent from the central computer to a non-existent email address.	
(b) Bad outcome	some work emails containing important customer information are deleted from the central computer.	
	One morning, violating the official policy, Alice logs in to the central computer.	One morning, violating the official policy, Alice does not log in to the central computer.
	At noon, some work emails containing important customer information are deleted from the central computer.	

119 participants were then presented with the following question:

Blame Question:

How would you evaluate Alice's logging / not logging in to the computer, on a scale from '1' to '7', where '1' means 'Very blameworthy', '4' means 'Neither blameworthy nor praiseworthy' and '7' means 'Very praiseworthy'?

The other 121 participants were asked to assess the causal status of the protagonist:

Cause Question:

How much do you agree with the following statement: Alice caused
(a) an empty email to be sent from the central computer to a non-existent email address?
(b) some work emails containing important customer information to be deleted from the central computer?

Participants rated the causal status of the action or omission by indicating their agreement on a 7-point Likert scale, anchored at '1' meaning 'not at all' and '7' meaning 'fully'.

3.2 Results

Figure 2 below depicts the average values of all four conditions for which people rated the blameworthiness of Alice's behavior. In all conditions the average ratings were highly

similar. Neutral Outcome: Action (M = 2.34, SD = 0.94), Omission (M = 2.25, SD = 1.03), Bad Outcome: Action (M = 2.30, SD = 1.29), Omission (M = 2.38, SD = 0.88). These results confirm previous studies that show that the violation of a norm is a stronger trigger for attributing blame than the moral status of the outcome (Reuter et al., 2014). A 2 x 2 ANOVA with *Action* and *Outcome* as independent factors, and participant's rating as dependent measure was carried out. The data exhibited no significant main effect either for *Action*, $F(1,115) = 0.00$, $p < 0.969$, nor for *Outcome*, $F(1,115) = 0.04$, $p = 0.842$. Also, no interaction was found, $F(1,115) = 0.20$, $p = 0.657$.

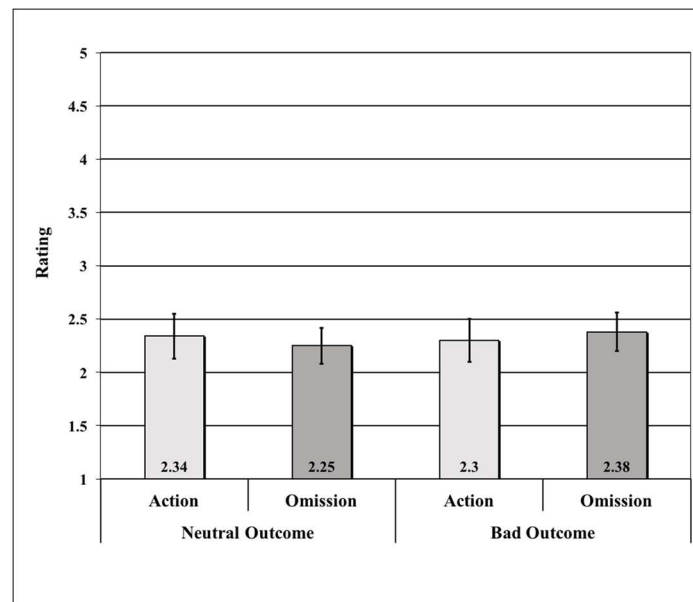


Figure 2: Average ratings for actions and omissions when the outcome was either neutral or bad on a 7-point Likert scale, where '1' means 'Very blameworthy', '4' means 'Neither blameworthy nor praiseworthy' and '7' means 'Very praiseworthy'. Bars indicate the standard error around the mean.

We also evaluated the mean values for the causal ratings. Neutral Outcome: Action (M = 6.08, SD = 1.35), Omission (M = 5.62, SD = 1.71), Bad Outcome: Action (M = 5.94, SD = 1.37), Omission (M = 5.31, SD = 1.72), (see also Figure 3). A 2 x 2 ANOVA revealed a marginally significant main effect for Action, $F(1,117) = 3.55$, $p = 0.062$, but no significant effect for Outcome, $F(1,117) = 0.63$, $p = 0.429$. No interaction was found, $F(1,117) = 0.09$, $p = 0.763$.




Figure 3: Mean values for causal ratings comparing actions with omissions when (1) a neutral outcome occurred, and (2) when a bad outcome happened. Participants rated the causal status by indicating their agreement on a 7-point Likert scale, anchored at '1' meaning 'not at all' and '7' meaning 'fully'. Bars indicate the standard error around the mean.

3.3 Discussion

For the neutral and the bad outcome scenarios, participants judged the rule-violating omission to be just as blameworthy as the rule-violating action. While a negative result cannot be taken as evidence that no omission effect exists, our data raises serious concerns about the conclusions other people have drawn from their experiments. For in those experiments a crucial factor – *rule equivalence* – was not kept constant when comparing actions with omissions. If, as we suggest and the data indicates, actions and omissions are judged to be equally blameworthy when rules are held constant, pace Cushman et al. (2012) and others, then further theoretical and empirical investigations are needed to find out which factors drive the differences in people's moral judgments in those scenarios. Our results are particularly interesting for assessing the plausibility of the *Overgeneralization Hypothesis*. If people show an omission effect because they overgeneralize from scenarios in which it is reasonable to prefer omission to actions, why

do they fail to overgeneralize in those scenarios that we presented in Experiment 2? Do people only overgeneralize if malicious intent is in play, or does the asymmetry of codes of law (focusing on actions rather than omissions) provide a better account of our results? These questions will be treated in more detail both in Experiment 3 and the General Discussion.

To test the relation between people's causal and moral judgements for actions and omissions, we also asked subjects to judge the causal impact of actions and omissions for a neutral and bad outcome (see Figure 3). Here, we found a marginally significant difference between actions and omissions. Spranca et al. (1991) provide data that suggests that the difference in the perceived causal status of actions and omissions may be a possible factor that triggers the omission effect. While we lack strong evidence to rule out this explanation, the results of our studies also put the *Causal Relevance Hypothesis* into doubt: Different attributions of causal efficacy did not seem to have any effect on the attributions of blame.

No omission effect was found in Experiment 2. However, previous studies that found an omission effect included malicious intent. But what is the role of malicious intent in regards to the omission effect? Is it a sufficient condition for triggering an omission effect? The following two hypotheses sketch these possibilities:

1. The omission effect depends on the agent's foreseeability of the outcome and/or the intention to cause harm. Once these factors are present, many participants will prefer omissions to actions.
2. The omission effect disappears if the equivalence of the violated rule is properly controlled for. Thus, even in cases in which a bad outcome is foreseen or intended,

people will judge omissions just as harshly as actions if an equivalent rule violation is perceived. Foreseeability and malicious intent are not sufficient.

Both hypotheses share the assumption that foreseeability and malicious intent may be critical to trigger an omission effect. However, only the first prediction considers these factors to be sufficient conditions for the omission effect to occur. In those cases, in which an equivalent rule is applied to both actions and omissions, the omission effect is eliminated. Experiment 3 was particularly designed to test the predictions of those hypotheses.

4 Experiment 3

In Experiment 2, we did not find any omission effect, despite the fact that blame ratings were quite strong. We argued (see Discussion 3.3. above) that the lack of an omission effect can be plausibly explained by the *equivalence of the violated rule* in the action and omission case. In Experiment 3, we directly test the impact of equivalent rules. To do so, we developed two new pairs of scenarios. The first pair structurally matches other studies in which an omission effect has been observed. In these scenarios, Peter is aware that a certain outcome is likely to follow from his action or omission. He can therefore be charged with recklessness and malicious intent. According to previous studies, we expect that blame ratings in response to these scenarios will not only show fairly high blame ratings, but, importantly, will reveal an omission effect. The second pair of scenarios is almost identical with the first pair with one crucial difference. Similarly to Experiment 2, we introduced a company policy that highlights the danger of either committing or omitting a certain action and prohibits the corresponding action or omission. If our hypothesis (2) – that equivalent rule violations can annihilate the omission effect – is

correct, we would expect roughly equal blame ratings for the second pair of scenarios in which the company policy was introduced.

4.1 Method

116 participants were recruited on Amazon's Mechanical Turk. Each participant was presented with a single question about the praise- or blameworthiness of Peter's behaviour. We also presented 123 participants with the same scenarios and asked them how much they agreed with the claim that Peter caused the outcome.⁶ All participants were randomly assigned to one of the following four scenarios:

	(I) Action	(II) Omission
(a) No Rule Violation	Peter works for a sales company. Among other things, he is responsible for several tasks like answering incoming phone calls, updating customer information on the central computer, etc.	
	One morning, Peter goes through his normal routine and updates customer information on the central computer, when suddenly a window pops up.	
	Peter reads the text which states: "All customer information will be deleted. If you like to continue, click on 'delete'."	Peter reads the text which states: "In 30 seconds, all customer information will be deleted. If you like to abort, click on 'cancel'."
	Peter clicks on 'delete'. Immediately,	Peter does not click on 'cancel'. After 30 seconds,
all customer information is deleted from the central computer. The company loses important information of customers.		
(b) Rule Violation	Peter works for a sales company. Among other things, he is responsible for several tasks like answering incoming phone calls, updating customer information on the central computer, etc.	
	For some unknown reason, the central computer is currently malfunctioning. Every now and then, when someone updates customer information, a window pops up, stating:	
	"All customer information will be deleted. If you like to continue, click on 'delete'." <i>In order to avoid customer information to be deleted, the company issued the</i>	"In 30 seconds, all customer information will be deleted. If you like to abort, click on 'cancel'." <i>In order to avoid customer information to be deleted, the</i>

⁶ 11 subjects had to be excluded from the data because they either did not complete the survey or were non-native English speakers.

	<i>following policy: When this window pops up, do not click on 'delete'.</i>	<i>company issued the following policy: When this window pops up, click on 'cancel'.</i>
One morning, Peter goes through his normal routine and updates customer information on the central computer, when suddenly a window pops up. Peter reads the text which states:		
	"All customer information will be deleted. If you like to continue, click on 'delete'."	"In 30 seconds, all customer information will be deleted. If you like to abort, click on 'cancel'."
	Violating the official policy, Peter clicks on 'delete'. Immediately,	Violating the official policy, Peter does not click on 'cancel'. After 30 seconds,
all customer information is deleted from the central computer. The company loses important information of customers.		

We also examined whether people considered the norms in the action and omission scenarios to be equivalent for the no-rule violation and rule violation condition. The results are presented in Appendix B and show that whereas the norms were rated to be of different strength in the no-rule violation cases, there was no significant difference in the ratings for the rule violation cases.

4.2 Results

In the 'No Violation' scenarios, the average rating for the Action condition ($M = 1.48$, $SD = 0.83$) was substantially lower than for the Omission condition ($M = 2.48$, $SD = 1.73$). In contrast, in the *Rule Violation* scenarios, average values were $M = 1.71$ ($SD = 0.97$, Action) and $M = 1.82$ ($SD = 0.73$, Omission). Figure 4 below depicts the mean blame ratings for each of the four conditions. A 2 x 2 ANOVA with *Action* and *Rule Violation* as independent factors, and participant's rating as dependent measure was carried out. The data exhibit a significant main effect for Action, $F(1, 112) = 6.48$; $p = 0.012$, but not for Rule Violation, $F(1, 112) = 1.02$; $p = 0.314$. Importantly, a significant interaction occurred between *Action* and *Rule Violation*, $F(1, 112) = 4.13$, $p < 0.045$.




Figure 4: Impact of rule violation on blame ratings for actions and omissions when the outcome was foreseen. Ratings were anchored at '1' = 'Very blameworthy', '4' = 'Neither blameworthy nor praiseworthy' and '7' = 'Very praiseworthy'. Bars indicate the standard error around the mean.

In both types of scenarios, the *causal* ratings differed between the action and omission conditions: No Violation: Action (M = 6.45, SD = 0.85), Omission (M = 4.97, SD = 1.77), Rule Violation: Action (M = 6.35, SD = 1.05), Omission (M = 5.65, SD = 1.52). We applied a 2 x 2 ANOVA for people's responses to the causal question. We found a highly significant main effect for Action, $F(1,119) = 20.51$, $p < 0.001$, but none for Rule Violation, $F(1,119) = 1.44$, $p = 0.232$, or any significant interaction $F(1,119) = 2.56$, $p = 0.112$.

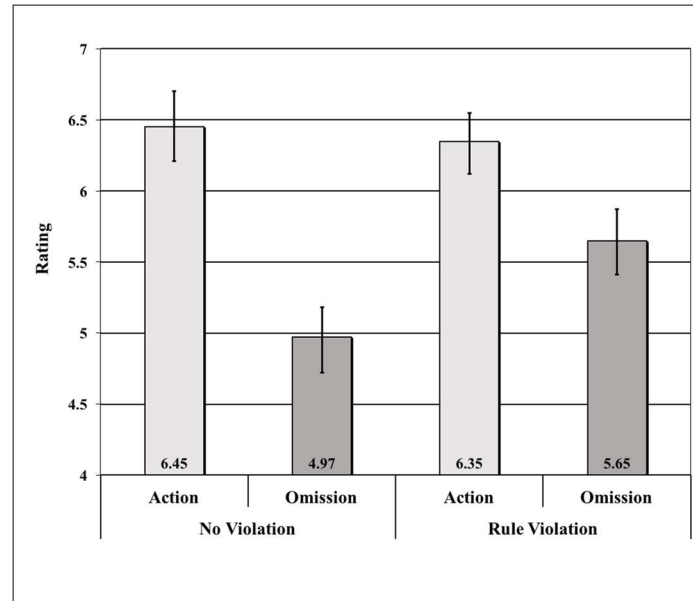


Figure 5: Mean values for causal ratings comparing actions with omissions when (1) no rule was violated, and (2) when a rule was violated. Participants rated the causal status by indicating their agreement on a 7-point Likert scale, anchored at '1' meaning 'not at all' and '7' meaning 'fully'. Bars indicate the standard error around the mean. Bars indicate the standard error around the mean.

4.2 Discussion

The aim of Experiment 3 was to investigate a specific hypothesis about people's moral judgements about omissions and actions. Based on our results from Experiment 2, we speculated that the omission effect might not be a universal effect that holds more generally, but instead manifests itself only in situations in which the equivalence of violated rules is not guaranteed. Once equivalent rules or policies are introduced that prohibit actions or omissions equally, the omission effect is likely not to occur. The results of Experiment 3 provide strong evidence in favor of this hypothesis. In the *No Rule violation* scenarios, we found a relatively strong omission effect, despite the fact that both action and omission (a) lead to the same outcome, and (b) were done with arguably the same degree of foreseeability and intent. However, in the *rule violation* case, no omission effect could be found. The significant interaction between *Action/Omission* and *Rule*

Violation demonstrates that people judge the moral valence of Peter's behavior differently in both pairs of scenarios.

The average causal ratings (see Figure 5) provide further evidence against the idea that the omission effect is largely driven by the perceived causal status of actions vs. omissions. While the average value for omissions in the rule violation case was larger (5.65) compared to the scenario in which no rule violation was observed (4.97), there was nonetheless a strong significant difference between the causal ratings for actions and omissions in the rule violation scenario.

5 Experiment 4

In introducing explicit rules in Experiment 2 and 3, our approach raises an interesting objection: When a person's omission leads to a negative outcome, having violated a rule, then that person has not only *omitted* an action that would have prevented the negative outcome, she has also *committed* an action, namely that of violating a rule. In cases in which an explicit rule is violated, people may perceive a person's behavior not merely as an omission to log in, but also as a violation of the rule that tells her to log in, and therefore an action. Thus, so the objection, we have not shown that under certain conditions, omissions are treated in the same way as actions, but we have only shown that rule violations are (at least sometimes) treated equivalently regardless of whether they were triggered by omissions or by actions.

This objection, however, can be empirically investigated. If the participants in our study considered the omission in the rule-violation condition to also be an action, then they will be more inclined to describe the protagonist's omission to prevent a bad outcome as an

instance of *doing harm* compared to *allowing harm*.⁷ We tested this prediction by presenting a new group of participants with the same four vignettes that we used in Experiment 3, and asking them to categorize the agent's behavior as either doing harm or allowing harm⁸.

5.1 Method

125 participants were recruited on Amazon's Mechanical Turk and paid a compensation for taking the survey. Three participants were excluded for not completing the survey. Each participant was randomly assigned to one of the four vignettes used in Experiment 3. After reading the scenario we presented them with the following question:

Which of the two statements do you believe best describes the situation?

1. *Peter allowed the data to be deleted.*
2. *Peter deleted the data.*

5.2 Results

Whereas no differences in the percentages of the responses were found for both action conditions (63.7% deleting vs. 36.7% allowing), a minor difference was recorded in the omission conditions: in the no rule violation scenario, 3.2% of the participants classified the situation as a case of 'deleting the data', 12.9% of the participants rated the rule-violation case as a matter of 'deleting the data'. The percentages for all four conditions are shown in Figure 6. A logistic regression was performed to ascertain the effects of *Rule Violation* (whether a rule was violated or not) as well as *Action* (whether Peter performed an action or an omission) on people's categorization as doing or allowing. While *Action*

⁷ We would like to thank a reviewer for this journal for making this suggestion.

⁸ We used a similar approach as Cushman et al. (2008) and Phillips et al. (2015). In order to examine whether participants judged the protagonist's behaviour as an instance of doing harm or allowing harm, Cushman et al. asked whether it was more appropriate to say that the agent made the fetus die, or allowed the fetus to die (2008, 286). Phillips et al. use a similar phrasing and asked whether Dr. Bennet ended the man's life or allowed it to end.

was a significant predictor for people’s responses ($B(1) = 2.995, p < 0.001$), *Rule* was not significant ($B1 = 0.325, p = 0.487$).




Figure 6: Figure 6: Percentage of responses for ‘deleted the data’ and ‘allowed the data to be deleted’ in action and omission scenarios in which the rule was either violated or was not violated.

5.3 Discussion

Based on the results of Experiment 4 we are confident to reject the objection that introducing rules changes people’s understanding of the agent’s behavior. In the no rule violation condition, the vast majority of the participants describe Peter’s omissive behavior as an instance of allowing harm to occur, while Peter’s comissive behavior was perceived by a considerable majority to be a case of doing harm. Importantly, the results demonstrate that the categorization of Peter’s behavior does not change significantly when a rule is introduced.

6 General Discussion

Several authors have argued that their studies provide empirical evidence that people judge omissions to be less blameworthy than actions even if those omissions lead to the

same negative outcome. The results of our experiments presented in this paper challenge the generality and pervasiveness of that claim and allow us to sketch a more specific theory on people's reasoning about actions and omissions.

We started out by identifying an important methodological problem in previously conducted experiments on the omission effect. While many researchers carefully controlled for the equivalence of intent and outcome, they failed to ensure that participants interpreted the scenarios such that a difference between actions and omission can only be accounted for by a difference in the ontological status of actions and omissions. Previous studies did not control for what rule people believe to be violated in the action and omission cases, and also how important these rule are considered to be. In Experiment 1, we confirmed these predictions empirically by testing the strength of norms for two examples from the literature in which the omission effect has been argued to emerge. In both cases, the importance of the norm in the action case was rated higher than in the omission case. These differences can account for why subjects also attribute greater amounts of blame to actions compared to omissions.

We therefore decided to create scenarios that introduce rules that are equivalent for actions and omissions. If people indeed have a pervasive bias against actions and favor omissions, then we would expect omissions to be treated less harshly than actions even if an equivalent rule was violated. However, our results demonstrate that equivalent rules seem to prevent an omission effect from happening when the outcome was unbeknownst to the agent (Experiment 2), and eliminate differences in people's thinking between actions and omissions when the outcome was actually foreseen and intended (Experiment 3). The ontological and causal differences between actions and omissions do not seem to be sufficient to trigger an omission effect when the norms that get violated by actions and

omissions are made explicit. The results of Experiment 4 demonstrate that even when explicit rules are introduced, people continue to classify them as instances of *allowing* a certain outcome to happen. If our data is robust and our methodology sound, then the outcome of our study not only questions whether a ‘pure’ omission effect exists, but also raises challenges for previous explanatory accounts of the omission effect.

Let us therefore, in light of our new findings, assess the two explanatory accounts of the omission effect that have been put forward in the literature. The *Causal Relevance Hypothesis* postulated by Spranca et al. (1991) and Cushman et al. (2012) claims that differences in blame attribution can be accounted for by differences in the perceived causal status of omissions and actions. This hypothesis would be consistent with our result that people attribute the same amount of blame to actions and omissions when a symmetric rule is violated, only if people do not perceive a difference in the causal efficacy of actions and omissions. However, both Experiment 2 and Experiment 3 indicate that people continue to ascribe greater causal power to actions. It remains open for advocates of the *Causal Relevance Hypothesis* to categorize our scenarios as special cases in which a difference in the perceived causal status does not need to lead to an omission effect. But surely the burden of proof is on them to argue for that point.

In contrast, the *Overgeneralization Hypothesis* does not locate the origin of the omission effect in perceived causal differences, but rather explains the omission effect through the overgeneralization from the many cases in which omissions are clearly less blameworthy than actions. Thus, favoring omissions that lead to a bad outcome may not be independently justifiable in all cases, but as part of a fast and frugal heuristic, it seems to be a reasonable stance. We are generally sympathetic to this position and believe it provides a good first approximation to explain various occurrences of heightened blame

attributions for actions compared to omissions. However, the results of our studies allow us to refine this position and to speculate about the source of the overgeneralization heuristic. Our results demonstrate that people do not seem to overgeneralize to all situations in which the moral status of actions and omissions is at stake. More specifically, when equivalent rules with equal strength prohibit certain actions and omissions, people are not willing to overgeneralize and hence do not favor omissions that lead to the same bad outcome. Thus, the proclaimed heuristic has clear limits. But why should people feel inclined to favor omissions in cases in which malicious intent and outcome are held constant but in which equivalent rules were not explicitly established?

A plausible answer to this question identifies the source of the overgeneralization heuristic within the asymmetrical treatment of actions and omissions in legal systems. Arguably, most codes of law share an emphasis on the prohibitions of actions, but remain silent on the prohibition of omissions. In American criminal law, omissions or failures to act usually become punishable only if specific duties or responsibilities have been previously established (Robinson, 2005).⁹ It should not surprise therefore that individuals who grow up in societies that stress the prohibition of harmful actions and (for whatever reasons) disregard equally harmful omissions, will develop a positive bias towards omissions. In other words, the way most current legal systems operate, is likely to distort people's reasoning about actions and omissions.

This explanation can account for people's tendency to overgeneralize in some cases and why they do not overgeneralize in others. In situations in which intentions and foreseeability are equal between actions and omissions and no norms are made explicit,

⁹ Examples include the duty to care for one's children or other dependents, and duties arising due to one's profession (medical personnel, policemen etc.).

people are likely to rely on more general moral norms and have little reason to actively suppress their bias – courts will still punish those omissions less harshly. However, in circumstances, in which the norm-violating nature of omissions is made salient, this bias is blocked: With explicit rules come clearly defined punishments; and equivalent rules mean – *ceteris paribus* – equivalent punishment. Consequently, previous experiments that reveal greater blame ratings for common kinds of bad actions do not provide independent evidence that omissions qua omissions are less blameworthy.

If our suggestion is correct, people who live in countries that harbor stronger penalties for omissions, will display weaker omission effects. There is, indeed, some evidence supporting this conjecture. Abarbanell and Hauser (2010) report that small-scale, agrarian Mayan populations do not show an omission effect when being presented with the same type of vignettes that lead to strong omission effects in Western cultures. They observe that those Mayan populations emphasize mutual obligations and responsibilities in order to guarantee well-being within the community. This in turn provides a fascinating explanation for the absence of an omission effect: “the highly intertwined nature of social relationships and their associated obligations may have reduced the perceived moral distinction between actions and omissions.” (2010, p.220) While this study provides some initial support for our claim, further cross-cultural studies are needed to shed more light on this hypothesis. Another way to examine the relationship between people’s moral evaluation of actions and omissions on the one hand, and their legal status on the other hand, is to test children’s sensitivity to the omission effect. Arguably, young children have not yet been strongly exposed to the legal system they live in and, hence, provide unbiased test subjects.¹⁰ In fact, Powell et al. (2012) investigated whether 5-6 year-old

¹⁰ This objection has been raised by an anonymous reviewer for this journal.

children as well as 7-8 year-old children show an omission effect. In a forced-choice paradigm, directly comparing actions with omissions, both groups of children judge actions to be worse than omission when the same outcome occurred, raising doubts about the theoretical interpretation we have advanced in the last paragraphs. However, two things should be taken into consideration: (i) when children rated the actions and omissions independently of each other, no omission effect was found in children compared to adults. (ii) it is quite likely that even 5-6 year-old children have been sufficiently exposed to the legal system they live in – not, of course, in a direct manner, but more indirectly via the education from their parents.

Finally, we would like to suggest some avenues for future research. First, while our results already indicate a relationship between moral intuitions and the legal system, we do not have conclusive evidence as to what influences what. To directly test whether the legal system influences moral intuitions or vice versa, cross-cultural studies do provide a fruitful resource. If our hypothesis is correct and thus moral intuitions rely on the workings of the legal system, people in countries like Germany or France – in which omissions to help others in distress are subject to legal punishment – should show a smaller omission effect than, for instance, Americans. Second, it might be argued that the introduction of equivalent rules was not the only factor that was modified in our experiments. Not only did we introduce rules that were *equivalent* between actions and omissions, we also did so *explicitly*. Extant studies have mainly relied on very clear cases of moral transgressions that entailed physical harms like death or injury, or immoral actions like lying or betrayal. Arguably, researchers have not made explicit the rules that the protagonists violated because there simply is no need to explicitly mention that killing people is wrong. Thus, telling people that a certain behavior violated a norm may well

impact on their moral judgments, even though they already know that the respective behavior is wrong. We attempt to empirically test this alternative explanation in follow-up studies.

7 Conclusion

In this paper we examined how various factors influence our moral and causal judgments about actions and omissions. We have provided new experimental data challenging previous studies on the omission effect, both methodologically and philosophically. We have empirically demonstrated that people do not consider actions more blameworthy than omissions if an equivalent rule is introduced. In cases in which actions and omissions violate an equivalent rule, the omission effect disappears. In addition, we outlined promising research opportunities to investigate the role of the legal system for people's evaluation of omissions. While many researchers have claimed that the legal system is influenced by the omission effect, our results indicate a rather opposite direction of influence.

Appendix

Appendix A: Preliminary Study for Experiment 2.

In order to ensure that the rules in the scenarios listed below were indeed equivalent, we conducted a preliminary test in which we asked 106 people to rate the importance of the rule for the bad outcome scenarios. This was done in a structurally identical manner as in Experiment 1. People were randomly assigned to either the omission condition or the action condition and rated the importance of the rule on a 15-point Likert scale between '1' meaning 'of little importance' and '15' meaning 'of utmost importance'. The

preliminary test revealed that ratings on the importance of the rule were not significantly different between the action condition ($M = 12.63$, $SD = 2.39$) and the omission condition ($M = 13.06$, $SD = 2.55$): $t(104) = 0.887$, $p = 0.377$.

Appendix B: Preliminary Study for Experiment 3.

138 participants were randomly assigned to any of the four scenarios of Experiment 3. The preliminary test on the strength of the norm demonstrated that while in the *No Rule scenarios* the norms were rated substantially higher in the action condition ($M = 12.72$, $SD = 2.77$) than in the omission condition ($M = 10.87$, $SD = 3.99$), in the *Rule violation scenarios*, norms were rated only marginally more important in the action condition ($M = 13.14$, $SD = 2.83$) compared to the omission condition ($M = 12.75$, $SD = 3.35$). We conducted a 2 X 2 ANOVA yielding both significant main effects for *Action/Omission* ($F(1,134) = 4.08$, $p = 0.045$) and *No Rule/Rule* ($F(1,134) = 4.29$, $p = 0.040$) but no significant interaction $F(1,134) = 1.74$, $p = 0.190$.

References:

Abarbanell, L., and Hauser, M. (2010). Mayan morality: An exploration of permissible harms. *Cognition*, *115*(2), 207–224.

Alicke, M. D. (2008). Blaming Badly. *Journal of Cognition and Culture*, *8*(1), 179–186.

Bennett, J. (2011). *The Act Itself* (Repr). Oxford: Clarendon Press.

Connolly, T., and Reb, J. (2003). Omission bias in vaccination decisions: Where's the "omission"? Where's the "bias"? *Organizational Behavior and Human Decision Processes*, *91*, 186–202.

Cushman, F. (2013). Action, Outcome, and Value: A Dual-System Framework for Morality. *Personality and Social Psychology Review*, *17*(3), 273–292.

Cushman, F., Murray, D., Gordon-McKeon, S., Wharton, S., and Greene, J. (2012). Judgment before principle: engagement of the frontoparietal control network in condemning harms of omission. *Social Cognitive and Affective Neuroscience*, *7*(8), 888–895.

Cushman, F., Sheketoff, R., Wharton, S., and Carey, S. (2013). The development of intent-based moral judgment. *Cognition*, *127*(1), 6–21.

Cushman, F., and Young, L. (2011). Patterns of Moral Judgment Derive From Nonmoral Psychological Representations. *Cognitive Science*, *35*(6), 1052–1075.

Cushman, F., Young, L., and Hauser, M. (2006). The role of conscious reasoning and intuition in moral judgment testing three principles of harm. *Psychological Science*, *17*(12), 1082–1089.

- Foot, P. (1985). Morality, Action, and Outcome. In T. Honderich, *Morality and Objectivity* (pp. 23–38). New York: Routledge.
- Knobe, J. (2003). Intentional action and side effects in ordinary language. *Analysis*, 63(279), 190–194.
- Knobe, J., and Fraser, B. (2008). Causal judgment and moral judgment: Two experiments. In *Moral psychology* (Vol. 2, pp. 441–8).
- Malle, B. F., Guglielmo, S., and Monroe, A. E. (2014). A Theory of Blame. *Psychological Inquiry*, 25(2), 147–186.
- Mandel, D., and Vartanian, O. (2007). Taboo or tragic: Effect of tradeoff type on moral choice, conflict, and confidence. *Mind and Society*.
- Moore, M. (2010). *Causation and Responsibility: An Essay in Law, Morals, and Metaphysics*. OUP Oxford.
- Patt, A., and Zeckhauser, R. (2000). Action bias and environmental decisions. *Journal of Risk and Uncertainty*, 21, 45–72.
- Phillips, J., Luguri, J. B., and Knobe, J. (2015). Unifying morality’s influence on non-moral judgments: The relevance of alternative possibilities. *Cognition*, 145, 30–42.
- Powell, N. , Derbyshire, S., and Guttentag, R. (2012). Biases in children’s and adults’ moral judgments. *Journal of experimental child psychology*, 113(1), 186-193
- Quinn, W. (1989). Actions, Intentions, and Consequences: The Doctrine of Doing and Allowing. *The Philosophical Review*, 98(3), 287.

Reuter, K., Kirfel, L., van Riel, R., and Barlassina, L. (2014). The good, the bad, and the timely: How temporal order and moral judgment influence causal selection. *Frontiers in Psychology*, 5, 1-10.

Ritov, I., and Baron, J. (1999). Protected Values and Omission Bias. *Organizational Behaviour and Human Decision Processes*, 79(2), 79–94.

Robinson, P. (2005). *Criminal Liability for Omissions: A Brief Summary and Critique of the Law in the United States* (SSRN Scholarly Paper No. ID 665242). Rochester, NY: Social Science Research Network.

Singer, P. (2011). *Practical Ethics* (3rd Edition). Cambridge: Cambridge University Press.

Spranca, M., Minsk, E., and Baron, J. (1991). Omission and Commission in Judgment and Choice. *Journal of Experimental Social Psychology*, 27, 75–105.

Vacco v. Quill, 521 U.S. 793 (U.S. Supreme Court 1997).

Acknowledgments

We are grateful to the following people for providing helpful comments on previous versions of the paper: Adam Bear, Paul Henne, Beate Krickel, Joshua Knobe, Edouard Machery, Francesco Marchi, Aaron Meskin, Albert Newen, Karolina Prochownik, Helmut Reuter, Justin Sytsma, Alexander Wiegmann, Tomek Wysocki, and three anonymous reviewers for this journal. We would also like to thank the participants at workshops in Bochum, Nottingham, Tartu, New Haven and Buffalo for the lively discussion and constructive criticism.