

Chapman University

Chapman University Digital Commons

Philosophy Faculty Articles and Research

Philosophy

7-2021

Asymmetry and Symmetry of Acts and Omissions in Punishment, Norms, and Judged Causality

Toby Handfield
Monash University

John Thrasher
Chapman University, thrasheriv@chapman.edu

Andrew Corcoran
Monash University

Shaun Nichols
Cornell University

Follow this and additional works at: https://digitalcommons.chapman.edu/philosophy_articles



Part of the [Other Philosophy Commons](#)

Recommended Citation

Handfield, T., Thrasher, J., Corcoran, A., & Nichols, S. (2021). Asymmetry and symmetry of acts and omissions in punishment, norms, and judged causality. *Judgment and Decision Making*, 16(4), 796-822. <http://journal.sjdm.org/21/210210/jdm210210.html>

This Article is brought to you for free and open access by the Philosophy at Chapman University Digital Commons. It has been accepted for inclusion in Philosophy Faculty Articles and Research by an authorized administrator of Chapman University Digital Commons. For more information, please contact laughtin@chapman.edu.

Asymmetry and Symmetry of Acts and Omissions in Punishment, Norms, and Judged Causality

Comments

This article was originally published in *Judgment and Decision Making*, volume 16, issue 4, in 2021.

Creative Commons License



This work is licensed under a [Creative Commons Attribution 4.0 License](https://creativecommons.org/licenses/by/4.0/).

Copyright

The authors

Asymmetry and symmetry of acts and omissions in punishment, norms, and judged causality

Toby Handfield^a John Thrasher[†] Andrew Corcoran[‡] Shaun Nichols[§]

Abstract

Harmful acts are punished more often and more harshly than harmful omissions. This asymmetry has variously been ascribed to differences in how individuals perceive the causal responsibility of acts versus omissions and to social norms that tend to proscribe acts more frequently than omissions. This paper examines both of these hypotheses, in conjunction with a new hypothesis: that acts are punished more than omissions because it is usually more efficient to do so. In typical settings, harms occur as a result of relatively few harmful actions, but many individuals may have had the opportunity to prevent or rectify the harm. Penalising actors therefore requires relatively few punishment events compared to punishing omitters. We employ a novel group paradigm in which harm occurs only if both actors and omitters contribute to the harm. Subjects play a repeated economic game in fixed groups involving a social dilemma (total N = 580): on each round self-interest favours harmful actions (taking from another) and harmful omissions (failing to repair the victim's loss), but the group payoff is maximized if individuals refrain from these behaviors. In one treatment harm occurs as a result of one action and two omissions; in the other, it is the result of two actions and one omission. In the second treatment, the more efficient strategy to maximize group benefit is to punish omissions. We find that subjects continue to prefer to punish acts rather than omissions, with two important caveats. There is still a substantial level of punishment of omissions, and there is also evidence of some responsiveness to the opportunity to enforce a more efficient rule. Further analysis addresses whether the omission effect is associated with asymmetric norm-based attitudes: a substantial proportion of subjects regard it as equally fair to punish harmful acts and omissions, while another portion endorse an asymmetry; and punishment behavior correlates with these attitudes in both groups.

Keywords: action, omission, act-omission, asymmetry, punishment, fairness

^aPhilosophy Department, SOPHIS, Monash University. Email: toby.handfield@monash.edu. ORCID: 0000-0003-2995-2067.

[†]Department of Philosophy and Smith Institute for Political Economy and Philosophy, Chapman University. ORCID: 0000-0001-9094-5283.

[‡]Cognition and Philosophy Lab, Monash University ORCID: 0000-0002-0449-4883.

[§]Sage School of Philosophy, Cornell University. ORCID: 0000-0002-9447-5618.

1 Introduction

There are several ways in which actions are treated differently than omissions in general, and in our moral psychology in particular (Baron & Ritov, 2004; Feldman, Kutscher & Yay, 2020; Ritov & Baron, 1990). Doing harm tends to be judged more wrongful than allowing harm to happen (e.g., Cushman, Young & Hauser, 2006). Actions that produce bad outcomes are more likely to attract condemnation and punishment than omissions that produce equally bad outcomes (Robinson & Darley, 1995; Spranca, Minsk & Baron, 1991). And if given a choice between harming through commission or omission, people will typically choose omission (e.g., Asch et al., 1994; Meszaros et al., 1996). Of these three asymmetries – in punishment, judgment, and choice – our focus in the present study is the asymmetry of punishment, though there are likely to be close relationships between all these phenomena.

There are three prominent explanations for the asymmetric moral status of acts versus omissions, and all of them can be applied to the asymmetry of punishment in particular:

1. *Causation*. An action which leads to a particular outcome is judged to be in some sense more causally responsible for that outcome than an omission that leads to the same or an equivalent outcome – for instance actions can be physically connected to outcomes in a way that omissions cannot (Dowe, 2004); and people prefer to punish those whom they hold more causally responsible (see e.g., Baron & Ritov, 2009; Cushman & Young, 2011; Greene et al., 2009; Jamison, Yay & Feldman, 2020; Royzman & Baron, 2002; Spranca et al., 1991).
2. *Intention*. People judge harmful outcomes that come about via omission to be less intentional than equivalently harmful actions (Hayashi, 2015), and prefer to punish those who are judged to have manifestly harmful intentions (Cushman & Young, 2011; Jamison et al., 2020; but see Ritov & Baron, 1990; Royzman & Baron, 2002).
3. *Social norms*. Punishment is generally reserved for transgressions of injunctive norms – lest it be considered sheer aggression – and societies tend to have a greater number of injunctive norms proscribing actions than norms that proscribe omission (Willemsen & Reuter, 2016; Willemsen & Kirfel, 2019; Robinson & Darley, 1995). Relatedly, legal institutions tend to hold individuals criminally liable for harmful omissions only in narrowly defined circumstances, often requiring particular relationships to the persons harmed (Ashworth, 2015; Kadish, Schulhofer & Barkow, 2016).

This research was supported by the Australian Research Council (DP170102834). We thank Jon Baron, David Rojo Arjona, Chet Pager, Erik Kimbrough, Bart Wilson, and Vernon Smith for many helpful suggestions and comments.

Preregistrations, data, and software for replication purposes, are available at <https://osf.io/v6u3q>.

Copyright: © 2021. The authors license this article under the terms of the Creative Commons Attribution 3.0 License.

There is evidence that all three of these factors play some role in explaining the asymmetry. Matters are further complicated in that there appears to be significant interaction between cognition of norms and of causation.

First, there is evidence that norm violations are more likely to be perceived as causes than norm-conforming behaviors (Alicke, 1992; Knobe & Fraser, 2008; Kominsky, Phillips, Gerstenberg, Lagnado & Knobe, 2015). For example, if two agents engage in physically identical actions (taking pens from a cup) that are both causally necessary for the outcome (the subsequent absence of pens), the agent who was breaking a rule in order to take the pens is judged more causally responsible for the absence of pens. So if social norms tend to prohibit actions more frequently than they prohibit omissions, actions will more typically be norm transgressions, and this may explain why actions are typically judged more causally responsible than omissions.

Second, causation by omission raises paradoxes about which omissions to count as causes, given there will often be innumerable many candidates. (Some plants died because no one watered them. Since no one watered them, *everyone* omitted to water them. Is everyone's failure to water them therefore a cause of their death?) There is evidence that norms are involved in selecting particular omissions as causes – whereby again, norm-violating behaviors are more liable to be regarded as causes (Henne, Pinillos & De Brigard, 2017). Suppose Bob's failure to water the plants broke his promise to water them, whereas your failure did not: people appear to judge Bob's omission more causal because it violated a norm of promise keeping.

Both these considerations suggest that the explanation in terms of norms may be a common cause of both causal perception and an omission effect in punishment. But there are several limitations of extant work on norms and the omission bias¹, and we designed our experiments to try to meet some of these limitations. First, rather than using vignette studies, we explore whether the omission bias will emerge in punishment behavior in an economic experiment with real monetary incentives. We use monetary incentives because real punishment decisions involve substantial risks and costs (Boehm, 1999; Nikiforakis, 2008; Wiessner, 2005).

A second feature of our studies is that subjects simultaneously judge whether to punish both actors and omitters. In typical psychological experiments on the punishment asymmetry, subjects read vignettes describing harmful behavior by one individual (an action or omission) and are asked to report the degree to which they believe the individual should be punished. These vignettes typically frame the bad outcome as the result of only one person's behavior, but this is a substantial simplification. There are always many causal factors in the past of a given event, so there will always be concerns that, in changing a scenario so as to manipulate whether an act or omission is the alleged sole cause of the harm, we may have introduced other assumptions about the causal history of the event that are morally

¹We use the term "bias" simply to denote an asymmetry in the treatment of harmful acts versus omissions. We do not intend to imply any normative judgment as to whether this asymmetry is justified or not.

relevant (Asscher, 2007). We therefore study settings in which there are multiple potential transgressors: some who harm by commission, some who harm by omission, to see if there is an asymmetry in preference to punish across these individuals. Rather than looking for different patterns of blame across two different scenarios, we can look at patterns within a single scenario, to get a tightly controlled test.

A third feature of our experiment requires a bit more explanation. Our studies examined an economic consideration that may contribute to explaining the punishment asymmetry: the relative *efficiency* of punishing acts versus omissions. In typical circumstances, it will often be less costly to enforce a rule that proscribes harmful actions than to enforce either a rule that proscribes harmful omissions, or a rule that proscribes both. Consider an example of a community that wishes to preserve a pristine beach environment. If a small amount of litter is left on the beach, this is likely to be the result of just one action (dropping some litter), and also the result of several omissions (every person who was near the beach at the time could have picked it up). Punishment could be used to reduce the frequency of this outcome in different ways. The bad outcome could be prevented by threatening to punish:

- those who drop litter (i.e., enforcing a rule that proscribes actions),
- those who fail to pick up litter (enforcing a rule that proscribes omissions), or
- those who do either of the above (enforcing a rule that demands certain outcomes be achieved)

Provided rates of enforcement are high enough, any one of these policies would suffice to ensure the beach remains pristine. But it is very likely that the cost of administering the second and third policies will be much greater than the cost of administering the first, simply because there will usually be multiple individuals who could have picked up the litter, but relatively few who dropped the litter.²

It should be noted that we focus on injunctive norms, i.e., norms that proscribe or prescribe certain actions (e.g., a norm against littering or a norm for paying taxes), as a possible determinant of the omission bias. Earlier work on perception of causes has found omissions are more likely to be judged causal if they violate either injunctive norms or descriptive norms – if they are immoral or abnormal (Hitchcock & Knobe, 2009; Knobe & Fraser, 2008; McGrath, 2005). Because we are focusing on punishment, however, we take it to be implausible that subjects will judge omissions as more deserving of punishment simply on account of their rarity (though see Eriksson et al., 2015 for evidence suggesting that subjects may conflate the two types of norm). Consequently, we assume that, if an

²Even if there were perfect compliance with a rule proscribing failure to pick up rubbish – in which case there would be no associated punishment costs – if the benefit to those who litter is of lesser magnitude than the inconvenience to those who remove it, then this arrangement will be less efficient than one in which there is perfect compliance with a prohibition on dropping litter. This is a second reason that a rule proscribing actions will tend to be more efficient. In our experimental setting we make the benefit of the harmful action exactly equal to the cost of preventing the harm, thus removing this possible basis for preferring to enforce an action rule.

asymmetry of norms explains the asymmetry of punishment of acts and omissions, it is an asymmetry of injunctive norms.

Two final methodological features of our studies are worth noting. First, punishment decisions are made collectively, by the victim and an unaffected bystander; and second, behavior is observed in a repeated game with fixed groups. We use a repeated game with fixed groups in order to give subjects the opportunity to learn about the social preferences of their groupmates, because social norms are very likely to be a significant factor in any explanation of punishment behavior in these settings (Ramalingam, Morales & Walker, 2019). We use collective punishment decisions for two reasons: there is evidence that it is important that punishment be perceived as legitimate if it is to reliably promote cooperation (Herrmann, Thöni & Gächter, 2008; Sylwester, Mitchell, Lowe, & Bryson, n.d.) and punishment decisions that are made collectively, especially if they involve an unaffected third party, are more likely to be perceived as justified. And secondly, several models suggest coordination of punishment may be especially conducive to the evolution of cooperative behavior (Boyd, Gintis & Bowles, 2010; Gavrilets, 2012; Hilbe et al., 2014).

To the best of our knowledge, no previous studies of the act–omission asymmetry have had this combination of features. While recent studies have demonstrated the act–omission asymmetry in punishment behavior with monetary incentives (Cox, Servátka & Vadovič, 2017; Ramalingam et al., 2019), they have not used collective punishment. Punishment involving an impartial spectator has been examined in social dilemmas (Croson & Konow, 2009; Tan & Xiao, 2012), but not in an environment where there is a meaningful act–omission distinction. One finding from the existing literature on collective punishment is that it can lead to suboptimal levels of punishment (Tan & Xiao, 2012): because impartial spectators are less likely to punish, they may not punish with enough frequency to deter free riders. This entails our design is likely to be a particularly stringent test for the existence of an act–omission asymmetry in punishment, which may be more evident if punishment were decided by individuals with a direct interest.

A primary aim of our studies is to evaluate the norm-based account of the omission effect regarding punishment. That account predicts that we will find an omission bias only if there is an asymmetry in the normative judgments about the propriety of punishing actors versus omitters. Humans may instead possess a heuristic of punishing actions rather than omissions, regardless of their causal or normative judgments, because this heuristic is computationally cheap and will be adaptive in the majority of naturally occurring scenarios (Baron & Ritov, 2004; Gigerenzer, 2008). A further corollary of the norms view is that individuals who regard punishing omitters and punishing actors as morally equivalent, will punish both to similar degrees.

To anticipate, we find a substantial minority of subjects demonstrate an omission effect, preferring to punish those who harm by their actions, and this finding remains robust across differences in the efficiency of punishment. In keeping with the norm theory, the omission bias we find is associated with asymmetric judgments about the fairness of punishing actors

and omitters. Although there is some omission effect – i.e., people are more likely to punish actors than omitters – a substantial proportion of subjects also punish both. Consonant with the norm-based view, this behavior is most pronounced in subjects who judge that it is equally fair to punish actors and omitters. We also find some modest evidence for our hypothesized preference to punish in a relatively efficient manner: a preference to punish actions when relatively few actors contribute to the harm, and to punish omitters when relatively few omitters contribute to the harm. But the magnitude of this effect is notably smaller than the magnitude of the average preference to punish actors rather than omitters.

Below we present our main experiment and a follow up. In addition to these experiments, we report two preliminary experiments, which obtained similar findings, in the Supplement, which also describes all preregistered hypotheses, and any deviations from those hypotheses.

2 Study 1

2.1 Methods

We recruited 210 subjects (mean age 22.0 years, $sd = 3.7$) from the Monash Laboratory for Experimental Economics subject pool to participate in a laboratory experiment involving an economic game. Subjects from our preliminary studies (see Supplement) were excluded. Subjects were randomly assigned to 42 groups of 5 to engage in an economic task with a structure that might be described as an asymmetric public goods game: individuals have the opportunity to engage in behavior that is personally costly but beneficial to the group, or to engage in personally beneficial behavior at cost to the group (Ledyard, 1997). (Our task is asymmetric because, unlike in a standard public goods game, on each round only some subjects have the opportunity to make selfish/altruistic decisions – though over the course of the experiment, all subjects have equal expectations of opportunities to cooperate/defect.)

Interactions were anonymous, via computer software, implemented using oTree (Chen, Schonger & Wickens, 2016). Subjects were paid in private at the end of the session. All the experimental procedures were common knowledge to the subjects, and there was no deception.

We had two treatments, differing in the causal roles of acts and omissions in sustaining the public good. Subjects undertook both treatments, in random order. In both treatments, subjects are randomly placed into groups of 5 players to play a repeated game over 12 or 15 rounds.³

2.1.1 20 treatment (“Two omissions”)

First, all players in each group are assigned an initial endowment of 180, 200, or 220 experimental currency units (ECU), with each endowment chosen independently, at random.

³In some early lab sessions, we conducted the game over 15 rounds, but found that this took too long to complete the experiment in the advertised time. All rounds are included in our analysis below.

(Each ECU was worth AUD0.12, which at the time was worth approximately USD0.09.) The randomization of initial endowments introduces some initial inequalities between individuals, generating more heterogeneous conditions across the experiment, and avoiding any strong implication of strict equality as a normative standard.

Second, players are randomly assigned to various roles for the given round: 1 x Red button presser, 2 x Green button pressers, 1 x Button-affected individual, and 1 x Unaffected individual. These were the terms used in all instructions to the subjects. For ease of exposition, however, we henceforth refer to the button-affected individual as the “victim” and the unaffected individual as the “bystander”.

The Red button presser is an individual who can choose to press a red button, for a gain of 40 ECU. Pressing the button, however, will cause the victim to lose 120 ECU. The Green button pressers can then choose whether or not to press their buttons, which will have the effect of undoing the harm to the victim. If a green button is pressed, the victim will have his/her initial endowment restored, but the green button pressers lose 40 ECU (if both press, the cost is shared, if only one presses, the cost is borne by the sole presser); no change occurs to the red presser when green is pressed. See Figure 1 for an overview of these stages of the task.




FIGURE 1: Overview of initial stages of 2O (Two omissions) treatment. Note that green button pressing decisions are made simultaneously, so in deciding to press green, an individual does not know whether they will have to shoulder the entire burden of restoring victim, or if it will be shared.

Both pressing the red button, and refraining from pressing a green button, are harmful, in the sense that they not only have negative consequences for the victim but the aggregate payoff to the group is also lower by 80 ECU. Because the game is repeated, and roles are randomly allocated each round, the diachronic extended game resembles a public goods problem. The aggregate group payoff is maximized either if the red button is never pressed, or if, whenever the red button is pressed, at least one of the green buttons is also pressed. The first of these possibilities requires the red button pressers to forgo an opportunity for personal gain, while the second requires the green button pressers to suffer a loss. If the group can enforce a norm of refraining from pressing the red button and/or always pressing green buttons, the public good will be preserved, to the expected benefit of all over the entire game.⁴

After decisions on button pressing are made, a punishment stage occurs. If the red button was pressed, the bystander is given the option to propose (costly) punishment of behaviors that contributed to the harm. The bystander is asked, e.g., “How many ECU do you wish to deduct from the red button presser?” and “How many ECU do you wish to deduct from each of the green button non-pressers (if any)?” By answering these questions, the bystander nominates punishment levels separately for both red pressers (harmful actors) and, if relevant, for green refusers (harmful omitters). The victim is informed of the proposed punishment levels and can choose to accept or reject the proposal. If it is accepted, the punishment is enacted, otherwise it has no effect. Punishment jointly costs the punishers 1 ECU per 3 ECU to be deducted from the punished party. If for instance the bystander nominates 9 ECU be deducted from each green omitter, and 6 ECU be deducted from the red presser, this will have the effect of reducing the relevant endowments by those amounts, but will also cost $((9 \times 2) + 6)/3 = 12$ ECU, and this cost is divided equally between the bystander and the victim. The maximum amount that can be deducted from an individual is 30 ECU.

Although punishment is a net destructive activity, if it sustains a norm of not pressing red, or always pressing green, it may be to the group’s overall benefit to punish, especially on earlier rounds (Gächter, Renner, & Sefton, 2008). Subjects receive full feedback on all decisions made in their group at the end of each round. Because we randomly shuffle identifiers between rounds, there is little danger to punishers that they will suffer retaliation on subsequent rounds, thus maximizing the potentially prosocial motivation for punishment.

⁴This game thus contains elements resembling “take-some” public goods games in the sense that the red button pushers can take additional resources, reducing the group’s aggregate payoff; and elements resembling the opposite, “give-some” framing in that the green button pushers can forgo resources, thereby increase the group’s aggregate payoff (Ledyard, 1997; Fehr & Gächter, 2000; Chaudhuri, 2011). In simple, two-person games it is common to see more cooperative, prosocial behavior in take-some framings, rather than in give-some framings (van Dijk & Wilke, 2000). These framing differences, however, have not been observed in more complex, multi-player games (Haesevoets, Hiel, Assche, Folmer & Bostyn, 2019). Further, if framing of this sort was important in our game, we should expect to see relatively more selfish behavior by green button subjects than red, but, as we report below, we find the opposite.

2.1.2 2A treatment (“Two actions”)

In the 2O treatment, harm occurs only if one action and two omissions are performed, hence we refer to that treatment as “Two omissions”: **2O**. Our other treatment involves reversing the causal structure, such that two actions and one omission is necessary for harm to occur, hence the treatment is referred to as “Two actions”: **2A**. Otherwise, the games are the same. See Figure 2 for a comparison.




FIGURE 2: The parallel causal structures of the two treatments. Normal arrowheads indicate causal promotion, circular heads indicate prevention. Nodes shaded grey indicate non-occurrence of the relevant event. In 2O, one omission is insufficient for harm (a), and correspondingly, in 2A, one action is insufficient for harm (d). In 2O, two omissions won't cause harm, provided no action occurs (b); and correspondingly in 2A, two actions won't cause harm, provided no omission occurs (e); finally, in 2O, two omissions cause harm, when a harmful action also occurs (c); and in 2A, two actions cause harm, when a harmful omission occurs also (f).

Our principal measure of interest is the amount of punishment proposed on rounds where the harm eventuates: where all available red buttons are pressed and no one presses a green button. On such rounds, there are three potential targets for punishment, and we separately record the bystander's proposed punishment for actors and for omitters.

Our design therefore addresses the possible interaction of two variables in determining amounts of punishment (i) whether the relevant subject performed an *action* or *omission* that was causally relevant, and (ii) whether the action or omission is the only event (*solely*

responsible) of that type required for the harmful outcome or if multiple events of that type were required (*jointly* responsible).

Because our principal measure is the amount of punishment proposed on a round where harm actually occurs, there was the possibility of obtaining insufficient data if harm did not often occur. As a precaution against this, on one round we collected subject responses using the “strategy method” (Selten, 1967). The relevant round was chosen at random in each lab session (to control for any order effects). The strategy method involves subjects indicating their choices hypothetically, for all possible contingencies, in advance of knowing which role they have been allocated to or what choices other subjects will make. These choices are then combined with allocation to roles to determine an outcome to the round, using the same payoff structure as in normal rounds. Observations from this round are excluded from our principal measure, but are consistent with our main results (see Supplement).

After the punishment decisions were made, subjects were asked to judge the propriety of punishing actors and omitters in the most recent treatment they completed. In particular, subjects were asked to answer the following questions on a scale of 1 (not at all fair) to 7 (completely fair):

To what extent do you think it would be fair to be punished for pushing the Red button?

To what extent do you think it would be fair to be punished for not pushing the Green button?

Subjects also completed a demographic survey. Subjects earned a fee of AUD10 for participation, plus their earnings from one round of the experiment, chosen at random. (Payment for a single randomly chosen round avoids any wealth effects that are liable to occur if earnings aggregate throughout the experiment.) Average earnings were AUD34.38 (approximately USD27), $sd = 7.44$.

2.2 Results

Both harmful actions and harmful omissions were relatively frequent, with the average subject pressing red on 68% of opportunities and refusing to press green on 61% of opportunities.

Across all experimental sessions, there were 296 rounds in which the group had an opportunity to punish those who contributed to a harmful outcome. We examine both the frequency with which punishment is proposed and the amount of punishment proposed. Punishment was proposed by the bystander on 160 of those rounds (54%). The distribution of these punishment proposals is summarized in Table 1 below. In both treatments, a large proportion of respondents proposed to punish both actions and omissions. Some degree of omission effect appears evident, however, in that there are more occasions when actions alone are punished than omissions alone. Across treatments, the distribution of

punishment proposals is relatively similar, suggesting little or no effect of our causal structure manipulation, at the level of punishment frequency.⁵

TABLE 1: Summary distribution of punishment proposals made across the two treatments. To control for repeated observations across groups, we use a Cochran-Mantel-Haenszel Chi-Squared test, with continuity correction. The test does not quite reject the null hypothesis that the probability of punishing acts and omissions is the same in the 2A treatment ($\chi^2(1)=3.54$, $p=0.060$, 95% CI for the odds ratio 1.08, 19.33), but rejects that null hypothesis in the 2O treatment ($\chi^2(1)=18.52$, $p<0.001$, 95% CI for the odds ratio 4.46, 839.21). Evidently both conditions show a very similar pattern.

Treatment	Punish actions only	Punish omissions only	Punish both	Punish neither
2A (n = 163)	11%	5%	36%	48%
2O (n = 106)	15%	5%	35%	45%

Examining the amounts of punishment proposed for acts versus omissions, we observed a consistent pattern whereby actions were punished more harshly than omissions, and solely responsible harmful behaviors are punished more than jointly responsible behaviors. We calculate the asymmetry in punishment towards acts by subtracting each group’s average proposed punishment of omissions from its average proposed punishment of actions. In 2A, 32 of 37 groups had positive bias towards punishing actions; in 2O, 27 of 35 groups were biased towards punishing actions. Table 2 presents the mean punishment levels proposed for each harmful behavior type (note that because these means include cases where zero punishment was proposed for both actors and omitters, it arguably underestimates the degree of act–omission asymmetry).

TABLE 2: Summary distribution of mean punishment levels proposed for jointly/solely responsible acts and omissions, conditional on any punishment being proposed.

	Actions	Omissions
Jointly responsible	6.95 (sd = 2.94, n = 83)	6.10 (sd = 3.62, n = 49)
Solely responsible	7.94 (sd = 3.08, n = 64)	6.44 (sd = 3.50, n = 72)

Although there was a clear pattern of preference to punish actions more harshly, we were struck by the relatively high proportion of our subjects who were willing to punish omissions. Across the 2A and 2O treatments, there were 179 subjects who had at least one opportunity to propose punishments of both actors and omitters. Some form of punishment was proposed by 110 of those subjects (61%). Of those subjects, 86 proposed punishing

⁵Because punishment will not be enacted without the approval of the victim, we focus on punishment proposed as a more informative measure of punishment intentions. For completeness, in the supplement, we also describe the distribution of actual punishment dispensed, which is broadly similar.

omitters at least once. In other words, 48% of all subjects who had the chance to punish, and 78% of subjects who ever actually proposed punishment, on at least one occasion suggested that some punishment should be inflicted on omitters. These figures were much higher than our prior expectations. (See Table S4 in the Supplement for further details, including confidence intervals.)

2.2.1 Fairness

We turn now to the descriptive results on the judgments of fairness. Figure 3 summarizes the results from the survey in which we asked subjects about their perception of the fairness of punishing harmful acts and omissions. We assume subjects will regard it as fair to punish a behavior if and only if that behavior is proscribed by social norms. Accordingly, these responses give us some indication of whether there is an asymmetry in (injunctive) normative attitudes to harmful omissions and actions in this setting.

Evidently, it is less controversial whether it is fair to punish a harmful action, whereas there is significant disagreement about the fairness of punishing a harmful omission. A within-subjects test confirms that the answers to these questions are significantly different at the group level ($z = 3.881$, $p = 0.0001$, Wilcoxon sign-rank test).

We then examined within subjects the difference between their answers to these two questions (Figure 4). A positive score indicates they agreed more with the fairness of punishing a harmful action than punishing a harmful omission, zero indicates that they regarded both as equally fair (or equally unfair), and a negative score indicates they regarded it as more fair to punish harmful omissions.

Only 19% of subjects rate it as more fair to punish omissions than actions. 30% rate the two as equally fair, and the remaining 51% of subjects rate it as more fair to punish actions than omissions. Comparing these results at the group level across the two possible treatment orders suggests no significant difference ($z = 0.556$, $p = 0.578$, Wilcoxon test).

2.2.2 Correlating asymmetries of punishment and fairness judgments

To initially assess the association of fairness judgments and punishment behavior, we generate a variable *Punishment bias* which is, for each round, the amount of punishment proposed for harmful actions divided by the total amount of punishment proposed for acts and omissions, less one half (thus omitting cases in which no punishment is proposed). This variable thus ranges from -0.5 for a subject who punishes only omissions, to $+0.5$ for a subject who punishes only actions, and is zero whenever acts and omissions are punished equally. Taking the group averages of punishment bias and fairness asymmetry, we observe there is a significant correlation between them (Kendall's tau = 0.30, $p = 0.037$): suggesting that the asymmetry of punishment is associated with asymmetric fairness judgments.⁶

⁶An alternative measure of punishment asymmetry, which includes cases of zero proposed punishment, is to subtract punishment of omissions from punishment of actions. Conducting a similar correlation test with fairness attitudes, using this measure, we obtain a Kendall's tau of 0.27, $p = 0.062$.




FIGURE 3: Histograms of individual answers to question asking about perceived fairness of punishing an agent for a harmful act (pressing red, left figure) or omission (failing to press green, right figure), by treatment (solid grey bars 2A, hollow red bars 2O). Answers relate to a 7-item Likert scale, ranging from “Completely fair” to “Not at all fair”.




FIGURE 4: Histogram of individual differences in perceived fairness of punishing an action versus punishing an omission. Responses range from -6 (maximally more fair to punish omissions) to $+6$ (maximally more fair to punish actions).

We further interrogate this data in a series of linear mixed-effects regression models to predict the amount of punishment proposed, presented in Table 3. The principal independent variables used in the various models are: *act/omission*, an indicator variable for whether the targeted behavior is an action (pressing a red button) or an omission (refraining from pressing green); *jointly responsible*, an indicator for whether the behavior type (act/omission) is one of two instances that contributed to the harm. So in the 2A treatment, for instance, actions are jointly responsible and omissions are not jointly responsible, and vice versa in the 2O treatment; *treatment order* is a dummy for the sequence in which the subject engaged in each treatment; *Asymmetry of fairness judgments* is a standardized continuous variable, reporting the difference between the subject's reported judgment of the fairness of punishing harmful acts and punishing harmful omissions.

Model (1) finds both the act–omission distinction and joint-versus-sole responsibility predict punishment levels, with actions being punished more and jointly responsible contributions being punished less. In other words, the model estimates that on average, solely responsible actions will attract the most punishment, and jointly responsible omissions will attract the least. The difference between these extreme cases is approximately 1.7 ECU in punishment expenditure, which has three times that magnitude difference in terms of effect on those punished. Model (2) adds additional covariates for gender of subjects and number

of experiments completed in the past, and finds similar results. Model (3) tests for a possible interaction between act/omission and joint responsibility, and does not find one.

TABLE 3: Regression models for punishment level proposed by bystander. Standard errors in parentheses. A coefficient of 1 means that an increment of the associated variable predicts an additional 1 ECU expenditure on punishment (which has an effect of 3 extra ECU being deducted from the punished party). All models include random effects for individual and continuous, fixed effects for round, coefficients not shown here. Additional covariates are for gender and number of experiments done in the past. Dummy variables other than gender are mean centered. Note that N is lower for some models because not all subjects completed the post-experiment survey items in the allotted time for the lab session.

	(1)	(2)	(3)	(4)	(5)	(6)
Action (not omission)	1.139*** (0.251)	1.070*** (0.256)	1.139*** (0.253)	1.097*** (0.269)	1.065*** (0.268)	1.065*** (0.268)
Jointly responsible (not solely)	-0.541* (0.251)	-0.625* (0.256)	-0.541* (0.253)	-0.649* (0.269)	-0.620* (0.268)	-0.620* (0.268)
Treatment order	0.764 (0.556)	0.752 (0.543)	0.820 (0.532)	0.568 (0.576)	0.567 (0.576)	0.461 (0.571)
Action * Jointly responsible			0.222 (0.687)			
Asymmetry of fairness judgments				-0.068 (0.280)	0.068 (0.280)	0.029 (0.279)
Action * Asymmetry of fairness judgments					0.519* (0.260)	0.519* (0.260)
Additional covariates	No	Yes	No	No	No	Yes
N	538	522	538	492	492	492

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Model (4) includes subjects' opinions about the fairness of punishing those who commit harmful actions versus harmful omissions. *Normative asymmetry* is positive when a subject rated it more fair to punish acts than omissions, and zero if they regard these punishment options equally fair. This variable is not itself a significant predictor of punishment levels, but the interaction of this variable with action/omission is significant (models (5) and (6)). Figure 5 illustrates the interaction: the act–omission distinction only makes a difference for subjects who have an asymmetry in their fairness attitudes towards punishment of acts and omissions.

This result suggests that it is primarily subjects holding this asymmetric normative attitude who punish acts more than omissions. For instance, model (6) estimates that, other




FIGURE 5: Interaction plot of normative judgment asymmetry and act/omission, with 95% confidence intervals, derived from model (6). Note that because the normative judgment asymmetry variable has been standardized for the purposes of mean-centering, zero on the horizontal axis does not mean a perfectly symmetrical attitude toward acts and omissions. The strictly symmetrical attitude is indicated with a dashed vertical line.

things being equal, an individual who has the maximally asymmetric attitude will spend 2.6 ECU more to punish an action than an omission, whereas an individual who has a symmetric attitude will spend only 0.6 ECU more on punishing actions than omissions.

More generally, we observe that the act–omission distinction, or its interaction with fairness, is a significant predictor in all models. We also observe that there is a tendency for subjects to spend approximately 0.6 ECU less on the punishment of jointly responsible individuals. This suggests some individuals may be adjusting their punishment to align with the more efficient punishment strategy (which is to forgo punishing jointly responsible individuals), but it is a relatively modest effect: about half the magnitude of the predicted increase in punishment for actions versus omissions in general.

The regression models reported above suggest that asymmetric judgments of fairness of punishing acts and omissions predict asymmetric punishment behavior. We cannot be confident about the causal direction – for instance, fairness judgments might be post hoc rationalizations of earlier behavior. However, this finding does conform to the predictions of the norms account.

As a robustness check, we performed similar regressions, including groups who failed our pre-registered comprehension criterion. The results were broadly similar. See the Supplement.

2.3 Discussion

Our first study confirmed that there is a robust act–omission distinction in the novel setting of our experiment, where punishment involved substantial personal cost for both punisher and punished. It also showed that, on average, subjects will be biased in their punishment decisions to punish acts more than omissions, even when the structure of the task is such that it is more cost-effective to punish harmful omissions.

Notwithstanding that we found a substantial act–omission distinction on average, we also found that a substantial portion of our subjects punished both acts and omissions equally. This is in line with earlier work on the omission effect with regard to choice and blame, which found substantial individual differences (Spranca et al., 1991). On closer analysis, this study suggests that the variation is associated with a difference in fairness attitudes towards the punishment of actions and omissions. A substantial proportion of our subjects thought it equally fair to punish a harmful actor or ommitter, and this attitude was highly predictive of punishment behavior. Though we cannot, with the present design, establish whether these attitudes are causally upstream or downstream of the punishment asymmetry, the finding fits well with a norm-based explanation of the punishment asymmetry. To more directly compare this explanation to a major competitor – a causal account – in the next study we added survey questions assessing the subjects’ perceptions of the causal contributions of actions and omissions to the overall outcome.

3 Study 2

3.1 Methods

In the previous study, we used an asymmetric structure where harm eventuated when two agents performed harmful acts/omissions, and only one agent performed the opposite type of behavior. Because of this, agents proposing punishment always faced twice the marginal cost to punish one type of behavior or another. Although this asymmetry was counterbalanced across the treatments, we decided to conduct a second study with a symmetrical structure, to enable a more direct comparison of willingness to punish acts versus omissions in a single round.

In Study 2, we therefore generalized the structure of our previous treatments to use a paradigm where multiple actions and multiple omissions are required for harm to eventuate. We refer to this as the 2A2O (“two act, two omission”) treatment. In this setting, there are six subjects in each group: two red button pressers, two green button pressers, a “victim” and a “bystander”. Both of the red buttons must be pressed, and neither of two green buttons must be pressed, for harm to eventuate to the victim. As in the previous study, the benefit of pressing red (+40) is fixed and will be shared evenly if both buttons are pressed. Similarly, the cost of pressing green (–40) is fixed and will be shared evenly if both green buttons are pressed. As in the previous study, punishment decisions are proposed by the bystander,

and must be approved by the victim in order to be enacted. Punishment costs the punishers 1 ECU to inflict 3 ECU loss on a punished party. The costs of any punishment imposed are shared equally between the two punishers. The maximum punishment that can be inflicted is 30 ECU per subject.

(The opportunity to punish potentially harmful behaviors is provided on rounds where no harm occurs – if, for instance, one or more red buttons is pressed and one or more green buttons is pressed. Our main analyses concern only punishment decisions where harm did occur, and thus could be ascribed to either red button pressing or omission to press a green button – but for completeness we also describe punishment data from no-harm events in the supplement.)

In addition to this structural change, we added two survey questions asking subjects to evaluate the causal influence of the relevant acts and omissions on the eventual outcome (see Box). By including these responses in our regression models, we could estimate the degree to which an asymmetry in perceptions of causal responsibility might explain the preference to punish acts more than omissions.

Box. Wording of questions regarding causal influence of acts and omissions on eventual harm.

Earlier in this experiment, you participated in one task in which there were **two people** who had the opportunity to press a **red button** and **two people** who had the opportunity to press a **green button**.

Imagine a scenario in which **both red buttons were pressed** by participants C and B, and **both green buttons were not pressed** by participants A and E. As a result, participant D – the button-affected participant – lost 120 ECU.

Indicate your agreement with the following:

- Participant C's pressing Red was one of the causes of the button-affected participant losing 120 ECU.

(7 point Likert scale ranging from 1 = Strongly disagree to 7 = Strongly agree.)

- Participant A's not pressing Green was one of the causes of the button-affected participant losing 120 ECU.

(7 point Likert scale ranging from 1 = Strongly disagree to 7 = Strongly agree.)

We conducted this experiment with 138 subjects in 23 groups of 6 (mean age 22.7 years, $sd = 3.4$) recruited from the same source as our earlier studies. As before, subjects interacted anonymously via software implemented using oTree, and were paid in cash, in private at the end of each session. Average earnings were AUD34.21, $sd = 4.58$. Subjects from our earlier experiments were excluded from participating.

3.2 Results

Consistent with Study 1, we observed that acts are more likely to be punished than omissions (Table 4), and gross expenditure on punishment of acts exceeds that on punishment of omissions. In all 23 groups, the average proposed punishment of harmful actions (mean 4.17, SD = 2.28) was greater than the average proposed punishment of harmful omissions (mean 1.17, SD = 1.07). Also consistent with Study 1, we observed high rates of punishing omitters. On 48% of rounds where punishment was possible, it was proposed that omitters receive some punishment. Conditional on any punishment being proposed at all, 22/33 punishment proposals included punishment of omitters.

TABLE 4: Contingency table of punishment proposal types, where both actors and omitters are jointly responsible. A Cochran-Mantel-Haenszel Chi-Squared test rejects the null hypothesis that the probabilities of punishing acts and omissions are the same ($\chi(1) = 4.88$, $p = 0.02721$).

	Punish omissions	Don't punish omissions
Punish actions	33%	28%
Don't punish actions	5%	34%

3.2.1 Causal judgments

Mean agreement with the proposition that pressing red was a cause of a subject being harmed was 5.3 on a 7-point Likert scale ($sd = 1.67$, 95% CI 5.02–5.59). For the equivalent proposition that failure to press green was a cause of harm, mean agreement was 4.4 ($sd = 1.97$, 95% CI 4.09–4.75). Figure 6 summarises the distribution of the within-subject differences on these items. The mean difference is 0.88 ($sd = 2.39$, 95% CI 0.48–1.29).

We attempted to assess whether causal judgment asymmetry might explain a preference to punish actions, by conducting individual level correlations of individual bias in punishment proposals (individual average proposed amounts to punish harmful actions, less proposed punishments of harmful omissions) with both causal judgment asymmetry and normative judgment asymmetry (Table 5). With only 23 groups, however, this study was insufficiently powered to detect any association.⁷

3.3 Discussion

Study 2 confirmed the main result of our main study: that individuals prefer to punish actions rather than omissions. However, as in Study 1, we also found a substantial amount

⁷Results are broadly similar if we use the alternative, absolute measure of punishment asymmetry, whereby proposed punishment of omissions is subtracted from proposed punishment of actions: correlations are positive, but modest, and not statistically significant.




FIGURE 6: Distribution of within-subject differences in judgments of causal responsibility for harmful acts and omissions. Responses range from -6 (maximal asymmetry favouring omission causation) to $+6$ (maximal asymmetry favouring action causation).

TABLE 5: Correlations of group level asymmetries in punishment proposals, causal judgment, and fairness attitudes regarding punishment ($n = 23$ groups; Kendall's tau; p-values in parentheses).

	Causal judgment	Fairness
Punishment	0.21 (0.26)	0.19 (0.29)
Causal judgment	–	0.05 (0.80)

of punishment of omitters. Compared to our first Study, this experiment demonstrated this result in a setting where the causal roles of actions and omissions were symmetric.

This study also directly asked subjects about their perception of the causal role of acts and omissions in the relevant scenario. There was no significant correlation between perceiving acts to be more causal and punishing acts more than omissions.

Our failure to find a significant association between causal perception and punishment behavior may of course be due to a lack of power, but is prima facie at odds with several findings which suggest that the preference to commit harm by omission is at least partly explained by differences in causal perception of acts versus omissions (Baron & Ritov, 2009; Cushman & Young, 2011; Greene et al., 2009). There are several possible explanations for why we failed to find this effect. We used a question about causal responsibility that was relatively broad and inclusive in the causal language used (roughly: do you agree that the pushing/non-pushing was “**a cause** of the outcome” [emphasis added]), and this may

have therefore prompted subjects to consider counterfactual notions of causation which are notoriously inclusive (Lewis, 2000, pp. 195–197), rather than notions of “direct” versus “indirect” causation, which have been argued to be important in this regard (Baron & Ritov, 2009). It is also possible that our findings diverge because our principal measure is how subjects *punish* acts versus omissions, rather than on the actual *performance* of act versus omission, but we find it unlikely that this difference explains the divergence.

4 General Discussion

Across a total of four experiments (two of which are reported in the supplementary materials) involving five treatments and 580 subjects in 123 groups of varying size, we observed a widespread tendency for subjects to punish harmful acts more than harmful omissions. The economic games across these studies shared a basic structure. Some players were given an opportunity to take valuable resources from another player, other players were given an opportunity to rectify the victim’s losses, and other players made decisions about whether to punish those who took resources (“actors”), those who failed to rectify victims (“omitters”), or both. In each study subjects were on average significantly more likely to punish actors than omitters, and they punished actors more harshly.

One motivation for our studies was to assess whether the act–omission bias in punishment reflects a choice to implement a relatively efficient punishment strategy. In typical settings, a single actor suffices to cause a bad outcome (e.g., litter being on the ground) but many people could potentially rectify this bad outcome. There is therefore an asymmetry between the number of people that you would need to punish if you enforced a prohibition on harmful action versus the number of people you would need to punish if you enforced an obligation to rectify harm. Other things being equal, it should thus be more efficient to punish actions rather than omissions in typical settings. Our experiment was designed to manipulate the typical situation, to assess whether the act–omission bias in punishment reflects a strategic choice to implement an efficient punishment strategy.

Our studies provided only limited support for the efficiency hypothesis. In Study 1, we manipulated that asymmetry by varying the number of actors and omitters causally implicated in causing harm. One condition had a single potential actor and two potential omitters. In this case the bad outcome would occur if the actor pushed the red button and both omitters failed to push the green button. Another condition flipped this around: there were two potential actors and one potential omitter. In this case, the bad outcome would occur if both actors pushed the red button and the omitters failed to push the green button. In this second condition, it is more efficient to punish the omitter than the actors. If agents adopted the strategy of punishing the smallest group of individuals necessary to deter the loss of public goods, then they would always have punished the solely responsible category. There was some evidence that the efficiency makes a modest difference, but the magnitude of this effect was substantially smaller than the effect of the act/omission distinction.

Our efficiency hypothesis was motivated by the thought that punishment of actions is less burdensome than punishment of omissions, because there will typically be fewer culpable actors than omitters to punish. Another way in which punishment of omissions is more burdensome than punishment of actions is that omissions tend to leave no physical traces of having occurred, hence it is difficult to have conclusive evidence that an individual culpably omitted to do something. Consider an air traffic controller who fails to notify a pilot of an imminent collision. Did the air traffic controller have the opportunity to notify the pilot, or was she unconscious at the time? There is unlikely to be any physical trace that would verify either hypothesis. By contrast, an air traffic controller who deliberately gave a misleading message that no collision was imminent (performed a culpable action) will leave a trace on the aircraft's voice recorder. DeScioli and colleagues (DeScioli, Bruening, et al., 2011) argue that this difference leads – in normal circumstances – to a preference to condemn harmful actions rather than harmful omissions. This in turn leads to a strategic preference on the part of individuals who wish to perform self-serving, harmful behaviors to do so by omission rather than commission (DeScioli, Christner & Kurzban, 2011). But they find that these omission effects – in punishment and in harmful choices – are reduced substantially when the evidential asymmetry is removed: omissions that leave a physical trace of having occurred are condemned more than omissions which come about from mere failure to act in time, and this holds even when controlling for equivalent, malevolent intentions in the two cases. In our setting, there was arguably no evidential asymmetry between acts and omissions: it was transparent to subjects when harmful omissions occurred, and omissions could not occur through a subject simply failing to make a decision in time. The account of DeScioli and colleagues, however, may contribute to explaining the smaller than expected magnitude of the omission effect that we observed.

A second motivation for our studies was to investigate whether the asymmetry of punishment behavior is likely the result of differences in normative attitudes towards acts and omissions. In both our studies, there was significant heterogeneity in subjects' perceptions of the relative fairness of punishing acts versus omissions. A substantial proportion of our subjects thought it equally fair to punish omissions and actions, while roughly half thought it more fair to punish acts than omissions. This heterogeneity in judgments of fairness turned out to be a significant predictor of the choice of punishment policy, which is at least consistent with the hypothesis that the act–omission bias is explained by an asymmetry in judgments of fairness.

An alternative hypothesis is that subjects punish actions more than omissions because they perceive actions as more causally responsible for the eventual harms. We examined this hypothesis in Study 2, where a substantial fraction of subjects reported in a post-experimental survey that they believe pressing the red button was more causally responsible for the bad outcome than refraining from pressing the green button. We could not draw any strong conclusions in this regard – there was no significant association between causal judgment asymmetry and punishment asymmetry.

One limitation of our design is that, although we made significant efforts to make acts and omissions causally equivalent in their contributions to the harm, in all our treatments, the harmful acts were temporally prior to the harmful omissions. This may have influenced perception of the causal role of the respective acts and omissions. In some legal doctrines regarding causation, for instance, an intervening action can render earlier events non-causal, but intervening omissions tend not to be regarded in this way (Moore, 2019). If this sort of consideration were relevant here, then in an experiment in which harmful omissions (like a failure to vaccinate) precede harmful actions, we might expect a stronger asymmetry of causal perception, and perhaps also a stronger asymmetry of punishment. But this speculation remains to be tested.

Although we found a robust tendency to punish acts more than omissions, we were surprised by the large proportion of subjects who chose to punish both acts and omissions. In both Studies 1 and 2, as well as the two additional studies reported in the supplemental materials, many subjects punished those who failed to exhibit beneficence, i.e., positive acts of kindness or generosity. The omitters had a chance to restore funds to a victim, and if they fail to do so, many subjects apparently thought it appropriate to punish them. This was perhaps the most interesting finding from our studies. It runs against the moral view suggested by Adam Smith, who asserts the following proposition:

Beneficence is always free, it cannot be extorted by force, the mere want of it exposes to no punishment; because the mere want of beneficence tends to do no real positive evil.

A recent paper by Vernon Smith and Bart Wilson (2017) presents empirical evidence in support of Smith's proposition in the context of trust games. Our results run counter to this claim in the context of social dilemmas. Indeed, among those who punished, a majority of them punished individuals who failed to exhibit beneficence.

What makes these results particularly surprising is that most of the explicit laws and moral rules we have are act-based (Ritov & Baron, 1990, 1999). It's against the law to steal, but not against the law to refrain from rectifying the victim of theft. By and large, littering is proscribed, but not failing to pick up litter. ("Bad Samaritan" laws do exist in some jurisdictions, but are a relative rarity (Robinson & Darley, 1995 chap. 2).) One is required to keep one's promises, but not to ensure that others keep their promises. Our results suggest that the prevalence of act-based rules is not due to some widespread psychological aversion at punishing those who fail to be beneficent. This accords with the suggestion due to Willemssen and Reuter (2016) that an act-omission bias only emerges when there is a clear norm that proscribes actions in particular. In our case – a relatively novel, artificial game in a laboratory setting – it is perhaps unsurprising that there may be some disagreement among subjects as to the relevant norms.

Another factor that may explain the high rate of punishment of omissions is that our paradigm may have created a social role of responsibility for the bystander. There is evidence that, at least in some cultures (Baron & Miller, 2000), social roles of responsibility

reduce the act/omission asymmetry, generating similar levels of condemnation for harms by omission or commission (Haidt & Baron, 1996). This mechanism, if present, is presumably compatible with a norm-based account of the asymmetry, since social roles are in large part defined by the norms that are associated with those roles.

Although we interpret our results as supporting a norm-based account, it should be emphasized that we did not directly manipulate the beliefs of the subjects about the fairness of punishing omissions and thus cannot draw any strong causal inferences. We welcome future research to address this limitation, though we note there is a potential difficulty in manipulating beliefs about what is morally required, given these are beliefs which are – unlike social norms – theorized as relatively insensitive to social cues (Bicchieri, 2005).

References

- Alicke, M. D. (1992). Culpable causation. *Journal of Personality and Social Psychology*, 63(3), 368.
- Asch, D. A., Baron, J., Hershey, J. C., Kunreuther, H., Meszaros, J., Ritov, I., & Spranca, M. (1994). Omission bias and pertussis vaccination. *Medical Decision Making*, 14(2), 118–123. <https://doi.org/10.1177/0272989X9401400204>.
- Ashworth, A. (2015). *Positive obligations in criminal law*. London, UNITED KINGDOM: Bloomsbury Publishing Plc. Retrieved from <http://ebookcentral.proquest.com/lib/monash/detail.action?docID=1724926>.
- Asscher, J. (2007). Killing and letting die: The similarity criterion. *Journal of Applied Philosophy*, 24(3), 271–282. <https://doi.org/10.1111/j.1468-5930.2007.00381.x>.
- Baron, J., & Miller, J. G. (2000). Limiting the scope of moral obligations to help: A cross-cultural investigation. *Journal of Cross-Cultural Psychology*, 31(6), 703–725. <https://doi.org/10.1177/0022022100031006003>.
- Baron, J., & Ritov, I. (2004). Omission bias, individual differences, and normality. *Organizational Behavior and Human Decision Processes*, 94(2), 74–85.
- Baron, J., & Ritov, I. (2009). Protected Values and Omission Bias as Deontological Judgments. In B. H. Ross (Ed.), *Psychology of Learning and Motivation* (Vol. 50, pp. 133–167). Academic Press. [https://doi.org/10.1016/S0079-7421\(08\)00404-0](https://doi.org/10.1016/S0079-7421(08)00404-0).
- Bicchieri, C. (2005). *The grammar of society: The nature and dynamics of social norms*. Cambridge University Press.
- Boehm, C. (1999). *Hierarchy in the forest: The evolution of egalitarian behavior*. Cambridge, MA: Harvard University Press.
- Boyd, R., Gintis, H., & Bowles, S. (2010). Coordinated punishment of defectors sustains cooperation and can proliferate when rare. *Science*, 328(5978), 617–620.
- Chaudhuri, A. (2011). Sustaining cooperation in laboratory public goods experiments: A selective survey of the literature. *Experimental Economics*, 14(1), 47–83. <https://doi.org/10.1007/s10683-010-9257-1>.

- Chen, D. L., Schonger, M., & Wickens, C. (2016). oTree—An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, *9*, 88–97. <https://doi.org/10.1016/j.jbef.2015.12.001>.
- Cox, J. C., Servátka, M., & Vadovič, R. (2017). Status quo effects in fairness games: Reciprocal responses to acts of commission versus acts of omission. *Experimental Economics*, *20*(1), 1–18.
- Croson, R., & Konow, J. (2009). Social preferences and moral biases. *Journal of Economic Behavior & Organization*, *69*(3), 201–212. <https://doi.org/10.1016/j.jebo.2008.10.007>.
- Cushman, F., & Young, L. (2011). Patterns of moral judgment derive from nonmoral psychological representations. *Cognitive Science*, *35*(6), 1052–1075. <https://doi.org/10.1111/j.1551-6709.2010.01167.x>.
- Cushman, F., Young, L., & Hauser, M. (2006). The role of conscious reasoning and intuition in moral judgment. *Psychological Science*, *17*(12), 1082–1089.
- DeScioli, P., Bruening, R., & Kurzban, R. (2011). The omission effect in moral cognition: Toward a functional explanation. *Evolution and Human Behavior*, *32*(3), 204–215. <https://doi.org/10.1016/j.evolhumbehav.2011.01.003>.
- Dowe, P. (2004). Causes are physically connected to their effects: Why preventers and omissions are not causes. In C. Hitchcock (Ed.), *Contemporary Debates in Philosophy of Science* (pp. 189–196). Malden, MA: Blackwell.
- Eriksson, K., Strimling, P., & Coultas, J. C. (2015). Bidirectional associations between descriptive and injunctive norms. *Organizational Behavior and Human Decision Processes*, *129*, 59–69.
- Fehr, E., & Gächter, S. (2000). Cooperation and punishment in public goods experiments. *American Economic Review*, *90*(4), 980–994.
- Feldman, G., Kutscher, L., & Yay, T. (2020). Omission and commission in judgment and decision making: Understanding and linking action-inaction effects using the concept of normality. *Social and Personality Psychology Compass*, *14*(8), e12557. <https://doi.org/10.1111/spc3.12557>.
- Gächter, S., Renner, E., & Sefton, M. (2008). The long-run benefits of punishment. *Science*, *322*(5907), 1510–1510.
- Gavrilets, S. (2012). On the evolutionary origins of the egalitarian syndrome. *Proceedings of the National Academy of Sciences of the USA*, *109*(35), 14069–14074.
- Gigerenzer, G. (2008). Moral intuition = fast and frugal heuristics? In *The Cognitive Science of Morality: Intuition and diversity*. MIT Press. Retrieved from <http://psycnet.apa.org/psycinfo/2007-14533-001>.
- Greene, J. D., Cushman, F. A., Stewart, L. E., Lowenberg, K., Nystrom, L. E., & Cohen, J. D. (2009). Pushing moral buttons: The interaction between personal force and intention in moral judgment. *Cognition*, *111*(3), 364–371. <https://doi.org/10.1016/j.cognition.2009.02.001>.
- Haesevoets, T., Hiel, A. V., Assche, J. V., Folmer, C. R., & Bostyn, D. H. (2019). An

- exploration of the motivational basis of take-some and give-some games. *Judgment and Decision Making*, 14(5), 534–546.
- Haidt, J., & Baron, J. (1996). Social roles and the moral judgement of acts and omissions. *European Journal of Social Psychology*. Retrieved from [http://onlinelibrary.wiley.com/doi/10.1002/\(SICI\)1099-0992\(199603\)26:2%3C201::AID-EJSP745%3E3.0.CO;2-J/full](http://onlinelibrary.wiley.com/doi/10.1002/(SICI)1099-0992(199603)26:2%3C201::AID-EJSP745%3E3.0.CO;2-J/full)
- Hayashi, H. (2015). Omission bias and perceived intention in children and adults. *British Journal of Developmental Psychology*, 33(2), 237–251. <https://doi.org/10.1111/bjdp.12082>.
- Henne, P., Pinillos, Á., & De Brigard, F. (2017). Cause by omission and norm: Not watering plants. *Australasian Journal of Philosophy*, 95(2), 270–283. <https://doi.org/10.1080/00048402.2016.1182567>.
- Herrmann, B., Thöni, C., & Gächter, S. (2008). Antisocial punishment across societies. *Science*, 319(5868), 1362–1367. <https://doi.org/10.1126/science.1153808>.
- Hilbe, C., Hilbe, C., Traulsen, A., Röhl, T., Röhl, T., & Milinski, M. (2014). Democratic decisions establish stable authorities that overcome the paradox of second-order punishment. *Proceedings of the National Academy of Sciences of the USA*, 111(2), 752–756. <https://doi.org/10.1073/pnas.1315273111>.
- Hitchcock, C., & Knobe, J. (2009). Cause and norm. *Journal of Philosophy*, 11, 587–612.
- Jamison, J., Yay, T., & Feldman, G. (2020). Action-inaction asymmetries in moral scenarios: Replication of the omission bias examining morality and blame with extensions linking to causality, intent, and regret. *Journal of Experimental Social Psychology*, 89, 103977. <https://doi.org/10.1016/j.jesp.2020.103977>.
- Kadish, S. H., Schulhofer, S. J., & Barkow, R. E. (2016). *Criminal law and its processes: Cases and materials*. Wolters Kluwer Law & Business.
- Knobe, J., & Fraser, B. (2008). Causal judgment and moral judgment: Two experiments. In W. Sinnott-Armstrong (Ed.), *Moral psychology* (Vol. 2, pp. 441–448). Cambridge, MA: MIT Press.
- Kominsky, J. F., Phillips, J., Gerstenberg, T., Lagnado, D., & Knobe, J. (2015). Causal superseding. *Cognition*, 137, 196–209. <https://doi.org/10.1016/j.cognition.2015.01.013>.
- Ledyard, J. (1997). Public goods: A survey of experimental research. In J. Kagel & A. E. Roth (Eds.), *The Handbook of Experimental Economics* (pp. 111–194). Princeton: Princeton University Press.
- McGrath, S. (2005). Causation by omission: A dilemma. *Philosophical Studies*, 123, 125–148. <https://doi.org/10.1007/s11098-004-5216-z>.
- Meszaros, J. R., Asch, D. A., Baron, J., Hershey, J. C., Kunreuther, H., & Schwartz-Buzaglo, J. (1996). Cognitive processes and the decisions of some parents to forego pertussis vaccination for their children. *Journal of Clinical Epidemiology*, 49(6), 697–703. [https://doi.org/10.1016/0895-4356\(96\)00007-8](https://doi.org/10.1016/0895-4356(96)00007-8).
- Moore, M. (2019). Causation in the law. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2019). Metaphysics Research Lab, Stanford University. Retrieved

- from <https://plato.stanford.edu/archives/win2019/entries/causation-law/>.
- Nikiforakis, N. (2008). Punishment and counter-punishment in public good games: Can we really govern ourselves? *Journal of Public Economics*, 92(1), 91–112.
- Ramalingam, A., Morales, A. J., & Walker, J. M. (2019). Peer punishment of acts of omission versus acts of commission in give and take social dilemmas. *Journal of Economic Behavior & Organization*, 164, 133–147. <https://doi.org/10.1016/j.jebo.2019.05.033>.
- Ritov, I., & Baron, J. (1990). Reluctance to vaccinate: Omission bias and ambiguity. *Journal of Behavioral Decision Making*, 3(4), 263–277. <https://doi.org/10.1002/bdm.3960030404>.
- Ritov, I., & Baron, J. (1999). Protected values and omission bias. *Organizational Behavior and Human Decision Processes*, 79(2), 79–94.
- Robinson, P. H., & Darley, J. M. (1995). *Justice, liability, and blame: Community views and the criminal law*. Boulder, Co.: Westview Press.
- Royzman, E. B., & Baron, J. (2002). The preference for indirect harm. *Social Justice Research*, 15(2), 165–184. <https://doi.org/10.1023/A:1019923923537>.
- Selten, R. (1967). Die Strategiemethode zur Erforschung des eingeschränkt rationalen Verhaltens im Rahmen eines Oligopolexperiments. In H. Saueremann (Ed.), *Beiträge zur experimentellen Wirtschafts-Forschung*. (pp. 136–168). Tübingen: Mohr.
- Smith, V. L., & Wilson, B. J. (2017). Sentiments, conduct, and trust in the laboratory. *Social Philosophy and Policy*, 34(1), 25–55. <https://doi.org/10.1017/S0265052517000024>.
- Spranca, M., Minsk, E., & Baron, J. (1991). Omission and commission in judgment and choice. *Journal of Experimental Social Psychology*. Retrieved from <http://www.sciencedirect.com/science/article/pii/002210319190011T>.
- Sylwester, K., Mitchell, J., Lowe, W., & Bryson, J. J. (n.d.). *Uses and consequences of costly punishment across populations*.
- Tan, F., & Xiao, E. (2012). Peer punishment with third-party approval in a social dilemma game. *Economics Letters*. Retrieved from <http://www.sciencedirect.com/science/article/pii/S0165176512004375>.
- van Dijk, E., & Wilke, H. (2000). Decision-induced focusing in social dilemmas: Give-some, keep-some, take-some, and leave-some dilemmas. *Journal of Personality and Social Psychology*, 78(1), 92–104. <https://doi.org/10.1037/0022-3514.78.1.92>.
- Wiessner, P. (2005). Norm enforcement among the Ju/'hoansi Bushmen. *Human Nature*, 16(2), 115–145. <https://doi.org/10.1007/s12110-005-1000-9>.
- Willemsen, P., & Kirfel, L. (2019). Recent empirical work on the relationship between causal judgements and norms. *Philosophy Compass*, 14(1), e12562. <https://doi.org/10.1111/phc3.12562>.
- Willemsen, P., & Reuter, K. (2016). Is there really an omission effect? *Philosophical Psychology*, 29(8), 1142–1159. <https://doi.org/10.1080/09515089.2016.1225194>.