

Diffusion of Being Pivotal and Immoral Outcomes*

Armin Falk[†] Nora Szech[‡]

December 2017

Abstract

We study how the diffusion of being pivotal affects immoral outcomes. In a first set of experiments, subjects decide about agreeing to kill mice and receiving money versus objecting to kill mice and foregoing the monetary amount. In a baseline condition, subjects decide individually about the life of one mouse. In the main treatment, subjects are organized into groups of eight and decide simultaneously. Eight mice are killed if at least one subject supports the killing. The fraction of subjects agreeing to kill is significantly higher in the main condition compared to the baseline condition. In the second set of experiments, we run the same baseline and main conditions but use a charity context and additionally study sequential decisions. We replicate our main finding from the mouse paradigm and additionally show that in the sequential treatment, prosocial behavior is even less pronounced. We further show that the observed effects increase with experience, i.e., when we repeat the experiment for a second time. Finally, we report evidence on beliefs, elicited in our main experiments but also from a treatment of noninvolved observers, and show that beliefs about being pivotal are a main driver of our results.

Keywords: Committees, diffusion of being pivotal, group decisions, morality, replacement logic

JEL Codes: C91, D01, D03, D23, D63

*We thank K. Albrecht, S. Altmann, R. Bénabou, T. Dohmen, J. Engel, D. Engelmann, S. Gächter, P. Heidhues, D. Huffman, S. Jäger, F. Kosse, B. Kőszegi, F. Krämer, G. Loewenstein, T. Neuber, F. Rosar, J. Sobel, N. Schweizer, F. Zimmermann and participants at various seminars for helpful comments. We thank M. Antony, T. Graeber, T. Wölk, and, in particular, T. Neuber, for excellent support. Falk acknowledges financial support by the German Science Foundation (DFG) through the Leibniz Program. The study was approved by Ethical Committee of the University of Bonn, Reference Number: 066/12.

[†]briq and University of Bonn, armin.falk@uni-bonn.de

[‡]Chair of Political Economy, Institute of Economics, Karlsruhe Institute of Technology, nora.szech@kit.edu

1 Introduction

This paper studies how groups favor moral transgression in diffusing responsibility and notions of being pivotal. Intuitively, acting in groups provides an excuse for acting immorally simply because an individual may perceive himself as not or only partly responsible for promoting a particular outcome. A striking example is the practice of firing squads, which typically consist of a group of executors rather than a single person. From an individual member’s perspective, being pivotal is diffused, as many people shooting at the same time implies that the killing is likely to happen, regardless of whether a particular member does or does not fire his gun. Moreover, members of firing squads are often randomly issued a gun containing a blank cartridge. This additionally diffuses being pivotal: Even if a member of the squad shoots his gun, he remains uncertain whether he can effectively cause the killing at all. Apparently, these features reduce feelings of responsibility and hence facilitate participating in executions.

To investigate the consequences of group settings which diffuse being pivotal, we ran two sets of experiments varying the choice environment and contrast environments where subjects are fully pivotal with contexts where being pivotal is exogenously diffused. In the latter, subjects are organized into groups and individual decisions are aggregated such that the individual can easily believe that his decision is unlikely to be pivotal. Organizing people into groups and implementing a decision rule that does not require the support of all members for immoral action enables a simple “replacement logic” (Sobel 2010). It allows each single actor to believe that even if he does not agree to engage in a morally questionable activity, others will, arguing that the immoral outcome happens anyway. This diffusion of being pivotal is pervasive at various levels of social interaction, such as firms, organizations, and markets.

We study two different choice paradigms. In the first set of experiments, the paradigm involves the trade-off between life and money. Subjects decide between receiving money and agreeing to kill mice versus not receiving money and objecting to the killing.¹ Importantly, mice used in the experiment are so-called “surplus” mice, which would have all

1. The study was approved by the ethics committee of the University of Bonn.

been killed without our intervention (see Section 2). Subjects learn about this default in a post-experimental debriefing. The paradigm is informed by the widely-held view that harming others in an unjustified and intentional way is considered immoral. While there exists no universal consensus about how to define the content of morality, avoiding and preventing harm is a central element according to most notions of morality.² We contrast two treatments: Baseline implements a simple binary choice where subjects either receive €0 for saving a mouse (Option A) or €10 for killing the mouse (Option B). In Baseline, subjects are hence fully pivotal. This condition serves as a comparison benchmark for the main treatment (Simultaneous). In the latter, eight subjects simultaneously decide between Option A and Option B. As in Baseline, a subject receives no money for choosing Option A and €10 for choosing Option B, irrespective of the other subjects' choices. If at least one subject chooses Option B, however, eight mice are killed. Thus, if a subject believes that at least one other subject is likely to choose Option B, he may no longer consider himself pivotal. From a utilitarian perspective, a low chance of being pivotal provides an excuse to choose Option B, as choosing Option B is unlikely to change the outcome but guarantees a payoff of €10. In line with this argument, we find that the fraction of subjects choosing Option B is significantly higher in Simultaneous than in Baseline, despite the fact that – upon being pivotal – killing causes the death of eight mice rather than one mouse. Moreover, the likelihood that a subject chooses to kill mice is decreasing in his belief of being pivotal. At the aggregate level, all mice are killed in Simultaneous.

Our second choice paradigm involves the binary decision to either keep money for oneself or to donate a higher amount to a charity. In particular, subjects choose between receiving €10 or donating €15 to a charity that supports children suffering from cancer. The purpose of this second set of experiments is three-fold. First, we study whether our main finding from the first experiment replicates using a different choice paradigm. Second, we additionally study a dynamic setting of diffusion of responsibility, i.e., a

2. See, e.g., Gert (2012, Section 1) on “The Definition of Morality,” The Stanford Encyclopedia of Philosophy: “In this descriptive sense, although avoiding and preventing harm is common to all, ‘morality’ can refer to codes of conduct of different societies with widely differing content, and still be used unambiguously.”

context where decisions are not made simultaneously but sequentially. The latter is particularly relevant in light of many real-world applications where people operate in a line of action, observing previous choices but being uncertain how individuals down the line will decide. Finally, we investigate experience effects, i.e., whether the observed effects get larger if subjects repeat the same experiment one more time.

We ran three treatment conditions. Besides using a different choice paradigm, these conditions are as similar as possible to the mouse conditions (including wording of the instructions, framing of decisions in terms of Option A vs. Option B, etc.). In *BaselineC*, (C for charity) subjects either donate €15 (Option A) or keep €10 for themselves (Option B). In this condition, subjects are fully pivotal. In *SimultaneousC*, a group of eight subjects is endowed with €120 ($8 \times €15$, analogously to the mouse condition). As in *Simultaneous*, subjects simultaneously decide to either donate or to keep the money. Choosing Option A implies no additional earnings while Option B guarantees €10, irrespective of the decisions of others. If, however, at least one of the group members chooses Option B, the whole donation is “destroyed,” i.e., subjects are pivotal only if all other group members choose to donate. In *SequentialC*, the decision context is exactly as in *SimultaneousC*, but subjects choose sequentially. The position is randomly determined such that one subject moves first, another second and so on, up to position 8. Every subject is informed about the previous choice history, i.e., he learns his position and how many subjects have previously chosen Option A or Option B, respectively. Unless Option B has already been chosen, this set-up creates diffusion of responsibility for a subject considering to choose Option A because it remains unclear what outcome will materialize, given that there is a line of others who might still choose Option B. This treatment delivers additional insights. For example, an individual’s choice may not only depend on expectations about behavior down the line but could also be affected by social learning depending on previous choices. All three treatments using the charity paradigm involve an identical second round, which comes at a surprise to subjects. The rationale for repetition is that, in contrast to *BaselineC*, many subjects in the first round of the group treatments may hold mistaken beliefs about their chance of being pivotal. Learn-

ing the outcome of the first round was therefore expected to affect behavior in the group treatments but not in BaselineC. Repeating choices provides useful additional information about the role of groups and about the implied diffusion of responsibility relative to choice contexts where individuals are fully pivotal.

The charity experiment replicates the main effect from the mouse conditions. The share of subjects choosing the selfish Option B is significantly higher in both group conditions than in the baseline treatment. Moreover, choosing a second time in BaselineC does on average not affect the likelihood of donations but – as expected – induces more selfish choices in both SimultaneousC and SequentialC. In the latter, we additionally find that previous history matters for behavior. In particular, learning that Option B has already been chosen basically eradicates the choice of Option A further down the line. In this respect it is intuitive that the choice of the first group member (position 1) is greatly affecting group behavior and explains why subjects in this role display a similar likelihood of choosing Option A as in BaselineC.

Perceptions of being pivotal are central to the mechanism under study and hinge critically on beliefs about the behavior of others. This is why, in both experiments, we elicit beliefs and confirm that choices are strongly associated with the perceived likelihood of being pivotal. A potential concern in eliciting beliefs of active players, however, is that stated beliefs may be used to “justify” behavior (Epley and Gilovich 2016; Gino, Norton, and Weber 2016). This may pollute findings even when beliefs are incentivized, as is the case in our experiment. Given the critical role of beliefs, we therefore ran a further experiment with noninvolved subjects. In this condition subjects read the instructions of all three treatments implemented with the Charity paradigm and are asked to predict the results from those experiments. They are paid for accuracy. These independently elicited beliefs of spectators corroborate our above-mentioned findings. In particular, we find that beliefs of spectators are very similar to those of subjects actually making a decision.

Taken together, the results from both the mouse and the charity paradigm show that an exogenously imposed diffusion of being pivotal facilitates moral transgression. In

particular, they show the power of organizational design to causally promote immoral behavior and outcomes. Our findings thus contribute to the understanding of the sources of malleability of moral behavior and of why “ordinary” people endowed with given moral values may engage in activities they would generally object to.

Our paper is related to work on contextual factors affecting fair outcomes in the context of simple dictator, bargaining, or allocation games. While we focus on the role of beliefs about being pivotal, other mechanisms that have been identified to favor “unfair” outcomes are delegation or exploiting moral “wriggle rooms,” as discussed, e.g., in Bartling and Fischbacher (2012), Hamman, Loewenstein, and Weber (2010), and Dana, Weber, and Kuang (2007).³ Falk and Szech (2013) analyze malleability of moral outcomes in bilateral and multilateral market situations and Falk (2017) studies the role of status inequality. Another related literature in social psychology concerns the so-called bystander effect (see, e.g., Latané and Darley 1968, and for a recent overview Fischer et al. 2011). Typical bystander experiments study helping behavior in response to a staged emergency (e.g., the experimenter becomes injured). What sets our simultaneous treatments apart is that even if a subject opts for the moral outcome, he remains uncertain about whether the moral outcome is implemented or not, similar, e.g., to firing squads. In bystander experiments on the other hand this uncertainty does not exist. If a subject opts for helping, the person in need receives help. Furthermore, in a typical bystander experiment, while deliberating about helping or not, subjects observe that others do not help either. In our simultaneous-move set-up, this type of social learning is ruled out. When deciding to kill a mouse or not to donate, respectively, subjects do not know whether other subjects opt for the selfish option as well. The dynamic properties of observing others, however, are explicitly studied in our sequential treatment. In addition, in a bystander experiment, participants need to realize that their help is required (and that it is better to step in than to hope that some other, say, more able helper will step in), while in our set-up, consequences of decisions are straightforward. Note also that in our experiment, consequences are real, incentives are exactly specified, and the mechanism (beliefs about

3. On the effects of institutions on values, see also Bowles (1998). On the role of authority, see Milgram (2009) [1974].

being pivotal) is explicitly measured.

The remainder of the paper is organized as follows. Section 2 describes the research paradigms, treatments, procedural details, and hypotheses. Results are presented in Section 3, and Section 4 concludes.

2 Design and Hypotheses

In this section, we first discuss and motivate the mouse experiment, including the paradigm, treatments, and procedural details. We then present the charity experiment, which was run to check replicability of our findings from the mouse experiment, to study sequential in addition to simultaneous decision-making, and to investigate the role of experience. Next, we describe the experiment on uninvolved subjects, run to elicit unbiased beliefs about being pivotal. Finally, we state our hypotheses.

2.1 Mouse experiment

Avoiding and preventing unjustified harm is central to most notions of morality. It is this notion that informs our “mouse paradigm,” which involves the trade-off between killing a mouse and receiving money versus saving a mouse life and receiving no money (Falk and Szech 2013).⁴ Subjects are explicitly informed that each mouse is a young and healthy mouse, which will live for about two years if saved. For illustrative purposes, we present subjects the picture of a mouse on an instruction screen. We guarantee subjects that mice, if saved, live in an appropriate, enriched environment, jointly with a few other mice. Hence, in case subjects decided to save mice, these mice were kept alive in an enriched environment, with good feed and comfortable nesting material, precisely as stated in the instructions.

4. Deckers et al. (2016) provide convergent and discriminatory validity of the mouse paradigm as a measure for morality. Killing is negatively related to agreeableness, one of the Big Five facets, which describes a tendency to be compassionate and cooperative rather than suspicious and antagonistic towards others, and positively related to Machiavellianism, measuring a person’s tendency to be unemotional, and detached from conventional morality. Moreover, killing is not related to disposable income, whether students are professionally involved with animal research or animal experiments or have a simple preference for animals, as expressed by having a pet at home.

Subjects are also informed in detail about the killing process. In the instructions, they read the following passage: “The mouse is gassed. The gas flows slowly into the hermetically sealed cage. The gas leads to breathing arrest. At the point at which the mouse is not visibly breathing anymore, it remains in the cage for another 10 minutes. It will then be removed.” To further rule out uncertainty about the decision context, subjects are shown a short demonstration video of the killing process. In the video, four mice first move vividly in the cage, then they successively slow down as more and more gas enters the cage. Eventually they die, with their hearts beating visibly heavy and slow.

It is important to stress that the mice used in the experiment were so-called “surplus” mice: These mice were bred for animal experiments, but turned out to be unsuited for scientific research. They were perfectly healthy, but keeping them alive would have been costly. It is common practice in laboratories conducting animal experiments to gas such mice. Thus, as a consequence of our experiment, many mice that would otherwise all have died were saved. Subjects were informed about this default in a post-experimental debriefing.⁵

Mouse treatments. We study the role of diffusion of being pivotal in contrasting two decision environments, one where subjects are fully pivotal (Baseline) and one where being pivotal is diffused by organizing subjects into groups (Simultaneous). The two decision contexts differ in how likely it is that any given subject is pivotal, keeping overall moral and financial consequences identical. In Baseline, each subject decides about the life of one mouse. Subjects face a simple binary choice between Option A and Option B. Option A implies that the mouse will survive and that the subject receives no money. Option B implies the killing of the mouse and receiving €10. The Baseline treatment informs us about the share of subjects who are willing to kill the mouse for €10 when obviously being pivotal.

5. While perceptions of the situation may have changed due to this information, consequences were exactly the same and as stated in the instructions. In future research, it would be interesting to explore whether using an alternative framing would affect decisions in response to institutional changes differently (compare evidence on the so-called omission-commission bias, e.g., in Spranca, Minsk, and Baron 1991).

In Simultaneous, subjects decide in groups of eight and are endowed with eight mice. As in Baseline, each subject faces an individual binary choice between Option A and Option B. Option A implies that a subject receives no money. If a subject chooses Option B, he receives €10. Individual monetary consequences are independent of other subjects' decisions. All subjects choose simultaneously. They know that if at least one subject chooses Option B, all eight mice are killed. Furthermore, they know that they will not receive feedback on whether the mice are finally killed or not (though it is obvious for a subject that the mice die if he chooses Option B). Note that we chose to endow a group with eight mice to keep the number of mice at the aggregate level identical to Baseline. Of course we do not know whether the valuation of mice lives is proportional to the number of saved mice, but keeping numbers identical at the aggregate level allows for a clean comparison of the overall impact of group vs. individual decision making. From an individual perspective, however, we could also have endowed groups with only one mouse. In this case, an individual considering himself as being pivotal would have faced the exact same consequences in Baseline and Simultaneous. Under the plausible assumption that subjects value eight mice lives at least as much as one mouse life, our treatment comparison is therefore conservative, in the sense of biasing against finding treatment effects.

In Simultaneous, right after subjects have made their decision, we elicit beliefs about being pivotal. Subjects are asked to estimate how many other subjects in their group have chosen Option B. They can enter any number from 0 to 7 and are paid €1 for a correct estimate (*belief_B*). We also ask subjects to indicate the probability that all other seven group members have chosen Option A. Subjects are asked to enter an integer percentage number, i.e., higher percentages indicate a higher perceived likelihood of being pivotal (*belief_pivotal*). Both types of beliefs (*belief_B* and *belief_pivotal*) are significantly correlated (Spearman rank correlation: -0.63 , $p < 0.001$).

Mouse procedures. 252 subjects, mainly undergraduate university students from all majors, took part in the experiment, 124 subjects in Baseline and 128 in Simultaneous. Each subject participated only in one treatment condition. We used z-Tree as the exper-

imental software (Fischbacher 2007). Subjects were recruited using the software ORSEE (Greiner 2004). At the beginning of an experimental session, participants received detailed information about the rules and the structure of the experiment. In all treatments, the experiment started only after all participants had answered several control questions correctly.

To reduce possible communication between subjects across sessions, the experiment was run on two consecutive days in six different rooms at the *Beethovenhalle*, the largest concert hall in Bonn. We set up six parallel, computerized labs in these rooms. Subjects received payments according to the rules of the experiment and an additional show-up fee of €20 to compensate for the remote location. In both treatments, subjects received their payments in a sealed envelope outside the room where the experiment had taken place. This way, neither other subjects nor the experimenters handing over the envelopes knew what a particular subject had earned. This procedure was explained in the instructions.

To ensure credibility, we stated right at the beginning that all statements made in the instructions were true, as is standard in economic experiments, and that all consequences of subjects' decisions would be implemented exactly as described in the instructions. We emphasized orally that the experimenters personally guarantee for the truthfulness of the instructions. Subjects were also invited to send us an email if they wanted to discuss the study.

2.2 Charity experiment

The charity treatments are essentially the same as in the mouse experiment, except that we use a different choice paradigm and study the role of experience as well as an additional sequential condition. As far as possible, we use the same design features, stake sizes (€10 for the selfish option), and wording and framing of choice options. At the beginning of the experiment, subjects are made familiar with the charity, which is devoted to support children who suffer from cancer. In particular, the charity is engaged in psychological assistance and in organizing leisure activities for children and their families, helps with follow-up care and school-related issues, and supports parents and siblings as well as

clinical research on cancer.

Charity treatments. To check replicability of our experimental results from the mouse paradigm, we study a baseline (BaselineC) and a simultaneous group condition (SimultaneousC), analogously to the mouse conditions. In BaselineC, subjects make the binary decision to either donate €15 (Option A) or to keep €10 for themselves (Option B).⁶ In SimultaneousC, subjects are in groups of eight and simultaneously choose either Option A or Option B, respectively. Choosing Option B implies receiving €10 and choosing Option A receiving no money, irrespective of the choices of other group members. A donation of €120 ($8 \times €15$) is initiated for the charity only if all group members choose Option A. If one group member or more choose(s) Option B, the donation of €120 is destroyed. To study how a dynamic setting affects diffusion of responsibility, we further run treatment SequentialC. This treatment is identical to SimultaneousC (including payments, donation, wording, etc.), except that subjects choose sequentially. It is randomly determined at which position a subject is asked to decide, one subject being first, another second, up to position 8. Prior to making the binary decision (Option A or Option B), subjects are informed about their position (1 to 8) and about the previous choice history, i.e., how many subjects have previously chosen A and how many have opted for B.

In both SimultaneousC and SequentialC, we also elicit beliefs analogously to Simultaneous in the mouse condition. Subjects are asked to estimate how many other subjects in their group have chosen Option B, with possible responses from 0 to 7 (*belief_B*). Correct answers are remunerated with €2. We also ask subjects to indicate the probability that all other seven group members have chosen Option A. Responses are given in percent using a slider, with higher percentages reflecting a higher perceived likelihood of being pivotal for the respective subject (*belief_pivotal*). Again, both types of beliefs (*belief_B* and *belief_pivotal*) are significantly correlated (Spearman rank correlation: -0.35 , $p < 0.001$ for SimultaneousC and -0.65 , $p < 0.001$ for SequentialC).⁷

6. Note that the design choice to donate €15 limits the plausibility of the argument that the €10 kept are spent on an alternative good cause.

7. Beliefs are elicited in the same way in SimultaneousC and SequentialC, but we note that in the latter, beliefs will depend on position and responses are affected by previous play, e.g., getting to know that Option B has already been chosen.

To measure potential experience effects, all three conditions include a second round, which came to subjects as a surprise.⁸ Subjects were told that they will make one more and final decision. In SimultaneousC and SequentialC, subjects learn whether at least one subject in their group has chosen Option B and thereby destroyed the donation, and that they will make the same decision in the same group of eight, as in the first round. In SequentialC, they also know that they act in the exact same order, i.e., each subject chooses at the same position as before. Payoffs and consequences are identical to the first round.

Charity procedures. 481 subjects, mainly undergraduate university students from all majors, took part in the experiments, 121 subjects in BaselineC, 120 in SimultaneousC and 240 in SequentialC (30 groups). Each subject participated only in one treatment condition. We used oTree as experimental software (Chen, Schonger, and Wickens 2016). Subjects were recruited using the software ORSEE (Greiner 2004). At the beginning of an experimental session, participants received detailed information about the rules and structure of the experiment. In all treatments, the experiment started only after all participants had answered several control questions correctly. The experiments were run at the BonnEconLab in March 2017. Subjects received a show-up fee of €10.

2.3 Belief experiment

A possible concern in interpreting beliefs is the potential endogeneity of beliefs due to motivated reasoning (Epley and Gilovich 2016; Gino, Norton, and Weber 2016). Eliciting and interpreting beliefs is notoriously difficult in this respect. To limit the problem, we incentivized beliefs about the number of other participants choosing Option B in the mouse and charity treatments, such that subjects could earn additional money for good estimates. However, to get an estimate of beliefs that is not biased in terms of justifying an action, we also ran an additional belief experiment. By comparing outcomes from

8. Of our 121 subjects who took part in BaselineC, only 79 took part in an experience condition, i.e., in a second round. For the first two sessions (with 42 subjects) we only ran one round. In the analysis, we therefore either use 121 observations (round 1) or 79 observations (round 2), respectively.

this experiment to beliefs elicited in the charity treatments, we can check whether or not beliefs in the latter are actually distorted. In the belief experiment, participants read the original instructions of BaselineC, SimultaneousC, and SequentialC (avoiding redundancies) and are then asked to answer questions on the respective treatments.

In both SimultaneousCB and SequentialCB (B for belief) we elicit the belief that a subject is in a group in which all other seven group members choose Option A (*belief_pivotal*). In addition we ask, assuming a group of eight, how many of the other seven group members choose Option B (*belief_B*). In answering the latter question for SequentialCB, subjects are asked to assume that they are the first mover.⁹

If the percentage answer (*belief_pivotal*) was correct within an interval of plus/minus five percentage points, subjects received €2. Likewise, if they estimated the correct number (*belief_B*), they received €2.¹⁰ 87 subjects participated in this condition, which was programmed with oTree (Chen, Schonger, and Wickens 2016) and run at the BonnEcon-Lab in March 2017.

2.4 Hypotheses

We expect that beliefs about the likelihood of being pivotal are crucial for morally relevant behavior. If the perceived likelihood of being pivotal is small enough, subjects will find it legitimate to opt for the morally problematic Option B. This replacement logic (Sobel 2010) can thus lead to a higher share of subjects opting to kill in the group treatments compared to Baseline, in which subjects know that they are pivotal for sure.

To fix ideas, suppose subjects value receiving €10 as $u > 0$ and attach zero utility to receiving €0. There is a moral cost of choosing Option B, $c(n) > 0$, where n indicates the group size, and thus, in terms of consequences, either the killing of $n = 1$ vs. $n = 8$ mice or destroying a donation of either $1 \times €15$ vs. $8 \times €15 = €120$, respectively. The expected disutility of killing n mice or destroying a donation of $n \times €15$ depends on the

9. This choice was made because uncertainty about all other seven group members' behavior only exists for the subject who moves first.

10. We also asked a set of position-dependent questions which we are not discussing further. Maximum earnings were €24 (twelve questions which were incentivized with €2 and one hypothetical question: "what would you choose...") plus a show-up fee of €5.

perceived likelihood of being pivotal, $p(n)$, where $p(1) = 1$ while $p(8) \in [0, 1]$. Assuming additivity of the arguments, it immediately follows that subjects choose Option B if $u - p(n)c(n) > 0$ and Option A otherwise.

We first note that the believed likelihood of being pivotal, $p(n)$, plays a crucial role for willingness to opt for the morally problematic Option B in the group contexts. In particular, Option B gives a higher utility under $n = 8$ than under $n = 1$ if and only if $p(8) < c(1)/c(8)$. This implies the following dichotomy result.

Belief-dependent dichotomy. *Assume that $c(1) < c(8)$. Then, the following dichotomy holds.*

(i) *For $p(8)$ small enough, opting for the morally problematic Option B becomes more attractive in the group contexts than in Baseline.*

(ii) *For $p(8)$ large enough, opting for the morally problematic Option B becomes less attractive in the group contexts than in Baseline.*

From this, we deduce our first Hypothesis.

Hypothesis 1 (Dichotomy). *In the group contexts, subjects with a high belief of being pivotal will be less likely to opt for the immoral Option B than subjects with a low belief of being pivotal.*

Further, if moral costs of agreeing to kill one versus eight mice are not that different,¹¹ opting for B should become easier for subjects in the group contexts than in Baseline, for a large range of beliefs of being pivotal, $p(8) \in [0, a]$, $a < 1$.

Hypothesis 2 (Baseline vs. group contexts). *More subjects will opt for B in the group contexts than in Baseline.*

Thus, in terms of treatment effects, we expect that there will be a substantial number of subjects for whom $u - p(8)c(8) > 0 > u - c(1)$, i.e., who choose Option A for $n = 1$ but prefer Option B for $n = 8$, reflecting a decrease in the perception of being pivotal

11. This seems rather plausible, compare Falk and Szech (2015) for individual willingness to kill one versus two versus three mice, respectively.

that exceeds the increase in moral consequences. Besides the notion of being pivotal, the relative valuation of killing eight mice vs. one mouse (or destroying a larger rather than a smaller donation, respectively) determines treatment differences. This implies that as long as $c(8) > c(1)$, we underestimate the role of being less pivotal in groups relative to Baseline in the following sense. We could have endowed groups only with one mouse or a donation of €15, respectively. In that case, we would expect even larger treatment effects. We opted for eight mice (€120, respectively) in order to keep the maximum possible extent of harm fixed at the aggregate level when comparing treatments.

In sum, smaller values of $p(8)$ imply stronger incentives for immoral behavior in groups, relative to the individual choice condition (Baseline). In addition, individual heterogeneity in $p(8)$ should translate into respective propensities to choose Option A or B. Hence, we expect that, on average, Option B is chosen more often in the group conditions than in Baseline, and that, at the individual level, the likelihood of choosing Option B is inversely related to perceptions of being pivotal. For SequentialC, this also implies that conditional on learning that Option B has already been chosen by another group member, $p(n)$ is zero, rendering Option B optimal. With respect to the repetition of SimultaneousC (and SequentialC), we further expect that if subjects overestimate being pivotal in the first round, they will update $p(n)$ downwards, yielding an increase in the share of subjects choosing Option B in the second round.¹² Since there is no such updating happening in BaselineC, we predict a similar share of subjects choosing Option B in rounds one and two.

The above arguments assume utilitarian reasoning. If, instead, a subject follows a deontological moral principle (e.g., the Kantian Categorical Imperative), he will stick to the morally preferred option, regardless of whether being pivotal is diffused or not. In other words, some subjects may not respond to treatment differences, opting for Option A even when $p(n)$ is perceived as low. The empirical relevance of these two moral conceptions, which have been the main combatants in occidental moral philosophy for the last

12. We thus assume here that equilibrium does not emerge directly, i.e., subjects may hold incorrect, heterogeneous beliefs, specifically in the first round of group decisions. For an equilibrium analysis of group decisions in morally relevant contexts, see Rothenhäusler, Schweizer, and Szech (2017).

centuries, has been empirically studied using the so-called trolley problem put forward by Philippa Foot (see also, e.g., Greene et al. 2004; Thomson 1976).¹³ Evidence from this literature suggests that both moral approaches are empirically relevant and that the extent to which people follow the one or the other largely depends on situational and emotional factors (e.g., framing the trolley problem as footbridge problem, which assigns actors a more active role, leads to more rule based behaviors). Yet in contrast to the Trolley evidence, which is using hypothetical outcomes, subjects in our experiments face real consequences. Moreover, they weigh a selfish benefit against a morally problematic outcome for third parties.

3 Results

We start in reporting our findings from the mouse experiment, before showing the results from the charity and the belief experiment.

3.1 Mouse results

Our main result from the mouse experiment is shown in Figure 1, where we compare the shares of subjects choosing to kill in Baseline and Simultaneous, respectively. In Baseline, 46.0 percent of subjects choose Option B. In Simultaneous the respective share is 58.6 percent, a difference of about 27 percent. This difference is significant ($p < 0.05$, two-sample test of proportions, two-sided). At the aggregate level, the group impact is striking. While in Baseline, 46 percent of mice are killed, *all* mice are killed in *all* groups in Simultaneous. It trivially follows that on top of being more pronounced, killing is also more “efficient” in Simultaneous: The average amount of money needed to kill a mouse is €10 in Baseline but only €5.86 in Simultaneous.¹⁴

Our results show that simple organizational rules can have a big effect on moral

13. The quandary to be resolved in this problem is to either follow the deontologically warranted option (and not to throw a switch that will divert a trolley and kill one person) or the option preferred from a consequentialist perspective (killing the person to save five others).

14. In other words, in Simultaneous it is possible to kill eight mice paying less than €80 (which would be needed in Baseline) since not all members choose to kill and yet, in the end, all mice are in fact killed.

outcomes. Put differently, an organization pursuing the goal of promoting socially responsible outcomes, should avoid delegating responsibility to groups where the support of only the “worst” member(s) is sufficient for selecting the immoral choice. Instead, it should attribute responsibility individually to its members.




Figure 1: Share of subjects choosing Option B in Baseline and Simultaneous. Error bars show 95% confidence intervals from an OLS estimation using heteroscedasticity-consistent standard errors.

We have argued above that individual perceptions of being pivotal are critical in driving the increase in selfish behavior in Simultaneous. Accordingly, we should observe that an individual’s willingness to choose Option B decreases in his belief of being pivotal. This is what we find. Recall that we asked subjects about the probability that all other group members had chosen Option A (*belief_pivotal*). Figure 2 displays the fraction of subjects choosing Option B depending on this belief. The four categories in Figure 2 are based on quartiles of the belief distribution with respective percentage values of 0–3.5, 3.5–10, 10–35, and 35–100. The figure shows a clear negative relation between subjective perceptions of being pivotal and the likelihood of choosing Option B (Spearman rank correlation: -0.54 , $p < 0.001$).




Figure 2: Share of subjects in Simultaneous choosing Option B depending on their belief of being pivotal. Error bars show 95% confidence intervals from an OLS estimation using heteroscedasticity-consistent standard errors.

OLS and probit regression results confirm the relationship between the belief of being pivotal and the choice of Option B (see columns 1 and 3 of Table 1). The respective coefficients are negative and significant. Assuming a linear relationship (column 1), the coefficient implies that a ten percentage point increase in the perception of being pivotal decreases the likelihood of killing mice by 8.4 percentage points. The importance of subjective perceptions of being pivotal can also be inferred from results reported in columns 2 and 4. Here we use subjects' responses to the question how many other subjects they think have chosen Option B (*belief_B*). We construct a dummy which takes value 1 if a subject stated the point belief that no other subject had chosen Option B, which would have rendered the subject's decision pivotal. 17 subjects (13.3 percent) stated this belief. Among these 17 subjects, only one subject chose Option B. The negative and significant coefficient of the dummy in column 2 implies that, relative to subjects with other point

beliefs, the likelihood of choosing Option B is about 60.8 percentage points lower for subjects who believe that they are pivotal.

	<i>Dependent variable: Option B</i>			
	OLS		Probit	
	(1)	(2)	(3)	(4)
belief_pivotal	-0.00840*** (0.00112)		-0.0252*** (0.00532)	
belief_B = 0		-0.608*** (0.0731)		-1.995*** (0.504)
Constant	0.791*** (0.0463)	0.667*** (0.0451)	0.823*** (0.155)	0.431*** (0.124)
Observations	128	128	128	128
R^2	0.273	0.175		
AIC	145.1	161.3	140.0	152.9

Table 1: Coefficient estimates, with binary choice option (Option B: kill mice vs. Option A: save mice) as dependent variable and heteroscedasticity-consistent standard errors in parentheses. Data come from the Simultaneous treatment. *belief_pivotal* is the belief that all other group members have chosen Option A (in percent). *belief_B = 0* is a dummy based on *belief_B* which takes value 1 for the point belief that all other subjects have chosen Option A. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

As briefly discussed above, notions of being pivotal provide an excuse only from an outcome-based or utilitarian perspective. In this respect, it is noteworthy that in Simultaneous, about 18 percent of subjects who held the belief that the chance of being pivotal is exactly zero chose Option A. From a utilitarian moral perspective, these subjects had no reason not to choose Option B and cash in €10. Possibly, these subjects have followed a deontological moral principle, sticking to their morally preferred option regardless of the outcome. This suggests that, in line with survey-based evidence (trolley problem), there is a co-existence of utilitarian and deontological moral conceptions, also in an incentivized choice task. The fraction of Kantian subjects, however, appears to be low.

3.2 Charity results

We now turn to our results based on the charity paradigm. The main findings are summarized in Figure 3, which displays the share of subjects choosing Option B (not to donate) in conditions BaselineC, SimultaneousC, and SequentialC, respectively. The dark bars show results from the first round, the light bars those of the second round (which was unexpected for subjects). Several observations can be made. First, we replicate the main result from the mouse experiment using a different choice paradigm. The share of subjects choosing Option B is significantly higher in SimultaneousC than in BaselineC with means of 58.3 percent and 39.7 percent, respectively ($p < 0.01$, two-sample test of proportions, two-sided). The increase in selfish behavior amounts to 47.0 percent, which is higher than the respective increase in the mouse condition.

Second, we find that, on average, selfish behavior is also more pronounced in the group setting where subjects choose sequentially rather than simultaneously. The overall share of participants choosing Option B in SequentialC is 72.1 percent, an increase of 81.7 percent relative to BaselineC. The difference between the two treatments is statistically significant ($p < 0.01$, two-sample test of proportions, two-sided).¹⁵ Hence, regardless of whether choosing simultaneously or sequentially, groups diffuse being pivotal and favor selfish behavior. At the aggregate level, no single group in SimultaneousC effectively donated and only two out of the 30 groups in SequentialC did not destroy the donation of €120. In Table 4 in Appendix C, we show regression results confirming these treatment differences controlling for various personal characteristics such as gender, cognitive skills (math grade) and personality (Big Five).

Third, the detrimental effects of group decision making on prosocial outcomes seem to increase with experience. Comparing results between periods one and two reveals an increase in the likelihood of immoral choices upon learning the previous outcome. The increases in SimultaneousC and SequentialC amount to 12.5 and 14.2 percentage points, respectively. These increases are statistically significant ($p = 0.033$ and $p = 0.058$, see

15. In running this test, for SequentialC we use the means of the 30 independent groups (of eight subjects) as observations. The difference between SimultaneousC and SequentialC is not statistically significant ($p = 0.17$, two-sample test of proportions, two-sided).




Figure 3: Share of subjects choosing option B in BaselineC, SimultaneousC, and SequentialC, per round. Error bars show 95% confidence intervals from an OLS estimation, where standard errors are clustered on the group level for the second round of SimultaneousC and for both rounds of SequentialC (395 clusters in total: 121 (BaselineC, round 1) + 79 (BaselineC, round 2) + 120 (SimultaneousC, round 1) + 15 (SimultaneousC, round 2) + 30 (SequentialC, round 1) + 30 (SequentialC, round 2) = 395).

OLS estimates shown in column 1 of Table 5 in Appendix C). In sharp contrast, moral behavior is not vulnerable to repetition in BaselineC, with an increase of Option B below one percentage point.¹⁶

Analogously to the mouse experiment, we find that the association between the belief about being pivotal and choosing Option B is negative and statistically significant for both SimultaneousC and SequentialC. This relationship is shown for both treatments in Figure 4 (panels A and B), where we display the share of subjects choosing Option B depending on *belief_pivotal* and in Table 2, which is constructed analogously to Table

16. The latter finding suggests that, on average, subjects neither display notions of moral licensing or moral cleansing nor of conscience accounting (see, e.g., Engel and Szech 2017; Gneezy, Imas, and Madarász 2014; Monin and Miller 2001). We cannot rule out these effects, however. In fact, despite no difference between rounds one and two on average, 12.7 percent of subjects switch from A to B and 12.7 percent switch from B to A.

1. When we regress the likelihood of choosing Option B on the belief of being pivotal as shown in Table 1 (columns 1 and 3), we get negative and significant coefficients of -0.006 and -0.009 , respectively. Hence, an increase in the perception of being pivotal of ten percentage points reduces the likelihood of not donating by six percentage points in SimultaneousC and by nine percentage points in SequentialC, similar to the effects observed in the mouse condition. Columns 2 and 4 show this effect for *belief_B*. In SimultaneousC, ten subjects (eight percent) stated the belief that no other subject had chosen Option B. Thus, they assumed to be pivotal. Among these, only one subject chose Option B. When we regress the choice of Option B in SimultaneousC on a dummy that takes value 1 for those who state this belief, the respective coefficient implies that relative to subjects with other point beliefs, the likelihood of choosing Option B is 53 percentage points lower (column 2). We report a similar finding for SequentialC in column 4, with an effect size of 77 percentage points. Finally, note that even among those in SimultaneousC who believe they are not pivotal (estimated likelihood of being pivotal of zero percent), 18 percent (three out of 17) of subjects choose Option A, presumably reflecting a Kantian kind of moral reasoning. Similarly, in SequentialC, of the 153 individuals for whom the group donation was already destroyed before, eight subjects (5.2 percent) nevertheless choose Option A and thus appear to follow a deontological moral rule.

As in BaselineC, we see switching in choices in both directions also in SimultaneousC and SequentialC. There is, however, a higher likelihood of switching from Option A to Option B: 11.7 percent in SimultaneousC and 7.9 percent in SequentialC, respectively, change from choosing Option B to Option A. However, 24.2 and 22.1 percent, respectively, switch from Option A to B. Intuitively, the prevalence of switching to the selfish option should depend on subjects' changes in beliefs about being pivotal. When we regress the choice in round 2 on the choice in period 1 and the change in the belief about being pivotal, there is a significant effect in the expected direction. Subjects who consider themselves less pivotal in period 2 than in round 1 become more likely to choose Option B in round 2 (see Table 6 in Appendix C).

To summarize, we replicate the main finding from the mouse condition in that subjects




Figure 4: Share of subjects choosing Option B in the first round depending on the belief of being pivotal in SimultaneousC (panel A) and SequentialC (panel B). Error bars show 95% confidence intervals from an OLS estimation using heteroscedasticity-consistent standard errors, which for SequentialC are clustered at the group level.

are less likely to choose the morally desired action in SimultaneousC than in BaselineC, and we observe a similar effect when subjects are choosing sequentially rather than simultaneously. In addition, we document that selfish outcomes in groups tend to increase with experience in contrast to individual decisions. Beliefs about being pivotal seem to be critical, both in a simultaneous and a sequential choice context, again very similar to what we have seen in the mouse experiment.

We now turn to a closer inspection of the dynamics of decision making in SequentialC. In contrast to the simultaneous choice context, participants in SequentialC know the choice history up to the point where they make their decision. This means that they can condition their choice on learning whether or not they are pivotal. Moreover, they act in a chain, rendering the specific position in the chain potentially relevant.

In Table 3 we explore the role of position and choice history in a simple panel regression framework using both rounds 1 and 2. In columns 1 and 4, we regress a participant's choice of Option B on his position. Not surprisingly, the respective coefficient is positive, suggesting that as play evolves, a previous choice of Option B by another subject is

	<i>Dependent variable: Option B</i>			
	SimultaneousC		SequentialC	
	(1)	(2)	(3)	(4)
belief_pivotal	-0.00574*** (0.00134)		-0.00879*** (0.000547)	
belief_B = 0		-0.527*** (0.106)		-0.769*** (0.0469)
Constant	0.774*** (0.0571)	0.627*** (0.0465)	0.954*** (0.0183)	0.846*** (0.0337)
Observations	120	120	240	240
R^2	0.139	0.087	0.448	0.400
Adjusted R^2	0.131	0.080	0.446	0.397

Table 2: OLS regression coefficient estimates, with binary choice option (Option B: destroy donation vs. Option A: donate) as dependent variable and heteroscedasticity-consistent standard errors in parentheses (clustered on the group level for SequentialC). Data come from round 1 of the SimultaneousC and SequentialC group treatments. *belief_pivotal* is the perceived chance that all other group members choose Option A (in percent) and *belief_B = 0* is a binary variable being either 0 or taking the value 1 for the point belief that all other group members choose Option A. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

getting increasingly likely. The effect is somewhat weaker in the second round (column 4), but we note that subjects are generally more likely to choose Option B in the second round (as reflected by the higher constant). In columns 2 and 5, we regress Option B on a dummy indicating that no other group member has yet chosen Option B (“not destroyed”). Both in the first (column 2) and second round (column 5), subjects react strongly to being potentially pivotal, reflected in the negative and significant coefficients. The effect is weaker in round 2, consistent with the notion that subjects have learned that the chance of being pivotal is actually quite low.

In columns 3 and 6 we combine position and history and also include the interaction of the two. Turning to round 1 (column 3), both the coefficient for position as well as the one for the interaction are insignificant, and the coefficient indicating that Option B has not yet been chosen is basically identical in columns 2 and 3. This suggests that the position itself plays a minor role. Instead, behavior is mainly determined by notions of

being pivotal. In fact, the likelihood of choosing Option B given that at least one previous group member has chosen Option B is 94.8 percent (see constant in column 2). The fact that this is below 100 percent could either reflect lack of attention or understanding (which is unlikely given the control questions and the prominent display of previous play on the decision screen) or a non-utilitarian notion of rule based decision making. Interestingly, the likelihood of choosing Option B conditional on learning that no other group member has yet chosen Option B is only 32.2 percent (28 out of 87 subjects), which is actually less than in BaselineC. In this respect, we note that a choice history consisting only of Option A choices conveys two messages, an increased likelihood of being pivotal and a signal about the prevalence of moral types in a given group. The latter can be interpreted as social learning about existing social norms and may create some social pressure to comply with this (evolved) norm.

Turning to round 2 in column 6, we see that experience has some effect on subjects' behavior. During the first round (column 3), subjects react strongly to whether or not there does still exist a chance of being pivotal. If any chance exists, they are less likely to choose Option B but apparently do not take into account how large that chance is, given their position in the decision line. With experience (column 6), subjects also consider that, depending on their position, many subjects who decide down the line need to comply in order to render their behavior pivotal. The positive interaction term in column 6 suggests that subjects who decide later in the decision line, understand that the chance of being pivotal, if still existent, is higher and thus are more likely to choose Option A.

We conclude in specifically investigating behavior of participants choosing first, i.e., at position 1. These subjects face no choice history and therefore a similar situation as subjects in SimultaneousC. Given that their behavior is observed down the line, however, they know that choosing Option B renders all following seven group members' decisions non-pivotal, which may create a particularly strong "feeling of responsibility" or a desire to act as a prosocial role model (Gächter et al. 2012; Gächter, Nosenzo, and Sefton 2013). In fact, the share of the 30 first movers in SequentialC who choose Option B is only

	<i>Dependent variable: Option B</i>					
	Round 1			Round 2		
	(1)	(2)	(3)	(4)	(5)	(6)
Position (1–8)	0.0552*** (0.0139)		-0.0121 (0.0110)	0.0210* (0.0115)		0.00108 (0.00963)
Not destroyed		-0.626*** (0.0626)	-0.607*** (0.115)		-0.458*** (0.122)	-0.154 (0.126)
Interaction			-0.0170 (0.0278)			-0.126*** (0.0229)
Constant	0.473*** (0.0988)	0.948*** (0.0228)	1.013*** (0.0467)	0.768*** (0.0798)	0.968*** (0.0161)	0.962*** (0.0572)
Observations	240	240	240	240	240	240
R^2	0.079	0.450	0.458	0.020	0.313	0.435
Adjusted R^2	0.076	0.448	0.451	0.015	0.310	0.428

Table 3: OLS regression coefficient estimates, with binary choice option (Option B: destroy donation vs. Option A: donate) as dependent variable. Data come from the SequentialC treatment. *Position* is the position in the move order from 1–8, *Not destroyed* is a dummy that is 1 if all subjects in the respective group have chosen Option A thus far, and *Interaction* is the interaction of the two above variables. Standard errors in parentheses are clustered at the group level (30 groups). * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

43.3 percent which is lower than for SimultaneousC and not statistically different from BaselineC where all subjects are fully pivotal ($p = 0.71$, two-sample test of proportions, two-sided).¹⁷

3.3 Belief results

As pointed out above, stated beliefs in our main experiments are potentially polluted by self-serving motives. These problems should be mitigated by the fact that we incentivized the belief about the number of other players choosing B. However, given the critical role of beliefs in understanding decision making in groups, we ran an additional belief treatment with uninvolved spectators. The latter were incentivized to correctly estimate behavior

17. This holds despite the fact that first movers' perceived percentage of being pivotal is only 31.3 percent, i.e., way below to what holds in BaselineC (100 percent). Note that for the 30 subjects who decided on position 1 we also find that the belief of being pivotal and Option A are significantly correlated in the expected direction ($p < 0.001$ for *belief_pivotal* and $p < 0.001$ for *belief_B*).

in the charity experiments and had no incentive for stating self-serving beliefs.

The main purpose here is to check whether elicited beliefs are in fact biased. To this end, we compare the beliefs of active subjects with those from spectators. In addition, we compare beliefs to outcomes. Note that beliefs in SequentialC are more difficult to interpret and less informative for the question at hand, given that beliefs depend on the position and, in particular, on the choice history (e.g., rendering beliefs trivial in case Option B has been chosen before). We therefore restrict this section to our simultaneous treatments. If differences between beliefs from SimultaneousC and SimultaneousCB are small, we will argue that self-serving motives are not seriously affecting our main findings.

Figure 5 shows results concerning *belief_pivotal*, i.e., the probability that a subject is in a group with all other seven group members choosing Option A. The actual probability was zero percent, both in rounds 1 and 2. In no single group, there were more than six subjects choosing Option A. A different way to estimate the actual probability of being pivotal is to use the whole distribution of choices and to calculate the likelihood – given the probability for Option A (42 percent) – of randomly being matched with seven group members who all choose Option A, which is 0.002. This value is shown in the first bar and the analogous value of 0.0002 for round 2 in the second bar (the probability of Option A in the latter round is 29 percent). Bars 3 and 4 show subjects’ beliefs for rounds 1 and 2 (in SimultaneousC), respectively. It is obvious that subjects heavily overestimate how likely it is that they are pivotal. While the shown average beliefs hide a substantial amount of heterogeneity, almost all subjects perceive themselves as being pivotal with a higher likelihood than they actually are. Moving from round 1 to round 2, subjects adjust in the correct direction but still heavily overestimate their impact.

These observations are interesting for several reasons. First, the negative effects of groups are apparently smaller than they “should” be – if subjects were holding more accurate beliefs. In other words, upon learning how unlikely it actually is for them to be pivotal, we would expect the effects to be much larger. This intuition is in line with the increased share of subjects choosing Option B in round 2 relative to round 1. Second, overestimating one’s sense of being pivotal could point to a human tendency

to overestimate one’s impact in general. This may well extend to other (non-moral) contexts and seems worth further investigation, e.g., in voting contexts (Duffy and Tavits 2008). A possible reason for overestimating one’s impact could come from a desire for meaning, self-attribution and determination, as well as for motivating action in general.¹⁸ Third, in terms of motivated reasoning, there is no indication that in the present context subjects form self-servingly biased beliefs in an attempt to justify selfish behavior. This is confirmed by a comparison with beliefs of the spectators (fifth bar). The relevant beliefs are those from the first round, which are actually slightly higher than those of the spectators. While this difference is not statistically significant ($p = 0.59$, Mann–Whitney U test, two sided) it suggests that, if anything, active subjects tend to overestimate the likelihood of being pivotal rather than engaging in self-serving underestimation.




Figure 5: Likelihood of being pivotal, i.e. the probability that all other seven members of a given subject’s group choose Option A (in percent). Error bars show 95% confidence intervals from an OLS estimation, where standard errors are clustered for the second round.

Qualitatively, we find very similar results concerning *belief_B*, i.e., the belief about

18. For a desire for efficacy, see research on the so-called IKEA-effect (e.g., Norton, Mochon, and Ariely 2012).

how many of the other seven group members choose Option B. Figure 6 shows actual numbers (for SimultaneousC) for rounds 1 and 2 (first two bars), beliefs in rounds 1 and 2 (bars 3 and 4), as well as beliefs of spectators (fifth bar). The number of subjects choosing B increases from round 1 to 2, which is reflected in changes in the beliefs of subjects. In contrast to *belief_pivotal*, however, subjects are overall much more accurate about actual outcomes.¹⁹ Importantly, as for *belief_pivotal*, the beliefs of active subjects and spectators are not statistically significantly different (comparison of bars 3 and 5 in Figure 6; $p = 0.59$, Mann–Whitney U test, two sided).




Figure 6: Number of other group members choosing Option B (0–7). Error bars show 95% confidence intervals from an OLS estimation, where standard errors are clustered for the second round.

4 Concluding Remarks

This paper has documented the malleability of moral outcomes in response to an exogenous diffusion of being pivotal. Simple organizational changes from an individual decision

¹⁹ A possible explanation is that subjects found answering the question concerning absolute numbers easier than estimating a probability.

context to group conditions increase moral transgression on the individual and even more so on the aggregate level. Further, we have shown that low beliefs about being pivotal lead to less moral behavior. Indeed, if beliefs about being pivotal had been more realistic, the willingness to engage in selfish behavior may have been even more pronounced. In this sense, it is conceivable that repeated interactions with learning possibilities increase the likelihood of immoral outcomes even further, as we observe in the second round of our experiment using the charity paradigm.

Our findings are largely in line with utilitarian moral thinking. Subjects consistently respond to notions of being pivotal and only few subjects appear to follow a Kantian conception. In Simultaneous, 18 percent of subjects who hold the belief that the chance of being pivotal is exactly zero choose Option A. In SimultaneousC, the respective share is again 18 percent. Finally, in SequentialC, of the 153 individuals for whom the group donation was already destroyed before, eight subjects (5.2 percent) nevertheless choose Option A. These numbers suggest the existence of deontological reasoning but they are quite low.²⁰ Our findings question the relatively high fractions of Kantian types in survey data such as the Trolley problem, where consequences are hypothetical rather than real.

Replacement arguments help explaining outcomes in markets that are violating traders' own moral or fairness preferences (Sobel 2010). Here replacement prevails if traders prefer concluding a trade themselves to letting another trader perform the same transaction, even if trading creates unfair outcomes for traders themselves, or imposes negative externalities on others. In cases where buying decisions create negative externalities, a frequently made "excuse" is that "if I don't buy, another buyer will." Conversely, suppliers of potentially harmful goods might argue that market demand would be met with or without their involvement. An argument along these lines was invoked by British Secretary of State Boris Johnson in October 2016 after allegations about weapons exported to Saudi Arabia being used for war crimes in Yemen. Faced with a motion in the House of Commons to suspend sales, he retorted that the respective members of parliament should "be in no doubt that we would be vacating a space that would rapidly

20. Of course, it may also be the case that some subjects made mistakes.

be filled by other Western countries who would happily supply arms with nothing like the same compunctions or criteria or respect for humanitarian law” (Peck 2016; see also Bartling and Özdemir 2017). This is a refined version of the discussed argument in pointing at positive “side effects” associated with Britain taking an active role (see Glover and Scott-Taggart 1975, pp. 177). Yet we have shown increased moral transgression in a context without any such subtleties. Thus, the latter might often represent mere excuses rather than sound justifications.

While the focus of this paper is to highlight possible negative consequences of organizational design on moral behavior, the reverse inference is of course our main interest. Our findings suggest that organizations aiming at promoting morality should reduce diffusion of being pivotal, and instead attribute individual responsibility to their members.

References

- Bartling, Björn, and Urs Fischbacher. 2012. “Shifting the Blame: On Delegation and Responsibility.” *Review of Economic Studies* 79 (1): 67–87.
- Bartling, Björn, and Yagiz Özdemir. 2017. *The Limits to Moral Erosion in Markets: Social Norms and the Replacement Excuse*. Mimeo.
- Bowles, Samuel. 1998. “Endogenous Preferences: The Cultural Consequences of Markets and Other Economic Institutions.” *Journal of Economic Literature* 36 (1): 75–111.
- Chen, Daniel L., Martin Schonger, and Chris Wickens. 2016. “oTree—An open-source platform for laboratory, online, and field experiments.” *Journal of Behavioral and Experimental Finance* 9:88–97.
- Dana, Jason, Roberto A. Weber, and Jason Xi Kuang. 2007. “Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness.” *Economic Theory* 33 (1): 67–80.
- Deckers, Thomas, Armin Falk, Fabian Kosse, and Nora Szech. 2016. *Homo Moralis: Personal Characteristics, Institutions, and Moral Decision-Making*. Mimeo.
- Duffy, John, and Margit Tavits. 2008. “Beliefs and Voting Decisions: A Test of the Pivotal Voter Model.” *American Journal of Political Science* 52 (3): 603–618.
- Engel, Jannis, and Nora Szech. 2017. *A Little Good Is Good Enough*. WZB Discussion Paper SP II 2017-301. Berlin: WZB.
- Epley, Nicholas, and Thomas Gilovich. 2016. “The Mechanics of Motivated Reasoning.” *Journal of Economic Perspectives* 30 (3): 133–140.
- Falk, Armin. 2017. *Status Inequality, Moral Disengagement and Violence*. CESifo Working Paper 6588. Munich: CESifo.

- Falk, Armin, and Nora Szech. 2013. "Morals and Markets." *Science* 340 (6133): 707–711.
- . 2015. *Irrelevance of Numbers and Concave Moral Costs*. Mimeo.
- Fischbacher, Urs. 2007. "z-Tree: Zurich toolbox for ready-made economic experiments." *Experimental Economics* 10 (2): 171–178.
- Fischer, Peter, Joachim I. Krueger, Tobias Greitemeyer, Claudia Vogrincic, Andreas Kastenmüller, Dieter Frey, Moritz Heene, Magdalena Wicher, and Martina Kainbacher. 2011. "The bystander-effect: A meta-analytic review on bystander intervention in dangerous and non-dangerous emergencies." *Psychological Bulletin* 137 (4): 517–537.
- Gächter, Simon, Daniele Nosenzo, Elke Renner, and Martin Sefton. 2012. "Who Makes a Good Leader? Cooperativeness, Optimism and Leading-by-Example." *Economic Inquiry* 50 (4): 953–967.
- Gächter, Simon, Daniele Nosenzo, and Martin Sefton. 2013. "Peer Effects in Pro-Social Behavior." *Journal of the European Economic Association* 11 (3): 548–573.
- Gert, Bernard. 2012. "The Definition of Morality." In *Stanford Encyclopedia of Philosophy*, Fall 2012, edited by Edward N. Zalta. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/fall2012/entries/morality-definition/>.
- Gino, Francesca, Michael I. Norton, and Roberto A. Weber. 2016. "Motivated Bayesians: Feeling Moral While Acting Egoistically." *Journal of Economic Perspectives* 30 (3): 189–212.
- Glover, Jonathan, and M. Scott-Taggart. 1975. "It Makes no Difference Whether or Not I Do It." *Proceedings of the Aristotelian Society, Supplementary Volumes* 49:171–209.
- Gneezy, Uri, Alex Imas, and Kristóf Madarász. 2014. "Conscience Accounting: Emotion Dynamics and Social Behavior." *Management Science* 60 (11): 2645–2658.
- Greene, Joshua D., Leigh E. Nystrom, Andrew D. Engell, John M. Darley, and Jonathan D. Cohen. 2004. "The Neural Bases of Cognitive Conflict and Control in Moral Judgment." *Neuron* 44 (2): 389–400.
- Greiner, Ben. 2004. "An Online Recruitment System for Economic Experiments." In *Forschung und wissenschaftliches Rechnen*, edited by Kurt Kremer and Volker Macho, 63:79–93. GWDG-Bericht.
- Hamman, John R., George Loewenstein, and Roberto A. Weber. 2010. "Self-Interest through Delegation: An Additional Rationale for the Principal-Agent Relationship." *American Economic Review* 100 (4): 1826–1846.
- Latané, Bibb, and John M. Darley. 1968. "Group inhibition of bystander intervention in emergencies." *Journal of Personality and Social Psychology* 10 (3): 215–221.
- Milgram, Stanley. 2009. *Obedience to Authority: An Experimental View*. Reprint. Harper Perennial Modern Classics.
- Monin, Benoît, and Dale T. Miller. 2001. "Moral credentials and the expression of prejudice." *Journal of Personality and Social Psychology* 81 (1): 33–43.

- Norton, Michael I., Daniel Mochon, and Dan Ariely. 2012. "The IKEA effect: When labor leads to love." *Journal of Consumer Psychology* 22 (3): 453–460.
- Peck, Tom. 2016. "If we don't sell arms to Saudi Arabia, someone else will, says Boris Johnson." *The Independent* October 26. <http://www.independent.co.uk/news/uk/politics/if-we-dont-sell-arms-to-saudi-arabia-someone-else-will-says-boris-johnson-a7382126.html>.
- Rothenhäusler, Dominik, Nikolaus Schweizer, and Nora Szech. 2017. "Guilt in voting and public good games." *European Economic Review*: forthcoming.
- Sobel, Joel. 2010. *Do Markets Make People Selfish?* Mimeo.
- Spranca, Mark, Elisa Minsk, and Jonathan Baron. 1991. "Omission and commission in judgment and choice." *Journal of Experimental Social Psychology* 27 (1): 76–105.
- Thomson, Judith Jarvis. 1976. "Killing, Letting Die, and The Trolley Problem." *The Monist* 59 (2): 204–217.

Appendix A Instructions of Mouse Experiment

Instructions have been translated from German.

A.1 Baseline

Thank you very much for your participation!

For your participation you will in any case receive 20 euros. In the following you can earn an additional amount of money. At the end of the experiment you will receive your money in an envelope. Neither the other participants of the experiment nor the experimenter will be able to see how much money you have earned.

Please note: Throughout the whole experiment **communication between the participants is not allowed**. On the computer, please only use the functions intended to be used. If you have questions please raise your hand. Your question will then be answered at your cubicle!

Please note: **All statements made in these instructions are true**. This holds for all experiments carried out by the BonnEconLab, and also for this experiment. **In particular, all actions to be taken will be implemented exactly in the way they are described**. If you want to, you will be able to verify the correctness of all statements made in these instructions after the experiment.

In this experiment, there is a **Quiz A** and a **Quiz B**. Both, Quiz A and Quiz B, are simple trivia quizzes with questions from history, geography, sports, and so on. One example question could be: “Capital of Belgium?” There will, respectively, be four possible answers out of which one answer is correct. The posed questions in Quiz A and Quiz B are identical, that means, they are exactly the same regarding their difficulty. You will get three minutes to solve the quiz. The more questions you solve correctly, the more you can earn. **For each question that is answered correctly, you receive 5 cents**.

A.1.1 Details on the mouse



In this study, the life of a mouse is entrusted to your care. It is a healthy, young mouse, living with some other mice together in a small group. The expected lifetime of this mouse is approximately two years.

A.1.2 What is the difference between Quiz A and Quiz B?

Quiz A: In Quiz A, at the end of the experiment, you earn no additional money besides the **20 euros** for participation and the mouse stays alive.

Quiz B: In Quiz B, at the end of the experiment, you **get 10 euros in addition**. As another consequence, **the mouse will get killed**.

A.1.3 Details on the killing process

If you opt for the death of the mouse, the mouse is gassed. The gas flows slowly into the hermetically sealed cage. The gas leads to breathing arrest. As soon as the mouse is not visibly breathing anymore, it remains in the cage for another 10 minutes. It will then be removed.

A.1.4 Summary

In Quiz A you earn no additional money, and the mouse does not get killed. In Quiz B, you earn additionally 10 euros, and the mouse gets killed. The decision is yours. You make your decision on a decision screen that will be shown as soon as you have answered the control questions on the following screen.

A.1.5 Control questions

In case of Quiz A: How many euros do you receive in addition? ____

Will a mouse be killed? Yes No

In case of Quiz B: How many euros do you receive in addition? ____

Will a mouse be killed? Yes No

A.2 Simultaneous

Introduction as in Baseline

In this study, the life of eight mice is entrusted to your group's care. These are healthy, young mice, living with some other mice together in a small group. The expected lifetime of these mice is approximately two years.

A.2.1 What is the difference between Quiz A and Quiz B?

In the following we describe the consequences of choosing Quiz A and Quiz B. The choice options and consequences are identical for all eight group members.

Quiz A: In Quiz A, at the end of the experiment, you earn no additional money besides the **20 euros** for participation. This holds for all group members. Each group member who chooses Quiz A receives no additional money.

Quiz B: In Quiz B, at the end of the experiment, you **get 10 euros in addition**. This holds for all group members. Each group member who chooses Quiz B receives 10 euros in addition.

Another consequence is that **eight mice get killed if at least one member of your group of eight chooses Quiz B**. Thus if in total one member of the group, or two, three, four, five, six, seven or eight group members choose Quiz B, eight mice get killed. **Only if no member in your group of eight chooses Quiz B, the mice will not get killed.**

A.2.2 Details on the killing process

If your group opts for the death of the mice, these will be gassed. The gas flows slowly into the hermetically sealed cage. The gas leads to breathing arrest. As soon as the mice are not visibly breathing anymore, they remain in the cage for another 10 minutes. They will then be removed.

A.2.3 Summary

In Quiz A you earn no additional money. In Quiz B, you earn additionally 10 euros. Whether the mice get killed depends on whether at least one member of your group of eight has chosen Quiz B. You make your decision on a decision screen, which will be shown as soon as you have answered the control questions on the following screen.

Control questions and video

Appendix B Instructions of Charity Paradigm

Instructions have been translated from German.

B.1 BaselineC

Welcome and thank you very much for your interest in today's experiment!

This experiment is part of a research project of the *Bonner Laboratorium für experimentelle Wirtschaftsforschung (BonnEconLab)*.

For your participation you will in any case receive €10.00, which will be handed to you in cash today at the end of the experiment. During the experiment, you will make

decisions at the computer. Depending on how you decide, you can earn additional money.

During the experiment, it is not allowed to communicate with other participants. Also note that the curtain of your cubicle has to be shut throughout the entire experiment. Please now switch off your mobile phone, to make sure that other participants are not being disturbed. On the computer, please only use the functions intended to be used and make all inputs using either the mouse or the keyboard. If you have questions, please contact the conductor of the experiment. To do so, please stick your hand out of the cubicle.

All statements made in this experiment are true. This holds for all experiments carried out by the BonnEconLab, and also for this experiment. In particular, all actions to be taken will be implemented exactly in the way they are described. If you want to, you will be able to verify the correctness of all statements made in these instructions after the experiment.

In what follows, we will first ask you to answer a question regarding your mood. Subsequently, the decisions you will have to make will be explained in detail.

B.1.1 How is your current mood?

Please give an answer to this on the following scale from 0 to 10.

0 means that your mood is very bad.

10 means that your mood is very good.

You can choose any integer number on the scale from 0 to 10 to express your current mood.

B.1.2 The donation

This experiment is about a donation to the *Förderkreis für krebskranke Kinder und Jugendliche e.V.*, a regional charity from Bonn.

Every participant, that means also you, will first be entrusted with a donation which will be made to the *Förderkreis für krebskranke Kinder und Jugendliche e.V.* after today's experiment.

During the experiment, you will make decisions which affect this donation. Moreover, the information which follow are also relevant for your personal payoff from this experiment.

Therefore, please carefully read the following instructions. In particular, make sure that you understand all decisions you can make as well as their potential consequences.

B.1.3 Information about the *Förderkreis*

The *Förderkreis*. The *Förderkreis für krebskranke Kinder und Jugendliche e.V.* **supports young people suffering from cancer and their families comprehensively in dealing with the disease.** The society is committed to psychological support, to organizing free time activities, as well as to aftercare and to supporting children and adolescents with school. Moreover, indirectly affected individuals like parents and siblings are extensively supported. This takes, for example, the form of a specifically established home for parents and of pedagogic support. Moreover, the *Förderkreis* supports clinical research on cancer.

Projects and tasks of the *Förderkreis*.

- *Klassissimo* school project: offers participation in school lessons using Skype
- *Bärenstark*: support of families at home
- Psychosocial and psychooncological counseling of patients and relatives
- Pedagogic support at the hospital department
- Start-up financing for new positions and financing of specific training of departments' staff.
- Financing of hospital clowns and music therapy
- Aftercare
- Support of clinical research on cancer

B.1.4 Your decision

The donation. You are entrusted with a donation of €15.00, which is supposed to be made to the *Förderkreis für krebskranke Kinder und Jugendliche e. V.* following today's experiment. Whether this amount will in fact be transferred to the *Förderkreis* at the end of the experiment depends on the decisions that you will make.

Anonymity. No other participant in this experiment can see your decisions. The subsequent analysis of all data is done anonymously, such that all your decisions cannot be linked to your identity anymore.

You can choose between two options: **Option A** and **Option B**. Depending on which of both options you choose, you can earn different amounts of money. Additionally, depending on which option you choose, consequences differ for the donation of €15.00 which was described above.

In what follows, the consequences associated with choices of **Option A** and **Option B**, respectively, will be described.

Option A. If you choose **Option A**, besides €10.00 for participation you will receive no additional money at the end of the experiment.

Option B. If you choose **Option B**, you will additionally receive €10.00 at the end of the experiment.

As a further consequence, the previously described donation of €15.00 will be destroyed.

Summary. If you choose **Option A**, you do not receive an additional payment and the donation will not be destroyed. If you choose **Option B**, you additionally receive €10.00 and the donation is destroyed. The decision rests with you.

You make your decision on a decision screen, which will be shown as soon as you have answered the control questions on the following screen.

B.1.5 Control questions

For Option A. How many euros do you receive in addition? ____

Will the donation be destroyed? Yes No

For Option B. How many euros do you receive in addition? ____

Will the donation be destroyed? Yes No

B.1.6 Your decision

Please now choose between **Option A** and **Option B**.

I choose: **Option A** **Option B**

B.1.7 Result

If Option A was chosen: You have decided **not to destroy** the donation.

Therefore, a donation of €15.00 to the *Förderkreis für krebskranke Kinder und Jugendliche e.V.* will be made for you by the BonnEconLab.

If Option B was chosen: You have decided **to destroy** the donation.

Therefore, **no donation will be made**.

B.1.8 Experiment 2

Now follows a second experiment. This experiment is the last experiment. Your final payoff comprises of €10.00 for participation in the experiment, your decision in the first experiment, and, independently, on how you decide in the second experiment.

The decision in the second experiment is the same as in the first experiment. Thus, you can again choose between **Option A** and **Option B**, i.e., you can decide whether a donation will be destroyed or not. The donation is again a donation to the *Förderkreis für krebskranke Kinder und Jugendliche e.V.*

B.2 SimultaneousC

Introduction as in BaselineC

B.2.1 Your decision

Your group. You are together with 7 other participants of today's experiment in a group of 8 people. Your group members have been allotted to you at the beginning of the experiment. You will at no point learn which participant is in your group.

Note: **You are making all decisions within this experiment autonomously and independent of the other members of the group.** The consequences of your decisions can depend on decisions of other group members. On the following screens, all decisions, alternatives, and consequences will be introduced and explained in detail.

The donation. Your group is entrusted with a donation totaling €120.00, which is supposed to be made to the *Förderkreis für krebskranke Kinder und Jugendliche e.V.* following today's experiment. Whether this amount will in fact be transferred to the *Förderkreis* at the end of the experiment, depends on the decisions that you and the other members of your group will make.

Anonymity. No other participant in this experiment can see your decisions. This is also true for the other members of your group. The subsequent analysis of all data is done anonymously, such that all your decisions cannot be linked to your identity anymore.

You can choose between two options: **Option A** and **Option B**. Depending on which of both options you choose, you can earn different amounts of money. Additionally, depending on which option you choose and which options the other participants of your group choose independently, consequences differ for the donation of €120.00 which was described above.

In what follows, the consequences associated with choices of **Option A** and **Option B**, respectively, will be described. The choices and the consequences are the same for all 8 participants in your group.

Option A. If you choose **Option A**, besides €10.00 for participation you will receive **no** additional money at the end of the experiment.

This holds for all group members: Each group member who chooses **Option A** receives no additional money.

Option B. If you choose **Option B**, you will **additionally** receive €10.00 at the end of the experiment.

This holds for all group members: Each group member who chooses **Option B** additionally receives €10.00.

As a further consequence, **the previously described donation of €120.00 will be destroyed if at least one of the 8 members of your group chooses Option B.** Thus, if one group member, or if two, three, four, five, six, seven, or eight group members decide for **Option B**, the donation is destroyed. **Only if none of the 8 members of your group chooses Option B, the donation will not be destroyed.**

Summary. If you choose **Option A**, you do not receive an additional payment. If you choose **Option B**, you additionally receive €10.00. Whether the donation to the *Förderkreis* is destroyed depends on whether at least one of the 8 members of your group has chosen **Option B**.

B.2.2 Decisions of participants in your group

Note: The consequences of your choice do not just depend on you but also on the decisions of the other 7 members of your group. This holds in particular for the execution of the donation to the *Förderkreis für krebskranke Kinder und Jugendliche e.V.*: **Only if none of the members of your group has chosen Option B, the donation of €120.00 is made.**

You and the other 7 members of your group decide **simultaneously**. After all group members have made their decision, you learn whether the donation will be made.

At the end of today's experiment, you will also learn how many members of your group have in total chosen **Option A** and how many members of your group have in

total chosen **Option B**.

You make your decision on a decision screen, which will be shown as soon as you have answered the control questions on the following screen.

B.2.3 Control questions

Suppose, no/one other group member chooses / two/six other group members choose Option B.

You choose Option A: How many euros do you receive in addition? ____

Will the donation be destroyed? Yes No

You choose Option B: How many euros do you receive in addition? ____

Will the donation be destroyed? Yes No

B.2.4 Your decision

Please now choose between **Option A** and **Option B**.

I choose: **Option A** **Option B**

B.2.5 What do you estimate?

How likely is it in your opinion that all other group members have chosen **Option A**?

Please enter a probability (from 0 to 100 percent): [Slider]

What do you think, how many of the other 7 group members have chosen **Option B**? If you estimate the correct number, you will additionally receive €2. *Enter a number between 0 and 7:* ____

B.2.6 Result

If Option A was chosen: You have decided **not to destroy** the donation.

In your group, at least one participant has decided **to destroy** the donation. The donation over €120.00 from you and the other members of your group will therefore **not be made**.

You have not made a correct estimation and therefore do not receive any additional payoff.

B.2.7 Experiment 2

Now follows a second experiment. This experiment is the last experiment. Your final payoff comprises of €10.00 for participation in the experiment, your decision in the first experiment, and, independently, on how you decide in the second experiment.

The decision in the second experiment is the same as in the first experiment. Thus, you can again choose between **Option A** and **Option B**, i.e., you can decide whether a donation will be destroyed or not. The donation is again a donation to the *Förderkreis für krebskranke Kinder und Jugendliche e.V.*

Please note: You are in the same group of 8 participants as in the first experiment.

B.3 SequentialC

Introduction as in SimultaneousC

B.3.1 Decisions of participants in your group

Note: The consequences of your choice do not just depend on you but also on the decisions of the other 7 members of your group. This holds in particular for the execution of the donation to the *Förderkreis für krebskranke Kinder und Jugendliche e.V.*: **Only if none of the members of your group has chosen Option B, the donation of €120.00 is made.**

You and the other 7 members of your group decide **one after the other**. Your position is randomly determined by a computer.

When it is your turn, you will learn whether among the people who have decided before you, someone has already chosen Option B. You will also learn your position within the sequence. Moreover, you will learn how many members of your group have already chosen **Option A** and how many members of your group have already chosen **Option B**. At the end of today's experiment, you will also learn how

many members of your group have in total chosen **Option A** and how many members of your group have in total chosen **Option B**.

Please note: If another participant in your group has already decided for **Option B** before it was your task, this means that the donation has already been destroyed. Thus in this case, your decision has **no** effect any more on whether the donation is made.

Control questions as in SimultaneousC.

B.3.2 Your decision

You are on **position 1** in the order of your group. Consequently, no other member in your group has made a decision yet.

Or:

You are on **position 2** in the order of your group. Consequently, 1 group member has already made a decision.

Of the 1 group members who have decided before you, 1 has decided for **Option A** and 0 for **Option B**.

Or:

You are on **position 3** in the order of your group. Consequently, 2 group members have already made a decision.

Of the 2 group members who have decided before you, 1 has decided for **Option A** and 1 for **Option B**.

Thus, the donation has already been destroyed.

Please now choose between **Option A** and **Option B**.

I choose: **Option A** **Option B**

Remaining instructions as in SimultaneousC.

B.4 Belief experiment

In the belief experiment, participants read the original instructions (avoiding redundancies, however), learn how many subjects have taken part in the respective treatment, and

are then asked to answer the following questions.

BaselineCB:

- How likely do you think it is that a randomly chosen participant of the just described experiment decides for Option A, i.e., not to destroy the donation?

SimultaneousCB:

- How likely was it for a participant in the experiment to be in a group in which all other 7 group members choose Option A? (answer in percent)
- How likely do you think it is that the donation is not destroyed in such a group in the end, i.e., that all 8 group members choose Option A. (answer in percent)
- Please imagine you are in the new situation at the BonnEconlab which was just described. What do you think: How many of the other 7 members of your group have decided for Option B, i.e., to destroy the donation?

SequentialCB:

- How likely was it for a participant in this experiment to be in a group in which all other 7 group members choose Option A? (answer in percent)
- How likely do you think it is that the donation is not destroyed in such a group, i.e., that all 8 group members choose Option A? (answer in percent)

Please now imagine yourself in the situation of a participant in the described experiment at the BonnEconLab.

- Imagine, you decide first and choose Option A. How many of the other 7 group members do think also choose Option A, such that the donation is not destroyed? (answer in percent)
- Imagine, the member at position 1 in your group chooses Option A. You decide second and also choose Option A. How likely do you think it is that all further 6

people in the group also choose Option A, such that the donation is not destroyed?
(answer in percent)

- Imagine, the members at positions 1 to 3 in your group all choose Option A. You decide as the fourth and also choose Option A. How likely do you think it is that all further 4 people in the group also choose Option A, such that the donation is not destroyed? (answer in percent)
- You decide last, i.e., as the eighth. How likely do you think it is that all 7 before you have chosen Option A? (answer in percent)
- Please again imagine yourself in the situation of the described experiment at the BonnEconLab. You decide first. What do you think: How many of the 7 other members of your group decide for Option B, i.e., for destroying the donation?
- Now, please imagine that you decide last in your group, i.e., as the eighth. All 7 group members before you have chosen Option A. Would you then choose Option A or Option B? (unincentivized)
- How likely do you think it is that a participant in the just described situation – decided last, all group members before have chosen Option A – also has chosen Option A? (answer in percent)

Appendix C Robustness Checks

	<i>Dependent variable: Option B</i>		
	OLS (1)	Probit (2)	Logit (3)
SimultaneousC	0.171*** (0.0643)	0.438*** (0.166)	0.701*** (0.268)
SequentialC	0.317*** (0.0723)	0.841*** (0.206)	1.362*** (0.340)
Female	-0.0583 (0.0451)	-0.164 (0.125)	-0.267 (0.209)
Math grade	-0.0102 (0.00690)	-0.0287 (0.0194)	-0.0467 (0.0316)
Big 5 – openness	-0.0140** (0.00653)	-0.0396** (0.0180)	-0.0641** (0.0300)
Big 5 – conscientiousness	0.00729 (0.00725)	0.0204 (0.0205)	0.0333 (0.0333)
Big 5 – extraversion	0.00475 (0.00543)	0.0135 (0.0149)	0.0218 (0.0248)
Big 5 – agreeableness	-0.00399 (0.00788)	-0.0112 (0.0215)	-0.0178 (0.0356)
Big 5 – neuroticism	0.00433 (0.00608)	0.0118 (0.0165)	0.0198 (0.0275)
Constant	0.580** (0.242)	0.256 (0.661)	0.404 (1.095)
Observations	481	481	481
R^2	0.091		
AIC	650.9	621.2	621.3

Table 4: OLS estimates of treatment effects in Charity experiment. Standard errors are clustered at the group level for subjects in SequentialC (30 groups; in total 271 clusters). * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

	<i>Dependent variable: Option B</i>		
	OLS (1)	Probit (2)	Logit (3)
BaselineC, round 2	0.00837 (0.0712)	0.0217 (0.184)	0.0348 (0.295)
SimultaneousC, round 2	0.125** (0.0583)	0.338** (0.157)	0.551** (0.257)
SequentialC, round 2	0.142* (0.0746)	0.506* (0.276)	0.888* (0.492)
BaselineC	0.397*** (0.0447)	-0.262** (0.116)	-0.419** (0.186)
SimultaneousC	0.583*** (0.0452)	0.210* (0.115)	0.336* (0.185)
SequentialC	0.721*** (0.0573)	0.585*** (0.170)	0.949*** (0.284)
Observations	920	920	920
R^2	0.709		
AIC	1115.3	1065.5	1065.5

Table 5: Difference estimates for second round effects in the charity experiments within each treatment. Standard errors are clustered at the group level for the second round of SimultaneousC and for both rounds of SequentialC (in total 395 clusters). * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

	<i>Dependent variable: Option B in round 2</i>	
	SimultaneousC (1)	SequentialC (2)
Option B in round 1	0.324*** (0.0989)	0.351*** (0.0783)
Increase in belief_pivotal	-0.00398** (0.00169)	-0.00538*** (0.00140)
Constant	0.498*** (0.0880)	0.536*** (0.0960)
Observations	120	240
Clusters	15	30
R^2	0.131	0.268

Table 6: Standard errors are clustered at the group level. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.