I Z A Institute
of Labor Economics
Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 13383

Moral Transgressions by Groups: What Drives Individual Voting Behavior?

Eberhard Feess
Florian Kerzenmacher
Gerd Muehlheusser

JUNE 2020

I Z A Institute of Labor Economics

Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 13383

Moral Transgressions by Groups: What Drives Individual Voting Behavior?

Eberhard Feess
Victoria University of Wellington

Florian Kerzenmacher
University of Innsbruck

Gerd Muehlheusser
University of Hamburg, CESifo and IZA

JUNE 2020

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Moral Transgressions by Groups: What Drives Individual Voting Behavior?*

We conduct an experiment where subjects are matched in groups of three and vote on a moral transgression. Analyzing different voting rules, the frequency of votes for the moral transgression increases with the number of votes required for it. This effect persists when considering pivotal votes only, which eliminates opportunities to save on own moral costs and to rely instead on sufficiently many votes for the transgression by other group members. A series of novel treatments allows us to identify guilt sharing and preferences for consensual voting as empirically relevant and independent drivers of voting behavior.

JEL Classification: C92, D02, D63, D71

Keywords: group decisions, unethical behavior, experiment, voting, diffusion of responsibility, guilt sharing, donations

Corresponding author:
Gerd Muehlheusser
University of Hamburg
Department of Economics
Von-Melle-Park 5
20146 Hamburg
Germany

E-mail: gerd.muehlheusser@wiso.uni-hamburg.de

* We are grateful to Daniel Chen, Martin Dufwenberg, Paul Heidhues, Bernd Irlenbusch, Oliver Kirchkamp, Aniol Llorente-Saguer, Petra Nieken, Ronald Peeters, Timo Promann, Anna Ressi, Georg von Weizsäcker, and Florian Zimmermann for very valuable comments. We also benefited from comments at the ESA World Meeting (Vancouver, 2019) and at seminar presentations at the Universities of Auckland, Dunedin, Hamilton, and Paderborn. Financial support from the German Research Foundation (DFG, Grant 658143) and the Austrian Science Fund (FWF, SFB F63) is gratefully acknowledged.

1 Introduction

The decision to carry out a morally questionable (or even illegal) action is often taken by groups rather than individuals. For example, hiring committees may follow their own nepotistic agenda instead of hiring the best candidate. Public officials may act in concert to deliberately award procurement contracts to companies which are inferior but pay high kickbacks. A group of top managers may agree on the whitewashing of balance accounts or even to commit outright fraud as in many topical corporate scandals such as Enron, Worldcom or Volkswagen.

The underlying decision-making processes can be perceived as a voting game where each group member decides on whether or not to support the moral transgression, which is then carried out if the number of affirmative votes is above a given threshold. In reality, these thresholds vary. For example, proposals in the European Union (Crombez, Huysmans, and Van Gestel, 2017; Bouton, Llorente-Saguer, and Malherbe, 2018) and decisions by procurement agencies (OECD, 2008) require either simple majority or unanimity. In some contexts, even a single group member might be sufficient to carry out the transgression, e.g. one out of several potentially corrupt officers accepting a bribe in exchange for a favorable administrative decision. Also, in the experimental setup of Falk, Neuber, and Szech (2020) with groups of eight, the moral transgression can be triggered by a single subject.

This paper studies the voting behavior with respect to moral transgressions in group settings under different voting rules. To the best of our knowledge it is the first attempt to disentangle various motives which are empirically relevant in this context. We conducted a pre-registered online experiment in which subjects were matched in groups of three and had to decide on taking money that was designated for donation to a charity.[1]

Our contribution is twofold: First, we analyze how the voting behavior is affected by the *voting threshold* (i.e. the number k of affirmative votes required to carry out the moral transgression), where our triad setting allows to compare dictatorial ($k = 1$), majority ($k = 2$), and unanimity ($k = 3$) voting. In our first treatment *SIM*, subjects vote simultaneously and independently for or against taking the money (YES or NO). We find that the number of subjects voting YES increases in the threshold k. This result is in line with predictions from a recent game-theoretic model by Rothenhäusler, Schweizer, and Szech (2018), where individuals incur a moral cost if and only if they vote for the moral transgression and it is actually carried out in the course of the group decision.[2] Furthermore, they assume that these moral costs decline in the number of other affirmative votes for the moral transgression, which they refer to as guilt sharing.[3] For example, under unanimity voting, individuals can downplay their own responsibility as anyone

[1]We adopt a donation setting as we want to consider a decision that harms subjects in need outside of the laboratory (Kirchler et al., 2016), and because taking money already designated for donation can clearly be seen as a moral transgression (Ariely, Bracha, and Meier, 2009; Bénabou and Tirole, 2010; Casal, Fallucchi, and Quercia, 2019).

[2]Such type of moral costs are usually referred to as *consequentialist* (Sinnott-Armstrong, 1988). By contrast, individuals incur non-consequentialist moral costs whenever they vote for the moral transgression, irrespective of whether or not it is actually carried out (Alexander and Moore, 2016).

[3]In the literature the term "diffusion of responsibility" is used as an alternative to guilt sharing. We use both terms synonymously.

else could have vetoed the moral transgression.

However, treatment *SIM* does not allow us to identify the presence of guilt sharing as there are several other potential explanations for the observed threshold effects. First and foremost, voting thresholds have a direct impact on free-riding opportunities: In particular, an individual in favor of the transgression might prefer to vote NO (thereby saving the moral cost) and instead rely on sufficiently many YES votes from the other group members such that it is nevertheless carried out. This free-riding incentive is largest under dictatorial voting (as it requires just one YES vote from one of the fellow group members), becomes smaller under majority voting (where both fellow group members must vote YES) and it is fully eliminated under unanimity voting. Hence, the observed threshold effects might simply be driven by these differences in free-riding incentives.[4]

This leads us to our second, and arguably more important, contribution: We design a series of novel treatments that allow us to isolate several behavioral motives potentially underlying our findings for treatment *SIM*. In a first step, we eliminate free-riding opportunities by considering only pivotal votes throughout. We proceed as follows: Each groups consists of two *unconditional* voters and one *conditional* voter. Unconditional voters decide independently on their votes, just as in *SIM*. For conditional voters we apply the strategy method and elicit their vote for each possible number of other YES votes. We then only consider the decisions of conditional voters when being pivotal for the moral transgression, i.e. the case with $k-1$ other YES votes. Our first treatment without free-riding incentives, *P-BASE*, reveals that the number of pivotal conditional YES votes still increases with the voting threshold k. This suggests that free-riding cannot be the sole explanation for the observed threshold effects but that there must exist additional motives.

To illustrate these different motives, consider a conditional voter in either one of two groups: one deciding under the majority rule (group M) and one deciding under the unanimity rule (group U). As a member of group M (U), this conditional voter knows that she is pivotal when there is one (two) other YES vote(s). Why would her behavior differ in the two groups? We argue that there exist three motives that should be distinguished conceptually and disentangled empirically. The first motive is *guilt sharing* or the diffusion of responsibility. When voting in favor of the moral transgression, the conditional voter can share guilt with one (two) fellow member(s) in group M (U), which would lead to a higher willingness to vote YES in group U. The second motive is a preference to follow a social norm, which we refer to as *social conformity*. Suppose that the conditional voter's only information about the social norm is the voting behavior of her two fellow members. In this case, pivotality in group U conveys a stronger signal that the social norm is to carry out the moral transgression.

The third motive, which we refer to as *consensual voting*, describes a tendency to not act against the revealed preferences of one's fellows.[5] Under the pivotal decision in group U, the

[4] Rothenhäusler, Schweizer, and Szech (2018) also point out that their prediction does not require the presence of guilt sharing, but can also be derived in a model with constant internal moral costs.

[5] As one participant puts it: "I always prefer to take the bonus, but in this case there were two reasons why I didn't. One is the new year and I give regularly to charities and the biggest reason is this is one of the best charities out there and they make a real difference. The people in my group will probably be upset with me but

conditional voter knows that both fellow members are in favor of the moral transgression, while there is disagreement under the pivotal decision in group M. Thus, when the conditional voter exhibits a preference for consensual voting, she would be more willing to votes YES in group U.[6] Importantly, each of these three motives leads to the same prediction that the observed frequency of conditional pivotal votes in favor of the moral transgression increases with the voting threshold k. Hence, each of these motives (or a combination of them) could be responsible for the observed threshold effects in *P-BASE*.

Consequently, in two additional treatments *P-INFO* and *P-RAND*, we identify guilt sharing and consensual voting as independent and empirically important motives for the voting behavior. These two treatments differ only with respect to the presence of guilt sharing; in *P-INFO* this motive is present, in *P-RAND* it is not.[7] The elimination of guilt sharing is accomplished by replacing the two unconditional votes in each group by random votes (i.e. whether or not the money is taken is now determined by these two random votes and the respective vote of the conditional voter). Comparing *P-INFO* and *P-RAND* reveals that the elimination of guilt sharing indeed leads to a lower impact of the voting threshold k on the number of pivotal conditional YES votes. This indicates that guilt sharing does play a role in our context.

Treatment *P-RAND* is also useful to assess the impact of consensual voting: While the unconditional votes in each group are no longer relevant for the decision to take the money, we still elicit the conditional votes for each number of unconditional YES votes. This allows us to analyze, separately for each voting threshold k, whether these unconditional YES votes do have an impact on the decisions of conditional voters. We find that this is indeed the case: For each voting rule, the number of pivotal conditional YES votes increases with the number of (irrelevant) unconditional votes for it. This provides strong evidence that also consensual voting is a highly relevant motive here.

Finally, we investigate whether our list of motives (free-riding, social conformity, guilt sharing and consensual voting) is close to exhaustive. If so, then after keeping all of these motives constant, we should observe no more impact of the voting rule on the number of pivotal conditional YES votes. Our results suggest that this is indeed the case.

The remainder of the paper is organized as follows: Section 2 discusses the related literature. Section 3 explains the experimental design and the treatments from which the hypotheses are derived. The results are presented in Section 4. Section 5 concludes and points to further

there is no way I could vote to take the money."

[6]According to our definition, there is a subtle but important difference between social conformity and consensual voting: Assume there are 100 groups, all taking their decision under the unanimity rule. Now consider the conditional voter of one of these groups, who learns that her group is the only one where both fellow members voted YES, while the overwhelming majority of individuals in the other 99 groups voted NO. Then, it seems reasonable to define the social norm as the behavior of the large majority in the overall population. A preference for social conformity would then lead to a vote in line with the social norm (i.e. NO), while a preference for consensual voting would lead to a vote which caters to the (revealed) preferences of one's own fellow members (i.e. YES).

[7]Importantly, we keep the impact of social conformity constant in both treatments by informing conditional voters about the overall frequency of unconditional YES votes as observed in treatment *P-BASE*. Moreover, we avoid any risk of deception by informing unconditional voters that their votes may or may not count for the decision on the moral transgression.

research. Appendix A develops a formal framework underlying the hypotheses. Appendix B contains further robustness checks. The experimental instructions are shown in Appendix C.

2 Relation to the literature

Our first treatment *SIM*, where all voters are unconditional, tests a proposition developed by Rothenhäusler, Schweizer, and Szech (2018). The additional treatments designed to identify guilt sharing and preferences for consensual voting are related to the literature on the diffusion of responsibility through markets (Falk and Szech, 2013; Bartling, Weber, and Yao, 2015b). Falk, Neuber, and Szech (2020) compare individual to group-decision making for the decision to sacrifice 10 Euros each for saving mice as initially developed in Falk and Szech (2013). Each group consists of eight subjects and eight mice are killed if at least one is unwilling to sacrifice the 10 Euros. They hence consider dictatorial voting for the morally questionable outcome and find that less individuals sacrifice the 10 Euros in the group setting compared to individual decision making.[8] An important distinctive feature of their approach is that anyone voting for saving the mice loses their 10 Euros even when the mice are effectively killed as at least one subject votes for it. This is different to our approach where a group member who votes against the moral transgression still gets the money if sufficiently other group members vote YES.

In Irlenbusch and Saxler (2019), a seller and a buyer can gain a monetary benefit at the expense of a charity if both agree on it. This unanimity rule for the moral transgression also yields a higher willingness to trigger the negative externality compared to individual decision making. Their approach is closely related to Dana, Weber, and Kuang (2007) who compare the behavior in different variants of the dictator game. In one of their treatments, two dictators decide between sharing the money equally with a recipient and keeping the lion's share to themselves. With unanimity required for the moral transgression, the unfair outcome is chosen more often than in a standard dictator game.

All of the experiments just discussed suggest a diffusion of responsibility but, similar to our treatment *SIM*, do not allow to disentangle the impact of behavioral preferences from changes in the incentive structure. In Falk, Neuber, and Szech (2020), the incentive structure in the group setting is identical to the individual setting if and only if a voter assumes a probability of 7/8 that at least one of the seven other group members votes for killing the mice.[9] In Dana et al. (2007) and in Irlenbusch and Saxler (2019), the results may be due to a diffusion in responsibility but could also be driven by an aversion against being pivotal, by a preference for consensual voting, or by a preference for increasing the partner's payoff. In line with the latter point, Wiltermuth (2011) finds that subjects lie more often when the benefit is shared with dummy players, which is related to preferences for white lies in Erat and Gneezy (2012).

[8]This result is confirmed in a donation setting.

[9]Only subjects with non-consequentialist moral costs could potentially sacrifice the money in the group setting even when their subjective probability that others are unwilling to safe the mice converges to one. Beliefs have hence not only a psychological impact but also influence the incentive structure.

Guilt sharing has also been investigated in settings where principals can delegate decisions on the allocation of money, and where more equal distributions reduce the own share. Bartling and Fischbacher (2012) find that delegation reduces the perception of guilt, thereby increasing the inequality of allocations. Oexl and Grossman (2013) show that this extends to settings where the dictator delegates the decision to an intermediary who can only choose between two unfair outcomes.

In our treatments, we disentangle the impact of changes in free-riding incentives from diffusion of responsibility (guilt sharing) and preferences for consensual decisions. This relates our findings to the literature on the interdependency of moral preferences (Bénabou, Falk, and Tirole, 2020). Monroe et al. (2018) find that the assessment of moral behavior depends to a large degree on what the majority does (social conformity) and less so on strict moral rules. We cannot quantify the impact of social conformity itself, but we can exclude that social conformity considerations influence the comparison of pivotal votes for the three voting thresholds.[10]

Nook et al. (2016) find that subjects donate more after learning that other participants donated high amounts. Somewhat contradictory, Gächter, Gerhards, and Nosenzo (2017) find that, while the behavior of other subjects influences what is viewed as morally appropriate in a dictator game, it has only a limited impact on the actual own behavior.[11]

Several studies suggest that social conformity and updates on the behavior of subjects one is matched with have little impact in lying experiments: In Peeters, Vorsatz, and Walzl (2015), the own behavior is not significantly influenced by the belief about the lying frequency of subjects one is matched with. Dato, Feess, and Nieken (2019) find that this holds even in highly competitive situations as lying by competing contestants has no impact on the own lying frequency. Conrads et al. (2013) find that lying in teams is more frequent compared to individual lying even when the own marginal monetary benefit is lower. This, however, might not be triggered by guilt sharing but rather by creating a positive externality on the partner. Kocher, Schudy, and Spantig (2018) also find a higher lying frequency in groups but the incentive structure again differs from the one with individual decision making.

Finally, our work is related to papers analyzing voting rules in moral contexts but with a different focus. Gerber, Neitzel, and Wichardt (2013) find that unanimity rules for the provision of public goods maximize the number of contributing subjects as it eliminates free-riding opportunities. Huck and Konrad (2005) show theoretically that, with non-consequentialist moral costs, the probability of reaching a predetermined threshold decreases in the committee size due to a decline in the probability of being pivotal. An equilibrium in mixed strategies then requires that all voters reduce their probability of voting for the moral transgression as this increases the

[10] A clean identification of the impact of social conformity itself would require to give different conditional voters different frequencies of unconditional votes for the moral transgression which, however, would not have been possible without deception. An alternative could have been to induce anchoring (Abeler, Nosenzo, and Raymond, 2019) or to ask for beliefs, but those could have been biased due to rationalizing the own behavior.

[11] While all of these papers contrast individuals to groups as decision makers, Bartling, Valero, and Weber (2017) compare settings where the negative externality affects either just one other party or is diffused over several parties. The frequency of moral transgression is slightly higher in the latter case but the difference is small and not robustly significant.

probability of being pivotal. Bartling, Fischbacher, and Schudy (2015a) consider a sequential voting structure and find experimentally that those who are pivotal for a moral transgression are more severely punished compared to other subjects.

3 Experimental design, treatments, and hypotheses

In this section, we explain the basic setup of the experiment and the treatments we have conducted. The treatments are designed to disentangle various possible motives for individuals' voting behavior concerning a morally questionable decision. This leads to a total of seven hypotheses to be tested (see Appendix A for a formal framework in which these hypotheses can be derived).

3.1 Experimental design

Subjects were randomly assigned to groups of three. They voted (YES or NO) on whether or not the group should take an amount of \$1.50 originally designated for donation to the charity *Make-A-Wish*, which engages in fulfilling the wishes of children with a critical illness.[12] If sufficiently many group members voted for taking the money, the \$1.50 were split equally among them. We consider three voting rules that differ with respect to the threshold k of YES votes required for taking the money. Throughout, we refer to $k = 1$, $k = 2$, and $k = 3$ as the *dictatorial*, *majority*, and *unanimity* rule, respectively, and to *threshold effects* when the number of YES votes increases in k.

3.2 Treatments and hypotheses

We consider a total of four treatments (see Table 1 for an overview). As explained in more detail below, we have employed a between-subject design. Each subject was randomly allocated to one treatment and then played under one of the three voting rules.

Treatment *SIM* In this treatment all three group members are *unconditional* voters in the sense that none of them can condition her own vote on the votes of the other two group members. This treatment hence captures the case of simultaneous voting as considered in the theoretical framework by Rothenhäusler, Schweizer, and Szech (2018). They consider a group of $n \geq 1$ agents voting on a moral transgression that yields the same monetary payoff for each group member, and which is implemented if there are at least k votes in favor of it. In the main part of their paper, they assume that subjects face moral costs if and only if they vote for the moral transgression *and* the transgression is actually implemented.

[12] See e.g. `www.wish.org`. Prior to the actual experiment, we ran an online survey (with a different subject pool) and asked subjects to rank several well-known charities such as the American Red Cross, the Children's Scholarship Fund, Doctors without borders, Donate Life and UNICEF with respect to their reputation. Make-A-Wish was ranked highest among these.

Moral costs are divided into two components, *non-diffusive moral costs* and *guilt sharing costs*, where the latter decrease in the number of votes in favor of the moral transgression.[13] Rothenhäusler, Schweizer, and Szech (2018) analyze (symmetric Bayesian Nash) equilibria, where each subject votes for the moral transgression if and only if her (privately known) moral costs are below a critical value. This yields their Proposition 1, part (ii), which translates directly into our first hypothesis.

Hypothesis 1. *In treatment* SIM*, the frequency of YES votes for the moral transgression increases in the voting threshold* k*.*

Our experiment provides a first empirical test of this basic prediction. It is important to note that the presence of internal moral costs alone is sufficient to derive Hypothesis 1 as the three voting rules give rise to different free-riding opportunities: The higher the likelihood that sufficiently many other group members vote for the moral transgression, the higher is the incentive to save the own moral costs by voting against it. In equilibrium, this likelihood decreases in the threshold k and is zero under unanimity. Hence, while testing Hypothesis 1 is adequate for analyzing the overall effect of voting thresholds, it is not informative about the presence of additional behavioral motives such as social conformity, guilt sharing, and consensual voting. Investigating those motives in more detail is the aim of our subsequent treatments.

Treatment *P-BASE* This treatment addresses the issue of free-riding. In contrast to *SIM*, one of the three members in each group is randomly chosen to become a conditional voter. For these conditional voters, we use the strategy method and elicit their decision for all possible numbers of YES votes by their two fellow group members (zero, one, or two) who remain unconditional voters. We then analyze the voting behavior of conditional voters for those vote constellations in which they are pivotal, i.e. where the number of unconditional YES votes in her group equals $k-1$.[14] Throughout, we refer to YES votes by conditional voters when being pivotal as *pivotal conditional YES votes*. Confining attention to these votes allows us to study threshold effects in a setting where free-riding opportunities are eliminated. If threshold effects continue to exist, then this would hint at (some of) the other motives discussed above.[15]

Hypothesis 2. *In treatment* P-BASE *the frequency of pivotal conditional YES votes increases in the voting threshold* k*.*

To permanently eliminate free-riding incentives from the subsequent analysis, all further treatments also focus on pivotal conditional YES votes. Moreover, apart from the introduction

[13] Their model also includes an aversion against being pivotal. We neglect this cost type as it does not influence the impact of the voting threshold in their model. Furthermore, we consider only pivotal votes anyway in all of our subsequent treatments.

[14] Note that whether the money is taken still depends on the actual choices of all three group members. For example, consider the case k = 3 with (i) a conditional voter who votes YES for any number of unconditional YES votes, and (ii) zero unconditional YES votes. For our analysis, we use only the decision where the conditional voter is pivotal, i.e. where she is asked for her decision in case of two unconditional YES votes. However, the moral transgression is not implemented as there is only one actual YES vote in the group.

[15] Considering pivotal votes only also has the merit that it does not matter whether moral costs are consequentialist or non-consequentialist (see the discussion of these two concepts in the Introduction).

of conditional voters, there is also a second difference between treatment *SIM* on the one hand and *P-BASE* (and the two subsequent treatments *P-INFO* and *P-RAND*) on the other: Before asking for their vote, unconditional voters were informed that it might be the case that their votes do not count towards the decision to take the money. We did so to ensure that unconditional voters are exactly in the same situation in all of these three treatments where the focus is on conditional voters. Otherwise, comparing the conditional pivotal voters' decisions between these treatments could potentially be contaminated by their different perceptions of the unconditional voters' motives.[16]

Treatment *P-INFO* This treatment acts as a bridge between treatments *P-BASE* and *P-RAND* (discussed below). It is identical to *P-BASE* with the only difference that, before asking conditional voters for their decision strategy, they were informed about the average voting behavior of the unconditional voters in *P-BASE* (which was played beforehand) under the respective voting rule. For instance, under majority voting in *P-BASE*, 65% of the unconditional voters voted YES, and this information was then provided to all conditional voters assigned to the majority rule in *P-INFO*. Such information was meant to provide conditional voters with a signal about the prevailing social norm in the population. The reason why we provide this information is that the test of Hypothesis 5, which is key for the identification of guilt sharing, will be based on a comparison of the voting behavior in treatments *P-INFO* and *P-RAND*. To cleanly identify guilt sharing, we need to make sure that the perception on what constitutes the social norm is the same in both treatments.

For *P-INFO* itself, we presume that providing the percentages of unconditional votes for the moral transgression does not erase the threshold-effect already predicted in *P-BASE*:

Hypothesis 3. *In treatment* P-INFO *the frequency of pivotal conditional YES votes increases in the voting threshold k.*

Treatment *P-RAND* This final treatment is identical to *P-INFO* except for one feature: The decision for the moral transgression is not determined by the three votes in a given group, but by the group's conditional voter and two random votes. The probability of the random vote being YES is taken from the frequency of YES votes in *P-BASE* for the respective voting threshold in order to ensure that a conditional pivotal voter's decision is not distorted by assigning different probabilities to the situations she might end up in.[17]

Whether a conditional voter is pivotal is hence no longer determined by the unconditional votes in her group, but by the outcome of the random votes. In our design, unconditional votes do not count for the group decision to take the money. But when the money is taken, it is still

[16] In the empirical analysis, we do not utilize the unconditional votes from these treatments. Unconditional votes are only used in the analysis of treatment *SIM*, where all voters are unconditional and where they are certain that their votes do count with respect to the decision to carry out the transgression.

[17] For instance, a conditional voter who assumes that she is pivotal with a very small probability anyway might just vote NO due to image concerns.

Table 1: Overview over treatments and hypotheses

	SIM	*P-BASE*	*P-INFO*	*P-RAND*
Conditional versus unconditional voters	No	Yes	Yes	Yes
Unconditional voters know vote might not count	No	Yes	Yes	Yes
Signal about social norm	No	No	Yes	Yes
Random votes	No	No	No	Yes
Motive eliminated	None	Free-riding	Free-riding	Free-riding, Guilt sharing
Hypotheses	1	2	3 & 5	4 – 7

shared among all three group members. For example, suppose that $k = 3$ and that there are zero conditional YES votes, two random YES votes and that the conditional voter votes YES for any number of random YES votes. Then the money is taken by the group and all three members receive their according payoffs, albeit the two unconditional voters actually voted against it.

The purpose of substituting the actual votes of unconditional voters by random votes is to eliminate guilt-sharing as a potential motive for conditional voters: Arguably, conditional pivotal voters can no longer share the guilt for an immoral outcome with the other group members when their votes do not matter for the decision. The effect of eliminating guilt sharing, and hence the difference in YES votes between *P-INFO* and *P-RAND*, should be strongest under the unanimity rule, followed by the majority rule. For the dictatorial rule, guilt-sharing is not an issue as a conditional voter can decide alone on the moral transgression, irrespective of other votes (actual or random).

When eliciting the decisions of conditional voters via the strategy method, we take into account both unconditional and random votes, leading to $3^2 = 9$ decisions for conditional voters (one for each of zero, one, or two YES votes from the unconditional and random votes, respectively).

We first investigate the behavior of conditional voters when we substitute unconditional by random votes but keep everything else constant. We therefore consider the case where both the number of random YES and the number of unconditional YES votes in a group is $k - 1$ (i.e. we use only one out of the elicited nine decisions from conditional voters). Note that this procedure eliminates guilt sharing but keeps the potential motive for consensual voting in place as the number of actual YES votes of unconditional voters increases in k when a conditional voter is pivotal. Our next prediction is that consensual voting matters:

Hypothesis 4. *In treatment* P-RAND*, when also the number of (irrelevant) unconditional votes* *in a group is $k-1$, the frequency of pivotal conditional YES votes increases in the voting threshold k.*

While the presence of consensual voting can directly be tested by considering *P-RAND*, the presence of guilt sharing can be tested by comparing the conditional pivotal voters' behavior in *P-RAND* and in *P-INFO*. As guilt sharing is present in *P-INFO* but switched off in *P-RAND*, we expect that the threshold effects are smaller in *P-RAND* compared to *P-INFO*:

Hypothesis 5. *Comparing treatments* P-INFO *and* P-RAND*, the difference in the frequencies of pivotal conditional YES votes is increasing in the voting threshold k.*

Treatment *P-RAND* also allows us to examine the relevance of preferences for consensual voting in more detail. For any given voting rule, the percentage of pivotal conditional YES votes should increase in the number of unconditional YES votes. Thus, for a given number of $k-1$ random YES votes, we now consider all three possible values of unconditional YES votes (i.e. three out of the elicited nine decisions from each conditional voter). These unconditional votes matter neither for the outcome nor for determining pivotality of conditional voters, but they should still matter when preferences for consensual voting are relevant:

Hypothesis 6. *In treatment* P-RAND*, for each voting threshold k, the frequency of pivotal conditional YES votes increases in the number of unconditional YES votes in their group.*

Conversely to all previous hypotheses, the test of Hypothesis 6 is not based on comparisons of YES votes *between* different voting thresholds k, but on the comparison of YES votes *within* each threshold.

A final question is whether, in addition to free-riding, social conformity, guilt-sharing, and consensual voting, there exist additional motives that we cannot capture with our design. This issue can also be addressed using treatment *P-RAND*. In *P-RAND*, only consensual voting remains, and we can eliminate this by fixing the number of unconditional YES votes in a given group. The absence of a threshold effect would then provide strong evidence against the relevance of additional motives other than the ones considered so far. This leads to our final prediction:

Hypothesis 7. *In treatment* P-RAND*, when fixing the number of unconditional YES votes in each group, the frequency of pivotal conditional YES votes is independent of the voting threshold k.*

3.3 Implementation of the experiment

The experiment was pre-registered at the RCT Registry of the American Economic Association (AEARCTR-0004317) and conducted in January and February 2020 via MTurk with a total of 4680 subjects (52% female). On average, it took participants about four minutes to participate and they earned 1.14$ (including a show-up fee of 0.80$).

We employed a between-subjects design. Each subject participated only in one treatment and played under only one of the three voting rules. We conducted three sessions for treatment *SIM* and four sessions for all other treatments. Subjects were randomly assigned to voting rules and groups, and to roles within their group (unconditional or conditional voter). They then received instructions depending on their assigned role and voting rule. Instructions were identical for all subjects in treatment *SIM* playing under the same voting rule (see Appendix C).

Unconditional voters in all other treatments received the same (voting-rule specific) instructions, while the instructions for conditional voters varied across treatments. After they had read the instructions, subjects were asked some control questions. They could continue to the next page only after they had answered correctly. They were then asked for their vote. Unconditional voters were simply asked to vote for or against taking the money (YES or NO). For conditional voters, we used the strategy method to elicit their votes conditional on other voting decisions: In treatments *P-BASE* and *P-INFO*, there were three conditional votes (one for each possible number of YES votes by the unconditional voters in their group). In treatment *P-RAND*, there were nine conditional votes (one for each combination of unconditional and random YES votes).

After the voting stage, we elicited subjects' age and gender. In addition, we asked them about their beliefs concerning the voting behavior of the other participants and how trustworthy they considered the charity Make-A-Wish.[18] Finally, we elicited risk preferences by asking subjects for their willingness to take risks (measured on an ordinal scale from zero to ten; similar to Dohmen at al., 2011).

In treatment *SIM*, where all subjects are symmetric (all unconditional voters), we have 120 subjects for each of the three voting rules. In *P-BASE*, *P-INFO*, and *P-RAND*, where the focus is on the behavior of conditional voters, the number of subjects was tripled, which gives us 120 conditional voters for each voting rule.

4 Results

In this section, we report our results for all within- and between-treatment predictions from Section 3 (Hypotheses 1 – 7). For each treatment, Figure 1 shows the percentage of votes for taking the money, separated by the three voting thresholds. We test the hypotheses by estimating linear probability models. The regression results are shown in Tables 2 and 3, and most of our predictions are supported by the data. Our findings suggest that guilt sharing and consensual voting are both independent drivers of the voting behavior. As shown in Appendix B, the results remain robust when instead using a non-parametric approach and when considering a richer set of controls. Results are also robust when estimating probit or logit models instead of linear probability models.

[18] As these two variables were elicited ex post, we only use them as a robustness check (see Appendix B).

Figure 1: Percentage of YES votes by voting thresholds (k) across treatments

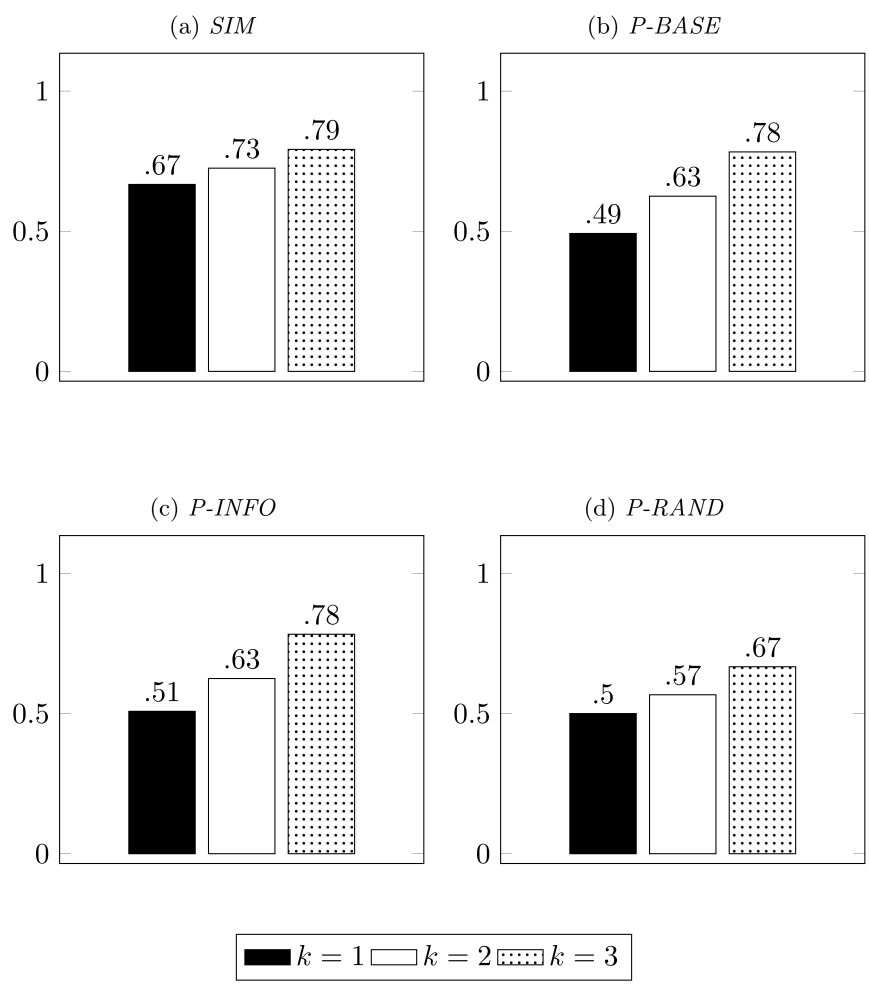

The impact of voting thresholds: Testing Hypothesis 1 We first consider the impact of the threshold k on voting behavior in treatment *SIM*. In this treatment, voters are fully symmetric and all three votes are taken into account. The percentage of votes in favor of taking the money is depicted in Figure 1(a) and increases with the voting threshold k. The respective regression results are shown in Table 2, column (1). The likelihood of YES votes is significantly higher at the 5%-level when moving from dictatorial ($k = 1$) to unanimity ($k = 3$) voting. Moreover, we reject the hypothesis of no joint impact of majority and unanimity voting compared to dictatorial voting. These findings provide support for Hypothesis 1 and the theory of Rothenhäusler, Schweizer, and Szech (2018). Recall that treatment *SIM* captures the overall impact of the voting threshold but does not allow us to disentangle different motives.

Eliminating free-riding: Testing Hypotheses 2 and 3 If results in treatment *SIM* were solely driven by free-riding, then the votes of subjects knowing that they are pivotal should not be influenced by the voting threshold. To test this, we next consider treatment *P-BASE* where one of three group members is asked for their vote conditional on the number of YES votes (zero, one, or two) by the other two group members. By taking only those vote constellations

Table 2: Regressions results for the tests of Hypotheses 1 to 5

	(1) *SIM* (H1)	(2) *P-BASE* (H2)	(3) *P-INFO* (H3)	(4) *P-RAND* (H4)	(5) *P-INFO + P-RAND* (H5)
Majority	0.0567	0.139**	0.117*	0.0614	0.117*
	(0.350)	(0.028)	(0.067)	(0.339)	(0.066)
Unanimity	0.118**	0.273***	0.287***	0.138**	0.290***
	(0.040)	(0.000)	(0.000)	(0.031)	(0.000)
P-RAND					0.00259
					(0.968)
Majority x P-RAND					-0.0586
					(0.516)
Unanimity x P-RAND					-0.150*
					(0.086)
Observations	360	360	360	360	720
Adjusted R^2	0.018	0.082	0.050	0.045	0.050
T1: Una. vs. Majo.	0.258	0.021	0.004	0.218	
T2: (Una. + Majo.) vs. Dict.	0.096	0.000	0.000	0.075	
T3: Treatment effect Majo.					0.373
T4: Treatment effect Una.					0.011
Controls Pers. Charact.	Yes	Yes	Yes	Yes	Yes
Session Dummies	Yes	Yes	Yes	Yes	No

Notes: Each column refers to a linear probability model with YES votes as the dependent variable. Regression (1) uses the observations from all three voters in a group; all other regressions use only pivotal conditional YES votes. p-values are reported in parentheses, and *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level (robust standard errors). The reference category is dictatorial voting. The bottom of the table shows p-values for the following Wald tests: (i) difference between unanimity and majority voting (T1), (ii) difference of unanimity and majority compared to dictatorial voting (T2), (iii) treatment effects for majority voting (T3) and unanimity voting (T4). All regressions contain controls for personal characteristics (gender, age, risk attitude). Session dummies are included in regressions (1) to (4), where within-treatment comparisons are conducted.

into account where the conditional voter is pivotal (i.e. where the moral transgression occurs if and only if the conditional voter votes YES), this removes free-riding incentives.

Figure 1(b) illustrates that the percentage of pivotal conditional YES votes increases in the voting threshold k. Regression results in Table 2, column (2), also show that, compared to dictatorial voting, the likelihood of YES votes is significantly higher for both majority and unanimity voting. Moreover, the Wald tests reveal at high levels of significance that (i) the impact of these two rules differs, and (ii) that there is a joint effect of these rules compared to dictatorial voting. All in all, these results are supportive of Hypothesis 2.

Very similar findings emerge in treatment *P-INFO* (see Figure 1(c)). Recall that this treatment is identical to *P-BASE* except that, before deciding on their voting strategy, conditional voters are informed about the percentage of unconditional YES votes in treatment *P-BASE*. This is meant as a signal about the prevailing social norm, as it refers to all 720 unconditional voters in *P-RAND*. Regression results for *P-INFO* are reported in Table 2, column (3). They show that the differences to dictatorial voting are again statistically significant for both majority and unanimity voting. The same is true for the comparison of the effects of unanimity and majority voting, as well as for their joint effect compared to dictatorial voting. These findings support Hypothesis 3. All in all, the evidence for Hypotheses 2 and 3 strongly suggest that the observed threshold effects are not solely driven by free-riding.

Relevance of consensual voting and guilt sharing: Testing Hypotheses 4 and 5 Treatment *P-RAND* plays a crucial role in the analysis of consensual voting and guilt sharing as potential drivers of threshold effects. Recall that *P-RAND* is identical to *P-INFO* with one exception: In *P-RAND*, whether a group takes the money is no longer determined by the three votes in a group, but by two random votes and the corresponding pivotal vote of the conditional voter. This removes guilt sharing as a potential motive for conditional voters, while keeping constant their information about the prevailing social norm (i.e. they still learn the overall behavior of unconditional voters in *P-BASE*). Hence, if threshold effects persist in *P-RAND*, then this would hint at the empirical relevance of consensual voting (Hypothesis 4).[19]

The behavior in treatment *P-RAND* is illustrated in Figure 1(d) and regression results are shown in Table 2, column (4). As can be seen, the results are qualitatively similar to those for *P-BASE* and *P-INFO*. In particular, threshold effects exist and the difference is statistically significant for unanimity voting. However, the coefficients in *P-RAND* become smaller compared to *P-INFO* (and *P-BASE*), so that also the Wald tests become less significant or even insignificant. Arguably, these findings are due to guilt sharing.

This issue can be analyzed in more detail by comparing *P-INFO* (where guilt sharing is present) and *P-RAND* (where it is not present). According to Hypothesis 5, for dictatorial voting ($k = 1$) where guilt cannot be shared, behavior should not differ between the two treatments. Conversely, for the majority and the unanimity rule, there should be fewer pivotal conditional YES votes in *P-RAND* than in *P-INFO* and the effect should be stronger under unanimity voting.

Comparing panels (c) and (d) of Figure 1 shows that the percentage of YES votes is indeed lower in *P-RAND* for majority and unanimity voting, while there is no difference for dictatorial voting. This is also confirmed in the regression analysis (see Table 2, column (5)): As predicted, there is no treatment effect for dictatorial voting. Furthermore, the interaction term of the unanimity rule with the treatment dummy reveals that the difference between *P-RAND* and *P-BASE* is significantly larger for unanimity compared to dictatorial voting. However, the predicted difference for the comparison of majority and dictatorial voting does not show up in the data. When considering treatment effects for each voting threshold separately, in line with Hypothesis 5, we find no treatment difference for dictatorial voting. Moreover, a Wald test reveals that there are significantly less pivotal conditional YES votes in *P-RAND* for unanimity voting. Also under majority voting, there are fewer pivotal conditional YES votes in *P-RAND*, but the effect is not strong enough to be statistically significant. Overall, our findings support the view that guilt sharing is an empirically relevant motive for the voting behavior in our context.

The role of consensual voting: Testing Hypothesis 6 Treatment *P-RAND* can also be used to test our remaining two hypotheses. Hypothesis 6 predicts a positive correlation between the pivotal conditional YES votes (as determined by the two random votes) and the number of YES votes of their two fellow group members (which do not count towards whether or not the

[19] Recall from above, that for the test of Hypothesis 4 we use the votes of conditional voters only for the case where the number of unconditional YES votes also equals $k-1$ (i.e. they are the same as the number of random YES votes).

Figure 2: Treatment *P-RAND*: Percentage of YES votes of conditional voters when pivotal, by voting thresholds and by number of unconditional YES votes in group (*GroupYes*)

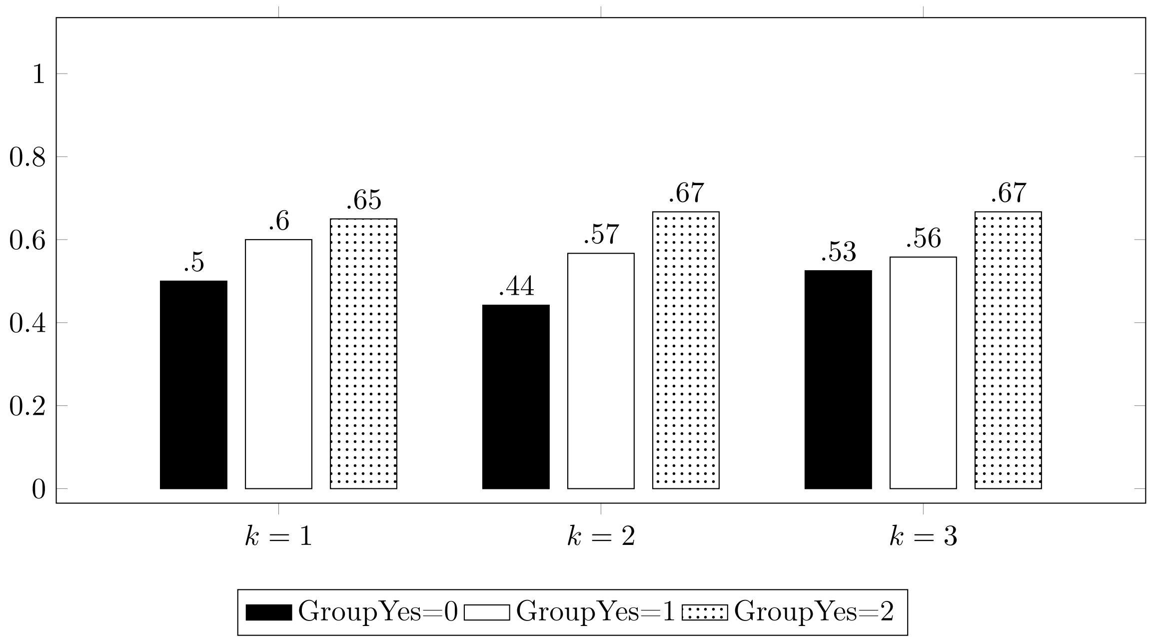

money is taken), measured by the variable *GroupYes* $\in \{0, 1, 2\}$. Such a positive relationship would provide evidence that consensual voting, defined as a preference for taking decisions in line with the other group members' preferences, matters in our context.

We test Hypothesis 6 separately for each of the three voting rules, utilizing the fact that we have elicited the conditional voters' decisions for all nine possible combinations of unconditional and random YES votes (see the treatment description in Section 3). Again, we focus on pivotal conditional YES votes, i.e. we still use only vote constellations with $k-1$ random YES votes. However, in contrast to the testing of Hypothesis 4, we here no longer confine attention to the case of $k-1$ unconditional YES votes. Instead, we use all three possible constellations of unconditional YES votes (zero, one, or two). As this yields three pivotal decisions per conditional voter, standard errors in the regressions are clustered at the subject level.

The results are illustrated in Figure 2 and the regression results are shown in Table 3, columns (1)–(3). For all three voting thresholds, the likelihood of pivotal conditional YES votes indeed increases in the number of YES votes by unconditional voters in their group. Five out of the six respective coefficients are statistically significant at high levels, thereby lending strong support for Hypothesis 6 and the empirical relevance of consensual voting in our context.

Relevance of further motives: Testing Hypothesis 7 The treatments discussed so far allowed us to address and isolate, in turn, several motives driving the observed threshold effects. In particular, we have *(i)* eliminated free-riding by considering only conditional voters in cases where they are pivotal (*P-BASE*, *P-INFO*, *P-RAND*), *(ii)* provided (and kept constant) information about the prevailing social norm (*P-INFO*, *P-RAND*), and *(iii)* isolated the effect

Table 3: Regressions results for the tests of Hypotheses 6 and 7

	(1)	(2)	(3)	(4)	(5)	(6)
	(H6)	(H6)	(H6)	(H7)	(H7)	(H7)
	Dictatorial	Majority	Unanimity	GroupYes=0	GroupYes=1	GroupYes=2
GroupYes=1	0.1000**	0.125***	0.0333			
	(0.024)	(0.003)	(0.163)			
GroupYes=2	0.150***	0.225***	0.142***			
	(0.001)	(0.000)	(0.001)			
Majority				-0.0589	-0.0406	0.0129
				(0.359)	(0.522)	(0.834)
Unanimity				0.000920	-0.0748	0.000926
				(0.989)	(0.242)	(0.988)
Observations	360	360	360	360	360	360
Adjusted R^2	0.036	0.091	0.109	0.031	0.029	0.011
T1: Both coeff. equal	0.113	0.003	0.005	0.349	0.588	0.847
T2: Sum both coeff. zero	0.004	0.000	0.003	0.606	0.297	0.898
Controls for Pers. Charact.	Yes	Yes	Yes	Yes	Yes	Yes
Session Dummies	Yes	Yes	Yes	Yes	Yes	Yes
Clust. Subj. Level	Yes	Yes	Yes	No	No	No

Notes: All regressions are based on treatment *P-RAND*. Each column refers to a linear probability model with pivotal conditional YES votes as the dependent variable. p-values are reported in parentheses, and *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level (robust standard errors). The reference category is dictatorial voting in regressions (1)-(3) and GroupYes=0 in regressions (4)-(6). The bottom of the table shows p-values for two Wald tests: one for the equality of the two coefficients in the respective columns (Test 1), and one for joint significance of both coefficients compared to the reference category (Test 2). All regressions include controls for personal characteristics (gender, age, risk attitude) as well as session dummies. Since each conditional voter is observed with three different decisions (one for each possible value of unconditional YES votes), standard errors are clustered at the subject level.

of guilt sharing and consensual voting (*P-RAND*).

This raises the question whether our design indeed captures the most relevant motives, or whether there exist further ones not considered so far. A straightforward way of addressing this issue is to examine if the threshold effects as observed in *P-RAND* vanish once we also neutralize consensual voting by fixing the number of YES votes by the unconditional voters in each group (Hypothesis 7).

Descriptive evidence for Hypothesis 7 can be observed from Figure 2 when comparing the numbers for a given level of unconditional YES votes (i.e. comparing the three black, white and dotted bars, respectively). As can be seen, all threshold effects indeed virtually disappear. This is also confirmed in the regression analysis as shown in Table 3, columns (4)–(6), where none of the six coefficients for the voting thresholds is even close to being statistically significant.[20] These findings provide strong support for Hypothesis 7, suggesting that our design indeed captures the most relevant motives.

[20] We have also conducted an additional treatment that is identical to *P-RAND* except that we do not inform conditional pivotal voters on the number of YES votes in their group. In this case, the voting threshold may still matter as, depending on the threshold, conditional voters may form different expectations on the number of YES votes by their group fellows. Consistent with such a reasoning, we find that the frequency of YES votes by pivotal voters is significantly lower under dictatorial voting (53.3%) compared to both the majority (65%) and unanimity (64.2%) voting.

5 Conclusion

Recent papers argue that moral transgressions are more likely when subjects decide in groups rather than on their own. This is often attributed to guilt sharing (also referred to as diffusion of responsibility) and supported by many experiments that compare individual decision making to situations where the members of a group act in concert (e.g. Dana, Weber, and Kuang, 2007; Kocher, Schudy, and Spantig, 2018) or vote on a final outcome (e.g. Falk, Neuber, and Szech, 2020). While it has been shown in this literature that the frequency of moral transgressions is higher in groups, we are not aware of any paper that disentangles guilt sharing from other motives. Identifying guilt sharing is nontrivial for at least two reasons: First, group decisions and voting rules differ from individual decision making with respect to the opportunities to free-ride on the (immoral) behavior of other participants. Second, guilt sharing is likely to overlap with social conformity considerations and preferences for consensual decisions. These motives are related but not identical.

In our triad setting, group members vote independently on whether to transfer money designated for donation to the own account. Comparing dictatorial, majority, and unanimity voting, we first consider a setting with only unconditional voters. Our experimental findings support the model prediction of Rothenhäusler, Schweizer, and Szech (2018) that, when moral concerns are consequentialist, the number of subjects voting for the moral transgression increases in the voting threshold. We then design a series of novel treatments in order to isolate the presence of guilt sharing and consensual voting. We find clear evidence for both motives. In addition, our data lend support to the hypothesis that our approach captures the most important behavioral motives: we do not observe any difference in conditional pivotal votes between the three voting rules after holding social conformity, guilt sharing, and consensual voting constant.

For future research, we aim at extending the analysis in two directions. First, we want to examine the robustness of our results by considering, instead of taking money designated for donation, a context of lying about the (privately observed) outcome of a lottery (see the seminal paper by Fischbacher and Föllmi-Heusi, 2013, and the meta study by Abeler, Nosenzo, and Raymond, 2019). Such an extension is potentially interesting as one might suspect that lying involves a larger non-consequentialist moral cost component as it is perceived as immoral even when the final outcome is unaffected. While this has no impact on the voting behavior of pivotal voters (as their votes always determine the final outcome by definition of pivotality), it may well influence the behavior of non-pivotal voters. In fact, the theoretical prediction of Rothenhäusler, Schweizer, and Szech (2018) that the frequency of YES votes increases in the voting threshold does not extend to non-consequentialist moral costs as there is a countervailing effect: A YES vote now bears the risk of incurring the non-consequentialist moral costs without a benefit (i.e. when the overall number of YES votes is below the respective threshold). When this "risk" is increasing in the voting threshold, this could lead to fewer YES votes accordingly. In a pilot for simultaneous voting with 864 subjects, we find that 14.3%, 27.5%, and 27.4% of participants lie under the dictatorial, majority, and unanimity voting rule, respectively.

Second, it would be interesting to extend the analysis to a sequential voting procedure, where each subject observes the votes of their predecessors. To illustrate, consider the second voter under majority voting. Would this voter be more inclined to vote for or against the moral transgression when the first voter has voted in favor of it? Arguably, this is unclear due to countervailing effects: On the one hand, preferences for consensual voting suggest to vote YES more often. On the other hand, however, the first voter's YES vote gives rise to free-riding opportunities, which disappear in case the first voter voted NO.

References

Abeler, J., D. Nosenzo, and C. Raymond (2019): "Preferences for truth-telling," *Econometrica*, 87, 1115–1153.

Alexander, L. and M. Moore (2016): "Deontological ethics," in *The Stanford Encyclopedia of Philosophy*, ed. by E. N. Zalta, Stanford University.

Ariely, D., A. Bracha, and S. Meier (2009): "Doing good or doing well? Image motivation and monetary incentives in behaving prosocially," *American Economic Review*, 99, 544–55.

Bartling, B. and U. Fischbacher (2012): "Shifting the blame: On delegation and responsibility," *Review of Economic Studies*, 79, 67–87.

Bartling, B., U. Fischbacher, and S. Schudy (2015a): "Pivotality and responsibility attribution in sequential voting," *Journal of Public Economics*, 128, 133–139.

Bartling, B., V. Valero, and R. Weber (2017): "On the scope of externalities in experimental markets," *Experimental Economics*, 22, 1–15.

Bartling, B., R. Weber, and L. Yao (2015b): "Do markets erode social responsibility?" *Quarterly Journal of Economics*, 130, 219–266.

Bénabou, R., A. Falk, and J. Tirole (2020): "Narratives, imperatives, and moral persuasion," *University of Bonn, mimeo.*

Bénabou, R. and J. Tirole (2010): "Individual and corporate social responsibility," *Economica*, 77, 1–19.

Bouton, L., A. Llorente-Saguer, and F. Malherbe (2018): "Get rid of unanimity rule: The superiority of majority rules with veto power," *Journal of Political Economy*, 126, 107–149.

Casal, S., F. Fallucchi, and S. Quercia (2019): "The role of morals in three-player ultimatum games," *Journal of Economic Psychology*, 70, 67–79.

Conrads, J., B. Irlenbusch, R. Rilke, and G. Walkowitz (2013): "Lying and team incentives," *Journal of Economic Psychology*, 34, 1–7.

Crombez, C., M. Huysmans, and W. Van Gestel (2017): "Choosing an informative agenda setter: The appointment of the Commission in the European Union," *European Union Politics*, 18, 145–167.

Dana, J., R. Weber, and J. X. Kuang (2007): "Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness," *Economic Theory*, 33, 67–80.

Dato, S., E. Feess, and P. Nieken (2019): "Lying and reciprocity," *Games and Economic Behavior*, 118, 193–218.

Dohmen, T., A. Falk, D. Huffman, U. Sunde, J. Schupp, and G. Wagner (2011): "Individual risk attitudes: Measurement, determinants, and behavioral consequences," *Journal of the European Economic Association*, 9, 522–550.

Erat, S. and U. Gneezy (2012): "White lies," *Management Science*, 58, 723–733.

Falk, A., T. Neuber, and N. Szech (2020): "Diffusion of being pivotal and immoral outcomes," *Review of Economic Studies, forthcoming.*

Falk, A. and N. Szech (2013): "Morals and markets," *Science*, 340, 707–711.

Fischbacher, U. and F. Föllmi-Heusi (2013): "Lies in disguise—an experimental study on cheating," *Journal of the European Economic Association*, 11, 525–547.

Gächter, S., L. Gerhards, and D. Nosenzo (2017): "The importance of peers for compliance with norms of fair sharing," *European Economic Review*, 97, 72–86.

Gerber, A., J. Neitzel, and P. C. Wichardt (2013): "Minimum participation rules for the provision of public goods," *European Economic Review*, 64, 209–222.

Huck, S. and K. A. Konrad (2005): "Moral cost, commitment, and committee size," *Journal of Institutional and Theoretical Economics*, 575–588.

Irlenbusch, B. and D. J. Saxler (2019): "The role of social information, market framing, and diffusion of responsibility as determinants of socially responsible behavior," *Journal of Behavioral and Experimental Economics*, 80, 141–161.

Kirchler, M., J. Huber, M. Stefan, and M. Sutter (2016): "Market design and moral behavior," *Management Science*, 62, 2615–2625.

Kocher, M., S. Schudy, and L. Spantig (2018): "I lie? We lie! Why? Experimental evidence on a dishonesty shift in groups," *Management Science*, 64, 3995–4008.

Monroe, A., K. Dillon, S. Guglielmo, and R. Baumeister (2018): "It's not what you do, but what everyone else does: On the role of descriptive norms and subjectivism in moral judgment," *Journal of Experimental Social Psychology*, 77, 1–10.

Nook, E., D. Ong, S. Morelli, J. Mitchell, and J. Zaki (2016): "Prosocial conformity: Prosocial norms generalize across behavior and empathy," *Personality and Social Psychology Bulletin*, 42, 1045–1062.

OECD (2008): "Compendium of country examples and lessons learned from applying the methodology for assessment of national procurement systems: Sharing experiences," *OECD Journal on Development*, 9.

Oexl, R. and Z. J. Grossman (2013): "Shifting the blame to a powerless intermediary," *Experimental Economics*, 16, 306–312.

Peeters, R., M. Vorsatz, and M. Walzl (2015): "Beliefs and truth-telling: A laboratory experiment," *Journal of Economic Behavior & Organization*, 113, 1–12.

Rothenhäusler, D., N. Schweizer, and N. Szech (2018): "Guilt in voting and public good games," *European Economic Review*, 101, 664–681.

Sinnott-Armstrong, W. (1988): *Moral dilemmas*, Blackwell Publishers.

Wiltermuth, S. S. (2011): "Cheating more when the spoils are split," *Organizational Behavior and Human Decision Processes*, 115, 157–168.

Appendix

A Theory

A.1 Voting behavior when voters are unconditional

The basic model where all voters are unconditional and vote simultaneously follows Rothenhäusler, Schweizer, and Szech (2018). Each member of a group of $n \geq 1$ agents votes on a moral transgression (YES or NO), which yields benefit $\pi > 0$ for each group member. The voting threshold k is defined as the minimum number of YES votes required for the transgression, where $1 \leq k \leq n$.

When the transgression occurs, each group member who has voted YES faces moral costs which consist of two components:[21] The first component is purely "internal" and hence independent of the number of YES votes in the group (sensitivity parameter $\alpha \geq 0$). The second component decreases in the overall number of YES votes in the group, thereby capturing the motive of *guilt sharing*. In particular, we consider a function $g(Y_i)$ in combination with a sensitivity parameter $\beta \geq 0$, where Y_i is the number of YES votes in the group other than voter i, and where $g(0) = 1$ and $g'(\cdot) < 0$. Both of these two cost components are weighed with a voter's moral type $\theta_i \in [0, \theta^{\max}]$, which is voter i's private information. Types are distributed with continuous density $f(\theta)$ with full support on $[0, \theta^{\max}]$, where $f(\theta)$ is common knowledge. The utility of agent i from voting YES (u_Y^i) and NO (u_N^i), respectively, is then given by

$$u_Y^i := (\pi - \theta_i (\alpha + \beta g(Y_{-i}))) \cdot \Pr(Y_{-i} \geq k - 1) \quad (1)$$

$$u_N^i := \pi \cdot \Pr(Y_{-i} \geq k). \quad (2)$$

Theoretical prediction for Hypothesis 1 Rothenhäusler, Schweizer, and Szech (2018) consider symmetric Bayesian Nash equilibria, and they prove existence of an equilibrium in which all types $\theta_i \leq \widehat{\theta}_{k,n}$ vote for the moral transgression and all types $\theta_i > \widehat{\theta}_{k,n}$ against it. Indifferent types $\widehat{\theta}_{k,n}$ are characterized by the set of solutions to

$$\pi b(k-1; n-1, p) = \widehat{\theta}_{k,n} \left(\sum_{j=k-1}^{n-1} (\alpha + \beta g(1+j))\, b(j; n-1, p) \right), \quad (3)$$

[21] Rothenhäusler, Schweizer, and Szech (2018) also consider a cost of being pivotal (which accrues only in case the transgression would not have occurred without the respective voter's YES votes) as a third component. We drop this cost type as it has no impact on Proposition 1 and because, in all other propositions and hypotheses, we consider only pivotal voters anyway, so that preferences against being pivotal cannot influence the impact of voting thresholds.

where

$$b(j; n-1, p) = \binom{n-1}{j} p^j (1-p)^{n-1-j}, \text{ and} \tag{4}$$

$$p = \Pr(u_y \geq u_n). \tag{5}$$

This yields the following result:

Proposition 1 (All voters unconditional)**.** *The indifferent type $\widehat{\theta}_{k,n}$, and hence the equilibrium frequency of YES votes, increases in the voting threshold k.*

Proof. See Rothenhäusler, Schweizer, and Szech (2018).

The proposition directly leads to Hypothesis 1 as spelled out in Section 3 above, and which is tested in treatment *SIM*, where all votes are unconditional. For our purposes, it is crucial to note that the proposition also holds when there is no guilt sharing ($\beta = 0$) but only internal moral costs ($\alpha > 0$). The reason is that lower voting thresholds provide higher free-riding opportunities: The higher the likelihood that sufficiently many other group members vote YES, the higher is the incentive to save moral costs by voting NO. In equilibrium, this likelihood decreases in the threshold k and is zero for the unanimity rule ($k = n$).

A.2 The voting behavior of a pivotal voter

To eliminate such free-riding incentives, from now on we focus on pivotal votes. Consider a voter who knows that she is pivotal for the moral transgression (i.e. there are $k-1$ other YES votes). The utility of the pivotal voter when voting YES is then

$$u_Y^P := \pi - \theta_i \left(\alpha + \beta g(W_{-i})\right), \tag{6}$$

where W_{-i} is the number of YES votes by other voters in the group counting towards the moral transgression, so that the pivotal voter can share guilt with them (and where, due to pivotality, $W_{-i} = k-1$ holds).

Now consider the case where the pivotal voter votes NO. Note first that in this case she has no chance of receiving the gain π, i.e. there are no free-riding opportunities for pivotal voters. Moreover, we introduce an additional cost when voting NO which is due to a preference for *consensual voting*. As explained in the main text, the underlying idea is that individuals suffer a disutility when not acting in line with the (revealed) preferences of their group fellows. When voting NO, the pivotal voter incurs a disutility which increases with the preference of her group fellows in favor of the transgression. We capture this by a function $c(V_{-i})$ in combination with a the sensitivity parameter $\phi \geq 0$, where V_{-i} denotes the number of YES votes by the group members other than i, and where $c(0) = 0$ and $c'(\cdot) > 0$.[22] In summary, when voting NO, the

[22]The assumption that the costs $c(V_{-i}) > 0$ arise only when voting against the transgression is without loss of generality. An equivalent approach would be to introduce a cost when voting against the stated preferences of

utility of a pivotal voter is given by

$$u_N^P := -\theta_i \phi c(V_{-i}), \tag{7}$$

where also $V_{-i} = k-1$ due to pivotality.

Note the subtle difference between W_{-i} and V_{-i}: In (6), W_{-i} is the number of group members other than i who vote YES *and* whose votes count for the moral transgression. Hence, W_{-i} is the relevant measure with respect to guilt sharing. By contrast, in (7), V_{-i} is the number of group members other than i who vote YES, irrespective of whether their votes actually count. Hence, V_{-i} is the relevant measure with respect to consensual voting. Recall that, in the experiment, in treatments *P-BASE* and *P-INFO* the decision concerning the moral transgression is based on the all three votes in the group, so that $W_{-i} = V_{-i}$ holds. By contrast, in treatment *P-RAND* the unconditional votes do not count for the decision concerning the transgression, i.e $W_{-i} \neq V_{-i}$.

Setting $u_Y^P = u_N^P$ then leads to the indifferent moral type of the pivotal voter, denoted by $\widehat{\theta}_k^P$, i.e.

$$\widehat{\theta}_k^P = \frac{\pi}{\alpha + \beta g(W_{-i}) - \phi c(V_{-i})}. \tag{8}$$

Theoretical prediction for Hypotheses 2 and 3 As explained above, treatments *P-BASE* and *P-INFO* differ only with respect to whether or not conditional voters receive information about the behavior of a population of unconditional voters outside their treatment which is not modeled here. Hence, the underlying proposition for the behavior of pivotal conditional voters is the same for both hypotheses. Taking into account that $W_{-i} = V_{-i} = k-1$, this leads to

Proposition 2 (Pivotal votes in treatments *P-BASE* and *P-INFO*)**.** *The indifferent moral type of the pivotal voter $\widehat{\theta}_k^P$, and hence the equilibrium frequency of YES votes, increases in the voting threshold k.*

Proof. Equating u_Y^P and u_N^P, we get the indifferent type as:

$$\widehat{\theta}_k^P = \frac{\pi}{\alpha + \beta g(k-1) - \phi c(k-1)}, \tag{9}$$

where

$$\frac{\partial \widehat{\theta}_k^P}{\partial k} = \frac{\pi(\phi c'(k-1) - \beta g'(k-1))}{[\alpha + \beta g(k-1) - \phi c(k-1)]^2} > 0 \tag{10}$$

as $g'(k) < 0$ and $c'(k) > 0$. ■

Intuitively, in case of pivotality, a higher voting threshold k also implies a higher number of other YES votes $k-1$. This reduces the moral costs of a YES vote for the pivotal voter due to guilt sharing. In addition, a higher value of $k-1$ YES votes indicates a stronger preference for the moral transgression among the other group members, so that the cost of voting against

one's fellow group members. All we need is that the cost difference between a vote for and against the transgression increases in the number of YES votes in the group.

it increases as well (consensual voting). Both effects make a YES vote more attractive for the pivotal voter.

Theoretical prediction for Hypothesis 4 This hypothesis applies to conditional voters in treatment *P-RAND* where both the number of random and unconditional YES votes is equal to $k-1$. Since only the former are relevant for the transgression, this mutes guilt sharing (i.e. $W_{-i}=0$) while the motive of consensual voting is still present (i.e. $V_{-i}=k-1$). This leads to

Proposition 3 (Pivotal votes in treatment *P-RAND*)**.** *Suppose $\phi > 0$. Then the indifferent moral type of the pivotal voter $\widehat{\theta}_k^P$, and hence the equilibrium frequency of YES votes, increases in the voting threshold k.*

Proof. Without the possibility of guilt sharing, we have $g(W_{-i}=0)=1$. The threshold type who is indifferent between voting YES or NO is given by:

$$\widehat{\theta}_k^P = \frac{\pi}{\alpha+\beta-\phi c(k-1)}, \tag{11}$$

where

$$\frac{\partial \widehat{\theta}_k^P}{\partial k} = \frac{\pi \phi c'(k-1)}{[\alpha+\beta-\phi c(k-1)]^2} > 0 \tag{12}$$

as $c'(V_{-i}) > 0$. ■

Intuitively, as $V_{-i}=k-1$ increases in k, so does the cost of voting NO due to consensual voting. Therefore, the number of pivotal conditional YES votes increases in the voting threshold even when guilt sharing is not present.

Theoretical prediction for Hypothesis 5 The impact of guilt sharing can be gauged by comparing the behavior of pivotal voters in treatments *P-INFO* and *P-RAND*. Denoting the indifferent types under *P-INFO* and *P-RAND* by $\widehat{\theta}_{1,k}^P$ and $\widehat{\theta}_{2,k}^P$, respectively, and defining the difference $\Delta\widehat{\theta}_k^P := \widehat{\theta}_{1,k}^P - \widehat{\theta}_{2,k}^P$, this leads to

Proposition 4 (Impact of guilt sharing)**.** *Suppose $\beta > 0$. Then, the indifferent moral type of the pivotal voter (and thus the equilibrium frequency of YES votes) is higher in treatment* P-INFO *than in* P-RAND *(i.e. $\Delta\widehat{\theta}_k^P \geq 0$). For $k=1$, there is no treatment difference ($\Delta\widehat{\theta}_1^P = 0$). For $k>1$, $\Delta\widehat{\theta}_k^P$ is strictly positive and increasing in k.*

Proof. Without the possibility for guilt sharing in treatment *P-RAND*, we have $Y_{-i}=0$ and $g(Y_{-i})=1$. Similarly, in treatment *P-INFO*, we have $Y_{-i}=k-1$ and $g(Y_{-i}) \leq 1$. The respective threshold types who are indifferent between voting YES or NO are then given by:

$$\widehat{\theta}_{1,k}^P = \frac{\pi}{\alpha+\beta g(k-1)-\phi c(k-1)} \tag{13}$$

and

$$\widehat{\theta}_{2,k}^P = \frac{\pi}{\alpha+\beta-\phi c(k-1)}, \tag{14}$$

where

$$\frac{\partial(\widehat{\theta}^P_{1,k} - \widehat{\theta}^P_{2,k})}{\partial k} = \pi \frac{\phi c'(k-1) - \beta g'(k-1)}{[\alpha + \beta g(k-1) - \phi c(k-1)]^2} - \pi \frac{\phi c'(k-1)}{[\alpha + \beta - \phi c(k-1)]^2} > 0 \tag{15}$$

as the numerator of the first fraction exceeds the numerator of the second fraction (because $c'(k-1) > 0$ and $g'(k-1) < 0$) and the denominator of the first fraction is (weakly) smaller than the denominator of the second fraction (because $g(k-1) \leq 1$). Note that $\widehat{\theta}^P_{1,k} = \widehat{\theta}^P_{2,k}$ if and only if $k = 1$. ■

Theoretical prediction for Hypothesis 6 The hypothesis concerns the impact of consensual votingand is again based on treatment *P-RAND* where guilt sharing is absent (i.e. $W_{-i} = 0$). However, we no longer fix $V_{-i} = k - 1$ but analyze, separately for each voting threshold k, how the behavior of pivotal voters varies in V_{-i}. This leads to

Proposition 5 (Impact of consensual voting)**.** *Suppose $W_{-i} = 0$ and $\phi > 0$. Then, for a given voting threshold k, the indifferent moral type of the pivotal voter $\widehat{\theta}^P_k$, and hence the equilibrium frequency of YES votes, increases in the number of unconditional YES V_{-i} in her group.*

Proof. With $W_{-i} = 0$, we have $g(0) = 1$. The threshold type who is indifferent between voting YES or NO is given by:

$$\widehat{\theta}^P_k = \frac{\pi}{\alpha + \beta - \phi c(V_{-i})}, \tag{16}$$

where

$$\frac{\partial \widehat{\theta}^P_k}{\partial V_{-i}} = \frac{\pi \phi c'(V_{-i})}{[\alpha + \beta - \phi c(V_{-i})]^2} > 0 \tag{17}$$

as $c'(V_{-i}) > 0$. ■

Intuitively, although the unconditional YES votes in a group do not count towards the moral transgression, they still matter for a (conditional) pivotal voter because of consensual voting. This effect size increases with the number of unconditional YES votes in a group.

Theoretical prediction for Hypothesis 7 According to this final hypothesis, other than for the ones considered so far, there are no further motives giving rise to threshold effects. The hypothesis is also tested in treatment *P-RAND* where, in addition to the absence of guilt sharing in this treatment, we also take out consensual voting by separately considering each value of V_{-i}. This leads to

Proposition 6 (No further motives)**.** *Suppose $W_{-i} = 0$ and $\phi > 0$. Then, for a given number of unconditional YES votes, $\overline{V}_{-i}$, the indifferent moral type of the pivotal voter $\widehat{\theta}^P_k$ is independent of the voting threshold k.*

Proof. Again, without the possibility of guilt sharing, we have $g(W_{-i} = 0) = 1$. The threshold type who is indifferent between voting YES or NO is given by:

$$\widehat{\theta}_k^P = \frac{\pi}{\alpha + \beta - \phi c(\overline{V}_{-i})}, \tag{18}$$

where $\overline{V}_{-i}$ is constant. Thus, $\frac{\partial \widehat{\theta}_k^P}{\partial k} = 0$. ■

B Robustness

In this appendix, we provide additional robustness checks for our results: First, we analyze whether threshold effects also arise in a non-parametric approach. Second, we check the robustness of the regressions results by considering a richer set of controls from the post-experimental questionnaire. Overall, these robustness checks do not lead to qualitatively different results than those presented in the main analysis.

Non-parametric approach Table 4 reports the p-values of the χ^2 tests for the comparison of the different thresholds in each treatment. This gives rise to results which are qualitatively very similar to those reported in Table 2. In particular, in each treatment there are significantly more YES votes under unanimity voting compared to dictatorial voting and also some of the other comparisons yield significant differences.

Table 4: Threshold effects in a non-parametric approach

	Dictatorial vs. Majority	Dictatorial vs. Unanimity	Majority vs. Unanimity
SIM	0.326	0.029**	0.228
P-BASE	0.038**	< 0.001***	0.007***
P-INFO	0.068*	< 0.001***	0.007***
P-RAND	0.301	0.009***	0.111

Notes: The entries in the table report the p-values for the respective χ^2 tests for differences of YES votes under the different thresholds in each treatment. As in the main text, in treatment *SIM* all votes are unconditional, and all of them are used in the analysis. In all other treatments, only pivotal conditional YES votes are used.

Additional controls Tables 5 and 6 replicate the regression analysis from the main text (Tables 2 and 3) but control in addition for the subjects' belief about the overall percentage of YES votes in the experiment and the degree of trust in the charity, measured on a ten-point scale. Both variables were elicited after subjects had taken their voting decisions.

Comparing Tables 5 and 2 shows that the results for the threshold effects are quite similar. In total, there are three cases where coefficients become insignificant when using the richer set of controls: (i) the comparison of unanimity and dictatorial voting in *SIM* (see column (1)),

Table 5: Robustness: Regressions results for the tests of Hypotheses 1 to 5

	(1) *SIM* (H1)	(2) *P-BASE* (H2)	(3) *P-INFO* (H3)	(4) *P-RAND* (H4)	(5) *P-INFO + P-RAND* (H5)
Majority	0.0561	0.159***	0.101*	0.0736	0.0975
	(0.285)	(0.004)	(0.092)	(0.199)	(0.102)
Unanimity	0.0410	0.312***	0.249***	0.142**	0.248***
	(0.387)	(0.000)	(0.000)	(0.016)	(0.000)
P-RAND					-0.0179
					(0.762)
Majority x P-RAND					-0.0265
					(0.748)
Unanimity x P-RAND					-0.110
					(0.171)
Belief	0.00866***	0.00820***	0.00783***	0.00801***	0.00797***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Trust	-0.0412***	-0.0273***	-0.0368***	-0.0516***	-0.0454***
	(0.000)	(0.003)	(0.000)	(0.000)	(0.000)
Observations	360	360	360	360	720
Adjusted R^2	0.282	0.307	0.167	0.211	0.195
T1: Una. vs. Majo.	0.750	0.002	0.007	0.240	
T2: (Una. + Majo.) vs. Dict.	0.270	0.000	0.001	0.032	
T3: Treatment effect Majo.					0.447
T4: Treatment effect Una.					0.019
Controls Pers. Charact.	Yes	Yes	Yes	Yes	Yes
Session Dummies	Yes	Yes	Yes	Yes	No

Notes: The structure of the regressions is the same as in Table 2 in the main text, so that also the explanation below Table 2 applies. The only difference is that the regressions reported here include additional controls for the reported beliefs about other participants' voting behavior (*Belief*) and the degree of trust in the charity (*Trust*).

(ii) the comparison of majority and dictatorial voting in *P-INFO* (see column (5)), and (iii) the interaction term for the treatment comparison between *P-INFO* and *P-RAND* (see column (5)). The comparison of Tables 6 and 3 (i.e. the test of Hypothesis 6 and Hypothesis 7) reveals that the two additional controls do not lead to any qualitative differences. In all regressions reported in Tables 5 and 6, *Belief* and *Trust* are highly significant in plausible directions. As these variables were elicited after the voting decisions were taken, the answers may well be influenced by rationalizations of the own behavior. It is hence not surprising that some of the effects become insignificant after controlling for *Belief* and *Trust*.

C Experimental instructions

Here, we provide the instructions for the exemplary case of majority voting ($k = 2$).

C.1 Introduction: all treatments

The instructions for all treatments started with the following text.

> Thank you for participating in this study. Please read the instructions carefully!
>
> You get \$0.80 for your participation and have the possibility to earn additional money.
>
> You are randomly paired with two other workers to form a **group of three**. Each member of your group votes separately on whether to take \$1.50 originally designated

Table 6: Robustness: Regressions results for the tests of Hypotheses 6 and 7

	(1)	(2)	(3)	(4)	(5)	(6)
	(H6)	(H6)	(H6)	(H7)	(H7)	(H7)
	Dictatorial	Majority	Unanimity	GroupYes=0	GroupYes=1	GroupYes=2
GroupYes=1	0.1000**	0.125***	0.0333			
	(0.025)	(0.003)	(0.164)			
GroupYes=2	0.150***	0.225***	0.142***			
	(0.001)	(0.000)	(0.001)			
Majority				-0.0485	-0.0293	0.0217
				(0.409)	(0.613)	(0.711)
Unanimity				0.0117	-0.0736	0.00460
				(0.846)	(0.219)	(0.938)
Belief	0.00720***	0.00419**	0.00716***	0.00573***	0.00778***	0.00570***
	(0.000)	(0.046)	(0.000)	(0.000)	(0.000)	(0.000)
Trust	-0.0490***	-0.0550***	-0.0316*	-0.0593***	-0.0433***	-0.0395***
	(0.000)	(0.000)	(0.054)	(0.000)	(0.000)	(0.000)
Observations	360	360	360	360	360	360
Adjusted R^2	0.187	0.197	0.193	0.169	0.167	0.107
T1: Both coeff. equal	0.114	0.003	0.005	0.321	0.453	0.772
T2: Sum both coeff. zero	0.004	0.000	0.003	0.718	0.312	0.796
Controls for Pers. Charact.	Yes	Yes	Yes	Yes	Yes	Yes
Session Dummies	Yes	Yes	Yes	Yes	Yes	Yes
Clust. Subj. Level	Yes	Yes	Yes	No	No	No

Notes: The structure of the regressions is the same as in Table 3 in the main text, so that also the explanation below Table 3 applies. The only difference is that the regressions reported here include additional controls for the reported beliefs about other participants' voting behavior (*Belief*) and the degree of trust in the charity (*Trust*).

by us for donation to the Make-A-Wish Foundation, a charity that grants wishes to children with a critical illness.

If there **are at least two votes** in favor of taking the money, then your group gets the $1.50 designated for donation (each of you gets $0.50 as a bonus payment). The charity Make-A-Wish gets nothing.

If there **are less than two votes** for taking the money, then your group does not get the $1.50. The money will instead be donated to Make-A-Wish.

C.2 Treatment *SIM*

After answering some control questions, participants in treatment *SIM* were then simply asked for their vote.

[Control questions]

Now it is your decision: What do you vote for?

- Take the money
- Do not take the money

C.3 Unconditional voters in treatments *P-BASE*, *P-INFO*, and *P-RAND*

All unconditional voters in treatments *P-BASE*, *P-INFO*, and *P-RAND* received the same instructions. After the introduction, their instructions continued as follows.

The timing is as follows: First, we ask you and one of the two other workers you are grouped with to vote independently from each other for or against taking the money. After the two of you have voted, we are going to ask the third worker for their vote. You then get the bonus payment if there are **at least two votes** for taking the money.

However, it may be that your actual votes do not matter for the transfer of the money. In some cases, your votes are instead substituted by the votes of two other randomly chosen workers from a similar study. The third worker will then be informed that these randomly chosen votes matter for the final outcome. The third worker's vote always counts and is never substituted.

As you do not know whether your vote actually counts for your own bonus payment or not you should make your decision assuming that it does count.

[Control questions]

Now it is your decision: What do you vote for?

- Take the money
- Do not take the money

C.4 Conditional voters in treatments *P-BASE* and *P-INFO*

After the introduction, conditional voters in treatments *P-BASE* and *P-INFO* received the same instructions except for the first paragraph, which was only shown to conditional voters in treatment *P-INFO*. Instead of just voting on whether or not to take the money, conditional voters in these treatments were asked for their vote depending on the votes of their group members.

[Only treatment *P-INFO* also included the following sentence:] In a similar study, **65%** of all workers who voted first voted for taking the money.

We have already collected the votes of the two workers you are grouped with. No one, one, or two of your group members may have voted for taking the money. We will ask for your decision for each of these three possibilities.

[Control questions]

Now it is your decision. Please vote for the following situations:

None of the workers in your group have voted for taking the money.

- Take the money
- Do not take the money

Exactly one of the workers in your group has voted for taking the money.

- Take the money

- Do not take the money

Both of the workers in your group have voted for taking the money.

- Take the money
- Do not take the money

C.5 Conditional voters in treatment *P-RAND*

After the introduction, conditional voters in treatment *P-RAND* received the following instructions. Instead of just voting on whether or not to take the money, conditional voters in this treatment were asked for their vote depending on the votes of their group members and the outcome of the random draws.

> We have already collected the votes of the two workers you are grouped with. No one, one, or both may have voted for taking the money. In a similar study, **65%** of all workers who voted first voted for taking the money.
>
> However, you were randomly assigned to a situation where the votes of the two workers you are grouped with do **NOT** matter for whether your group gets the money. Instead, the votes of your two group members are substituted by random draws. Each random draw has a probability of 65% for taking the money (as in the similar study).
>
> We will ask for your decision for all possible combinations of these random draws (which count for the outcome) and the actual votes of your group members (which do not count).
>
> [Control questions]
>
> We now ask you for your vote in the following situations:
>
> Suppose **none** of the two random draws were for taking the money. Please vote for the following situations:
>
> - **None** of the other two workers in your group (and whose votes do not count for the final outcome) have voted for taking the money.
>   - Take the money
>   - Do not take the money
> - **Exactly one** of the other two workers in your group (and whose votes do not count for the final outcome) has voted for taking the money.
>   - Take the money
>   - Do not take the money
> - **Both** of the other two workers in your group (and whose votes do not count for the final outcome) have voted for taking the money.

  - Take the money
  - Do not take the money

Suppose **exactly one** of the two random draws was for taking the money. Thus, the money will be taken only if you vote for it. Please vote for the following situations:

- **None** of the other two workers in your group (and whose votes do not count for the final outcome) have voted for taking the money.
  - Take the money
  - Do not take the money
- **Exactly one** of the other two workers in your group (and whose votes do not count for the final outcome) has voted for taking the money.
  - Take the money
  - Do not take the money
- **Both** of the other two workers in your group (and whose votes do not count for the final outcome) have voted for taking the money.
  - Take the money
  - Do not take the money

Suppose **both** of the two random draws were for taking the money. Please vote for the following situations:

- **None** of the other two workers in your group (and whose votes do not count for the final outcome) have voted for taking the money.
  - Take the money
  - Do not take the money
- **Exactly one** of the other two workers in your group (and whose votes do not count for the final outcome) has voted for taking the money.
  - Take the money
  - Do not take the money
- **Both** of the other two workers in your group (and whose votes do not count for the final outcome) have voted for taking the money.
  - Take the money
  - Do not take the money