



**University of
Zurich** ^{UZH}

University of Zurich
Department of Economics

Working Paper Series

ISSN 1664-7041 (print)
ISSN 1664-705X (online)

Working Paper No. 263

The Limits to Moral Erosion in Markets: Social Norms and the Replacement Excuse

Björn Bartling and Yagiz Özdemir

September 2017

The Limits to Moral Erosion in Markets: Social Norms and the Replacement Excuse*

Björn Bartling
University of Zurich

Yagiz Özdemir
University of Zurich

September 27, 2017

This paper studies the impact of a key feature of competitive markets on moral behavior: the possibility that a competitor will step in and conclude the deal if a conscientious market actor forgoes a profitable business opportunity for ethical reasons. We study experimentally whether people employ the argument “if I don’t do it, someone else will” to justify taking a narrowly self-interested action. Our data reveal a clear pattern. Subjects do not employ the “replacement excuse” if a social norm exists that classifies the selfish action as immoral. But if no social norm exists, subjects are more inclined to take a selfish action in situations where another subject can otherwise take it. By demonstrating the importance of social norms of moral behavior for limiting the power of the replacement excuse, our paper informs the long-standing debate on the effect of markets on morals.

Keywords: Replacement excuse, social norms, moral behavior, competition, markets, utilitarianism, deontological ethics

JEL classification: C92, D02, D63

* Corresponding author: bjoern.bartling@econ.uzh.ch, Department of Economics, Blümlisalpstrasse 10, 8006 Zurich, Switzerland. We thank Alexander Cappelen, Ernst Fehr, John List, Bertil Tungodden, Lise Vesterlund, and conference and seminar participants at Chicago, Dijon, Edinburgh, Heidelberg, Geneva, Karlsruhe, Lyon, Munich, NHH Bergen, San Diego, St Andrews, Stanford, Stavanger, Vienna, and Zurich for helpful comments. We are grateful to the Choice Lab at NHH Bergen for granting access to their AMT account and to Som Datye for excellent research assistance.

“If we want to stop the defence industry operating in this country we can do so, and the result incidentally will be that someone else supplies the arms that we supply.”—Tony Blair¹

1. Introduction

One of the founders of the Chicago school, Frank H. Knight, wrote in an early paper that “it must be conceded that the lines along which a competitive economic order tends to form character are often far from being ethically ideal” (1923, p. 591). The possibility that competitive markets erode morals has been debated ever since (for recent contributions see, e.g., Bowles 1998, Shleifer 2004, Satz 2010, Sandel 2012, Falk and Szech 2013, Bartling et al. 2015, Pigors and Rockenbach 2016, Kirchler et al. 2016). Data on the causal effect of markets on morals and, in particular, on the mechanisms by which markets might affect morals are however still scarce.

This paper studies the impact of a key feature of competitive markets on moral behavior: the possibility that a competitor will rush in and cut the deal if a conscientious market actor forgoes a profitable business opportunity for ethical reasons. Suppose the conscientious market actor considers refusing to sell arms to a dictatorial regime. The only consequence of the refusal might be that a less scrupulous competitor will grab the chance; hence, the regime still receives the arms and the profits are passed to the competitor. In anticipation of this outcome, the conscientious market actor might sell the arms himself. The *replacement excuse*, that is, the argument that “if I don’t do it, someone else will,” has intuitive appeal and might lead even the most conscientious market actors to ignore their moral standards.

The two main normative ethical theories disagree fundamentally on the justification the replacement excuse offers. From the viewpoint of consequentialist or utilitarian ethics, the replacement excuse provides justification for any action because outcomes alone matter for ethical assessment. Given that an unethical outcome will come about anyway (say, a dictatorial regime obtains new arms), taking the action that effectively implements the outcome (supplying the arms) does not change or worsen the outcome. The action is therefore not ethically wrong. Tony Blair, for instance, then UK prime minister, sought to justify an increase in UK arms exports in the introductory quote by claiming that a competitor would have stepped in had the UK refused to supply the arms. In contrast, deontological or duty ethics judges an action with respect to its adherence to a rule or principle. The fact that an unethical outcome will come about anyway does

¹ Tony Blair, July 25, 2002: <http://webarchive.nationalarchives.gov.uk/20100105053147/number10.gov.uk/page3000>

not render the action that effectively implements the outcome ethically right: “If we accept this as a justification, it is hard to see what acts, however otherwise wicked, could not be defended in the same way” (Glover and Scott-Taggart 1975).

This paper takes a positive approach and studies experimentally how people *actually* make decisions in situations where the replacement excuse is available. Do only outcomes matter for a market actor’s utility (utilitarianism) or do also his actions matter for his utility (deontological ethics)? If outcomes alone mattered for a market actor’s utility, he would always prefer to carry out a profitable but unethical business transaction himself rather than let someone else conclude the deal. Since the transaction will be implemented one way or another, own profits would be higher. This would imply that competitive markets erode moral standards by providing market actors with the replacement excuse. But if a market actor incurs a utility loss by taking an action that does not conform to the prevailing social norms of moral behavior (e.g., Elster 1989, Lindbeck 1997, López-Pérez 2008, Krupka and Weber 2013), then he might not conclude the unethical transaction—even if forgoing it will make no difference and profits are passed on to a competitor.² This would imply that competitive markets do not necessarily undermine morals, at least not by way of the replacement excuse.³

Our experimental design consists of a series of games, where players make binary choices between realizing a monetary gain (selfish outcome) and forgoing it for an altruistic reason (pro-social outcome). In each game, only one player can realize the monetary gain, and the pro-social outcome will only materialize if all players, who decide sequentially, forgo the monetary gain. Our treatments vary the number of players who can realize the monetary gain if the first player in the sequence decides to forgo it. That is, we exogenously vary the likelihood that one of the subsequent players’ selfish choice will undo the first player’s pro-social choice. Moreover, our games differ in the prevailing social norms. While our experimental measures show that the implementation of the selfish outcome is socially inappropriate and inconsistent with moral or proper social behavior in some of our games, no such norm exists in others.

² The source of the utility loss can be self-image or identity concerns (e.g., Akerlof and Kranton 2000, 2005; Bodner and Prelec 2003; Bénabou and Tirole 2004, 2006, 2011; Köszegi 2006, Mazar et al. 2008), social image concerns (e.g., Akerlof 1980, Bernheim 1994, Andreoni and Bernheim 2009, Ariely et al. 2009), or the “cold prickle,” rather than the “warm glow,” of taking an action that does not conform to the prevailing social norms (Andreoni 1989, 1990, 1995).

³ By studying the replacement excuse, our paper is related to a recent strand of papers on excuse driven behavior. The literature documents, for instance, that people employ self-serving interpretations of ambiguity and risk (Haisley and Weber 2010, Exley 2015), avoid information and exposure to sharing opportunities (Dana et al. 2006, 2007; Lazear et al. 2012, Andreoni et al. 2016), or use time to develop own excuses (Exley and Petrie 2016).

Our experimental design allows studying two questions. First, do players act more selfishly if they can employ the replacement excuse? That is, does the existence of subsequent players who could undo a first player's altruistic choice promote selfish behavior by the first player? This would always be the case if people's choices were guided by utilitarian ethics. The reason is that the opportunity cost of taking the pro-social choice (the monetary gain that is realized by implementing the selfish outcome) remains constant, while there is no guarantee that the pro-social outcome will ensue if a subsequent player's selfish choice can undo it.

Second, does the impact of the replacement excuse depend on the existence of a social norm of moral behavior? This would be the case if deontological ethics guided peoples' choices, but not if people were utilitarian.⁴ If the implementation of the selfish outcome is consistent with the prevailing social norm of moral behavior, no norm is broken if the selfish action is taken; hence, there is no utility loss from taking the selfish action, and the predictions for deontological and utilitarian players coincide. If, however, the implementation of the selfish outcome is inconsistent with the prevailing social norm of moral behavior, taking the selfish action constitutes a norm deviation. Importantly, pointing to the fact that the selfish outcome is likely to ensue anyway does *not* justify the norm deviation from the viewpoint of deontological ethics. Consequently, a deontological player will not implement the selfish outcome simply because a subsequent player is likely to do so.

Studying whether competitive markets and the replacement excuse lead people to behave in narrowly self-interested ways is of increasing importance because the "reach of markets, and market-oriented thinking, into aspects of life traditionally governed by nonmarket norms is one of the most significant developments of our time" (Sandel 2012, p. 7). Although a single ruthless competitor, or latent market entry of such a competitor, suffices to evoke the replacement excuse, behavior consistent with ethical conduct can be observed in many competitive markets. Financial services firms that are committed to ethical investment strategies—forgoing possibly more profitable investments in, say, arms manufacturing or tobacco—can serve as examples (e.g., Sparkes and Cowton 2004). With naturally occurring field data it is, however, difficult to separate true corporate social concern from reputational incentives or concerns about consumer demand

⁴ Note that social norms relate to actions, not outcomes. Elster (1989), for example, defines "social norm by the feature that they are *not outcome-oriented*. The simplest social norms are of the type: Do X, or: Don't do X" (p. 99). Similarly, Ostrom (2000) defines social norms as "shared understandings about *actions* that are obligatory, permitted, or forbidden" (p. 143-144, emphasis added).

(e.g., Aupperle et al. 1985, Porter and Kramer 2006, Riedl and Smeets 2017). Ethical business practices can thus be compatible with the goal of profit-maximization, which renders it difficult to study when the replacement excuse gives rise to narrowly self-interested behaviors in the field and when it does not. In contrast, while being stylized and subject to concerns regarding external validity, our experimental approach allows us to control for confounding factors and to study the causal effect of the availability of the replacement excuse on behavior in a clean way.

The remainder of the paper is organized as follows. Sections 2 to 4 lay out the details of our experimental games and present the results. We draw conclusions in Section 5.

2. Donation Game

2.1 Experimental Design

We study the effect of the availability of the replacement excuse on behavior in a simple donation game. Subjects make binary choices between receiving a payment of 20 CHF (about 21 USD) and having the experimenter transfer 60 CHF to FAIRMED, a Swiss charity, to finance the surgery of a leprosy patient in India.⁵

We conducted two experimental conditions of the donation game, as shown in Figure 1. The *baseline condition* is an individual decision task. A subject (player 1) decides to either take 20 CHF (selfish outcome) or having the experimenter finance the surgery (altruistic outcome). The *replacement condition* is a three-player game. Subjects are randomly placed into groups of three and assigned the role of either player 1, 2, or 3. At most one player can take the 20 CHF and at most one surgery is financed per group. Players decide sequentially whether to take the 20 CHF. If a player decides to take the 20 CHF, the surgery is not financed and the game ends. The surgery is financed only if first player 1, then player 2, and finally player 3 forgoes taking the 20 CHF.

The key feature of the experimental design is that player 1's decision to forgo the 20 CHF cannot be replaced in the *baseline condition*, while we endow player 1 with the replacement excuse in the *replacement condition*. In the latter condition, player 1 can argue that if he does not take the 20 CHF, player 2 or player 3 will do so. The design thus allows us to isolate the causal effect of

⁵The experimental instructions provided basic facts about leprosy and the use of the donation. Leprosy is an infectious disease that causes damage of nerve cells and blockage of blood vessels, which can lead to bodily disfigurement. Although the disease can be cured, many leprosy victims suffer from stigmatization due to disfigurement. Often small surgical interventions can significantly reduce the scope of disfigurement. Almost 60 percent of the global leprosy cases occur in India. Funding an operation, which costs about 60 CHF on average and would allow for a life in dignity, is not possible for most of the victims with disfigurement due to the prevalence of poverty.

the availability of the replacement excuse by comparing the take rate (the fraction of players 1 who take the money) in the *baseline condition* to the take rate in the *replacement condition*.

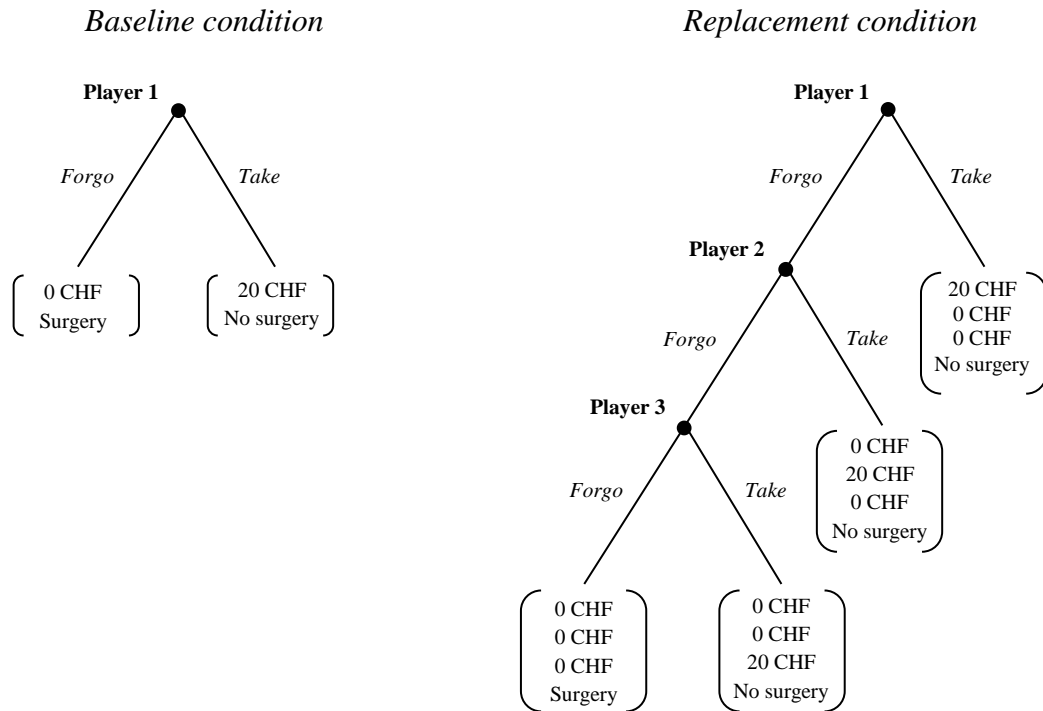


FIGURE 1. DONATION GAME

To check the effectiveness of our experimental manipulation we elicited beliefs about the replacement probability (the probability that either player 2 or 3 will take the money if player 1 does not do so) in the *replacement condition*. We implemented a between-subjects design to avoid the potential confound of ex-post rationalization (e.g., a subject stating a high replacement probability to justify having taken the money). The replacement probability is zero, by design, in the *baseline condition*. Our experimental manipulation is thus effective if subjects expect a positive replacement probability in the *replacement condition*. Subjects in the belief elicitation sessions received the instructions and control questions for the *replacement condition*. But instead of asking the subjects to choose between money and surgery, we asked them to state their beliefs about the replacement probability. Subjects had to enter integer numbers between 0 and 100 to indicate their beliefs in form of a percentage point and they earned 10 CHF if the stated belief was no further away than 5 percentage points from the true value.

Moreover, we measured the social norm regarding the decision to take the money, using a coordination game as measurement tool (Krupka and Weber 2013). We implemented a between-subjects design and asked subjects, either for the *baseline condition* or for the *replacement condition*, whether player 1's choice to take the 20 CHF would be rated by most people as "socially appropriate and consistent with moral or proper social behavior" or as "socially inappropriate and inconsistent with moral or proper social behavior" or as "neutral." Subjects had to indicate their beliefs by choosing one of five options: (i) very socially inappropriate, (ii) somewhat socially inappropriate, (iii) neutral: neither socially inappropriate nor appropriate, (iv) somewhat socially appropriate, or (v) very socially appropriate. Importantly, we did not ask subjects to provide the rating they believe to be right, but the rating they believe the other subjects in the session most frequently chose. The social norm is assumed to serve as the coordination device in this game. Subjects received the instructions and control questions for the respective condition of the donation game. But instead of asking the subjects to choose between money and surgery, we asked the subjects to provide their beliefs about the most frequently given response. Subjects received a bonus of 10 CHF if their answers matched the modal response.

2.2 Procedural Details

All sessions took place at the laboratory of the Department of Economics at the University of Zurich. We implemented the study with z-Tree and h-Root (Fischbacher 2007; Block et al. 2014). Subjects were students from the University of Zurich and the ETH Zurich. 432 subjects participated in total. We conducted eight sessions of the donation game, with 67 subjects in the *baseline condition* and 177 subjects, i.e., 59 in each role, in the *replacement condition*. We used the strategy method for players 2 and 3 in the *replacement condition*. Sessions lasted about 45 minutes and subjects earned either 15 CHF (the show-up fee) or 35 CHF (if they took the 20 CHF); average earnings were CHF 19.75. Subjects' choices resulted in 68 leprosy surgeries being financed, i.e., in a donation of 4080 CHF to FAIRMED. We elicited beliefs in the *replacement condition* from 64 subjects. 124 subjects participated in the measurement of the social norm; 62 in each of the two conditions. Subjects earned either 10 CHF (if their answers qualified for the bonus) or 0 CHF; average earnings were 4.10 CHF. These measurements took about 15 minutes and were added at the end of unrelated experiments. No subject participated in our experiments more than once. The experimental instructions for all games and measurements are in the online appendix.

2.3 Predictions

The behavioral predictions depend on the assumptions about subjects' moral motivations. If subjects are guided by utilitarian ethics, then we should expect that the take rate is higher in the *replacement condition* than in the *baseline condition*. But if deontological ethics guide subjects, and if taking the 20 CHF constitutes a norm violation, then the take rate will not necessarily be higher in the *replacement condition* than in the *baseline condition*.

To see this more formally, suppose first that subjects are utilitarian, so that utility is defined solely over outcomes. Let d and m denote the pro-social outcome (donation) and selfish outcome (money), respectively. A subject forgoes taking the money in the *baseline condition* if $U(d) > U(m)$. Let p denote the (belief about the) replacement probability in the *replacement condition*. A subject then takes the money if $(1 - p) \cdot U(d) < U(m) \Leftrightarrow p > (U(d) - U(m))/U(d)$, i.e., if p is sufficiently large. Subjects with only a slight preference for the donation over taking the money in the *baseline condition* are thus likely to take the money instead in the *replacement condition*.

Suppose next that subjects are deontological, so that utility is defined over outcomes and actions. Suppose further that taking the money in the donation game is socially inappropriate and inconsistent with moral behavior. Let c denote the utility cost incurred by deviating from a social norm of moral behavior. A subject in the *baseline condition* then chooses the donation and not the money if $U(d) > U(m) - c$ and a subject in the *replacement condition* forgoes taking the money if $(1 - p) \cdot U(d) > U(m) - c$. Consequently, if the utility cost of deviating from the applicable social norm of moral behavior is sufficiently large, $c > U(m)$, a subject will not even take the money if $p = 1$ in the *replacement condition*.

2.4 Results

The grey bars in Figure 2 show that the take rates of players 1 are almost identical in the *baseline condition* and in the *replacement condition*, 23.9 percent (16 of 67) and 25.4 percent (15 of 59), respectively ($p=0.50$, Fischer exact test, one-sided). We observe these take rates despite the fact the replacement probability increases from 0, by design, in the *baseline condition* to 0.51 in the *replacement condition*, indicated by the black diamonds in Figure 2. Also the take rates of players 2 and 3 in the *replacement condition* are very similar, 30.5 percent (18 of 59) and 28.8 percent (17 of 59), respectively. We cannot reject that the take rate is identical across both conditions and all types of players (Pearson's chi squared test, $p=0.83$). Note that players 2 could invoke the

replacement excuse as well, as their choice to forgo the 20 CHF might be replaced by player 3. In contrast, players 3 in the *replacement condition* are in the same strategic situation as players 1 in the *baseline condition* as their choice to forgo taking the money cannot be replaced. We summarize these findings in our first result.

Result 1: *Subjects do not use the replacement excuse in the donation game. The possibility that subsequent players can replace an earlier player’s decision not to take the money in the replacement condition does not result in take rates that are different from the baseline condition, where replacement is not possible.*




FIGURE 2. TAKE RATES AND REPLACEMENT PROBABILITIES IN THE DONATION GAMES

Notes: The grey bars show the take rates of all types of players in both conditions. The black diamonds indicate the respective replacement probabilities. The replacement probability is zero, by design, for players 1 in the *baseline condition* and for players 3 in the *replacement condition*; it is 0.51 and 0.29 for players 1 and 2 in the *replacement condition*, respectively. The figure illustrates that the take rates are very similar despite large differences in the replacement probabilities.

Our manipulation check excludes the possibility that our treatment is ineffective. We find that 75 percent of the subjects state a strictly positive replacement belief.⁶ Hence, the large majority of subjects in our belief elicitation task expect that player 1’s altruistic choice to forgo the money

⁶ 25 percent of the subjects stated beliefs of 5 percent, which could reflect a true belief of 0 percent. Subjects receive a bonus of 10 CHF if their stated belief is within 5 percentage points of the true value. The fact that no subject stated a belief lower than 5 percent is evidence that the subjects understood the monetary incentives in the elicitation task.

and enable the surgery in the *replacement condition* will be undone with positive probability by a subsequent player's selfish choice. However, subjects tend to underestimate the replacement probability; the average belief is 0.23, while the true value is 0.51.⁷

The finding that the replacement probability does not affect the take rate in the donation game is consistent with the possibility that deontological ethics guide subjects' choices. A necessary condition for this interpretation to be correct is the existence of a social norm that taking the 20 CHF is inconsistent with moral behavior. Figure 3 illustrates that our measurements confirm the existence of a clear social norm. Averaged over both conditions, 97.6 percent of the subjects (121 of 124) rate the decision to take the 20 CHF as either "very socially inappropriate" or "somewhat socially inappropriate." The former is the modal choice in both conditions. Not a single subject chose "somewhat appropriate" or "very appropriate" in either condition. The distribution of the ratings is virtually identical in both conditions ($p=0.80$, Mann Whitney U test). The latter finding indicates that the availability of the replacement excuse does not render it more appropriate to take the 20 CHF, which is in itself consistent with Result 1. We summarize the measurement of the social norm in our second result.

Result 2: *A strong social norm of moral behavior exists in the donation game. Almost all subjects rate the decision to take the 20 CHF, instead of enabling a leprosy surgery in India, as either "very socially inappropriate" or "somewhat socially inappropriate" in both the baseline and the replacement condition.*

In sum, we find that subjects do not use the replacement excuse in our donation games, where a strong social norm exists that taking the 20 CHF is inconsistent with moral or proper social behavior. These data suggest that competitive institutions and the associated replacement excuse do not necessarily lead to the erosion of moral behavior.

⁷ This observation is consistent with existing research showing that subjects systematically underestimate the probability of disjunctive events, such as drawing at least one "take" in a sequence of two; see, e.g., Bar-Hillel (1973).




FIGURE 3. SOCIAL NORMS IN THE DONATION GAMES

Notes: The bars show the measures of the social appropriateness of taking the 20 CHF, instead of enabling a leprosy surgery in India, in both conditions. The figure illustrates that vast majority of subjects believe that the choice to take the 20 CHF would be rated by most people as either *very inappropriate* or *somewhat inappropriate*. Almost no subject believes that most people would rate the choice to take the 20 CHF as *neutral* and no subject believes that most people would rate this choice as *appropriate* in either condition.

3. Take Games with and without Punishment

Our Result 1, the irrelevance of the availability of the replacement excuse in the donation game, challenges often-held intuitions. For example, Falk and Szech (2013) write: “This ‘replacement’ logic is a common feature of markets, and it is therefore not surprising that the rhetoric of traders often appeals to the phrase that ‘if I don’t buy or sell, someone else will’” (p. 710).

In the following, we analyze behavior in two different games, using a different subject pool and different stake sizes, to check the robustness of Result 1. We want to rule out the possibility that behavior in the donation game is an exception, driven, for instance, by an exceptionally strong social norm that taking the 20 CHF and thereby defeating the leprosy surgery in India is inconsistent with moral behavior.

3.1 Experimental Design

We conducted a series of simple take games to revisit the question whether the replacement excuse encourages narrowly self-interested behaviors. *Take Game 1 (TG-1)* consists of two players: A and B1, who both start with an endowment of 0.5 USD. Player B1 chooses to either take away 0.4 USD from player A or to refrain from doing so. Player A cannot take an action. If B1 takes the money,

player A's payoff is 0.1 USD and player B1's payoff is 0.9 USD; if B1 does not take the money both players receive their endowments of 0.5 USD. *Take Game 2 (TG-2)* is different from *TG-1* only in that it consists of three players: A, B1, and B2, who all have endowments of 0.5 USD. First, B1 can take away 0.4 USD from A. If B1 does not take the 0.4 USD, then B2 can do so. Finally, *Take Game 3 (TG-3)* is different only in that it consists of four players: A, B1, B2, and B3. If neither B1 nor B2 takes the money, B3 can finally do so.

The key feature of the experimental design is the exogenous variation of the number of players who can replace player B1's pro-social choice not to take away the money from player A. That is, we vary the replacement probability by treatment. Comparison of the take rates of players B1 in conditions *TG-1*, *TG-2*, and *TG-3* allows for a first robustness check of Result 1.

As a manipulation check, we measured the beliefs that player B1's decision to forgo taking the 0.4 USD from player A will be replaced by player B2 in *TG-2* or by players B2 or B3 in *TG-3*, respectively. We also elicited the social norm that applies to player B1's decision to take the 0.4 USD from player A in each of the three conditions. We used the same methods as in the donation game to measure replacement beliefs and social norms.

Moreover, we conducted three additional conditions of the take games that feature a punishment option for player A. The take games with punishment, *TG-P-1*, *TG-P-2*, and *TG-P-3*, are identical to *TG-1*, *TG-2*, and *TG-3*, respectively, except that player A can spend up to 5 USD cents of his payoff to punish a player B who took the money away. For each 1 USD cent player A spends, the targeted player B's payoff decreases by 10 USD cents. Players B who do not take money away from player A cannot be punished. Adding a punishment option for player A provides for a second robustness check of Result 1, in a game where the incentive to take the 0.4 USD from player A is reduced because it might result in monetary punishment. Moreover, the punishment data allow studying whether the replacement excuse provides justification for the act of taking the money from the viewpoint of the "victim." In *TG-P-3*, for example, do players A consider players B1, who are endowed with the replacement excuse, less worthy of punishment for taking the money than players B3, who are not endowed with the replacement excuse?

Table 1 provides an overview of our treatment conditions and additional measurements, and it shows the respective numbers of observation.⁸

⁸ We did not repeat the measurements of beliefs and norms in the conditions with punishment to economize on subjects.

TABLE 1. OVERVIEW OF TAKE GAME CONDITIONS AND NUMBER OF OBSERVATIONS

Condition	Role				Additional Measurements	
	A	B1	B2	B3	Social Norm	Replacement Belief
<i>TG-1</i>	126	126	-	-	126	-
<i>TG-2</i>	108	108	108	-	110	104
<i>TG-3</i>	108	108	108	108	100	104
<i>TG-P-1</i>	101	101	-	-	-	-
<i>TG-P-2</i>	112	112	112	-	-	-
<i>TG-P-3</i>	100	100	100	100	-	-

3.2 Procedural Details

Subjects were recruited on Amazon Mechanical Turk (AMT), an online labor market (see, e.g., Horton et al. 2011). Participation was restricted to U.S. workers with at least 500 completed assignments and a minimum approval rating of 95 percent. We implemented the experiment with the software o-Tree (Chen et al. 2016). We controlled the subjects' understanding of the instructions by asking a set of test questions. Subjects could not participate if they were not able to provide correct answers within two attempts. All experimental instructions and test questions are in the online appendix. Overall, 2486 workers participated. We randomized workers who clicked the link to our study into treatments and roles or into a norm or belief measurement task. We used the strategy method for players B2 and B3 in the conditions with and without punishment. We also used the strategy method for players A in the punishment conditions. For example, player A in *TG-P-3* made three punishment decisions, one for each player B who could have taken the money. We stopped the data collection once we had 100 observations in each cell. Due to the random assignment, we obtained more than 100 observations in some games and measurement tasks; see Table 1. Subjects could participate only once. Each subject received a fixed payment of 0.50 USD and earned a variable payoff on top. The variable payment depended on the choices of players B and, if applicable, the punishment behavior of players A, as explained in Section 3.1. Subjects in the belief and norm measurement tasks earned 3 additional USD if their norm ratings matched the modal response or if their beliefs of the replacement probability were not further away than 5 percentage points from the true value, respectively. On average, the workers received a total payment (fixed payment plus variable payment) of 1.07 USD and took about 6 minutes to complete the study, translating to an average hourly rate of about 10.70 USD.

3.3 Predictions

As in the donation game in Section 2, the behavioral predictions depend on the assumptions about subjects' moral motivations. If subjects are utilitarian, then we should expect that the take rates will be higher in the conditions with possible replacement, *TG-2* and *TG-3*, than in the condition without replacement, *TG-1*. The same holds for the respective conditions with punishment. But if subjects' choices are guided by deontological ethics and if taking the 0.4 USD from player A constitutes a norm violation, then the take rates in *TG-1* and *TG-P-1* will not necessarily be different from the take rates in the respective conditions with possible replacement.

3.4 Results

We report the data from the take game (without punishment) first. The grey bars in the left panel of Figure 4 show that the take rates of players B1 are similar in all three conditions, 68.3 percent (86 of 126) in *TG-1*, 63.9 percent (69 of 108) in *TG-2*, and 60.2 percent (65 of 108) in *TG-3*. We cannot reject that the take rates are identical ($p=0.44$, Pearson's chi-square test). We observe these take rates even though the replacement probability increases from 0, by design, in *TG-1* to 0.69 in *TG-2* and 0.89 in *TG-3*, indicated by the black diamonds in Figure 4. The take rates of players B2 in *TG-2* and of players B2 and B3 in *TG-3* are equally similar, 68.5 percent (74 of 108), 67.6 percent (73 of 108), and 64.8 percent (70 of 108), respectively. We cannot reject that the choices of all types of players B in all conditions originate from the same distribution (Pearson's chi square test, $p=0.36$).⁹ We summarize this finding in our next result.

Result 3: *Subjects do not use the replacement excuse in the take games. The possibility that subsequent players B can replace an earlier player B's decision not to take the money from player A does not affect take rates.*

⁹ Note that the number of players is not constant between our conditions with and without replacement, neither in the donation game nor in the take game. While this can be considered as a natural feature arising from the fact that the possibility of replacement comes with the presence of additional players, we want to exclude the possibility that simply adding players affects behavior. For example, if B1 is sensitive to the size of the "audience" in the take game, this could lead to *lower* take rates in *TG-2* and *TG-3*, where one and two other players B, respectively, are informed about B1's choice, than in *TG-1*, where no other player B is informed about B1's choice. To test for an "audience effect" (Andreoni and Bernheim, 2009), we conducted a control condition on AMT that is identical to *TG-1*, except that it is common knowledge that two passive players are added, acting as an "audience," each receiving 0.5 USD and getting informed about B1's decision. We collected 107 novel observations for B1 in *TG-1* and 109 observations for B1 in the new condition with spectators; hence, a total of 650 workers participated in the control study. The experimental instructions are in the online appendix. We find that the take rates of B1s are not different in the treatments with and without an audience, 56.9 percent (62 of 109) and 61.7 percent (66 of 107), respectively ($p=0.28$, Fischer exact test, one-sided). We conclude that "audience effects" do not explain the irrelevance of the replacement excuse in our games.

The manipulation check reveals that 98.1 percent (102 of 104) and 96.0 percent (96 of 100) of the subjects believe that the replacement probability for player B1 is strictly positive in *TG-2* and *TG-3*, respectively.¹⁰ The average belief in *TG-2*, 0.61, is close to the true value, 0.69. The average belief in *TG-3* is 0.65, while the true value is 0.89. That is, as in the *replacement condition* of the donation game, subjects underestimate on average the replacement probability when more than one subject can replace the first player’s pro-social choice. Importantly, however, even though replacement beliefs are much higher in the take games than in the donation game (0.61 and 0.65, compared to 0.23), our main result that subjects do not make use of the replacement excuse is replicated in the take games.




FIGURE 4. TAKE RATES AND REPLACEMENT PROBABILITIES IN THE TAKE GAMES

Notes: The grey bars show the take rates in the take game (left panel) and the take game with punishment (right panel). The black diamonds indicate the respective replacement probabilities. The figure illustrates that our main result on the irrelevance of the replacement excuse for the frequency of narrowly self-interested choices in the donation game replicates in both take games. The take rates within each game are very similar, despite large differences in the replacement probabilities.

Result 3 is again consistent with the possibility that deontological ethics guide subjects’ choices. As in the donation game, a necessary condition for this interpretation to be correct is the existence of a clear social norm that taking the 0.4 USD from player A is inconsistent with moral

¹⁰ Four subjects stated a belief of 5 percent, which could reflect a true belief of 0 percent. Only 2 of 204 subjects stated a belief of less than 5 percent, suggesting that the vast majority of the subjects understood the monetary incentives.

behavior. Figure 5 illustrates that our measurements confirm the existence of a clear social norm. In total, 82 percent of the subjects rate player B1’s decision to take the money from player A as either “very” or “somewhat socially inappropriate.” The latter is the modal response, chosen by 46 percent of the subjects. No significant differences in the norm ratings exist between *TG-1*, *TG-2*, and *TG-3* (Kruskal-Wallis test, $H(2)=3.91$, $p=0.14$). This shows that the availability of the replacement excuse does not render it more appropriate to take the 0.4 USD from player A, which is in itself consistent with Result 3. We summarize these findings in the following.

Result 4: *A clear social norm exists in the take game. The large majority of subjects rate player B1’s decision to take 0.4 USD from player A as either “somewhat socially inappropriate” or “very socially inappropriate” in all three conditions.*




FIGURE 5. SOCIAL NORM IN THE TAKE GAME

Notes: The bars show the measures of the social appropriateness of taking the 0.4 USD from player A in all three conditions. The figure illustrates that large majority of subjects believe that the choice to take the 0.4 USD from player A would be rated by most people as either *very inappropriate* or *somewhat inappropriate* in all three conditions.

We next turn to the take games with punishment. The grey bars in the right panel of Figure 4 show that the take rates of players B1 are very similar in all three conditions, 36.6 percent (37 of 101) in *TG-P-1*, 32.1 percent (36 of 112) in *TG-P-2*, and 25.0 percent (25 of 100) in *TG-P-3*. We cannot reject that these take rates are identical (Pearson’s chi-square test, $p=0.20$). We obtain this result even though the replacement probability increases from 0, by design, in *TG-P-1* to 0.29 in

TG-P-2, and to 0.47 in *TG-P-3*, as shown by the black diamonds. The take rates of players B2 in *TG-P-2* and of players B2 and B3 in *TG-P-3* are equally similar, 29.5 percent (33 of 112), 28.0 percent (28 of 100), and 27.0 percent (27 of 100), respectively. We cannot reject that the choices of all players B in all three conditions originate from the same distribution (Pearson's chi square test, $p=0.24$).¹¹ We summarize this result as follows.

Result 5: *Subjects do not use the replacement excuse in the take games with punishment. The take rates are identical in all three conditions.*

In sum, the behavior observed in the take games replicates the results obtained in the donation game. In two versions of the take game, with and without punishment, we find that subjects do not make use of the replacement excuse, a context where a clear social norm exists (measured in the take game without punishment) that taking the 0.4 USD is inconsistent with moral or proper social behavior.

We finally report the punishment data, which allows studying whether players A accept the replacement excuse. Do players A who lost 0.4 USD, i.e., the “victims,” punish less when a player B's pro-social choice not to take the money could have been replaced? Our data show that this is not the case. Players B1 and B2 in condition *TG-P-2* were punished at virtually identical levels if they took the 0.4 USD from player A, leading to an average punishment of 0.26 USD for both types of players B. Also the average punishment of players B1, B2, and B3 is very similar when they took the 0.4 USD from player A in condition *TG-P-3*, 0.19, 0.19, and 0.21 USD, respectively, and not significantly different (Friedman test, $p=0.83$). These findings are consistent with Result 3.¹² We summarize the punishment data as follows.

Result 6: *From the point of view of players A, possible replacement of the pro-social choice is not an acceptable excuse for making the selfish choice in the take games. Within our conditions, Players A punish all types of players B equally for taking the money, independently of their positions in the sequence.*

¹¹ It is noteworthy that the threat of punishment reduces the take rate significantly compared to the conditions without punishment (Fisher exact test, one sided, $p<0.01$ for all possible bilateral comparisons). While this is not surprising in itself, it shows that workers on AMT react sensibly to incentives. This suggests that the absence of the replacement effect in the take games is not driven by a general insensitivity of AMT workers to experimental variations.

¹² Average punishment of players B1 in *TG-P-1* for taking the 0.4 USD from player A is 0.29 USD. Punishment thus appears to be lower in conditions with a larger numbers of players B.

4. Ultimatum Games with Responder Competition

We demonstrated in three different games in Sections 2 and 3 that subjects do not employ the replacement excuse. The games share the social norm that implementing the respective selfish outcome is inconsistent with proper social behavior. Our interpretation of these empirical results was that subjects are guided by deontological ethics, where the argument that the selfish outcome is likely to come about anyway does not justify deviations from a social norm. To further test this interpretation, we turn to the study of ultimatum games with responder competition in this section.

It is a well-known finding that even very low offers are frequently accepted in ultimatum games with responder competition, where multiple responders can accept an offer. The same low offers are, in contrast, mostly rejected in standard ultimatum games with a single responder (e.g., Güth et al. 1998, Grosskopf 2003, Fischbacher et al. 2009).¹³ For example, the rejection rate of a 10-percent-offer in the ultimatum games in Fischbacher et al. (2009) is 0.80 in the standard game with a single responder but only 0.46 in a game with two responders and 0.26 in a game with five responders (leading to expected rejection probabilities of 0.21 and 0.001, respectively). Rejecting a low offer in a standard ultimatum game ensures that the proposer will be punished, e.g., with intent to enforce a fairness norm (Fehr and Gächter 2002). Rejecting a low offer in an ultimatum game with responder competition, however, will not ensure that the proposer is punished because an offer will only be rejected if every single responder rejects. Apparently, many responders in games with competition do not want to pay the cost of punishing the proposer, which is incurred by rejecting a low offer. A possible explanation is that responders invoke the replacement excuse in these games and argue that “if I don’t accept the low offer, another responder will do so.”

How can our results in the donation and take games be reconciled with the existing results on responder behavior in ultimatum games? A necessary condition for our interpretation of the findings in Sections 2 and 3 to be correct and to generalize to other games is the *absence* of a social norm that accepting a low offer in an ultimatum game is inconsistent with moral or proper social behavior. In this case, the prediction for deontological players coincides with that for utilitarian players because no social norm exists that guides the decision whether to accept or reject a low offer. Consequently, irrespective of whether responders follow a utilitarian or a deontological

¹³ Multiple responders decide simultaneously in these games. One of the accepting responders is randomly determined to receive the offer if more than one responder accepts. An offer is rejected only if all responders reject.

fairness view, they will more frequently accept a low offer if the only consequence of rejecting the offer is that other responders get the chance to accept it.

4.1 Experimental Design

We conducted a series of ultimatum games in which we vary the numbers of responders. *Ultimatum Game 1 (UG-1)* is a standard, binary ultimatum game with two players, one proposer (player A) and one responder (player B1). Player A can either split 1 USD equally or make a low offer of 0.1 USD to player B1 to keep 0.9 USD for himself. The equal split is automatically implemented (for simplicity) but player B1 can decide whether to accept the low offer, in which case the unequal split is implemented, or to reject it, in which case both players receive nothing. *Ultimatum Game 2 (UG-2)* is different from *UG-1* only in that it consists of three players: A, B1, and B2. If player A proposes the equal split, A and B1 each receive 0.5 USD and B2 receives nothing. If player A makes the low offer, players B1 and B2 decide sequentially whether to accept it. If B1 rejects the low offer, B2 can accept it. The player B who accepts the low offer receives the 0.1 USD, the respectively other player B receives nothing. The low offer is finally rejected only if both players B reject it. *Ultimatum Game 3 (UG-3)* is different from *UG-2* only in that a third responder, B3, is added. If first B1 and then B2 reject the low offer, B3 can finally accept it.¹⁴

The key feature of the experimental design is the exogenous variation of the number of players who can replace player B1's choice to reject the low offer. Comparison of the rejection rates of players B1 in *UG-1*, *UG-2*, and *UG-3* allows studying if the existing results on responder behavior in ultimatum games with responder competition replicate in our setting.

As a manipulation check, we measured the beliefs that player B1's decision to reject the low offer will be replaced by that of player B2 in *UG-2* or by that of players B2 or B3 in *UG-3*, respectively. Finally, we elicited the social norm that applies to player B1's decision to accept the low offer in each of the three conditions. We used the same methods as in the donation and take games to measure replacement beliefs and social norms.

Table 2 provides an overview of our treatment conditions and additional measurements, and it shows the respective numbers of observation.

¹⁴ We implemented a sequential version of ultimatum games with responder competition to minimize the differences with respect to our take games. Also the stake size and possible payoff differences among players are held constant.

TABLE 2. OVERVIEW OF ULTIMATUM GAME CONDITIONS AND NUMBER OF OBSERVATIONS

Condition	Role				Additional Measurements	
	A	B1	B2	B3	Social Norm	Replacement Belief
<i>UG-1</i>	109	109	-	-	111	-
<i>UG-2</i>	106	106	106	-	102	104
<i>UG-3</i>	101	101	101	101	100	102

4.2 Procedural Details

Subjects were recruited on Amazon Mechanical Turk (AMT). Participation restrictions and other methods were exactly as in the take games. All instructions and test questions are in the online appendix. Overall, 1459 workers participated. Each subject received a fixed payment of 0.50 USD and earned a variable payoff on top. The variable payment depends on the choices of players A and, in case of a low offer, the acceptance decision of players B, as explained in Section 4.1. Subjects in the belief and norm measurement tasks could earn additional 3 USD, exactly as in the take games. On average, the workers received a total payment of 1.13 USD and took about 6 minutes to complete the study, translating to an average hourly rate of about 11.30 USD.

4.3 Predictions

We expect our implementation of ultimatum games with responder competition to replicate the results reported in the literature. That is, we expect players B1 to more frequently accept the low offer in *UG-2* than in *UG-1* and more frequently in *UG-3* than in *UG-2*. Moreover, we hypothesize that no social norm exists that accepting the low offer is inconsistent with moral or proper social behavior.

4.4 Results

The grey bars in Figure 6 show the acceptance rates of the low offer by players B1 in all three conditions. 51 percent of players B1 accept the low offer in the standard ultimatum game *UG-1*. The acceptance rate is significantly higher, 74 percent, in condition *UG-2*, where player B2 can accept the low offer if player B1 rejects it ($p < 0.01$, Fisher exact test). The acceptance rate increases further, to 85 percent, in condition *UG-3*, where players B2 and, if not B2, then B3 can accept the low offer in case B1 rejects it (*UG-3* vs. *UG-1*, $p < 0.01$; *UG-3* vs. *UG-2*, $p = 0.06$; Fisher exact

tests).¹⁵ The black diamonds in Figure 6 indicate the replacement probabilities; it is zero, by design, in *UG-1* but increases to 0.80 and 0.98 in *UG-2* and *UG-3*, respectively. We summarize these findings in the following.

Result 7: *The possibility that subsequent responders can accept and thereby replace the first responder’s decision to reject a low offer in the ultimatum game significantly increases the rate at which the first responder accepts the low offer.*




FIGURE 6. ACCEPTANCE RATES AND REPLACEMENT PROBABILITIES IN THE ULTIMATUM GAMES

Notes: The grey bars show the acceptance rates of the first responder in the three conditions. The black diamonds indicate the replacement probabilities. The replacement probability is zero, by design, for the responder in *UG-1*; it is 0.80 and 0.98 for the first responder in *UG-2* and *UG-3*, respectively. The figure illustrates that the acceptance rates increase as the replacement probability increases.

Our replication of the result that responders accept low offers more frequently in ultimatum games with responder competition than in ultimatum games with a single responder already indicates that our treatment manipulation was effective. Our manipulation check revealing that the vast majority of subjects expect a positive replacement probability (but underestimate it on

¹⁵ We focus on the acceptance decisions by players B1 because the sequential order of moves of the responders implies that the low offer is the kindest action that is available to player A from the point of view of players B2 and B3; they receive nothing if player A splits equally with player B1. The acceptance rates of players B2 and B3 are accordingly high; 0.8, 0.86, and 0.82 for B2 in *UG-2*, B2 and B3 in *UG-3*, respectively.

average) confirms this.¹⁶ Moreover, we cannot reject that the beliefs in the ultimatum games and the take games, which are equal in terms of stake sizes and subject pool and thus lend themselves to being compared, are identical (*UG-2* vs. *TG-2*, $p=0.14$; *UG-3* vs. *TG-3*, $p=0.82$; Mann Whitney U tests). This indicates that differences in beliefs do not explain why the replacement excuse is effective in the ultimatum games but not in the take games.

The finding that the replacement probability affects the acceptance rate in the ultimatum games would be consistent with the possibility that subjects' choices are guided by deontological ethics—as our results in the donation and take games suggest—if no social norm existed that accepting a low offer is inconsistent with moral behavior. This possibility is supported by the data. Figure 7 shows that the modal response in all three condition is that accepting the low offer is “neutral: neither socially inappropriate nor appropriate,” chosen by 33.5 percent of all subjects. 38.0 percent of the subjects rate player B1's decision to accept the low offer as either “very” or “somewhat socially appropriate,” and only 28.4 percent choose “very” or “somewhat socially inappropriate.” There is no significant difference in the norm ratings between *UG-1*, *UG-2*, and *UG-3* (Kruskal-Wallis test, $H(2)= 1.03$, $p=0.60$). These data clearly indicate that there is no common understanding that accepting a low offer is socially inappropriate. We summarize the measurement of the social norm as follows.

Result 8: *No social norm exists that accepting a low offer in the ultimatum game is socially inappropriate and inconsistent with moral or proper social behavior.*

In sum, we replicate the existing result that responders accept low offers at a higher rate in ultimatum games with responder competition than in standard games without competition. Importantly, our data indicate that no social norm exists that governs the decision whether to accept or reject a low offer in an ultimatum game. Hence, preferences over outcomes solely guide even deontological responders' choices, and if the only consequence of rejecting a low offer is that another responder will be able to accept the offer, most responders prefer accepting the low offer right away. The responders' behavior in ultimatum games with multiple responders is thus consistent with our interpretation of the behavioral patterns in the donation and take game.

¹⁶ We find that 97.1 percent of the subjects believe that the replacement probability is strictly positive in *UG-2* (101 of 104) and *UG-3* (99 of 102). As in the take games, four subjects stated a belief of 5 percent, which could reflect a true belief of 0 percent, given the incentive structure. Only two subjects stated a belief of less than 5 percent. The average beliefs are 0.64 and 0.62 in *UG-2* and *UG-3*, respectively.




FIGURE 7. SOCIAL NORM IN THE ULTIMATUM GAMES

Notes: The bars show the measures of the social appropriateness of accepting the low offer in all three conditions. The figure illustrates that the modal response in all three conditions is that subjects believe that the first responder's choice to accept the low offer would be rated by most people as *neutral*.

4.5 Additional norm measurement in the laboratory

As a final test, we measure the social norms regarding responder behavior in ultimatum games in the laboratory, using the experimental design by Fischbacher et al. (2009). We hypothesize that we can replicate Result 8, which is based on data obtained on AMT, with data obtained in a standard laboratory subject pool. Fischbacher et al.'s study was run in the laboratory of the Department of Economics at the University of Zurich, where we also ran our measurement of the applicable social norm. We provided the subjects with Fischbacher et al.'s original instructions and test questions from either the ultimatum game with a single responder or the ultimatum game with two competing responders. We then asked the subjects to state their beliefs about how most people would rate a responder's choice to accept a 10-percent-offer in the respective ultimatum game. The same five options were available that we used to measure social norms throughout the paper. Subjects received a bonus of 10 CHF if their answers matched the modal response. We elicited beliefs in the standard ultimatum game from 63 subjects and in the ultimatum game with two competing responders from 69 subjects. The measurement of the social norms was added at the end of unrelated experiments; it took about 15 minutes and subjects earned 3.27 CHF on average.

We find that the modal response in the standard ultimatum game is that accepting the low offer is “neutral: neither socially inappropriate nor appropriate,” chosen by 31.7 percent of the subjects. A similar fraction of subjects, 30.4 percent, chose the “neutral” option in the ultimatum game with two responders. We cannot reject that the norm ratings in the two conditions are identical ($p=0.69$, Mann Whitney U test). Overall, only 27.3 percent rate the decision to accept the low offer as either “very” or “somewhat socially inappropriate.” These ratings are very similar to the ratings obtained on AMT. We cannot reject that the ratings elicited in the laboratory (both conditions pooled) and the ratings elicited on AMT (all three conditions pooled) are identical ($p=0.43$, Mann Whitney U test).¹⁷ The measurement of the social norms regarding responder behavior in the laboratory thus confirm the measurements on AMT, which are summarized in Result 8.

5. Conclusion

We studied experimentally if people make use of the “replacement excuse”—the argument “if I don’t do it, someone else will”—to justify narrowly self-interested behaviors. A key feature of competitive markets is that a competitor can step in if a responsible market actor forgoes a business opportunity for ethical reasons. Our leading example was a market actor who considered refusing to sell arms to a dictatorial regime. If the only consequence of refusing to sell the arms was that a competitor could conclude the deal, the refusal might make no difference except that the profits that accrue from selling the arms are passed on to the competitor. The forces of competition would undermine the moral standards of the responsible market actor if, in anticipation of this outcome, he sells the arms himself.

It is a long-standing hypothesis that market interaction affects moral behavior. One theory, originating in the eighteenth century, asserts that markets are “civilizing.” An example is Marquis de Condorcet, who claimed that “Manners have become more gentle [...] through the influence of the spirit of commerce and industry” (1795; as quoted in Hirschman, 1982, p. 1465). An opposed theory, of which Marx is a prominent representative, asserts that the “modes of production and of exchange” have “left remaining no other nexus between man and man than naked self-interest”

¹⁷ We obtain the same measures of the social norm despite differences in the design of the respective ultimatum games. In particular, proposers in Fischbacher et al.’s design could offer any integer percentage amount of their endowment. In contrast, we implemented a binary version where the proposer could offer either the equal split or 10 percent of the endowment. Moreover, multiple responders decided simultaneously whether or not to accept an offer in Fischbacher et al., whereas responders made their acceptance decisions in a sequential order in our design.

(Marx and Engels, 1848, Ch. 1). Causal evidence on the effect of markets on morals and, in particular, on the mechanisms by which markets might affect morals is however still scarce. The contribution of our paper is to provide clean experimental evidence on one potentially powerful mechanism—the replacement excuse—by which competitive markets might erode morals.

Our data revealed a clear behavioral pattern, consistently found in a series of games and in two different subject pools. People do not use the argument “if I don’t do it, someone else will” to justify taking a narrowly self-interested action if a social norm exists classifying the selfish action as immoral—even when another subject is likely to replace the omission of the action. But if no social norm exists classifying the selfish action as immoral, subjects more often take the selfish action when an omission of the action can be replaced by another subject.

By demonstrating that social norms can outweigh the forces of competition, our paper informs the long-standing debate on the effect of markets on morals. While one should be careful about drawing general conclusions from the results of a single study, our data suggest that competitive markets do not necessarily undermine morals—at least not by way of the replacement excuse.

References

- Akerlof, George. 1980. “A Theory of Social Custom, of Which Unemployment May Be One Consequence.” *Quarterly Journal of Economics*, 94(4): 749–75.
- Akerlof, George, and Rachel Kranton. 2000. “Economics and Identity.” *Quarterly Journal of Economics*, 115(3): 715–53.
- Akerlof, George, and Rachel Kranton. 2005. “Identity and the Economics of Organizations.” *Journal of Economic Perspectives*, 19(1): 9–32.
- Andreoni, James. 1989. “Giving with Impure Altruism: Applications to Charity and Ricardian Equivalence.” *Journal of Political Economy*, 97(6): 1447–58.
- Andreoni, James. 1990. “Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving.” *Economic Journal*, 100(401): 464–77.
- Andreoni, James. 1995. “Warm-Glow versus Cold-Prickle: The Effects of Positive and Negative Framing on Cooperation in Experiments.” *Quarterly Journal of Economics*, 110(1): 1–21.
- Andreoni, James, and Douglas Bernheim. 2009. “Social Image and the 50-50 Norm: A Theoretical and Experimental Analysis of Audience Effects.” *Econometrica*, 77(5): 1607–36.

- Andreoni, James, Justin Rao, and Hannah Trachtman. 2016. "Avoiding the Ask: A Field Experiment on Altruism, Empathy, and Charitable Giving." *Journal of Political Economy*, forthcoming.
- Ariely, Dan, Anat Bracha, and Stephan Meier. 2009. "Doing Good or Doing Well?" *American Economic Review*, 99(1): 544-55.
- Aupperle, Kenneth, Archie Carroll, and John Hatfield. 1985. "An Empirical Examination of the Relationship between Corporate Social Responsibility and Profitability." *Academy of Management Journal*, 28(2): 446-63.
- Bar-Hillel, Maya, 1973, "On the Subjective Probability of Compound Events." *Organizational Behavior and Human Performance*, 9(3): 396-406.
- Bartling, Björn, Roberto Weber, and Lan Yao. 2015. "Do Markets Erode Social Responsibility?" *Quarterly Journal of Economics*, 130(1): 219-66.
- Bénabou, Roland, and Jean Tirole. 2004. "Willpower and Personal Rules." *Journal of Political Economy*, 112(4): 848-86.
- Bénabou, Roland, and Jean Tirole. 2006. "Incentives and Prosocial Behavior." *American Economic Review*, 96(5): 1652-78.
- Bénabou, Roland, and Jean Tirole. 2011. "Identity, Morals and Taboos: Beliefs as Assets." *Quarterly Journal of Economics*, 126(2): 805-55.
- Bernheim, Douglas. 1994. "A Theory of Conformity." *Journal of Political Economy*, 102(5): 841-77.
- Bock, Olaf, Ingmar Baetge, and Andreas Nicklisch. 2014. "Hroot: Hamburg Registration and Organization Online Tool." *European Economic Review*, 71: 117-20.
- Bodner, Ronit, and Drazen Prelec. 2003. "Self-Signaling and Diagnostic Utility in Everyday Decision Making." In: *The Psychology of Economic Decisions*, Vol. 1, pp. 105-26, Isabelle Brocas and Juan Carrillo (eds), Oxford University Press.
- Bowles, Samuel. 1998. "Endogenous Preferences: The Cultural Consequences of Markets and Other Economic Institutions." *Journal of Economic Literature*, 36(1): 75-111.
- Chen, Daniel, 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, Daylian Cain, and Robyn Dawes. 2006. "What You Don't Know Won't Hurt Me: Costly (but Quiet) Exit in Dictator Games." *Organizational Behavior and Human Decision Processes*, 100(2): 193-201.

- Dana, Jason, Roberto Weber, and Jason Kuang. 2007. "Exploiting Moral Wiggle Room: Experiments Demonstrating an Illusory Preference for Fairness." *Economic Theory*, 33(1): 67–80.
- Elster, Jon. 1989. "Social Norms and Economic Theory." *Journal of Economic Perspectives*, 3(4): 99–117.
- Exley, Christine. 2015. "Excusing Selfishness in Charitable Giving: The Role of Risk." *Review of Economic Studies*, 83(2): 587–628.
- Exley, Christine and Ragan Petrie. 2016. "Finding Excuses to Decline the Ask." HBS Working Paper, No. 16–101.
- Falk, Armin, and Nora Szech. 2013. "Morals and Markets." *Science*, 340(6133): 707–11.
- Fehr, Ernst, and Simon Gächter. 2002. "Altruistic Punishment in Humans." *Nature*, 415(6868): 137–40.
- Fischbacher, Urs. 2007. "Z-Tree: Zurich Toolbox for Ready-Made Economic Experiments." *Experimental Economics*, 10(2): 171–78.
- Fischbacher, Urs, Christina Fong, and Ernst Fehr. 2009. "Fairness, Errors and the Power of Competition." *Journal of Economic Behavior & Organization*, 72(1): 527–45.
- Glover, Jonathan and M. J. Scott-Taggart. 1975. "It Makes No Difference Whether or Not I Do It." *Proceedings of the Aristotelian Society, Supplementary Volumes*, 49(1): 171–209.
- Grosskopf, Brit. 2003. "Reinforcement and Directional Learning in the Ultimatum Game with Responder Competition." *Experimental Economics*, 6(2): 141–58.
- Güth, Werner, Nadege Marchand, and Jean-Louis Rullière. 1998. "Equilibration and Context Dependency: An Experimental Investigation of the Ultimatum Bargaining Game." *Revue Economique*, 49(3), 785–94.
- Haisley, Emily and Roberto Weber. 2010. "Self-serving interpretations of ambiguity in other-regarding behavior." *Games and Economic Behavior*, 68(2): 634–45.
- Hirschman, Albert. 1982. "Rival Interpretations of Market Society: Civilizing, Destructive, or Feeble?" *Journal of Economic Literature*, 20(4), 1463–84.
- Horton, John, David Rand, and Richard Zeckhauser. 2011. "The Online Laboratory: Conducting Experiments in a Real Labor Market." *Experimental Economics*, 14(3): 399–425.
- Kirchler, Michael, Jürgen Huber, Matthias Stefan, and Matthias Sutter. 2016. "Market Design and Moral Behavior." *Management Science*, 62(9), 2615–25.
- Knight, Frank. 1923. "The Ethics of Competition." *Quarterly Journal of Economics*, 37(4): 579–624.

- Köszegi, Botond. 2006. "Ego Utility, Overconfidence, and Task Choice." *Journal of the European Economic Association*, 4(4): 673–707.
- Krupka, Erin, and Roberto Weber. 2013. "Identifying Social Norms Using Coordination Games: Why Does Dictator Game Sharing Vary?" *Journal of the European Economic Association*, 11(3): 495–524.
- Lazear, Edward, Ulrike Malmendier, and Roberto Weber. 2012. "Sorting in Experiments with Application to Social Preferences." *American Economic Journal: Applied Economics*, 4(1): 136–63.
- Lindbeck, Assar. 1997. "Incentives and Social Norms in Household Behavior." *American Economic Review*, 87(2): 370–77.
- López-Pérez, Raúl. 2008. "Aversion to Norm-Breaking: A Model." *Games and Economic Behavior*, 64(1): 237–67.
- Marx, Karl, and Friedrich Engels. 1848, "Manifesto of the Communist Party." *Marx/Engels Selected Works*, Vol. 1, 98–137, Progress Publishers: Moscow, 1969.
- Mazar, Nina, On Amir, and Dan Ariely. 2008. "The Dishonesty of Honest People: A Theory of Self-Concept Maintenance." *Journal of Marketing Research*, 45(6): 633–44.
- Ostrom, Elinor. 2000. "Collective Action and the Evolution of Social Norms." *Journal of Economic Perspectives*, 14(3), 137-58.
- Pigors, Mark and Bettina Rockenbach. 2016. "Consumer Social Responsibility." *Management Science*, 62(11), 3123-37.
- Porter, Michael, and Mark Kramer. 2006. "Strategy and Society: The Link between Competitive Advantage and Corporate Social Responsibility." *Harvard Business Review*, 84(12): 42–56.
- Riedl, Arno and Paul Smeets. 2017. "Why Do Investors Hold Socially Responsible Mutual Funds?" *Journal of Finance*, forthcoming.
- Sandel, Michael. 2012. *What Money Can't Buy. The Moral Limits of Markets*. London: Alan Lane.
- Satz, Debra. 2010. *Why Some Things Should Not Be for Sale: The Moral Limits of Markets*. Oxford University Press.
- Shleifer, Andrei. 2004. "Does Competition Destroy Ethical Behavior?" *American Economic Review, Papers and Proceedings*, 94(2): 414–18.
- Sparkes, Russell, and Christopher Cowton. 2004. "The Maturing of Socially Responsible Investment: A Review of the Developing Link with Corporate Social Responsibility." *Journal of Business Ethics*, 52(1): 45–57.