

Morals in multi-unit markets*

Andreas Ziegler, Giorgia Romagnoli, Theo Offerman

University of Amsterdam

February 9, 2021

Abstract

We examine how the erosion of morals in markets depends on the market power of individual traders. Previously studied single-unit markets provide market power to individual traders by limiting the roles of two forces: (i) the replacement logic, whereby immoral trading is justified by the belief that others would trade otherwise; (ii) market selection, by which the least moral trader determines quantities. In an experiment, we compare single-unit to (more common) multi-unit markets which may activate these forces. We find that, in contrast to single-unit markets, multi-unit markets show full erosion of morals. Especially the replacement logic drives this finding. In addition, we find that (i) market experience leads to biased social learning, whereby subjects believe that others are less moral than they actually are; (ii) erosion of morals persists partially after multi-unit markets; (iii) changes in social norms are not driving these results.

Keywords replacement logic, social responsibility, markets, externalities, competition, charity

JEL Codes C91, C92, D62

*We thank Björn Bartling, Sourav Bhattacharya, Guillaume Fréchette, Hartmut Kliemt, Florian H. Schneider, Jeroen van de Ven, Joël van der Weele and seminar and conference audiences at the CREED-ESE meeting, TI Workshop on Communication in Markets and Organizations, EEA Virtual 2020, 2020 ESA Global, CERGE-EI Prague and CeDEX Nottingham for helpful suggestions and comments. Financial support from the Research Priority Area Behavioral Economics of the University of Amsterdam is gratefully acknowledged. This experiment is preregistered with RCT ID AEARCTR-0004603. Andreas Ziegler: a.g.b.ziegler@uva.nl, Giorgia Romagnoli: g.romagnoli@uva.nl, Theo Offerman: t.j.s.offerma@uva.nl.

1 Introduction

People’s morals may easily take a backseat in markets. Consider the market for air travel. Passengers may think that even if they refrain from buying a cheap ticket, someone else will, leaving total emissions unchanged. This reasoning, the so called replacement logic, may excuse frequent flying also for environmentally conscious consumers. Barr, Shaw & Coles (2011) find that many people combine a sustainable lifestyle at home with the use of short-haul air travel for leisure, instead of the more fuel efficient modes of transport like trains and coaches. Airlines themselves may excuse their offering of flights by arguing that a competitor will offer an additional flight if they decide to withdraw a connection. At the same time, the choices of a minority of consumers can have a disproportionate impact on aggregate outcomes. Gössling & Humpe (2020) find that in the US, 12% of adults account for 68% of all trips. As a result, aggregate behavior may not reflect the average person’s concern for environmental damages. Anecdotal evidence suggests that resorting to the replacement logic to excuse morally questionable behavior as well as disproportional activity of few irresponsible actors are also features of other markets, such as the market for opioids, the shipping industry and weapons trade.¹

Recent laboratory experiments have investigated the extent to which morals are eroded in single-unit markets, which are markets where each participant is restricted to trade at most one unit. In a seminal paper, Falk & Szech (2013) find that while 45.9% of subjects are willing to kill a mouse for €10 in individual decision-making, 75.9% do so in single-unit markets.² However, key results of Falk & Szech (2013) are contested. Market prices can decline also without moral erosion (Sutter, Huber, Kirchler, Stefan & Walzl 2020). Moreover, while Falk & Szech (2013) compare a single decision in individual decision-making with repeated decisions in a market, Bartling, Fehr & Özdemir (forthcoming) show that the partial erosion in markets

¹In the opioids market, a spokes-woman for McKesson, which was the largest distributor in the US from 2006 to 2012, stated: “Any suggestion that McKesson influenced the volume of opioids prescribed or consumed in this country would reflect a misunderstanding of our role as a distributor” (<https://apnews.com/98963bb70e0f462295ccc02fe9c68e71>). In 2017, the number of Americans dying from an overdose of opioids (47,600) surpassed the number dying from car accidents (Scholl, Seth, Kariisa, Wilson & Baldwin 2019). Empirically, Vuillemeay (2020) documents an erosion of standards in the shipping industry, where jurisdictions compete to register additional ships by relaxing regulatory requirements. In the market for weapon trading, both UK prime minister Tony Blair (in 2002) and British Secretary of State Boris Johnson (in 2016) made the argument that they could stop the defense industry operating in their country, but that then someone else would step in to supply the arms that they supplied (Bartling & Özdemir 2017, Falk, Neuber & Szech 2020).

²Engelmann, Friedrichsen & Kübler (2018) show that the morality of behavior in laboratory markets correlates with the type of choice they are intended to capture outside of the laboratory. The follow-up literature also discusses how different market forces, such as anonymity, market framing, or joint decision-making affect traders’ morals in markets (Kirchler, Huber, Stefan & Sutter 2016, Irlenbusch & Saxler 2019).

disappears under repetition of both environments. So far, the evidence that people's morals are eroded in markets is inconclusive.

We think that the single-unit markets studied so far provide a poor approximation of most real world markets.³ In single-unit markets, each trader possesses substantial market power. For each pair of traders who trade a unit, at least one of them is pivotal. The total quantity traded would be reduced if this trader would have refrained from trading. Buyers and sellers can usually trade more than one unit and traders are rarely endowed with the ability to affect aggregate market outcomes.

In this paper, we investigate how the erosion of morals depends on the ability of traders to affect aggregate market outcomes. In a laboratory experiment, we compare single-unit to multi-unit markets. Within multi-unit markets, we vary individuals' market power by removing traders' pivotality for aggregate outcomes. Our main contribution is twofold. First, we investigate whether market outcomes reflect participants' concern towards causing negative externalities and the extent to which this is affected by individuals' market power. Second, we disentangle how much of this effect is due to either of two forces: (i) *market selection* and (ii) *replacement logic*.

Market selection predicts that trade is mainly carried out by the traders that are least concerned with causing negative externalities. The repeated decision whether to trade an additional unit can be made by these traders only in multi-unit markets, as only in these markets they will not be prevented from trading again after doing so once. This means that market quantities and prices can be fully determined by the buyer and seller caring the least about causing an externality. Market selection does not necessarily imply that participants behave any differently in markets. Instead, markets may weigh the preferences of less moral market participants more heavily when aggregating preferences and determining outcomes. Under this force, outcomes are more self-interested even when market participants' preferences are unchanged.

We further decompose two dimensions along which the market selection effect may operate. First, there may be initial heterogeneity in how different people care about causing a particular negative externality. Lower moral costs imply, *ceteris paribus*, a competitive advantage that could lead traders with low moral costs to serve a larger share of the demand or supply. Second, there may be a desensitization effect, which occurs when perceived marginal moral costs of causing a negative externality decrease in the number of units that a participant trades. Due to the large scale of their operations, a relatively small number of traders can produce large negative externalities with low average moral costs.

³Also within experimental economics, markets were extensively studied in multi-unit rather than single-unit settings (e.g. Smith 1962, Ketcham, Smith & Williams 1984, Plott 1983).

Aside from market selection, the replacement logic may be particularly strong in multi-unit markets. According to this principle, participants may decide to trade as they realize that their individual actions do not affect aggregate outcomes. They then feel justified in trading and reaping the profits for themselves (Sobel 2007). In the single-unit markets studied so far, a market participant on at least one side of the market is pivotal for the aggregate volume traded and thus cannot use the replacement logic argument. In contrast, in multi-unit markets, traders on both sides of the markets can excuse their trading with the argument that if they would not have traded someone else would have taken advantage of the opportunity,

We test whether markets erode morals with a laboratory experiment based on four between-subject treatments, including three multi-lateral market treatments. In all market treatments, participants can earn monetary payments by trading. However, trading causes a negative externality. By agreeing on a trade, participants decrease a monetary amount donated to UNICEF, which provide measles vaccines. We will refer to taking money for oneself instead of donating money as cancelling donations. We follow the convention in the literature and conclude that if participants are more willing to cancel donations in some treatments compared to others, there is an erosion of morals (Falk & Szech 2013, Kirchler et al. 2016). We separately establish that cancelling donations to UNICEF in exchange for the moderate and small rewards utilized in the experiment is perceived as immoral by our subjects.⁴

Our first market, treatment SINGLE, is a single-unit market. This treatment is comparable to the markets studied in the current literature. We further run two types of multi-unit markets. In MULTI, we implement a scaled-up version of SINGLE, where instead of one unit per participant, three units can be traded. Each trader is similarly pivotal as in SINGLE. Additional opportunities for social learning may erode morals further in MULTI compared to SINGLE, as more precise information from repeated interaction in a market is available. MULTI also serves as a control treatment for the third market treatment, FULL. In FULL, we remove pivotality, as each trader is now able to serve the entire market by herself. This activates both the replacement logic as well as the market selection effect. We use the identity of each participant to disentangle the two forces. We can do so because we begin by eliciting subjects' preferences in individual decision-making tasks in all treatments.

We measure whether these market forces erode morals by comparing behavior between market treatments. As in the previous literature, we also run an individual decision-making control treatment, IDM. We employ price lists to elicit participants'

⁴Using the elicitation method by Krupka & Weber (2013), we find that 331 out of 381 participants rate taking €1 as a payment to one-self instead of donating €1.50 in an individual decision-making task as "(somewhat) socially inappropriate".

preferences for canceling a donation for measles vaccines. In this context, we show that participants are willing to forgo significant payments to ensure that donations are made. Importantly, their preferences towards this donation opportunity are heterogeneous and characterized by desensitization. We repeat individual decision-making as we do for all markets. This allows us to control for a potential erosive effect of repetition.

We provide conclusive evidence that markets erode morals.⁵ We find partial erosion of morals in single-unit markets when comparing SINGLE and IDM. In IDM, we detect erosion when measuring morals in repeated decision making, in line with Bartling et al. (forthcoming). However, a partial erosion of morals comparing IDM and SINGLE persists even when controlling for repetition.

We find additional erosion in multi-unit markets, compared to behavior in single-unit markets. Additional erosion is already visible in the comparison of SINGLE and MULTI. Strikingly and more importantly, morals in FULL appear to be completely eroded, in the sense that quantities traded are statistically indistinguishable from selfish competitive equilibrium predictions.

We further show that this behavior is largely driven by the widespread use of the replacement logic. In FULL, 83% of market participants attempt to trade units yielding minuscule gains and comparatively large negative externalities whereas only 16% of participants in SINGLE and 32% in MULTI attempt to trade these units. Furthermore, the quantities traded in the FULL markets largely exceed the predicted quantities based on the preferences of each market's participants obtained in individual decision-making. Therefore, the observed market outcomes are inconsistent with selfish outcomes driven by market selection.

We complement the main analysis with additional findings that shed new light on the effects of and mechanisms behind moral erosion in markets. First, we show that there are detrimental spillover effects of trading in multi-unit markets that carry over to outside of the market itself. After all markets have taken place, we repeat individual decision-making. While we observe partial recovering of morals in this environment, part of the moral erosion attributable to market experience persists.

⁵Following Samuelson & Nordhaus (2005, p. 26), we define a market as a mechanism through which buyers and sellers interact to determine prices and exchange goods and services. It has been argued that the measurement of preferences in price lists may exhibit market features, such as the sharing of responsibility between the subject and the experimenter when a subject chooses money instead of a donation (e.g., Bartling et al. (forthcoming)). To the extent that such considerations actually play a role, this may lead to an underestimation of the true effect of markets on erosion in the comparison of IDM and SINGLE. Our key comparison is between the three markets treatments, which all share a similar structure and between which substantial moral erosion is observed. Notice further that with a price list, subjects do not act as "price-takers", because their decisions affect the expected price in case of a transaction.

This erosion is stronger for the multi-unit markets, and is strongest in FULL. Second, after the markets, we elicit participants' beliefs about the median subjects' morals exhibited in individual decision-making at the start of the experiment. We find that multi-unit markets lead to strongly biased social learning. Subjects in these markets are overly pessimistic about their fellow traders' morals. This points to subjects' beliefs not accounting well for how much market selection and the replacement logic in such markets may shape behavior. Third, we show that our findings are not driven by differences in social norms. The average norm deems trading socially inappropriate equally in each of our market treatments. Still, frantic trading ensues in FULL, against the prevailing norm.

A separate literature studies specific market structures that allow markets to partially sustain pro-social behavior. For instance, Bartling, Weber & Yao (2015) show that allowing for sorting of buyers and sellers into externality-free market segments can be such a force. When intensifying competition in their setup, by increasing the relative number of sellers, the share of pro-social transactions increases.⁶ Other examples in which competition and pro-social behavior can be mutually reinforcing are provided by Byambadalai, Ma & Wiesen (2019) and van Leeuwen, Offerman & Schram (2020).

Our paper contributes to a literature that investigates how the replacement logic and diffusion of pivotality affect behavior in non-market games. Roth, Prasnikar, Okuno-Fujiwara & Zamir (1991), Prasnikar & Roth (1992), Fischbacher, Fong & Fehr (2009) study ultimatum games with proposer or responder competition, which lead to the side with competition receiving almost nothing of the endowment. Bartling & Özdemir (2017) find support for the replacement logic if and only if the selfish action is not perceived as immoral. The games used in Bartling & Özdemir (2017) abstain from a market framing, in addition the sequential nature of their game may remove the possibility to use a replacement logic argument for the last player, and for others by backward induction. In group decision games, also absent market framing, Falk et al. (2020) find support for the replacement logic, whereas Brütt, Schram & Sonnemans (2020) find mixed evidence for the effect of decreased pivotality.

Our takeaway from this literature is that whether or not the replacement logic is activated depends on the setting. To know if it matters in markets, it has to be studied in markets. Notice further that differently from the existing literature, our paper studies the replacement logic vis-à-vis other forces activated by markets such as market selection. Our paper contributes by showing that an unrestrained activation

⁶Bartling et al. (2015)'s findings are robust to different specifications of the externalities (Bartling, Valero & Weber 2019). In a similar spirit, Schneider, Brun & Weber (2020) document an endogenously arising wage premium, and associated sorting, for morally questionable tasks.

of the replacement logic in markets leads to trade of even the most unprofitable units, for which the gains from trade do not compensate the moral costs of even the least moral traders.

In the following, we begin by describing the experimental design. We present the novel features of the experimental markets in detail. We continue by presenting our hypotheses and by describing our results. We conclude by discussing the implications of our findings.

2 Experimental design

The experiment consisted of three main parts. Parts 1 and 3 were identical to each other, and they were the same in each treatment. In these parts, subjects faced an individual decision-making task which elicited their willingness to accept (WTA) to cancel donations towards UNICEF for varying stakes. In Section 2.3, we give more details on the donation opportunity.

For the individual decision-making task, we employed multiple price lists where subjects chose between varying amounts of money and donations to UNICEF. Monetary amounts ranged between €0 and twice the monetary amount of the donation under consideration, with a total of 21 steps in each list. In the experiment, we elicited subjects' WTA for multiple number of units, consistent with potential market trades. Each subject faced separate price lists for 1, 2, 3, 5, 7, 10 and 15 units of donation, in increasing order. We restricted participants to switch at most once in each price list. In our analysis, we set a subject's moral costs equal to the payment at which the subject switched.⁷ We set the moral costs of subjects who never choose to cancel a donation equal to the upper bound of the multiple price list.

Part 2 varied in the four between-subject treatments. In our control treatment (individual decision-making, or IDM), part 2 presented a repetition of the task of part 1 for four times. In the three market treatments, four market periods with negative externalities were implemented. Part 3 repeated the individual decision-making task of part 1 in each treatment.

2.1 Markets

We implemented the market as a double auction with induced values and costs, with some special features. Each market consisted of five buyers and five sellers interacting repeatedly and anonymously. Buyers posted bids, sellers asks, and all traders could

⁷We do this to match behavior in the markets, where we can only infer that a subject's moral costs is at most equal to the profit margin of a submitted or accepted offer.

accept an offer of the other market side. If accepted, a trade was implemented at the price of the accepted offer, where the buyer received a payment of the induced value minus the price and the seller received a payment of the price minus the induced costs. For every unit traded, a donation to UNICEF which costs approximately €1.50 was cancelled.

We ran three market treatments: SINGLE, MULTI and FULL. In each market treatment, traders first participated in a practice market where no externality was present, to make them familiar with the market environment. Afterwards, we implemented four market periods in which every trade caused an externality through the cancelled donations.

The single-unit market treatment, SINGLE, allows for most forces of erosion considered in the current literature. In SINGLE, each trader was restricted to trade at most one unit. In total, up to five units could be traded in a market.

The multi-unit market, MULTI, was implemented identically to SINGLE, with the exception that each trader could trade up to three units. This implies that in each market, 15 units could be traded. We also scaled up induced values and costs exactly proportionally. Doing so, MULTI only differs from SINGLE in the scale of an otherwise identical market.

We allowed each trader to cater to the entire market in the unrestricted market, FULL. This is the market that reduces the market power of individual traders and allows for full replacement and market selection. Treatment FULL was identical to MULTI apart from one key aspect: We removed the capacity constraints of each trader. This means that each participant was able to trade up to 15 units. As a consequence, one buyer and one seller were able to serve the entire market.

To evaluate the effect of different market forces, we designed the experimental market closely following the literature, with two key innovations: (i) the double auction was implemented *sequentially*, (ii) we induced a *common* supply- and demand-schedule.

We implemented a sequential double auction, in which buyers and sellers moved in turns, trading unit by unit. Traders were informed of the costs and values of all units that could be traded in the market.

In each market period, one side of the market – i.e., the buyers or the sellers – was randomly determined to move first. The starting side had the opportunity to submit offers to the second movers within a time constraint of 14 seconds. We restricted all offer submissions to yield non-negative profits for both market sides. Afterwards, the second movers could either accept the most favorable standing offer, or decide to submit a counter offer. A counter offer had to improve upon all preexisting offers.

If no trader accepted an offer, the most favorable counter offer was presented to the original starting side, and traders could again decide whether to accept the most favorable offer or improve upon the best offer they had submitted so far.

If both market sides did not accept or submit an improved offer at least twice, the market period ended and no further units could be traded. Participants were shown a reminder of this feature after neither side had been active once. Whenever an offer was accepted and the 14 seconds time limit had elapsed for all traders currently moving, the trade was implemented for the two agreeing traders. If more than one trader accepted an offer, or if multiple offers were equally favorable, one randomly determined buyer and one randomly determined seller traded, irrespective of the exact time at which an offer was made or accepted.

After a unit had been traded, all pre-existing offers were removed and the previous second-movers were first to propose new offers for the subsequent unit.

These design features have two key advantages: (i) the responding market side has most bargaining power, as they only observe the most favorable offer of the proposers, therefore we obtain relatively tight bounds on the profits proposers deem acceptable; (ii) the speed at which subjects act is removed as a confounding variable, and we generate observations on the willingness to trade for all active traders, so beyond what is normally observed in a traditional double auction where trade is implemented immediately after agreement.

We also innovate in the way values and costs are induced for market participants. We implemented a common schedule for values and costs. The essential difference with a standard double auction is that traders can take advantage of the trading opportunities left by others at the exact same conditions. If one trader foregoes the possibility to trade a unit, another trader can step in and take advantage of this opportunity. In a standard double auction all traders have a private schedule, which implies that if one trader refrains from trading her unit, no other trader can take advantage of that particular trading opportunity. We think that our common schedule matches the markets that are claimed to suffer from replacement arguments more closely.

A consequence of the common schedule is that a buyer's values and a seller's costs at each point only depended on the aggregate number of units traded, and not on the number of units traded by the particular market participant. In our design, costs and values each trader faces were identical and known to all traders. In Figure 1, we plot the costs and values we induced using the common schedule in treatment SINGLE on the left, and treatments MULTI and FULL on the right. The first units were designed such that trade is efficient: The surplus available to traders is larger

than the associated costs to UNICEF by trading these units (surpluses of € 3.80 and € 2.40 compared to a cost of donating of € 1.50). Profitability decreased progressively with subsequent units where market participants could only split € 0.60, € 0.40 and € 0.20.

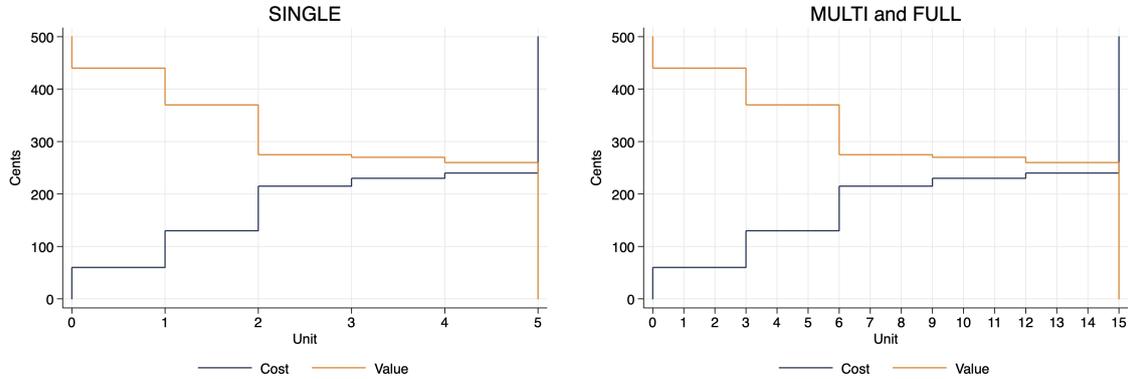


Figure 1: Induced common costs and values

In SINGLE and MULTI, each participant is pivotal for the maximum number of units that they could trade: The entire market could only be served if all traders were active for all units assigned to them. Here, by foregoing a trade, they unilaterally reduced the maximum quantity to be traded by one. We removed these capacity constraints in FULL, so each participant could serve the entire market if she wished to do so.

FULL activates two market forces: the replacement logic and market selection. For the former, if a subject believes that there is at least one other trader on her side who continues to be active, this subject concludes that her actions do not affect aggregate outcomes. This is the case only in FULL. For the latter, unconstrained trading implies that the marginal trades can all be executed by traders with the smallest concerns about causing the externality.

Subjects' trading in the practice market without externalities allows us to see if our special features lead to different market outcomes than previously established in the literature. In all groups, the competitive equilibrium outcome arose in the absence of externalities. Thus, our market produces the same aggregate outcomes as a standard double auction.

2.2 Additional elicitations

To shed light on the mechanisms and guide the interpretation of the results, we included some measurements of subjects' views and attitudes after part 3. We elicited:

(i) beliefs about the median trader’s WTA to cancel donations; (ii) social norms about behavior in individual decision-making and markets; (iii) risk preferences. For the beliefs, subjects were asked to fill in a multiple price list reporting what they “think the average participant did” in the first list of part 1. If their belief matched the choice of the median participant, they received €1. To elicit subjects’ perception of the social norms for canceling donations in either individual decision-making or the market, we followed the procedure by Krupka & Weber (2013) and asked subjects to state whether scenarios described to them were considered “socially appropriate” and “consistent with moral or proper social behavior” on a 4-point scale from “very socially inappropriate”, to “somewhat socially (in)appropriate” and “very socially appropriate”. For one randomly picked scenario, subjects received €2 if their choice matched the modal choice in their session. Among the scenarios described were “[Individual] 1 chooses to receive 1 Euro instead of making a donation of 4 doses of measles vaccine to UNICEF” and “[Individual] 2 decides to accept an offer which allows him to earn 1 EURO”. For the full list of scenarios, see the Appendix Section A.4. We also elicited risk attitudes using the method introduced by Eckel & Grossman (2002).

2.3 Experimental procedures

The computerized laboratory experiment was run in 28 sessions in September and October 2019, at the CREED laboratory of the University of Amsterdam. We preregistered the experiment (Offerman, Romagnoli & Ziegler 2019). In total, 381 subjects participated. 47% were female, with an average age of 21. We had 100 participants per market treatment and 81 participants in IDM. Sessions lasted on average 1.5 hours, with average payments of €19 per subject, besides payments to UNICEF.

Subjects knew that they were paid for only one randomly selected part from the first three parts. All subjects within a session were paid for the same part. If individual decision-making was selected, one decision from one of the multiple price lists was randomly chosen and paid for each subject. If one of the markets was selected, the sum of earnings in two out of the four market periods and the practice market was paid. Additionally, subjects received a show-up fee of €7, all earnings from the three additional elicitation at the end of the experiment as well as an unannounced lump-sum payment of €9 if the markets were selected for payment, to guarantee sufficient minimum earnings.

Subjects read the computerized instructions at their own pace and separately for each part of the experiment (see Appendix Section C). They also received handouts with summaries of the instructions. Subjects were required to complete a set of test

questions before they could proceed. Subjects were paid in cash and in private at the end of the experiment.

In the experiment, several choices affected donations to UNICEF. As in Kirchner et al. (2016), Sutter et al. (2020), donations were intended for measles vaccine. We used a text of UNICEF to inform subjects about the consequences of measles.⁸ One dose of measles vaccine through UNICEF costs approximately €0.375, and two doses are required to vaccinate one person. In the experiment, one unit was chosen to consist of four doses, corresponding to a donation of €1.50. This amount was communicated to subjects in the instructions and the handout.⁹ In the instructions, subjects were presented with sample receipts of such a donation to UNICEF.¹⁰ At the end of each experimental session, the donation was immediately implemented by the experimenter. Subjects were presented with the UNICEF receipt for their session (i) immediately in the experimental interface, jointly with their experimental earnings; (ii) when receiving their earnings in cash; (iii) via email if subjects so desired. These emails were collected on separate handouts and thus could not be linked to specific subjects or choices in the experiment. Subjects were made aware of this procedure at the start of the experiment. In total, approximately €889 was donated to UNICEF as a result of subjects' choices.

3 Hypotheses on how markets erode morals

To evaluate whether and which market forces may contribute to moral erosion, we compare behavior between our four treatments: individual decision-making IDM, and especially the three market treatments SINGLE, MULTI and FULL.

In the first part of each treatment, we elicited subjects' willingness to accept to cancel a donation to UNICEF for seven different stake sizes. We use this data to construct a "moral cost curve" for each subject, by fitting a second-order polynomial to elicited moral costs and subsequently obtain predicted moral costs.

More precisely, for moral costs $\Theta_i(q)$ for subject i and q units, we use the moral costs we had elicited for $q \in \{1, 2, 3, 5, 7, 10, 15\}$ to estimate α_i, β_i in i 's moral cost curve using OLS:

⁸"Measles are highly infectious and very often deadly. Each day hundreds of children become victims of this disease. The survivors often suffer consequences for their whole life, like blindness or brain damages. This, even though protecting the children would be so easy. Measles kills more than 160,000 children worldwide each year."

⁹This particular donation was only available in packs of 40 doses, excess donations were made over to UNICEF as a generic donation, which subjects were aware of and could verify as well.

¹⁰At the time of the study, this donation is available at <https://market.unicef.org.uk/inspired-gifts/measles-vaccines-to-protect-20-children/S359163X/>, which we also communicated to subjects.

$$\Theta_i(q) = \alpha_i q + \beta_i q^2 + \epsilon_{j,q}$$

Where $\epsilon_{j,q}$ is an individual-unit error. We can use this to predict moral costs, based on estimated $\hat{\alpha}_i, \hat{\beta}_i$. Then we construct marginal, for unit j , moral costs $\theta_i(j)$, where $\Theta_i(q) = \sum_{j=1}^q \theta_i(j)$. These are predicted marginal moral costs $\widehat{\theta}_i(q)$ to cancel donations for all 15 units, $q \in \{1, 2, \dots, 15\}$, for each subject i .

In the market treatments, we use the moral cost curves to predict market outcomes under the assumption that markets do not erode morals. That is, we use traders' estimated moral costs $\widehat{\theta}_i(q)$ and predict how many units we would expect to be traded if $\widehat{\theta}_i(q)$ is *not* affected by moving to our market setup, given the market rules of the treatment subjects are participating in. As the literature finds repetition to be a force behind erosion (Bartling et al. forthcoming), we correct moral costs in each period by the average erosion we find in the corresponding repetition of the treatment IDM.

For this, starting from the first unit, we randomly draw a buyer b , with marginal moral costs to trade an additional unit of $\widehat{\theta}_b(q)$, and a seller s , with marginal moral costs of $\widehat{\theta}_s(q')$. If the sum of the two moral costs do not exceed the available gains from trade, given by the difference in induced values and costs, this pair of traders is designated to trade. Afterwards, we proceed to the next unit, and repeat the procedure. If the marginal moral costs of the pair (b, s) exceed the gains from trade, we attempt to find 200 times a pair for whom trading is feasible. In drawing random pairs, we keep track of the number of units previously traded, which may affect marginal moral costs or individual capacity constraints. At the point where no further pair can be found, the predicted quantity is the last unit which can be traded. Our predictions are the average outcome of 10,000 simulations, to account for differences in drawing random buyer-seller pairs.¹¹ For a more detailed description of the implemented analysis, see the Appendix Section A.2.

We call the outcome of this exercise the “competitive equilibrium with moral costs” or “moral competitive equilibrium”. Note that this exercise is only possible in a design such as ours, where we observe participants both in individual decision-making and in a market environment. The benchmark of the moral competitive equilibrium allows us to: (i) disentangle whether observed market outcomes can be reconciled with the preferences of market participants or whether markets do erode morals; (ii) carry out counterfactual simulations to highlight the role of market selection.

¹¹In order to focus on the most relevant equilibria, we keep those trades observed in the experiment which are consistent with traders' moral costs in our simulations.

Regarding (i), we compare the degree of moral erosion by ranking the extent to which observed quantities exceed predicted quantities in the moral competitive equilibrium between treatments. This is of particular interest in treatment FULL: due to market selection, the least moral traders can determine quantities by themselves. If preferences are heterogeneous, or additionally if marginal moral costs are strongly decreasing for some of the traders, predicted quantities in the moral competitive equilibrium are higher in FULL than in MULTI or SINGLE. Under market selection, aggregate market outcomes in FULL appear to be more selfish than we would expect on the basis of homogeneous traders having the same median preferences. However, this does not imply that moral costs have eroded in markets, it just represents the fact that the traders least concerned with causing an externality are setting quantities. These traders might not be representative of the median trader. In the analysis, we will use each trader's estimated moral costs to verify whether her trading behavior is consistent with her stance outside of markets.

The possibility to run counterfactual simulations, in (ii), provides another important advantage of the moral competitive equilibrium. In predicting quantities, we use the estimated moral cost curve. In this process, both initial heterogeneity in how subjects care about the externality as well as heterogeneity generated through decreasing marginal moral costs enter in the prediction. As a counterfactual scenario, we simulate a homogeneous moral cost equilibrium, where we set the moral costs for all units equal to the median subject's moral cost for the first unit, within their market. The difference between the two benchmarks provides a measure of how severe market selection is, or, how well markets are predicted to reflect the preferences of an average market participant.

If market quantities exceed the predictions of the moral competitive equilibrium, markets do erode morals in the sense that traders care less about the externality they cause in a market than outside of a market. The literature so far studied a number of forces that might explain such erosion. For example, in contrast to individual decision-making, in the market participants share guilt when they decide to trade a unit. Also, people may respond differently when they face a market frame. These forces are present in all market treatments, and our main interest is in the question whether there is moral erosion in multi-unit markets beyond what is found in the comparison between SINGLE and IDM.

In particular, in MULTI and FULL, we study two additional forces for moral erosion. First, compared to SINGLE, participants interact repeatedly in a market, as each participant can now trade multiple units which enhances the scope for learning from the actions of others. Here, social learning might be biased, as frequently

trading participants are salient in this procedure; whereas it is more difficult to notice that some traders are actually refraining from trading at low profits while many units are being exchanged. If participants do not recognize this feature of multi-unit markets, and they for example tend to imitate others' behavior, we would expect additional erosion of moral costs in MULTI compared to SINGLE. Second, moving from MULTI to FULL, individuals no longer have any ability to affect aggregate quantities as long as other participants continue trading. This activates the replacement logic: participants might decide to reap the benefits from trading themselves, given that some other participant will trade otherwise and the corresponding negative externality cannot be prevented anyways.¹²

In the Results Section, we test a number of preregistered hypotheses (Offerman et al. 2019). We summarize our hypotheses below and also state our expected findings.

H1. Shared social responsibility and market frame do not matter.

For hypothesis H1, we compare behavior between IDM and SINGLE. We expected to reject H1.

H2. Social learning does not matter.

For hypothesis H2, we compare the normalized market outcomes between SINGLE and MULTI, as quantities relative to competitive equilibrium and moral competitive equilibrium benchmarks. We expected to reject H2.

H3. Replacement logic and market selection do not matter.

For hypothesis H3, we compare the normalized market outcomes between MULTI and FULL, as quantities relative to competitive equilibrium and moral competitive equilibrium benchmarks. To decompose the effect due to replacement logic from the market selection effect, we study the identity of traders trading the final unprofitable units in the market. If market selection drives the effect, marginal trades are executed by traders with the lowest moral costs, but at profits still consistent with their moral costs. If however the replacement logic is responsible for the effect, more traders are active, irrespective of their moral costs. We do so by comparing what share of traders are how active at the final units, splitting traders in above and below median moral costs within their market. We expected to reject H3.

¹²We also analyze erosion of morals on an individual level. For this, we use submitted offers and acceptances during the market to approximate a market moral cost curve on the profits participants made on offers and acceptances. However, these are likely an overestimation of the underlying moral costs, as profits in this exercise are assumed to just cover moral costs. We reproduce this analysis in the Appendix Section A.3. Results for SINGLE and FULL are in line with the analysis presented in the main part, we discuss results for MULTI in Result 2.

H4. The erosion of moral costs is a temporary phenomenon.

For hypothesis H4, we compare estimated moral costs before and after market participation between treatments. We expected not to reject H4.

H5. Social norms are not affected by market participation: there are no treatment differences in social norms.

We expected to reject hypothesis H5.

H6. Errors in the beliefs about the median moral cost are constant between treatments.

We compare the correctness of beliefs about the median participants' moral cost between treatments. We did not have a directional prediction for H6.

H7. Markets without externalities converge to the competitive equilibrium.

We confirmed H7 in all markets, and will not analyze this data further.

4 Results

We now present the results of the experiment. For all market outcomes, we perform tests based on averages on matching-group levels, which yields 10 observations per market treatment (10 groups with 10 participants each per treatment). For all hypotheses on the individual level (H4 to H6), for which participants do not interact, we study individual level data. In graphs, we bootstrap confidence intervals, as we frequently plot proportions close to 0% or 100%, so observe skewed distributions. We plot bias-corrected confidence intervals, allow for clustering on a matching group level (the market group for market treatments and the session for IDM) and use 10,000 replications.

4.1 Can market selection matter?

Market selection means that at each moment in a market, the market participants who care least about causing an externality will be deciding on whether additional units will be traded. This force is unrestricted in the FULL market, as each market participant can serve the entire market, which cannot be prevented by other traders.

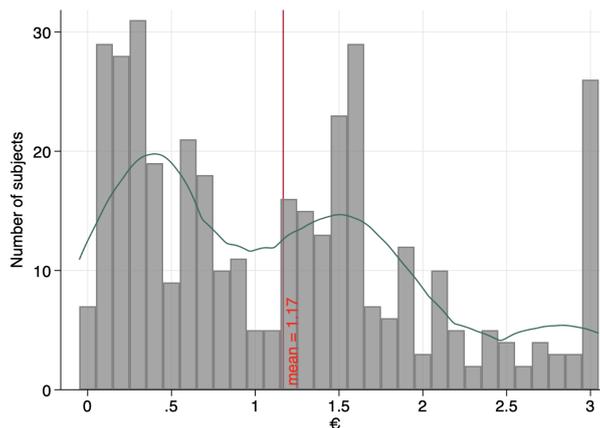
Market selection in multi-unit markets can matter in two dimensions: (i) initial heterogeneity in how traders value donations, and (ii) decreasing marginal moral

costs when causing an externality. Potentially, the two dimensions may interact. For all treatments, we obtain estimates of moral costs of the subjects, and check whether these two features are indeed present in the data.

From choice data obtained in multiple price lists for 1, 2, 3, 5, 7, 10 and 15 units of donation to UNICEF, we fit a quadratic polynomial for each subject. We then predict the marginal moral costs for units 1 to 15.

In Figure 2, we provide a histogram of the predicted per-unit moral costs of all subjects in part 1 of the experiment, averaged at the subject level. We estimate for what payment to herself a subject is willing to cancel a donation to UNICEF, which costs €1.50. Evidently, there is substantial heterogeneity in how subjects value the opportunity to donate to UNICEF. A minority of subjects hardly cares about donating to UNICEF. There is also a remarkable share of subjects whose moral costs are estimated to be above €1.50, implying that they value donating *more* than the corresponding monetary value.¹³

Figure 2: Heterogeneity in valuations of donations



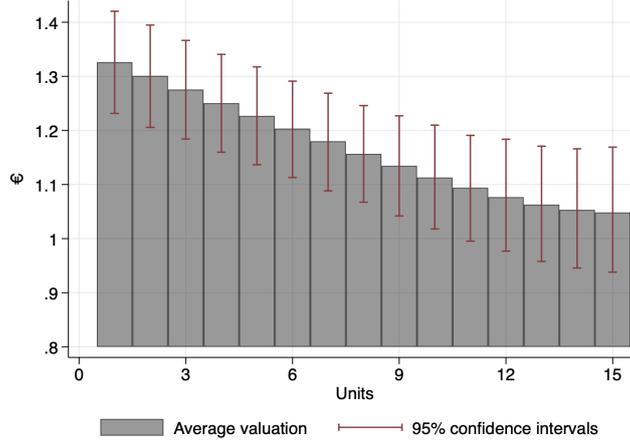
Notes: Histogram of average estimated per-unit moral costs for 1 to 15 units of donation with a value of €1.5 each, by subject. Kernel density in green, mean in red.

In Figure 3, we provide evidence of a desensitization effect; per-unit moral costs decrease over units. At larger stakes, subjects need to be paid less such that they are willing to cancel an additional unit of donation. The effect is quite strong: For the first unit, subjects on average reported moral costs of €1.33, this decreases to €1.05 for the fifteenth unit.¹⁴

¹³Bénabou, Falk, Henkel & Tirole (2020) show that elicited moral costs can be affected by the method of elicitation, when using either direct elicitation or multiple price lists, as image motives are differentially affected by these methods. In our data, we find only few “observationally deontological” subjects, those who never cancel a donation across all price lists, as only 10 out of 381 subjects do so across part 1, compared to 26% of subjects who do not cancel the donation for any monetary amount in Bénabou et al. (2020).

¹⁴Regression evidence supports decreasing marginal moral costs. For a more in-depth analysis of the desensitization effect, we refer to Appendix Section A.5.

Figure 3: Decreasing marginal moral costs



Notes: Average predicted per-unit moral costs for 1 to 15 units of donation with a value of €1.5 each, by unit.

Given these data, there is a clear possibility for market selection to play an important role. We will discuss this theme in the following Section.

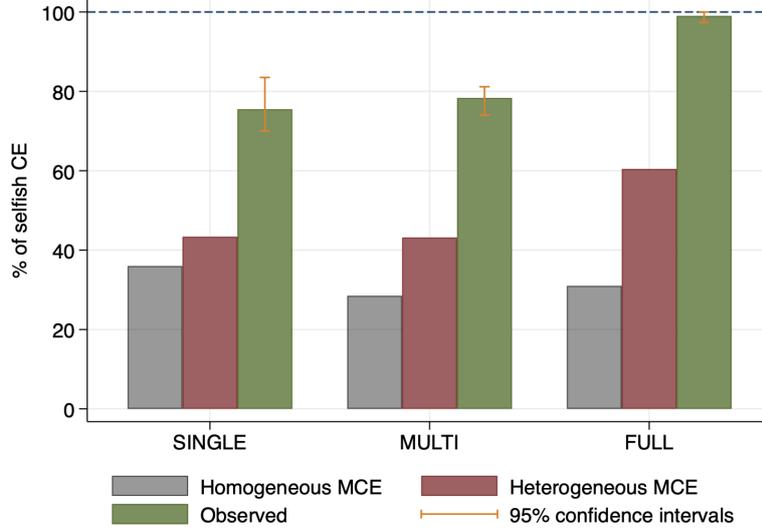
4.2 Market outcomes and moral competitive equilibria

In this Section, we start with a comparison of market outcomes and the moral competitive equilibria, before we proceed to test the key hypotheses of the paper. In Figure 4, we plot the observed market quantities in combination with the quantities predicted in the competitive equilibria with moral costs, as introduced in Section 3. All quantities are relative to the selfish competitive equilibrium outcome. In it, 5 units will be traded in SINGLE, and 15 units in MULTI and FULL, as the associated externality is ignored by traders in this benchmark.

The first bar for each treatment, in grey, shows the predicted quantity in the competitive equilibrium with homogeneous and constant moral cost. For each market, we use the median trader's moral costs for the first unit to simulate how many units will be traded on average. Average quantities are between 28.5% and 36%. These differences between treatments are purely driven by initial heterogeneity of subjects, and are not related to underlying market features.¹⁵ As it turns out, our subjects valued donations to UNICEF somewhat higher in FULL and MULTI than in SINGLE.

¹⁵Note that average moral costs between treatments are quite similar. However, for this exercise, we rely on distributions of the median, where we continue to see some variability between treatments.

Figure 4: Market outcomes and competitive equilibria



Notes: Average quantities relative to selfish competitive equilibrium for two moral competitive equilibrium (“MCE”) benchmarks and observed quantities. MCE use participants’ moral costs elicited in individual decision-making to predict market quantities. Heterogeneous MCE are based on actual moral costs, homogeneous MCE are based on the median trader’s moral cost for the first unit within the matching group.

The second bar, in red, shows predicted quantities given the actual, heterogeneous moral costs of market participants. These are higher quantities in all treatments than in the homogeneous moral competitive equilibrium. As expected, the differences are largest in FULL. The difference between the two equilibria can be attributed to market selection: the least moral traders in FULL are no longer constrained, thus they can expand the size of the market. This market force increases quantities by 29.4 percentage points. In MULTI and SINGLE participants’ heterogeneity has a smaller impact on traded quantities. Whereas in SINGLE the increase is only 7.4 percentage points, this increases to 14.7 percentage points in MULTI.

The third bar, in green, shows observed quantities across the three treatments. We see that there is erosion of moral costs in all treatments. We observe partial erosion of morals in SINGLE and MULTI. In FULL, market outcomes are fully selfish. Compared to the competitive equilibrium with heterogeneous moral costs, quantities increase in SINGLE. They increase stronger in MULTI, and by even more in FULL.

Consistent with the partial erosion of morals measured in SINGLE when comparing quantities to the moral competitive equilibria, we find erosion in single-unit markets based on the measures used in the literature. We compare the share of subjects who are willing to cancel a donation for at most its cost, €1.50, in IDM

to the share of traders who conclude a trade for which they receive a payment of at most €1.50 in SINGLE. In the first period of markets, this applies to 62% of traders, whereas in IDM only 49.4% of subjects cancel one unit of donation when paid €1.50 or less. If we compare the lowest reported valuation in the entire part 2, we find evidence for erosion in markets as well as for erosion through repetition: 59.3% of subjects in IDM cancel a donation for less than its value at least once across the four periods, while 89% of subjects trade with at most such profits. Summarizing, we find evidence for both a partial erosion of morals in markets (Falk & Szech 2013, Bartling et al. 2015) as well as erosion when measured by a subject cancelling a donation at least once in a repeated task (Bartling et al. forthcoming). We provide a more detailed analysis of this result in the Appendix Section A.6.

Result 1 *We reject hypothesis H1, and find partial erosion of morals in single-unit markets. Market forces such as shared social responsibility or market framing contribute to moral erosion.*

Our key hypotheses are on behavior in multi-unit markets. We want to establish whether there is an erosion in this market, *in excess* of the erosion we find in single-unit markets. For this, we return to observed quantities in markets. In Table 1, we summarize market quantities relative to the selfish competitive equilibrium quantities together with p -values of Mann-Whitney U-tests (10 observations per treatment) of quantity comparisons between treatments.¹⁶

Table 1: Treatment effects

| | SINGLE | MULTI | FULL |
|------------------------|--------|-------|-------|
| Quantity in % | 75.5 | 78.3 | 99 |
| p -values vs. SINGLE | - | .378 | .0005 |
| vs. MULTI | - | - | .0001 |

Notes: Average quantities relative to selfish competitive equilibrium. Mann-Whitney U-tests, on matching group averages, 10 observations per treatment.

A comparison of SINGLE and MULTI sheds light on the question whether the enhanced possibility for social learning in MULTI contributes to moral erosion. Quantities in SINGLE and MULTI are quite similar and not statistically distinguishable. Also, in comparison to the moral competitive equilibrium benchmark, we find only some additional erosion between SINGLE and MULTI. However, moral costs curves

¹⁶These treatment differences also arise when regressing quantities on treatment indicators, with and without controlling for period indicators, moral costs (average, median and minimum within matching group), as well as risk measures; see the Appendix Section A.1 for results.

estimated on offer data exhibit erosion when comparing SINGLE to MULTI, see the Appendix Section A.3. Even though we do not find an effect of social learning on the aggregate quantity traded, subjects appear to be somewhat more eager to trade in MULTI.

Result 2 *We find mixed evidence against hypothesis H2. The enhanced possibility for social learning in MULTI leads to an at most modest effect on moral erosion.*

In FULL, however, quantities are substantially increased and statistically indistinguishable from the selfish competitive equilibrium, as almost all units are being traded. Strikingly, morals appear to be entirely eroded in unrestricted multi-unit markets. Note that this includes units where participants can split €0.20, while causing negative externalities of €1.50. The difference between MULTI and FULL is significant at conventional levels.

Moral erosion in FULL is particularly strong, even though differences between observed and heterogeneous moral competitive equilibria might appear not too different between MULTI and FULL at first sight in Figure 4.¹⁷ Erosion is much stronger in FULL as additional units traded are causing larger negative externalities the more units have already been traded, relative to the potential gains from trade. This is the case as the induced gains from trade are decreasing at higher quantities, while damages stay constant. Below 40%, trading is efficient, as the damage to UNICEF is less than the associated payments to market participants. An increase from 40% to 60% leads to additional negative externalities of €4.50, whereas traders receive €1.80. The gains relative to damages to UNICEF decrease further, and an increase from 80% to 100% also yields damages of €4.50, however traders only receive total payments of €0.60. To quantify the size of the erosion, we summarize how many additional units compared to the moral competitive equilibrium benchmark are traded in each treatment in Table 2. We also show what damages to the donation traders are willing to accept for an additional payment of €1 per additional unit that is traded. Damages, and the associated erosion of moral costs, are highest in FULL. Consistent with this analysis we also find stronger erosion in FULL than in MULTI based on the moral cost curves estimated on offer data, see Appendix Section A.3.

¹⁷The severity of erosion compared to the moral competitive equilibrium benchmark in FULL also arises when making the very conservative assumption that moral costs are halved in markets compared to individual decision-making. This would account for example for traders sharing guilt equally with their trading partner, and does not require a comparison to our treatment MULTI to identify erosion. Predicted average quantities in the moral competitive equilibria with halved moral costs are 11.1 units, an increase from the 9.1 units in the baseline simulation. This quantity still cannot match the observed traded quantities of 14.9 units.

Table 2: The size of erosion in markets

| | SINGLE | MULTI | FULL |
|-----------------------------|--------|-------|------|
| Normalized units | 4.8 | 5.3 | 5.8 |
| Damage per €1 gain per unit | 3.2 | 3.2 | 4.9 |

Notes: Number of units traded beyond heterogeneous moral competitive equilibrium as well as damages to UNICEF on average per additional unit, normalized across treatments. Damage per unit is fixed at €1.50, gains from trade vary between €0.20 and €3.80.

Result 3 *We reject hypothesis H3. We find strong erosion of morals in FULL, compared to MULTI. The replacement logic and market selection are market forces that play important roles.*

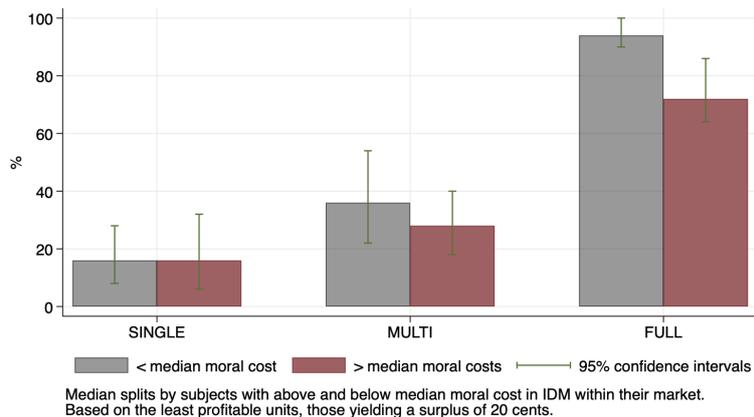
4.3 Market selection versus replacement logic

The preceding Section presented evidence for a complete erosion of morals only in FULL markets. In the competitive equilibria with moral costs, we observe that the final units in the market are not predicted to be traded, yet they are almost always traded in the experimental markets. This is inconsistent with the selfish outcomes in FULL being driven by market selection. In this case, we would be able to approximate market outcomes by sampling the traders with lowest moral costs. However, we see that even those traders are not willing to trade units with gains from trade of €0.40 or €0.20 when causing an externality of €1.50. Instead, the replacement logic seems to be a major force in driving up market quantities.

To study the robustness of this result, we use a second measure of whether traders fall prey to the replacement logic. We study whether traders are active in the sense of submitting or accepting offers at the final units, the units which yield gains from trade of €0.20. To evaluate what type of trader is active we split the sample into those with below and above median moral costs in part 1.

In Figure 5, we plot the share of traders who are active at least once at these least profitable units. We see that in both SINGLE and MULTI, both groups of market participants are similarly active. However, this share is much higher in FULL, where 94% of traders with below median moral costs are active, but even 72% of traders with above median moral costs are active. This points to a strong erosion attributable to the replacement logic. Traders with above median moral costs clearly are not predicted to be as active at these units, unless they employ the replacement logic to excuse their behavior.

Figure 5: Share of traders active at the least profitable units

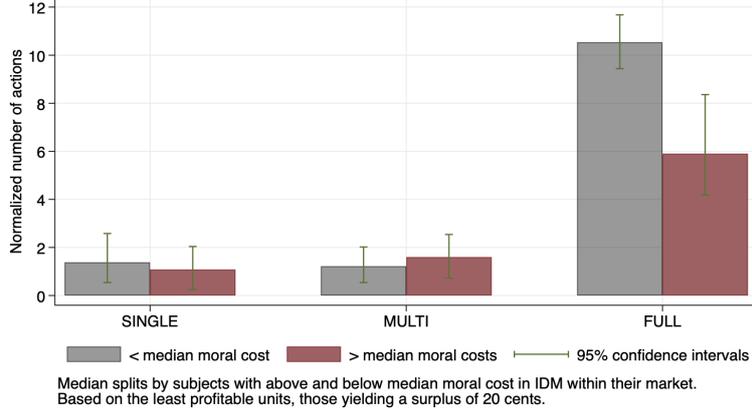


Notes: Share of traders who submit or accept an offer at the final units, which yield gains from trade of €0.20 in exchange for an externality of €1.50. Median splits based on predicted moral costs within matching group.

In Figure 6, we also show the intensive margin of this phenomenon: how many offers and acceptances do we observe from traders? To normalize the number of actions per trader across treatments, so to account for the smaller total market size in SINGLE, we multiply the observed number of actions in SINGLE by 3. We again observe that erosion due to replacement logic appears to matter most. Frequent trading of both types of traders is observed in FULL, with 8.2 actions per trader observed on average, whereas in SINGLE and MULTI only 1.2 and 1.4 actions per traders are observed on average. In the Appendix Section A.7, we also show this statistic when we divide traders in whether or not they reported average moral costs of higher than €1.50 in part 1. This captures traders who revealed not to follow consequentialistic reasoning, as their moral costs exceed the underlying value of the donation. Even those participants are active in FULL, whereas they are much less active in SINGLE and MULTI.

Result 4 *In unrestricted multi-unit markets, market outcomes converge to selfish competitive equilibrium. This is mostly driven by the replacement logic, and less so by market selection.*

Figure 6: Number of acceptances and offers at the least profitable units



Notes: Average number of offer submissions or acceptances per trader at the final units, which yield gains from trade of €0.20 in exchange for an externality of €1.50. Median splits based on predicted moral costs within matching group.

4.4 Effects of market participation

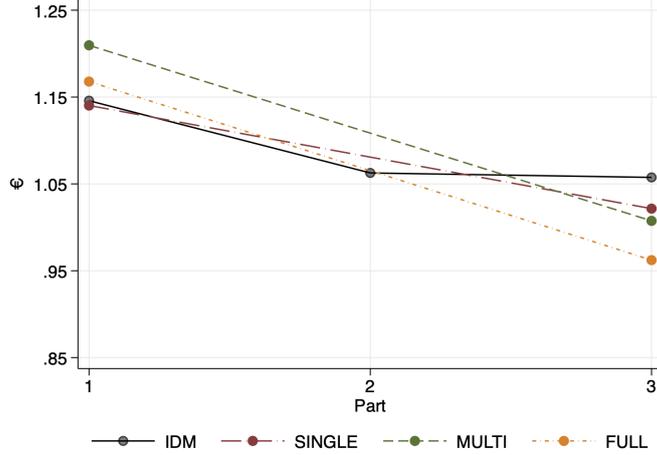
In the following, we explore hypotheses H4 to H6, which all concern effects of market experience on choices made outside the market itself.

As a first question, we hypothesized that any erosion of moral costs we find within markets is temporary. In all treatments, we elicited subjects' valuations of donations in part 3, identically to the start of the experiment. Treatment IDM allow us to study whether repetition by itself is eroding morals. Comparing this erosion to erosion after experiencing markets in treatments SINGLE, MULTI and FULL allows us to determine whether the erosion in markets does have an effect outside the immediate market environment. In addition, we can evaluate whether specific market features lead to stronger erosion outside the market.

In Figure 7, we plot the average predicted moral costs per treatment, by parts. In treatment IDM, we elicit moral costs in parts 1, 2 and 3. In the market treatments, we use individual decision-making only in parts 1 and 3.

We do observe that moral costs are decreasing over time. In IDM, average per-unit moral costs in part 3 decrease by 8.8 cents (relative to a donation of €1.50), compared to the moral costs in part 1. This change slightly increases in the markets, in SINGLE it amounts to 11.9 cents. In the multi-unit markets MULTI and FULL, erosion is most drastic, with decreases of moral costs of 20.2 cents and 20.6 cents, respectively, after market exposure. Across all treatments, this decrease is significant (Wilcoxon signed-rank tests, 100 observations per market treatment, 81 in IDM, p -values of IDM=.0051, SINGLE=.0553, MULTI=.000, FULL=.0003). Comparing the decrease

Figure 7: Persistence of erosion



Notes: Average predicted moral costs in individual decision-making, for €1.50 donations, by part. In part 2, only IDM employs individual decision-making.

between treatments, we do not find significant differences between IDM and SINGLE (Mann-Whitney U-test, 81/100 observations, p -value=.5513). We do find that multi-unit markets in turn show stronger erosion, as the decrease in MULTI compared to SINGLE is significant (MWU, 100 observations per treatment, p -value=.0107), while the decrease between MULTI and FULL is similar (MWU, 100 observations per treatment, p -value=.3022). This indicates that, to our surprise, erosion of morals does seem to persist outside of markets, especially so in multi-unit markets. Repetition does seem to contribute to erosion as well, but its role appears to be less pronounced than that of multi-unit market exposure.

Result 5 *We reject hypothesis H4: erosion of moral costs persists outside of the market, most so in multi-unit markets.*

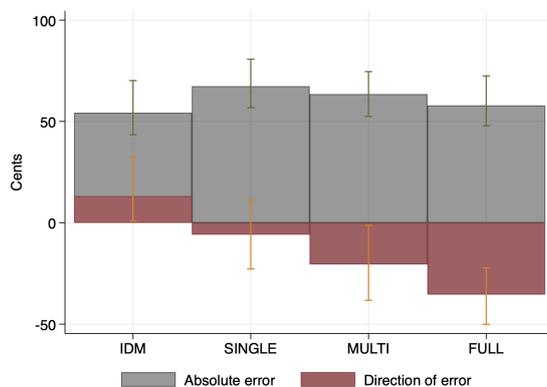
As a second question, we investigated whether participation in our experimental markets affects traders' beliefs about the median trader's moral costs. For this, subjects were paid if they correctly estimated the median participant's choices in the first multiple price list, for the first unit in the first part of the experiment, within their session.

In Figure 8, we report the mean difference from prediction to observed moral cost of the median trader in red, as well as the absolute prediction error in grey. At first sight, observing their fellow peers does not help participants to improve their estimate: the absolute error is not decreasing in the markets compared to IDM. Also, there do not appear to be strong differences between the market treatments.

Interestingly, the direction of the error does change systematically between treatments. If anything, participants in IDM slightly overestimate how much the median participant values a donation to UNICEF. While there is a slight decrease in SINGLE, the multi-unit markets MULTI and FULL lead to systematic errors: participants strongly underestimate how much their participants care about donations for the measles vaccine.¹⁸

There is biased social learning in the sense that participants believe that their peers are more selfish than they truly are.¹⁹ Participants do not sufficiently take into account that other traders' behavior in the market is to a large extent shaped by market forces. This is also consistent with multi-unit markets complicating inference about the moral costs of traders who are less active. Market participants observe frequent trading, but do not comprehend that this does not necessarily reflect the preferences of an average participant outside of the market. In the Appendix Section A.8, we provide some correlational evidence that those traders who are most strongly underestimating their peers' morals are those who adjust their moral cost most strongly downward in part 3, compared to part 1. Therefore, biased social learning and any resulting biased imitation may be a channel of sustaining erosion of morals outside of markets.

Figure 8: Errors in beliefs about median subject's moral cost



Notes: Average error in estimating the session's median subject's moral cost for cancelling one unit of donation of €1.50 in part 1 of the experiment. In grey the absolute distance between prediction and target, in red the difference between prediction and target.

¹⁸Within treatment, we find no evidence of excuse-driven belief reports of beliefs, see Appendix Section A.9.

¹⁹Regressing (381 observations, clustering standard errors on matching group level) subjects' absolute errors on treatment fixed effect shows insignificant dummies (all p -values $> .1$). Regressing the error on treatment fixed effects shows differences in fixed effects, compared to the IDM baseline, for SINGLE (estimate of -19.0, p -value=.112), MULTI (estimate of -33.7, p -value=.009) and FULL (estimate of -48.5, p -value $< .001$).

Result 6 *We reject hypothesis H6. Social learning is differentially biased in multi-unit markets, as participants become overly pessimistic about their fellows' morals.*

As a third question, we hypothesized that market participation might affect social norms about what constitutes behavior that is “consistent with moral or proper social behavior”. To investigate this, we elicited social norms using the method introduced by Krupka & Weber (2013). Subjects were incentivized to report what they believed was their session’s modal answer on a 4-point scale from “very socially inappropriate”, to “somewhat socially (in)appropriate” and “very socially appropriate” in response to scenarios in which a participant in an experiment chose to cancel donations of € 1.50 when paid € 1 either in individual decision-making or in an experimental market.

In Figure 9, we report the mean report of the identical norm question, in which a participant cancels donation in individual decision-making in the left panel, and in a market in the right panel. We observe that across all treatments, cancelling a donation for this monetary gain is rated on average at most as “somewhat socially inappropriate”. Thus, there does exist a clear social norm that trading is not appropriate, contrary to our observed trading behavior especially in FULL.

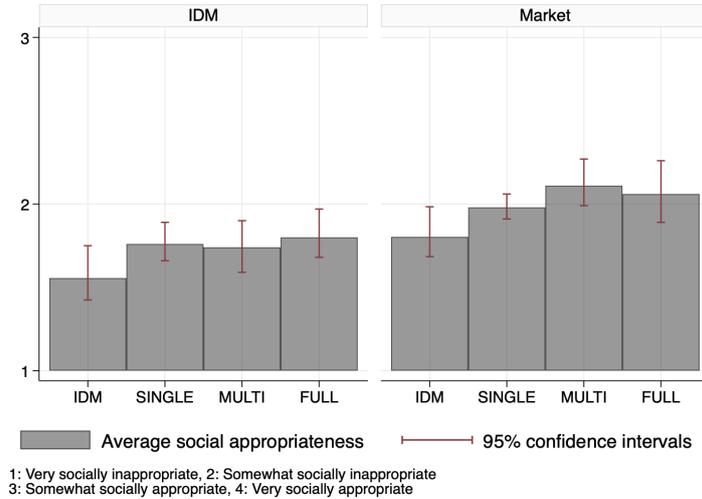
In accordance with our expected finding, causing an externality in a market is perceived as less inappropriate as the same choice outside of a market frame (Wilcoxon signed-rank, 381 observations, p -value=.000). Somewhat surprisingly, this does not apply to differences between treatments. It does not seem to be the case that the more selfish behavior we found in the multi-unit markets, especially in FULL, is also reflected in changes in social norms.²⁰ We cannot reject equality of norms in markets comparing IDM and SINGLE (Mann-Whitney U-test, 81/100 observations, p -value=.111), between SINGLE and MULTI (MWU, 100 observations per treatment, p -value=.238) and between MULTI and FULL (MWU, 100 observations per treatment, p -value=.705).²¹ We report additional descriptive statistics for other scenarios in the Appendix Section A.4, which yield similar conclusions.

Result 7 *We do not reject hypothesis H5. Social norms are not differentially affected by the market treatment. The finding that market outcomes are most selfish in FULL cannot be explained by the prevailing social norm.*

²⁰As for beliefs, we find no evidence of excuse-driven norm reports, see Appendix Section A.10.

²¹Results are similar when regressing subjects’ norms (2 elicitation for 381 subjects, so 762 observations) on treatment fixed effects, a dummy for the market scenario and interactions of this dummy with the treatment fixed effects, clustering standard errors at the matching group. Significant is the dummy for the market scenario (p -value < 0.001), but none of the interactions is significant (all p -values > .1). This confirms that there is not a specific treatment effect on social norms in markets.

Figure 9: Social norms in individual decision-making and in markets



Notes: Average social norm in response to cancelling one donation of €1.50 when paid €1 in individual decision-making (left panel) and in the experimental market (right panel). A rating of 2 corresponds to “somewhat socially inappropriate”.

5 Discussion

In this paper, we study market forces that can lead to a widespread erosion of morals and selfish market outcomes. As market power is reduced by allowing each trader to serve the entire market by herself, we show that aggregate outcomes as well as the behavior of a large share of market participants change dramatically.

Our paper provides conclusive evidence that markets can erode morals. We start by documenting that markets which retain pivotality of individual traders lead to a partial erosion of morals, as we observe trading of units not justified by the individual valuation of the negative externality obtained in price lists. These results support Falk & Szech (2013)’s and Bartling et al. (2015)’s conclusion that single-unit markets partially erode morals.

We then expand the analysis of markets by introducing multi-unit trading and removing pivotality. These changes lead to a full erosion of morals. Participants appear to entirely disregard their moral concerns towards preventing negative externalities in markets. Meanwhile, they are willing to forgo substantial amounts of money before and after markets in an individual decision-making task.

In markets without capacity constraints, we further investigate the relative role played by market selection and the replacement logic in deteriorating market outcomes. We show that market selection is indeed predicted to increase the number of units that will be traded. However, the identity of the active traders reveals the

replacement logic as the key force at play. Trading may become justifiable if traders believe another market participant will otherwise make use of trading opportunities in a self-reinforcing cycle. This logic seems to apply to a very large share of traders: 83% of subjects are willing to trade when they can share gains from trade of €0.2, whereas only 9% of these same subjects are willing to cancel the first donation when each is paid €0.2 in individual decision-making, averaged on part 1 and 3-data.

It is particularly interesting and worrisome to see how far the replacement logic can push trade in markets without pivotality. In principle, moral subjects prone to replacement logic arguments could coordinate to trade only as long as the underlying reported valuation of the least moral subjects justifies trading. If this was the case, the identity of the traders may change, but aggregate market outcomes would not reflect any erosion of morals. Under this rule, trading would stop at the moment where further trade produces gains from trade that are too small to convince even the least moral subjects to engage in trading.

However, the replacement logic can and does push trade further. Now, units with gains from trade that are minuscule are being traded, which is inconsistent with almost all participants' moral costs. Our data suggests that it may be the absence of pivotality that leads to an erosion of morals in markets, rather than a generic effect of participating in a market.

Conceptually, the replacement logic can be used to justify arbitrary trading volumes under two conditions in our setting: (i) enough subjects (i.e., at least 2 on each side of the market) use consequentialistic reasoning, in that they suffer from moral costs only if they are pivotal (Falk et al. 2020); (ii) these subjects hold the self-reinforcing belief that others are willing to trade a particular unit. If these conditions are met, any quantity between the moral competitive equilibrium and the selfish competitive equilibrium can be traded in our FULL markets.

This leads to the important empirical question which of these equilibria is most likely to emerge in each environment. As it turns out, in our market setting almost always the equilibrium associated with the aggregate quantity corresponding to full selfishness was selected. Absent pivotality, large shares of participants engage in frenzied trade of units which cause large damages compared to the available gains from trade. Strikingly, this frenzied trading contrasts with the social norm. Subjects do not perceive it to be socially appropriate to trade these units. This led to widespread frustration among subjects, some of whom spontaneously wrote down their thoughts after the experiment. One subject commented that “The level of selfishness displayed on market 2 has almost made me cry during the experiment. Today, my faith in humanity has taken a giant blow.”.

The large erosion of morals we detect has also implications for our understanding of markets as aggregators of preferences. Using market outcomes to infer individuals' preferences regarding damages to third parties is complicated by key market design features. Participants can behave very selfishly and quite generously depending on specific features of the market structure. A poor understanding of the forces that apply in a given environment might fundamentally lead to a misrepresentation of individuals' preferences. In this sense, markets may not aggregate preferences in a straightforward way. Aside from concerning economists attempting to estimate preferences, this inference problem affects market participants themselves: Our subjects strongly underestimate how much their peers care about the donation to UNICEF after having participated in multi-unit markets. This brings up another potential danger of inference from market outcomes: We might be systematically underestimating by how much fellow members of our society would actually want to prevent the externalities they cause. This in turn could lead to biased social learning and biased imitation (which we find suggestive evidence for in Appendix Section A.8) and an underestimation of the demand for regulating such markets.

References

- Barr, S., Shaw, G. & Coles, T. (2011), ‘Times for (Un)sustainability? Challenges and opportunities for developing behaviour change policy. A case-study of consumers at home and away’, *Global Environmental Change* **21**(4), 1234–1244.
- Bartling, B., Fehr, E. & Özdemir, Y. (forthcoming), ‘Does market interaction erode moral values?’, *The Review of Economics and Statistics* .
- Bartling, B. & Özdemir, Y. (2017), ‘The limits to moral erosion in markets: Social norms and the replacement excuse’.
- Bartling, B., Valero, V. & Weber, R. (2019), ‘On the scope of externalities in experimental markets’, *Experimental Economics* **22**(3), 610–624.
- Bartling, B., Weber, R. A. & Yao, L. (2015), ‘Do markets erode social responsibility?’, *The Quarterly Journal of Economics* **130**(1), 219–266.
- Bénabou, R., Falk, A., Henkel, L. & Tirole, J. (2020), ‘Eliciting moral preferences: Theory and experiment’.
- Brütt, K., Schram, A. & Sonnemans, J. (2020), ‘Endogenous group formation and responsibility diffusion: An experimental study’, *Games and Economic Behavior* **121**, 1–31.
- Buser, T., Niederle, M. & Oosterbeek, H. (2020), ‘Can competitiveness predict education and labor market outcomes? Evidence from incentivized choices and validated survey measures’, *Tinbergen Institute Discussion Paper* **20-048/I**.
- Byambadalai, U., Ma, C.-t. A. & Wiesen, D. (2019), ‘Changing preferences: An experiment and estimation of market-incentive effects on altruism’, *Boston University - Department of Economics - Working Papers Series* **WP2019-11**.
- Eckel, C. C. & Grossman, P. J. (2002), ‘Sex differences and statistical stereotyping in attitudes toward financial risk’, *Evolution and human behavior* **23**(4), 281–295.
- Engelmann, D., Friedrichsen, J. & Kübler, D. (2018), ‘Fairness in markets and market experiments’.
- Falk, A., Neuber, T. & Szech, N. (2020), ‘Diffusion of Being Pivotal and Immoral Outcomes’, *The Review of Economic Studies* **87**(5), 2205–2229.
- Falk, A. & Szech, N. (2013), ‘Morals and markets’, *Science* **340**(6133), 707–711.

- Fischbacher, U., Fong, C. M. & Fehr, E. (2009), ‘Fairness, errors and the power of competition’, *Journal of Economic Behavior & Organization* **72**(1), 527–545.
- Gössling, S. & Humpe, A. (2020), ‘The global scale, distribution and growth of aviation: Implications for climate change’, *Global Environmental Change* **65**, 102194.
- Irlenbusch, B. & Saxler, D. J. (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.
- Ketcham, J., Smith, V. L. & Williams, A. W. (1984), ‘A comparison of posted-offer and double-auction pricing institutions’, *The Review of Economic Studies* **51**(4), 595–614.
- Kirchler, M., Huber, J., Stefan, M. & Sutter, M. (2016), ‘Market design and moral behavior’, *Management Science* **62**(9), 2615–2625.
- Krupka, E. L. & Weber, R. A. (2013), ‘Identifying social norms using coordination games: Why does dictator game sharing vary?’, *Journal of the European Economic Association* **11**(3), 495–524.
- Offerman, T., Romagnoli, G. & Ziegler, A. (2019), ‘Morals in multi-unit markets’, *AEA RCT Registry* .
URL: <https://doi.org/10.1257/rct.4603-1.0>
- Plott, C. R. (1983), ‘Externalities and corrective policies in experimental markets’, *The Economic Journal* **93**(369), 106–127.
- Prasnikar, V. & Roth, A. E. (1992), ‘Considerations of fairness and strategy: Experimental data from sequential games’, *The Quarterly Journal of Economics* **107**(3), 865–888.
- Roth, A. E., Prasnikar, V., Okuno-Fujiwara, M. & Zamir, S. (1991), ‘Bargaining and market behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An experimental study’, *The American Economic Review* **81**(5), 1068–1095.
- Samuelson, P. A. & Nordhaus, W. D. (2005), *Economics*, 18 edn, Columbus, OH: McGraw-Hill Irwin.
- Schneider, F., Brun, F. & Weber, R. A. (2020), ‘Sorting and wage premiums in immoral work’, *University of Zurich, Department of Economics, Working Paper* **353**.

- Scholl, L., Seth, P., Kariisa, M., Wilson, N. & Baldwin, G. (2019), ‘Drug and opioid-involved overdose deaths—united states, 2013–2017’, *Morbidity and Mortality Weekly Report* **67**(51-52), 1419.
- Smith, V. L. (1962), ‘An experimental study of competitive market behavior’, *Journal of Political Economy* **70**(2), 111–137.
- Sobel, J. (2007), ‘Do markets make people look selfish?’.
- Sutter, M., Huber, J., Kirchler, M., Stefan, M. & Walzl, M. (2020), ‘Where to look for the morals in markets?’, *Experimental Economics* **23**(1), 30–52.
- van Leeuwen, B., Offerman, T. & Schram, A. (2020), ‘Competition for status creates superstars: An experiment on public good provision and network formation’, *Journal of the European Economic Association* **18**(2), 666–707.
- Vuillemeys, G. (2020), ‘Evading corporate responsibilities: Evidence from the shipping industry’.

A Appendix

In this Appendix, we provide additional analyses of the data. Sections A.1 to A.4 were preregistered, whereas the remaining sections are additional exploratory analyses or based on comments we received.

A.1 Robustness of treatment effects

In Table A1, we regress the quantities on treatment indicators to verify robustness of our main results. Each market outcome provides one observation. Quantities in FULL differ significantly both in (1) and when including controls in (2).

Table A1: Treatment effects

| | (1) | (2) |
|--------------------|--------------------------|-----------------------|
| | Relative quantity traded | |
| (1 if MULTI) | 2.833 (3.980) | 4.040 (3.256) |
| (1 if FULL) | 23.500*** (3.604) | 23.315*** (3.506) |
| (1 if Period=2) | | -4.667** (1.795) |
| (1 if Period=3) | | -7.111*** (2.020) |
| (1 if Period=4) | | -4.889** (2.153) |
| Mean moral cost | | 0.040 (0.080) |
| Median moral cost | | -0.098 (0.065) |
| Minimum moral cost | | 0.108 (0.100) |
| Mean risk measure | | 3.834 (3.842) |
| Constant | 75.500*** (3.512) | 69.044*** (12.992) |
| Observations | 120 | 120 |
| Adjusted R^2 | 0.508 | 0.555 |

Note: Dependent variable is observed quantity relative to selfish competitive equilibrium. Mean, median and minimum moral cost are the mean, median and minimum of marginal moral costs, averaged on a subject level, in part 1 within a matching group. Mean risk is the average chosen lottery in the risk task per matching group. Standard errors clustered on matching group level in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A.2 Simulating moral competitive equilibria

In this section, we describe in more detail how competitive equilibria with moral costs are simulated. In this, we follow the analysis plan.

As described in the main text, we first estimate participants' moral costs from individual decision-making in part 1, and obtain predicted marginal moral costs $\widehat{\theta}_i(q)$ for individual i to cancel a q -th unit of donation, and impose $\widehat{\theta}_i(q) \geq 0$. To account for erosion through repetition, in treatment IDM, we repeat individual decision-making 4 times in part 2. We estimate moral costs each period, and rescale estimated moral costs for each market period with the average erosion found in IDM.

We take p as the price agreed between one buyer and one seller. For unit Q to be traded, $v(Q)$ are induced values, $c(Q)$ induced costs, which are common across all traders at this unit. $\widehat{\theta}_i(q)$ are predicted marginal moral costs for trader i to cancel a q -th unit of donation.

For heterogeneous moral competitive equilibria, we take the following steps, in each market period, where the simulation proceeds sequentially unit by unit:

1. We record individually traded quantities at every step, keeping track of which traders are constrained by capacity constraints (in SINGLE and MULTI) and what predicted marginal moral costs to trade one more unit are for each trader i : $\widehat{\theta}_i(q)$.
2. First, we verify whether any trade made in the experiment is consistent for both this buyer-seller pair we observe. That is, profits are larger than the moral costs if $\widehat{\theta}_s(q) \leq p - c(Q)$ for seller s and $\widehat{\theta}_b(q) \leq v(Q) - p$ for buyer b . We keep all trades which are consistent for this buyer and seller. By doing so, we keep those equilibria which are closest to observed trading behavior.
3. Second, we verify whether additional units can be traded, beyond the number of units kept in step 2. For each additional unit, we draw at most 200 times a random pair of buyer b and seller s . In drawing random traders, we incorporate that our market picked one buyer and one seller randomly among those who submitted an equally favorable offer and among those who accepted the offer in question. We check whether for a candidate pair of traders, their moral costs allow them to trade one more unit, compared to the available gains from trade. That is, we verify that the sum of marginal moral cost is at most the difference in induced values and costs: $\widehat{\theta}_s(q) + \widehat{\theta}_b(q) \leq v(Q) - c(Q)$. If moral costs satisfy this equation, the two traders can agree on a price p at which they are both willing to trade. For the first randomly drawn pair of traders for whom the equation is satisfied, we designate these two to trade the Q -th unit,

and continue to the $Q + 1$ -th unit. We continue to simulate additional units, up to the point where for all 200 randomly drawn pairs of traders, marginal moral costs are prohibitively high: $\widehat{\theta}_s(q) + \widehat{\theta}_b(q) > v(Q) - c(Q)$. At this point, trading stops, and the predicted quantity is the last unit which could be traded.

4. For each market and period, we repeat this procedure 10,000 times, as the order in which trader pairs are drawn potentially affects outcomes. Predictions shown in the main text are averages across all simulations and periods.

For homogeneous moral competitive equilibria, we adapt the above procedure only in the predicted marginal moral costs $\widehat{\theta}_i(q)$: for each market group, we use the median traders' moral costs for the first unit as the moral costs for all traders and all units. We thus remove both initial heterogeneity within a market and the desensitization effect from estimated moral costs. We then perform the above procedure, which again yields a predicted quantity to be traded.

A.3 Estimating moral costs in markets

Competitive equilibria with moral costs are one measure of the erosion of morals in markets. In addition, we have a second method of evaluating any such erosion. Our markets were designed such that traders are able to submit and accept offers sequentially, which ensures that we observe reservation prices of our subjects in many instances.

We can then use this data to reconstruct moral costs in a similar manner as we do for data obtained in multiple price lists. In a market, when a subject submits or accepts an offer, the margin (obtained from value less price, or price less costs, respectively) she would make when trading this unit needs to cover two components: her moral costs from causing the negative externality, as well as any profits she intends to make. Note that, in our experimental design especially the first units traded offer large surpluses, and are efficient to trade, as the gains from trade exceed the costs of the negative externality. This was also implemented to ensure sufficient earnings for our participants, as well as limit fatigue, which may arise if only little money is to be made in the entire part 2.

We use the margins obtained from a trade as an upper bound of measuring moral costs in markets. In order to focus on the moral cost-component of the margins, we construct a data set where, for each unit a subject is trading, we use the lowest margin we observe her accepting and offering.²² For this, we pool all data from part 2, irrespective of the period in which an offer was submitted. We then estimate

²²We verified that results are similar when using the two lowest offers per subject.

marginal, per unit, moral costs, allowing for changes in marginal moral costs as we do when estimating moral costs based on individual decision-making, so we estimate $\theta_i(q) = \gamma_i + \delta_i q + \epsilon_{j,q}$ for each i , where $\theta_i(q)$ now can also contain profits a subject makes beyond moral costs. We can directly observe marginal moral costs in markets, and compare these to the marginal moral costs inferred from the predicted total moral costs obtained in individual decision-making. Note that for all participants who are active only for their first unit, thus for example all participants in SINGLE, we are constrained to estimate constant marginal moral costs ($\delta_i = 0$).

From this estimated moral cost curve, we take the difference between the moral cost curve from part 1 (individual decision-making only) and part 2 (market data or individual decision-making in treatment IDM), for all units that a subject was observed to trade in the market. This means that, for each subject, differing number of units contribute to potential moral erosion in markets. To make the size of erosion comparable, we measure the erosion for the first three units in treatment IDM, which corresponds to the maximum number of units traded by a subject on average in the market treatments. Correspondingly for SINGLE, we scale up measured erosion by the factor three.²³ We then sum all measured erosion per subject, and afterwards take treatment averages.

In Table A2, we present the average erosion per subject by treatment.

Table A2: Moral erosion on an individual level

| | IDM | SINGLE | MULTI | FULL |
|--------------------------|-------|--------|--------|--------|
| Measured erosion | -16.5 | -65.0 | -205.9 | -370.2 |
| vs. no erosion | .3743 | .0367 | .0051 | .0051 |
| <i>p</i> -values vs. IDM | - | .0864 | .0002 | .0002 |
| vs. SINGLE | - | - | .0082 | .0003 |
| vs. MULTI | - | - | - | .0126 |

Notes: Normalized average erosion comparing part 1 and part 2, in Eurocents, testing vs. the null hypothesis of no erosion and between treatments (Diff-in-Diff). Mann-Whitney U-tests for between treatment and Wilcoxon signed-rank tests for within treatment tests, on matching group averages, 10 observations per treatment.

Confirming the analysis presented in the main text, moral costs are not strongly eroded in IDM through repetition. We detect some erosion in SINGLE compared to the no erosion benchmark and potentially compared to the erosion in IDM. In MULTI, erosion is stronger, and erosion is at it strongest in FULL. Based on this analysis we can reject hypotheses H1 to H3.

²³If anything, this may exaggerate the erosion in SINGLE, as marginal moral costs are typically decreasing, and this decrease is not measurable in this market.

A.4 Social norms

We elicited social norms using the method introduced by Krupka & Weber (2013). We described seven different scenarios in the experiment, where subjects evaluated whether they deemed the behavior as “socially appropriate” and “consistent with moral or proper social behavior” on a 4-point scale from “very socially inappropriate”, to “somewhat socially (in)appropriate” and “very socially appropriate”. In particular, we described four scenarios involving individual-decision making as well as three scenarios in an experimental market. In Appendix B, we reproduce the full instructions and interface.

Scenarios 1 to 4 mirror the individual decision-making task in the experiment, where Individual 1 makes the following choices (as a reminder, 4 doses cost approximately €1.5.):

1. “1 chooses to receive 1 Euro instead of making a donation of 4 doses of measles vaccine to UNICEF.”
2. “1 chooses to receive 2 Euro instead of making a donation of 4 doses of measles vaccine to UNICEF.”
3. “1 chooses to receive 3 Euro instead of making a donation of 12 doses of measles vaccine to UNICEF.”
4. “1 chooses to receive 6 Euro instead of making a donation of 12 doses of measles vaccine to UNICEF.”

Three scenarios with Individual 2 mirror the experimental markets, where trading canceled a donation of four doses of measles vaccine.

5. “2 decides to accept an offer which allows him to earn 1 Euro.”
6. “2 decides to accept an offer which allows him to earn 2 Euro.”
7. “2 makes an offer in the market. If a trade is concluded based on this offer, 2 would earn 1 Euro.”

In addition to the data presented in the main text, below are histograms of the responses of subjects for all scenarios across the four treatments.

Figure A1: Norms in individual decision-making

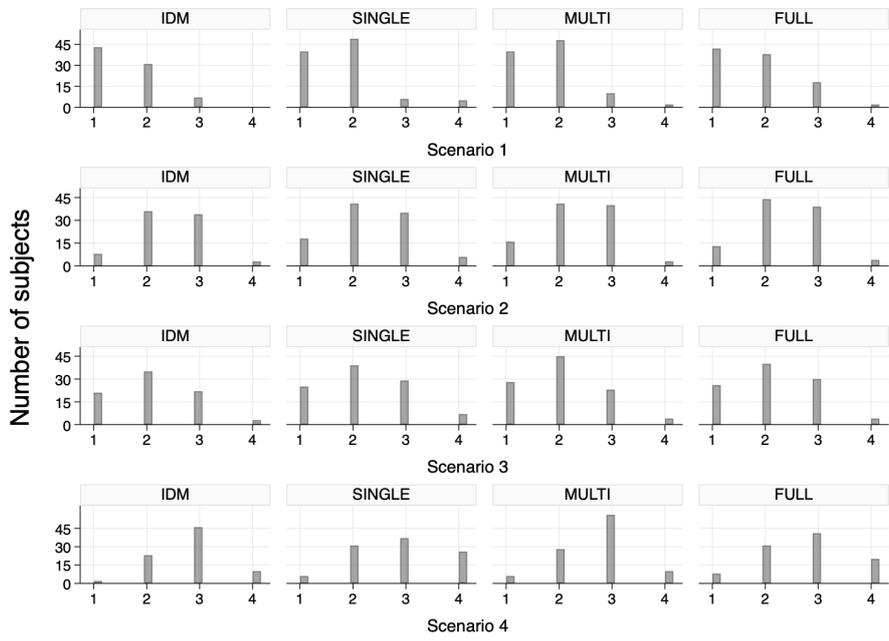
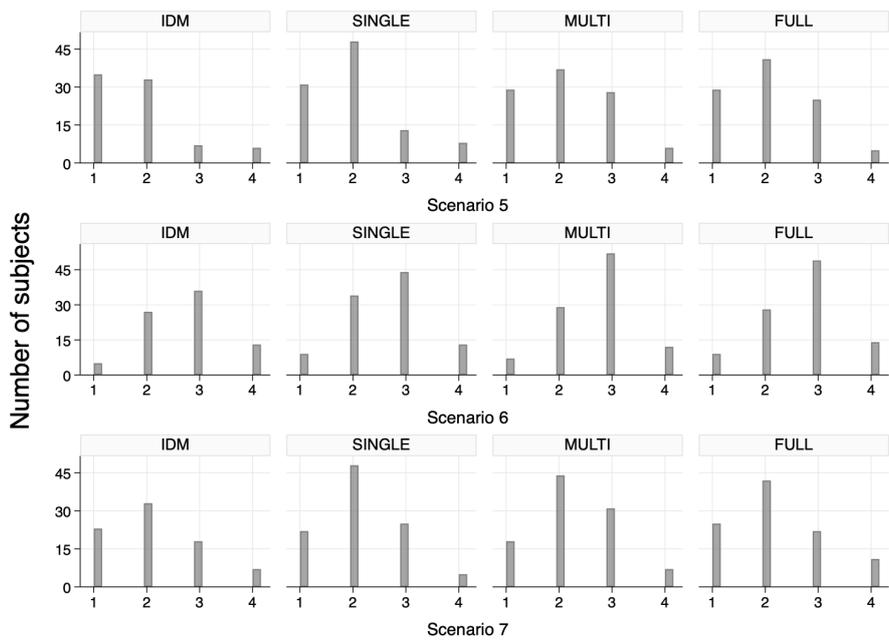


Figure A2: Norms in markets



A.5 Are marginal moral costs decreasing?

We estimate individual-level moral costs as follows. Using individual decision-making data, where subjects report the monetary equivalent of cancelling donations of different sizes (i.e., number of units), we estimate:

$$\Theta_i(q) = \alpha_i q + \beta_i q^2 + \epsilon_{j,q}$$

where $\Theta_i(q)$ is the monetary equivalent of a donation of q units for subject i . We then estimate marginal moral costs for the q -th unit as $\hat{\Theta}_i(q) - \hat{\Theta}_i(q - 1)$.

Empirically, 80.2% of subjects exhibit $\alpha_i \geq 0$ and $\beta_i \leq 0$, suggesting that marginal moral costs are decreasing for the vast majority of subjects. The decreasing pattern in marginal moral costs is also seen in the aggregate as shown in figure 3 in the main text.

The decreasing pattern of marginal moral costs is statistically significant. In an OLS regression with subjects fixed effects, we allow for changes in marginal moral costs as a function of the size of the donation (Unit). The results, presented in Table A3, show that this variable is empirically important. The estimate on Unit is negative and significant.

Table A3: Evidence for decreasing marginal moral costs

| | Marginal moral costs |
|----------------|-----------------------|
| Unit | -2.076*** (0.409) |
| Constant | 166.610*** (3.274) |
| Observations | 5715 |
| # of subjects | 381 |
| Subject FE | Yes |
| Adjusted R^2 | 0.859 |

Note: Dependent variable is estimated marginal moral costs. Unit captures the unit number running from 1 to 15. Subject fixed effects control for level differences in moral cost curves across subjects. Standard errors clustered on matching group level in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A.6 Moral erosion in single-unit markets

Our treatments IDM and SINGLE are similar to treatments which the literature so far has studied. We can use this treatment comparison to evaluate whether our experimental setup leads to result consistent with the current literature. This is es-

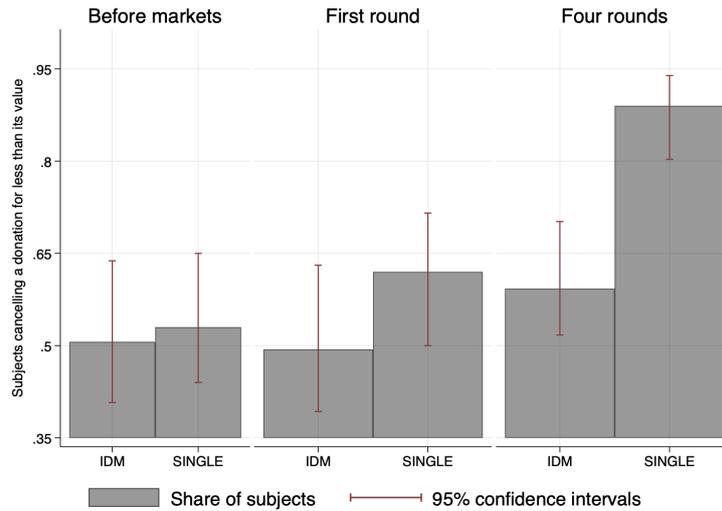
pecially important because there is a number of differences between our design and the literature so far. Compared to previous papers, we introduce a common schedule of induced costs and values, and employ a sequential double auction. With respect to Falk & Szech (2013), the most appropriate comparison in their paper are treatment differences between individual decision-making and their multi-lateral market. However, we use charity donations instead of mice, different monetary stakes and an equal number of traders per market-side. Also, we repeat individual decision-making, but have only four market periods. Close is also Kirchler et al. (2016), where the appropriate treatments are CL_BASE and DA_BASE. We use the identical donation opportunity, however at different stakes, and use only 4 instead of 10 periods. Similar is also Bartling et al. (forthcoming), however we study multi-lateral instead of bi-lateral markets, and decisions in IDM only affect payoffs to one subject and a charity, instead of payoffs to two subjects and a charity.

In Figure A3, we plot the share of subjects who cancel a donation for at most its value of € 1.50 in different environments and at different stages of the experiment, for treatments IDM and SINGLE. In the first two bars, we plot the share of subjects who cancel the first unit of donation already for a payment of at most € 1.50 in individual decision-making in part 1. These treatments are balanced in this dimension. The following two groups of bars compare behavior of these subjects either in repeated individual decision-making in IDM or in markets in SINGLE. For the markets, we study whether a trader concluded a trade for which she was paid at most € 1.50. In the middle panel, we compare behavior in the first period in part two. We observe that there appears to be slight erosion of morals in SINGLE, compared to IDM, this effect is however not particularly strong. In the right panel, we compare whether in the entire four periods of the experiment, a given subject at least once cancelled a donation for at most € 1.50. There does appear to be some erosion through repetition in IDM. Erosion is much stronger in SINGLE, as a large share of subjects can be observed to trade at least once. Therefore, our experimental design shows similar effects as the existing literature, and we find partial erosion in SINGLE compared to IDM.

A.7 Replacement logic and consequentialistic subjects

In the main text, we show that a large share of subjects engage in trading when these units only yield € 0.20 for two market participants, in exchange for causing a damage of € 1.50 to UNICEF. In the main text, part of the traders with above median moral costs are potentially consequentialistic subjects, who can use the replacement logic: given that the donation will in any case not go through, it may be legitimate to trade.

Figure A3: Erosion in single-unit markets



Notes: Share of participants who cancelled a donation for at most its value (€1.50) in individual decision-making and based on implemented trades in the market. The left panel shows take rates are similar in part 1 of the experiment, for treatment IDM and SINGLE. The middle panel compares behavior in the first round of part 2. The right panel states the share of participants who, in the four rounds of part 2, at least once cancelled a donation.

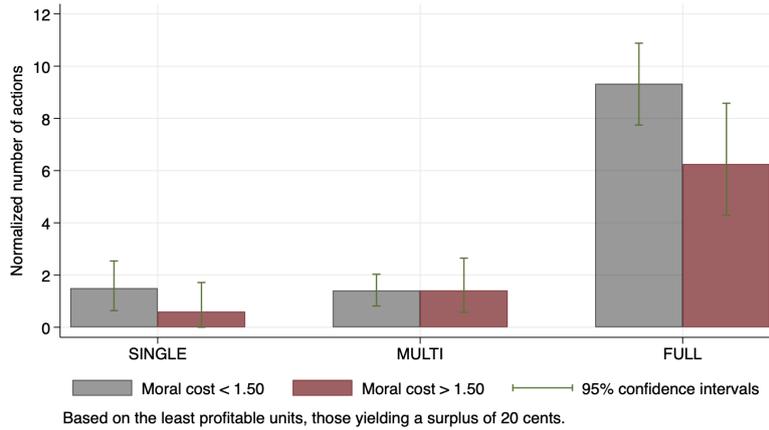
Interestingly, this activity also carries over to subject who likely do not use consequentialistic reasoning. In the first part of the experiment, we have a subset of participants who report moral costs above the corresponding value of the donation. This set of participants decided to forgo a higher payment in order not to cancel the donation, which they could have instead donated to UNICEF themselves. Approximately 34% of subjects report such preferences.²⁴

In Figure A4, we show what share of traders are active at the least profitable units in the markets, splitting them into subjects with moral costs below and above €1.50. While in SINGLE and MULTI, these subjects rarely are active, they are very active in FULL. For these subjects, it appears to be the case that their morals were eroded. This is the case as for these subjects, the replacement logic is hardly a justification to trade.

In Table A4, we show correlates of an indicator capturing whether a subject was active at the least profitable units, those with available gains from trade of €0.20. All statements in quotation marks are statements from the questionnaire, rated from

²⁴Note that this is unlikely to be driven by misunderstanding: regressing subjects' moral costs, or equivalently a dummy equal one if they report moral costs above €1.50, on the number of attempts this subject required to complete the practice questions for part 1 shows an insignificant correlation. Results are also similar when splitting subjects at even higher moral costs, such as at €1.70 or €2, which implies transaction costs are also unlikely to explain these results.

Figure A4: Replacement logic in non-consequentialistic subjects



Notes: Average number of offer submissions or acceptances per trader at the final units, which yield gains from trade of €0.20 in exchange for an externality of €1.50. Splits based on average predicted moral costs above and below €1.50, the cost of the donation.

1 to 7 whether subjects agreed with a given statement. What appears to matter are (1) initial moral costs of subjects, (2) leaning politically to the right, (3) using a statement modeled to fit the replacement logic: “I decided to trade in market 2 because I realized the units I traded would have been traded by others in any case.”. In (2), we report average marginal effects of the logistic regression in (1), as well as OLS estimates in (3).

Table A4: Who uses the replacement logic?

| Dep. variable: | (1) | (2) | (3) |
|---|-----------------------------|----------------------|---------------------|
| | (1 if active at last units) | | |
| Change in moral cost from part 1 to 3 | 0.267 (0.531) | 0.036 (0.071) | 0.033 (0.069) |
| Moral cost in part 1 | -0.829*** (0.279) | -0.111*** (0.037) | -0.094** (0.035) |
| (1 if male) | -0.914*** (0.311) | -0.123*** (0.040) | -0.111** (0.042) |
| (1 if international student) | 0.168 (0.378) | 0.023 (0.050) | 0.016 (0.053) |
| Risk measure | -0.013 (0.098) | -0.002 (0.013) | -0.003 (0.015) |
| Belief about median subject's moral cost | 0.068* (0.040) | 0.009* (0.005) | 0.009 (0.006) |
| Social norm in ind. dec.-making | 0.219 (0.306) | 0.029 (0.041) | 0.024 (0.044) |
| Social norm in market | -0.175 (0.218) | -0.023 (0.030) | -0.022 (0.033) |
| "I believe the donations for measles vaccines to UNICEF are helpful." | 0.221 (0.186) | 0.030 (0.025) | 0.030 (0.019) |
| "I believe measles vaccines save lifes." | -0.021 (0.157) | -0.003 (0.021) | -0.001 (0.021) |
| "When making a moral decision, I try to always follow a rule, instead of evaluating the consequences of each particular option every time." | -0.177 (0.111) | -0.024 (0.015) | -0.024 (0.016) |
| "When deciding on whether I should trade in market 2, I studied at what profits other traders were willing to trade." | 0.065 (0.136) | 0.009 (0.018) | -0.000 (0.018) |
| "I decided to trade in market 2 because I realized the units I traded would have been traded by others in any case." | 0.338*** (0.099) | 0.045*** (0.012) | 0.047*** (0.012) |
| "How competitive are you?" (1 not competitive, 7 very; Buser, Niederle & Oosterbeek (2020)) | -0.052 (0.148) | -0.007 (0.020) | -0.008 (0.021) |
| "Where do you see yourself in the left-right political spectrum?" (1 left, 7 right) | 0.325** (0.128) | 0.044*** (0.016) | 0.039** (0.015) |
| (1 if MULTI) | 0.987** (0.435) | 0.142** (0.061) | 0.146** (0.070) |
| (1 if FULL) | 4.434*** (0.511) | 0.694*** (0.047) | 0.686*** (0.053) |
| Constant | -5.227*** (1.792) | | -0.195 (0.219) |
| Study fixed effect | yes | yes | yes |
| Observations | 273 | 273 | 278 |
| Estimation | Logit | Avg. ME | OLS |

Note: Dependent variable is a dummy equal one if a subject submitted or accepted an offer at least once for units with gains from trade of €0.20. Change in moral cost is defined as moral costs in part 3 less moral costs in part 1 in Euro. Standard errors clustered on matching group level in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A.8 Evidence of biased imitation

In the main text, we show that multi-unit markets bias trader’s beliefs, in the sense that they become overly pessimistic about their fellow trader’s morality.

In Table A5, we regress the change in average marginal moral costs from part 1 to part 3 in cents on the mean error a subject makes in predicting the median trader’s moral cost, treatment indicators, interactions between errors and treatment fixed effects as well as their moral costs as reported in part 1. This allows us to infer whether subjects with strongly miscalibrated beliefs adjust their own reported moral costs more strongly.

Interestingly, we see that subjects whose values appear to erode most, as their decrease is strongest, are also those subjects who most strongly underestimate how much their peers care about the donation. Subjects who misunderstand in the market how much their fellow subjects actually value the donation are also the subjects who strongly adjust their valuation downwards.

This implies that incorrect updating from observed market behavior might actually be associated with more selfish future behavior of market participants. For example, imitation of others paired with biased social learning might lead markets to continue to erode participants’ morals through this channel, which we refer to as biased imitation.

Table A5: Evidence of biased imitation

| Change in moral cost from part 1 to part 3 | |
|--|-----------------------|
| Error in belief | 0.148*** (0.038) |
| (1 if MULTI) | -6.094 (6.711) |
| (1 if FULL) | -5.792 (8.375) |
| (1 if MULTI) × Error in belief | -0.067 (0.055) |
| (1 if FULL) × Error in belief | -0.058 (0.078) |
| Moral cost in part 1 | -20.953*** (4.299) |
| Constant | 12.886*** (4.527) |
| Observations | 300 |

Note: Change in moral cost is defined as average marginal moral costs in part 3 less average marginal moral costs in part 1 per subject in cents. Error in belief is the difference between reported belief and median subject’s moral cost within the session. Standard errors clustered on matching group level in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A.9 Beliefs as excuses

One potential concern regarding the findings on belief biases can be that subjects report beliefs in order to provide an excuse for their own selfish behavior in the markets. These excuses may be needed most in treatments MULTI and FULL, where we also find that subjects are pessimistic about their peers' morals.

To verify whether this might be driving our results, we report average beliefs, by treatment, for those traders who likely need the excuse most: those traders who we observe to be active at the least profitable units, those yielding profits of €0.20. In Table A6, we see that there are no meaningful patterns that would support such excuse-driven reporting of beliefs. Similarly, regressing beliefs (or errors in beliefs) on a dummy variable equal one if a trader was active at the least profitable units, with treatment fixed effects, yields insignificant, and for that matter positive, coefficients on the dummy variable capturing the need for an excuse. Therefore, it is unlikely that our findings on beliefs can be explained by participants' need to provide justification for their own selfish behavior.

Table A6: Average beliefs for (in)active traders at last units

| | SINGLE | MULTI | FULL |
|----------|--------|-------|------|
| Inactive | 10.20 | 10.07 | 8.24 |
| Active | 10.25 | 10.56 | 8.84 |

Notes: Average belief of median participant's switching point in the multiple price list for the first unit (11 corresponds to indifference between payments to self and UNICEF). Split by whether the subject was active at the final units, those with gains from trade of €0.20.

A.10 Norms as excuses

As for the beliefs, it may be the case that subjects report perceived norms to excuse their behavior in part 2 of the experiment. We report the same analysis for the elicited norms in Table A7.

There are no systematic patterns which suggest that norms are reported self-servingly. This is confirmed by regression evidence, similar to the analysis for beliefs. Regressing the reported norm in the experimental market on a dummy variable equal one if a trader was active at the least profitable units, with treatment fixed effects, yields insignificant coefficients on the dummy variable (p -value=.469) and on the treatment fixed effects (p -value=.222 for MULTI, p -value=.923 for FULL).

Table A7: Average norm report for (in)active traders at last units

| Norm in | SINGLE | | MULTI | | FULL | |
|----------|--------|--------|-------|--------|------|--------|
| | IDM | Market | IDM | Market | IDM | Market |
| Inactive | 1.76 | 1.99 | 1.68 | 2.01 | 1.65 | 2.12 |
| Active | 1.75 | 1.94 | 1.88 | 2.31 | 1.83 | 2.05 |

Notes: Average norm report for cancelling donations of €1.50 in return for a payment of €1 in individual decision-making (IDM) or with in an experimental market (Market). 2 corresponds to “somewhat socially inappropriate”. Split by whether the subject was active at the final units, those with gains from trade of €0.20.

B Experimental interface

Below is an example screenshot from the experimental markets.

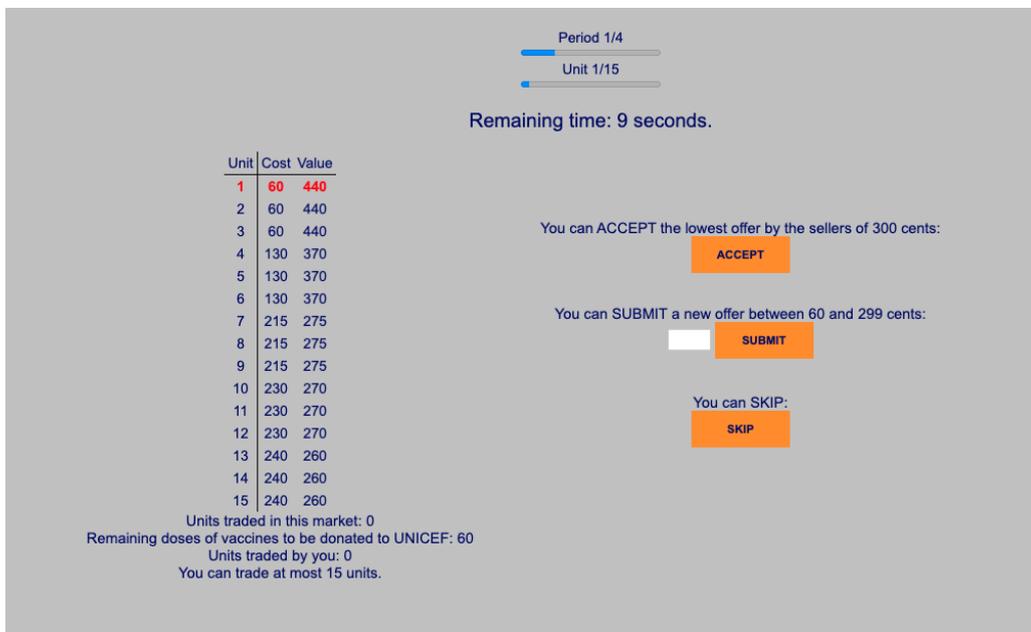


Figure A5: Experimental market interface

After each unit, traders received the following feedback:

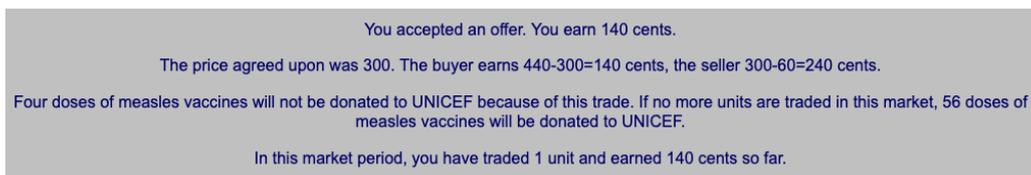


Figure A6: Feedback after each trade

Further, after the end of each market period, participants received this feedback:

The market for this trading period has closed, no more units can be traded.
In this period, you traded one unit. You earn 140 cents if this period will be randomly selected for payment.
In the entire market, 1 unit was traded.
If this period will be randomly selected for payment, 56 doses of vaccine will be donated to UNICEF. Donations for 4 doses have been cancelled in this market.

Figure A7: Feedback after a market period

Therefore, traders were reminded of the negative externality that was caused by trading continuously within the market, after each unit and at the end of each market period.

C ONLINE APPENDIX: Instructions part 1

C.1 Page 1

Welcome!

Welcome to this experiment. Please read the following instructions carefully.

Please do not communicate with other people and refrain from verbally reacting to events that occur during the experiment. The use of mobile phones or laptops is not allowed.

There are pen and paper on your table, you can use these during the experiment. We will also distribute a handout with some key facts about this experiment later.

If you have any questions, or need assistance at any time, please notify the experimenter by raising your hand. The experimenter will assist you privately.

C.2 Page 2

General information

This experiment consists of multiple parts. Your decisions in one part will not affect any of your choices or potential earnings in other parts. You will receive instructions for each part separately.

For your participation in this experiment, you will be paid 7 Euro. Additionally, you can earn money by your decisions in this experiment. These earnings will depend on your decisions and may depend on other participants' decisions. **One out of the first three parts will be randomly selected to be paid to you.** Additionally, you will be paid for three short tasks at the end of the experiment. Your earnings will be paid to you privately in cash at the end of today's session. **All your earnings will be denoted in cents** (100 Cents = 1 Euro).

C.3 Page 3

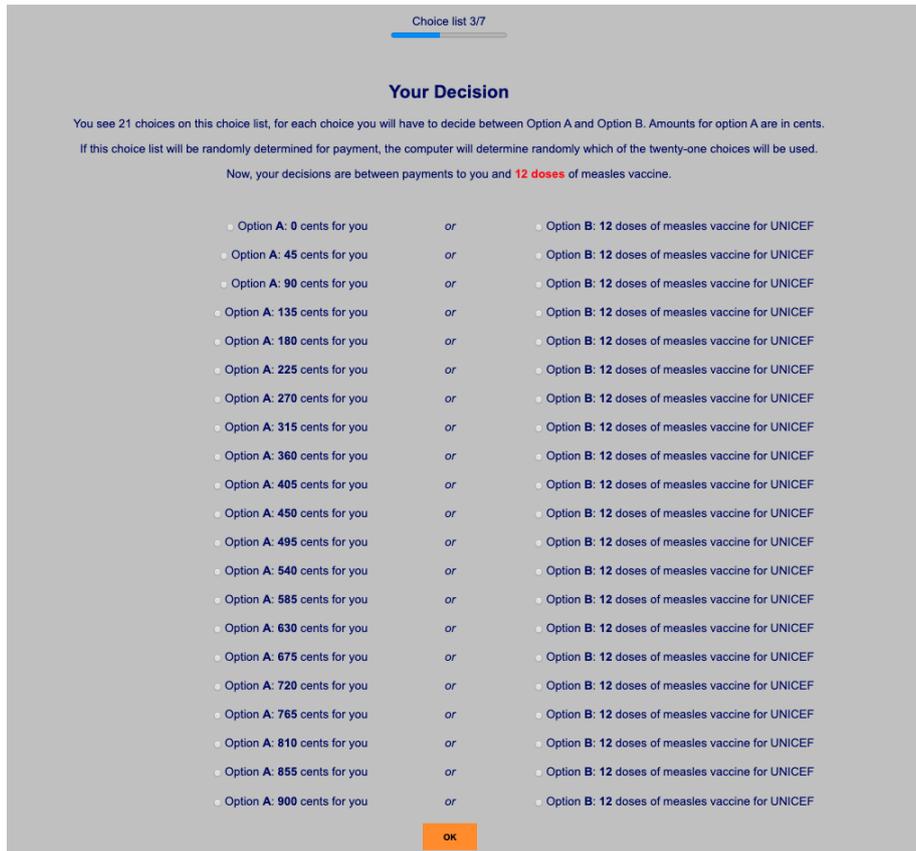
Part 1

In this first part, you will repeatedly choose between two options, A and B:

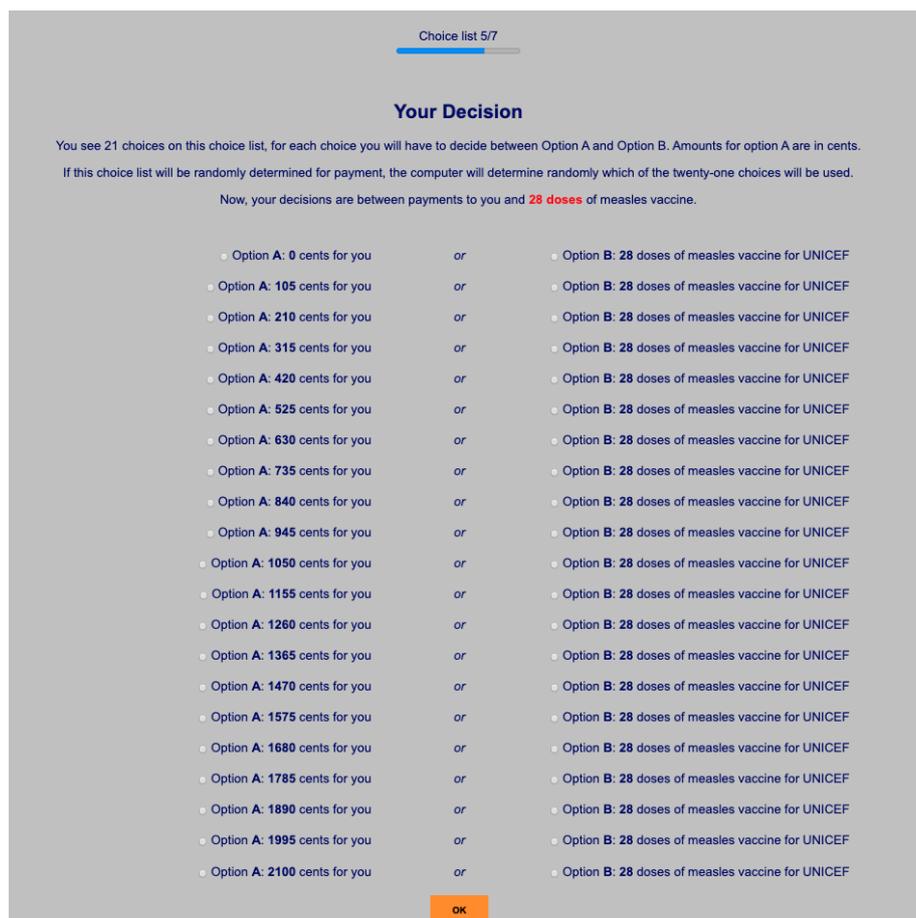
- **A:** This option will pay a certain amount of money to you.
- **B:** This option will donate a certain amount of money to UNICEF. With this donation, UNICEF will buy measles vaccines. With two doses of this vaccine, one child can be vaccinated against measles (details on the donation follow below).

A list of repeated choices between A and B on one screen is called a choice list.

Below is an example of a choice list. In this example, you choose between varying amounts of money paid to you on the left (option A) and 12 doses of measles vaccine on the right (option B). A donation of four doses costs approximately 1.5 Euro. Even though you will be asked to make multiple decisions, at most one of them will affect your earnings.



In the second example screenshot below, you see another choice list. Here, you choose between varying amounts of money paid to you in option A and 28 doses of measles vaccines in option B. Note that also the available payments in option A vary across choice lists.



You will face several choice lists like the one in the screenshots above. On each list, two things change. First, the number of doses of vaccine donated to UNICEF change, which are 12 and 28 in the two examples given here. Second, the available payments in option A change.

Within each choice list, only option A changes between choices. As you scroll down the list, the amount of measles vaccines donated to UNICEF stays the same. The money that would be paid to you if you choose not to donate to UNICEF is increasing on each choice list. To simplify the decision, as soon as you click on one choice, the computer will pre-fill subsequent choices automatically. If for a particular choice you chose A (money to you), then all choices on the choice list below this choice of A pay even more money to you than that choice while option B does not change. Then, the computer will pre-fill A for all these choices below. Similarly, if for a particular choice you chose B (donation to UNICEF), then all choices above this choice pay even less money to you than that choice, so then the computer will pre-fill B for all these choices above. Until you click on OK, you can always change your decision. The pre-filled choices will adjust automatically while you change your decision.

C.4 Page 4

Payment

If this part is randomly selected for payment, one of your decisions from this part will be randomly selected to be paid out at the end of the experiment. For the decision to be paid out, first, one of the seven choice lists you faced will randomly be chosen, with each choice list being equally likely. Second, within this chosen list, one decision will be randomly chosen for payment, with each decision being equally likely. If you chose option A for this decision, you will be paid the number of cents indicated for this choice. If you instead chose option B, the specified number of doses of measles vaccines will be donated to UNICEF at the end of the experiment.

Details about the donation Two doses of measles vaccine are sufficient to vaccinate one child (see the next page for more details) and can be bought with a donation of approximately 75 cents. Depending on your choice in the selected period a certain amount of money is donated to UNICEF by the experimenters. We will show you a donation receipt by UNICEF at the end of this experiment, right after we transferred the announced donation. As an example, below we show you how such a receipt for a previous donation looks like. A confirmation of the donation to UNICEF can also be sent to you via email, to allow you to verify the correctness of the statements made here. To do so, you can write your email address on the form on your table, which will be collected after the experiment. Your email address will not be linked to any other data in this experiment.

As UNICEF only allows us to donate in bundles of 40 doses, any excess donations in your session will be paid to UNICEF as a direct transfer, and this transfer will be included in the receipt we show you.

[DONATION RECEIPT EXAMPLES]

C.5 Page 5

Information on the measles vaccines

Measles are highly infectious and very often deadly. Each day hundreds of children become victims of this disease. The survivors often suffer consequences for their whole life, like blindness or brain damages. This, even though protecting the children would be easy. Measles kills more than 160,000 children worldwide each year.

Measles are extremely infectious and spread especially fast when many people live densely together, as in refugee camps. Especially with weakened children the

disease often ends deadly or leads to lasting physical or mental damages. Measles are one of the main causes for blindness among children and often become critical when no medical help is available. This, even though measles vaccination offers quick, reliable, and cheap protection. UNICEF conducts major vaccination campaigns, especially after natural disasters and in other emergency situations, to prevent the spreading of the disease. With a measles vaccination you do not only protect the children, but you also reduce the risk for all who get in contact with them.

Source text on measles vaccines by UNICEF: <https://unicef.at/shop/index.php/gesundheit-und-schutz/masern-impfstoff.htm> and <https://market.unicef.org.uk/inspired-gifts/measles-vaccines-to-protect-20-children/S359163X/>

Source pictures: <https://market.unicef.org.uk/inspired-gifts/measles-vaccines-to-protect-20-children/S359163X/>

C.6 Practice questions (page 6)

NOTE: ALL NUMBERS IN THE QUESTIONS ARE ARBITRARILY CHOSEN, AND ARE NOT RELEVANT FOR THE EXPERIMENT.

Please answer the following questions:

1. Your decisions in other parts do not affect your earnings in this part. Also, your choices in this part do not affect your earnings in other parts. [TRUE/false]
2. The following choice has been randomly selected for payment:

Option A: 360 cents for you or Option B: 12 doses of measles vaccine for UNICEF

In this choice, you have chosen option A, as indicated. How much will be paid to you? [FREE FORM: 360] cents

How many doses of measles vaccines will be donated to UNICEF? [FREE FORM: 0] doses

3. Now, the following choice has been randomly selected:

Option A: 1260 cents for you or Option B: 28 doses of measles vaccine for UNICEF

In this choice, you have chosen option B, as indicated. How much will be paid to you? [FREE FORM: 0] cents

How many doses of measles vaccines will be donated to UNICEF? [FREE FORM: 28] doses

4. At the end of this experiment, the promised donations will immediately be transferred by the experimenter. You can verify this with the receipt from UNICEF. [TRUE/false]

C.7 Page 7

End of instructions

You have reached the end of the instructions. When you are ready for the experiment, please push the button READY. When all participants have pushed READY, the experiment will start.

If you still have questions, please raise your hand, and the experimenter will assist you!

D ONLINE APPENDIX: Instructions part 2

{FOR IDM: Repeated instructions of part 1: These are the identical instructions as those you saw at the start of the experiment [see above]}

{FOR MARKET TREATMENTS [ONLY SELLER INSTRUCTIONS ARE REPRODUCED, BUYERS APPROPRIATELY ADJUSTED]}:

D.1 Page 1

Market instructions

In this part of the experiment you will repeatedly trade in a market. In the market, 5 sellers can trade with 5 buyers. You will be a SELLER in the entire experiment. You will trade in two markets, market 1 and market 2, which proceed according to similar rules. After market 1 is completed, you will receive additional instructions for market 2.

Market 1

Trading profits

In market 1, a total of 5 units can be traded. Each trader can trade at most {FULL: 5 units; MULTI: 2 units; SINGLE: 1 unit}. Trading will proceed unit by unit. For each unit, one buyer and one seller can conclude a trade by agreeing on a price for that unit.

If bought, each unit has a certain cost to the seller. This will be denoted in cents. Similarly, each unit sold will have a value in cents to the buyer. Earnings for the buyer and seller for concluding a trade are:

- The seller earns the difference between the price and the cost for this unit:
PRICE-COST
- The buyer earns the difference between the value and the price for this unit:
VALUE-PRICE

These costs and values are presented during the market, as in the screenshot below. In this example, the first unit is being traded, which is highlighted by the red first line in the table.

Unit 1/5

Remaining time: 8 seconds.

You can ACCEPT the highest offer by the buyers of 180 cents:

ACCEPT

| Unit | Cost | Value |
|------|------|-------|
| 1 | 60 | 440 |
| 2 | 130 | 370 |
| 3 | 215 | 275 |
| 4 | 230 | 270 |
| 5 | 240 | 260 |

You can SUBMIT a new offer between 181 and 199 cents:

180 SUBMIT

You can SKIP:

SKIP

Units traded in this market: 0
Units traded by you: 0

Example in the screenshot: The buyer has a value given by 440, the seller has a cost given by 60. You, as the seller, and one of the buyers agree on a price of 180. Then,

- You get: $PRICE - COST = 180 - 60 = 120$ cents.
- The buyers get: $VALUE - PRICE = 440 - 180 = 260$ cents.

In the screenshot, notice that the cost of the seller and the value of the buyer change with the unit transacted (e.g. for the first unit the cost for the seller is 60 cents and the value for the buyer is 440 cents, for the third unit the cost is 215 cents and the value is 275 cents and so on). **However, for each unit, they are the same for all buyers or sellers.** Costs and values only depend on the number of units traded up to that point in the entire market by any of the traders. That is, they do not depend on the number of units you yourself have traded previously.

D.2 Page 2

Trading protocol

To agree on a price, the side of the sellers and the side of the buyers submit and accept offers sequentially. This means that first one side of the market decides ("the active side"), afterwards this side will wait and the other side of the market decides. If trading continues, the first side of the market is allowed to decide again, and so forth.

While your side (the sellers' side) is active in the market, you have three available choices:

1. SUBMIT: Submit an offer to the buyers
2. ACCEPT: Accept an offer of the buyers
3. SKIP

You can see all three options available in the screenshot below:

Unit 1/5

Remaining time: 8 seconds.

You can ACCEPT the highest offer by the buyers of 180 cents:

ACCEPT

You can SUBMIT a new offer between 181 and 199 cents:

SUBMIT

You can SKIP:

SKIP

| Unit | Cost | Value |
|------|------|-------|
| 1 | 60 | 440 |
| 2 | 130 | 370 |
| 3 | 215 | 275 |
| 4 | 230 | 270 |
| 5 | 240 | 260 |

Units traded in this market: 0
Units traded by you: 0

Each of the options works according to these rules:

1. ACCEPT:
 - You will see the highest price offered by any of the buyers.
 - You can accept this highest offer. If you do so, a trade for one unit is concluded, the profits are calculated as explained before.

- If multiple sellers accept an offer, or if multiple offers are equally good, it will be randomly chosen which of the traders who wanted to can conclude this trade.
- Afterwards, trading of the next unit can begin, old offers are removed and new ones can be made.

2. SUBMIT:

- You can submit a new offer, which will be presented to the buyers as soon as they become active.
- A new offer has to improve upon previous offers. This means that a new offer needs to be above the lowest offer submitted by any of the other sellers. A new offer cannot be above the buyers' values, or below the highest offer by the buyers.

3. SKIP:

- If you skip, you immediately move to the waiting screen.
- As soon as all sellers are on the waiting screen, the buyers become active and can submit new offers or accept the lowest offer of the buyers. Clicking on skip can speed up the market. However, you will no longer be able to submit or accept an offer at that moment.
- If you do not submit or accept an offer within the trading time of 14 seconds, you will skip automatically.

End of trading

Trading ends if all available units are sold in the market.

Also, if no trader on both sides of the market chooses SUBMIT or ACCEPT, a warning sign will be shown. Then, each trader on both market sides can once again SUBMIT, ACCEPT or SKIP. If again no trader on either of the two sides chooses SUBMIT or ACCEPT, the market ends for this and all subsequent units. This means that you will not be allowed to trade additional units after this happens.

D.3 Page 3

Additional details

- At the start of the market for the first unit, it is randomly determined whether the side of the buyers or the side of the sellers first becomes active. For the next units, the active side for making the first offer is alternated.

- On the top of the trading screen you always see the remaining trading time. We will also show how many units you have traded. The specific moment at which you submit or accept offers does not matter, as long as you submit or accept within the 14 second trading time.
- No trader knows with whom in the room he or she has traded. That means that your anonymity is ensured.

Reminders

- At each moment, you can choose only one of the three options (SUBMIT, ACCEPT and SKIP). If trading continues and your side of the market becomes active again, you can again choose between these options.
- A maximum of five units can be traded in market 1; after the 5th unit is sold the market ends. Each trader can trade at most {FULL: 5 units; MULTI: 2 units; SINGLE: 1 unit}.
- Each unit is traded by one buyer and one seller, all other traders get a payoff of zero for that unit.

Payment

If this part and this market is selected for payment, for each trade a participant concluded, his or her payment is calculated with the rules described above. That is, for each unit, the seller will be paid the difference between the price and the cost for this unit. The buyer will be paid the difference between the value for this unit and the agreed upon price. {FULL/MULTI: The earnings for this market are then given by the sum of earnings for all units traded by each participant.}

In part 2, there will be a total of 5 markets. 2 out of the 5 markets will be randomly selected to be paid.

D.4 Practice questions (page 4)

NOTE: ALL NUMBERS IN THE QUESTIONS ARE ARBITRARILY CHOSEN, AND ARE NOT RELEVANT FOR THE EXPERIMENT.

Please answer the following questions:

1. Each seller pays the same costs as any of the other sellers to supply any unit, and each buyer values any unit equally as any of the other buyers. [TRUE/false]

2. If no buyer or seller submits an improved offer twice, the market for this period will end and no more units can be traded. [TRUE/false]
3. We will ask you several questions about the scenario below. Note that the behavior in this scenario is randomly determined, only for the purpose of asking these questions.
 - (a) The first unit is being traded in this market. This unit costs 60 cents to any of the sellers, and has a value of 440 cents to any of the buyers. The buyers were randomly selected to first submit offers.
 - (b) Buyer B1 decides to submit a price of 140 cents to the buyers and buyer B2 submits a price offer of 200 cents. The trading time of 14 seconds expires without any other buyer submitting an offer.
 - (c) Now the sellers become active. As buyer B2's offer is the highest offer, the sellers will only see buyer B2's offer of 200 cents.
 - (d) However, none of the sellers decides to accept this offer. Instead, seller S1 submits a new offer. This offer needs to be higher than 200 cents, as otherwise accepting buyer B2's offer is more favorable to seller S1. Seller S1 submits a new offer of 260 cents. Again, the trading time of 14 seconds expires without any other seller submitting or accepting an offer.
 - (e) Now, the buyers become active again. Seeing seller S1's offer of 260 cents, buyer B3 decides to accept this offer. The trading time of 14 seconds expires without any other buyer accepting this offer. This means that the first unit has been traded.
 - (f) Afterwards, bargaining about the second unit begins.

How many cents does buyer B1 earn from the first unit? [FREE FORM: 0] cents

How many cents does buyer B2 earn from the first unit? [FREE FORM: 0] cents

How many cents does buyer B3 earn from the first unit? [FREE FORM: 180] cents

How many cents does seller S1 earn from the first unit? [FREE FORM: 200] cents

[MARKET 1 TAKES PLACE, AFTERWARDS INSTRUCTIONS FOR MARKET 2 (with externality) FOLLOW]

D.5 Page 1

This concludes market 1. Now, trading in **market 2** begins.

Generally, the same rules apply in this market. We will therefore highlight here only the differences between the two markets:

- Trading behavior in this market determines an amount of money that will be donated to UNICEF, in addition to your own earnings. The number of units successfully traded in the market is used to calculate how many doses of measles vaccines will be donated to UNICEF. The maximum number of doses donated to UNICEF in one market period is {FULL/MULTI: 60 doses; SINGLE: 20 doses}. The more units are traded in the market, the less will be donated to UNICEF: for each unit that is traded in market 2 that is selected for payment, 4 doses of measles vaccines will be subtracted from the donation to UNICEF, which cost approximately 1.5 Euro. Recall that with two doses of measles vaccine, one child can be protected. UNICEF will be paid the donation amount at the end of the experiment. The following table summarizes how the number of traded units in the market translates into the number of MEASLE DONATIONS. For example, if at the end of the market, zero units have been traded, then a total of {FULL/MULTI: 60 doses; SINGLE: 20 doses} are donated to UNICEF for this market. If at the end of the market 3 units have been traded then in total {FULL/MULTI: 48 doses; SINGLE: 8 doses} doses are donated. Donations to UNICEF are only affected by the overall number of units traded in the market and not by whom these units are traded.

Final number of units traded and number of doses: [TREATMENT-SPECIFIC TABLE WITH COST/VALUES, SEE PAPER]

- Each unit traded has a VALUE and a COST according to the table below. These costs and values (in cents) will be the same in all markets of this experiment. [TABLE HERE, STATING NUMBER OF TRADED UNITS AND CORRESPONDING DONATIONS]
- While market 1 only lasted for 1 period, you will now be trading in a sequence of 4 market periods. Each market period is conducted in the same way. Your choices in one period have no consequences on any other period.
- {FULL/MULTI: While in market 1 a maximum of 5 units could be traded, now the maximum number of units tradeable in each market period is 15.; SINGLE: As in market 1, a maximum of 5 units can be traded.} Just like in

market 1, fewer than {FULL/MULTI: 15; SINGLE: 5} units will be traded if the traders no longer SUBMIT or ACCEPT after the warning sign. Moreover, each trader can trade at most {FULL: 15 units; MULTI: 3 units; SINGLE: 1 unit}. {MULTI/SINGLE: This means that if you decide not to trade one unit that you are allowed to trade, you reduce the number of units that can be traded by one, which would also reduce the corresponding damage to the donation to UNICEF.}

Payment

If this part is selected for payment, two of the market results are randomly selected for payment. It is equally likely that each one of the 4 market periods of market 2 or the one period in market 1 is selected for payment. Payment for participants are then calculated according to the same rules as in market 1.

If a market period of market 2 is selected, the trades in the selected period also determine the amount donated to UNICEF. At the end of the experiment, the experimenter will transfer this amount.

D.6 Page 2

[REPEATED INFORMATION ON UNICEF, SEE INSTRUCTIONS FOR PART 1]

D.7 Practice questions (Page 3)

NOTE: ALL NUMBERS IN THE QUESTIONS ARE ARBITRARILY CHOSEN, AND ARE NOT RELEVANT FOR THE EXPERIMENT.

Please answer the following questions:

1. If this part is selected for payment, two market results are randomly selected for payment. These can be market 1 or one of the market periods of market 2. {FULL/MULTI: Each trader earns the sum of cents generated by all of his or her trades} [TRUE/false]
2. For each unit that is traded, how many doses of measles vaccines will be subtracted from the donation to UNICEF? [FREE FORM: 4] doses
3. We will ask you several questions about the scenario below. Note that the behavior in this scenario is randomly determined, only for the purpose of asking these questions.

- (a) The first unit is being traded in the market. This first unit costs 60 cents to any of the sellers, and has a value of 440 cents to any of the buyers. The sellers are first to submit offers.
- (b) Seller S1 decides to submit a price of 290 cents to the buyers. Also, seller S2 submits a price offer, of 310 cents. The trading time of 14 seconds expires without any other seller submitting an offer.
- (c) Now the buyers become active. As seller S1's offer is the lowest offer, the buyers will only see seller S1's offer of 290 cents.
- (d) Buyer B1 and buyer B2 decide to accept this offer.
- (e) It is randomly determined that buyer B2 buys the first unit. This means that the first unit has been traded and that 4 fewer doses of measles vaccines will be donated to UNICEF.
- (f) Afterwards, bargaining about the second unit begins.

How many cents does seller S1 earn from the first unit? [FREE FORM: 230] cents

How many cents does seller S2 earn from the first unit? [FREE FORM: 0] cents

How many cents does buyer B1 earn from the first unit? [FREE FORM: 0] cents

How many cents does buyer B2 earn from the first unit? [FREE FORM: 150] cents

}

E ONLINE APPENDIX: Instructions part 3

E.1 Page 1

Part 3

You will now face a set of choices identical to the choices at the start of the experiment. As before, you have several choice lists, where each choice asks you to choose between points for yourself or varying doses of measles vaccine donated to UNICEF.

This part is conducted identically to the first part, and you will also be paid according to the same rules. On the next page, we reproduce the instructions from the start of the experiment in case you want to review them again.

Note that your earnings from your decisions in this part are not depending on any decision you have made up to now, or on any of your decisions you will make in the following set of questions.

E.2 Page 2

[SEE ABOVE FOR INSTRUCTIONS]

F ONLINE APPENDIX: Instructions for the three additional tasks

This is the end of the main parts of this experiment. In the remainder you will be able to make some additional money for three short tasks.

F.1 Instructions part 4 (belief elicitation)

Now, think of all other subjects who participate in this session today. The first task everyone completed in this experiment was a choice list where you could choose between an amount for yourselves and a donation of 4 doses of measles vaccines donated to UNICEF.

What do you think other participants chose on average in this choice list?

Please fill out this choice list **as you think the average participant did** in their first choice list. If your choice matches what the average participant did, you will earn an additional bonus of 100 cents.

Part 4

Below is a choice list which and all other participants faced as first task in this experiment, in part 1.

Please fill in this choice list as you think the average participants did.

- | | | |
|---|----|---|
| <input type="radio"/> Option A: 0 cents for you | or | <input type="radio"/> Option B: 4 doses of measles vaccine for UNICEF |
| <input type="radio"/> Option A: 15 cents for you | or | <input type="radio"/> Option B: 4 doses of measles vaccine for UNICEF |
| <input type="radio"/> Option A: 30 cents for you | or | <input type="radio"/> Option B: 4 doses of measles vaccine for UNICEF |
| <input type="radio"/> Option A: 45 cents for you | or | <input type="radio"/> Option B: 4 doses of measles vaccine for UNICEF |
| <input type="radio"/> Option A: 60 cents for you | or | <input type="radio"/> Option B: 4 doses of measles vaccine for UNICEF |
| <input type="radio"/> Option A: 75 cents for you | or | <input type="radio"/> Option B: 4 doses of measles vaccine for UNICEF |
| <input type="radio"/> Option A: 90 cents for you | or | <input type="radio"/> Option B: 4 doses of measles vaccine for UNICEF |
| <input type="radio"/> Option A: 105 cents for you | or | <input type="radio"/> Option B: 4 doses of measles vaccine for UNICEF |
| <input type="radio"/> Option A: 120 cents for you | or | <input type="radio"/> Option B: 4 doses of measles vaccine for UNICEF |
| <input type="radio"/> Option A: 135 cents for you | or | <input type="radio"/> Option B: 4 doses of measles vaccine for UNICEF |
| <input type="radio"/> Option A: 150 cents for you | or | <input type="radio"/> Option B: 4 doses of measles vaccine for UNICEF |
| <input type="radio"/> Option A: 165 cents for you | or | <input type="radio"/> Option B: 4 doses of measles vaccine for UNICEF |
| <input type="radio"/> Option A: 180 cents for you | or | <input type="radio"/> Option B: 4 doses of measles vaccine for UNICEF |
| <input type="radio"/> Option A: 195 cents for you | or | <input type="radio"/> Option B: 4 doses of measles vaccine for UNICEF |
| <input type="radio"/> Option A: 210 cents for you | or | <input type="radio"/> Option B: 4 doses of measles vaccine for UNICEF |
| <input type="radio"/> Option A: 225 cents for you | or | <input type="radio"/> Option B: 4 doses of measles vaccine for UNICEF |
| <input type="radio"/> Option A: 240 cents for you | or | <input type="radio"/> Option B: 4 doses of measles vaccine for UNICEF |
| <input type="radio"/> Option A: 255 cents for you | or | <input type="radio"/> Option B: 4 doses of measles vaccine for UNICEF |
| <input type="radio"/> Option A: 270 cents for you | or | <input type="radio"/> Option B: 4 doses of measles vaccine for UNICEF |
| <input type="radio"/> Option A: 285 cents for you | or | <input type="radio"/> Option B: 4 doses of measles vaccine for UNICEF |
| <input type="radio"/> Option A: 300 cents for you | or | <input type="radio"/> Option B: 4 doses of measles vaccine for UNICEF |

OK

If you have any questions, please raise your hand and one of us will come to your desk.

F.2 Instructions part 5 (norms)

On the following screens, you will read descriptions of a series of situations. These descriptions correspond to situations in which one person, Individual 1, must make a decision.

After you read the description of the decision, you will be asked to evaluate the different possible choices available to the person and to decide, for each of the possible actions, whether taking that action would be “socially appropriate” and “consistent with moral or proper social behavior” or “socially inappropriate” and “inconsistent with moral or proper social behavior”. By socially appropriate, we mean behavior that most people agree is the “correct” or “ethical” thing to do. Another way to think about what we mean is that if Individual 1 were to select a socially inappropriate choice, then someone else might be angry at Individual 1 for doing so.

In each of your responses, we would like you to answer as truthfully as possible, based on your opinions of what constitutes socially appropriate or socially inappropriate behavior.

At the end of the experiment today, we will randomly select one of the situations. For this situation, we will also randomly select one of the possible choices that Individual 1 could make. Thus, we will select both a situation and one possible choice at random. For the choice selected, we will determine which response was selected by most people participating in this experiment right now. If you give the same response as that most frequently given by other people, then you will receive an additional 200 cents. This means that you will earn most money if you select the response given most frequently by other participants.

Part 5

Individual 1 participates in an experiment. In this experiment, 1 repeatedly chooses between money paid to 1 and donations to UNICEF, which pay for measles vaccines. These are identical donations as you saw earlier in this experiment, a donation of two doses costs approximately 75 cents.

| Individual 1's choice | Very socially inappropriate | Somewhat socially inappropriate | Somewhat socially appropriate | Very socially appropriate |
|--|-----------------------------|---------------------------------|-------------------------------|---------------------------|
| 1 chooses to receive 1 Euro instead of making a donation of 4 doses of measles vaccine to UNICEF | • | • | • | • |
| 1 chooses to receive 2 Euro instead of making a donation of 4 doses of measles vaccine to UNICEF. | • | • | • | • |
| 1 chooses to receive 3 Euro instead of making a donation of 12 doses of measles vaccine to UNICEF. | • | • | • | • |
| 1 chooses to receive 6 Euro instead of making a donation of 12 doses of measles vaccine to UNICEF. | • | • | • | • |

Submit

As a reminder, you are asked to evaluate the different possible choices available to Individual 2 and to decide, for each of the possible actions, whether taking that action would be "socially appropriate" and "consistent with moral or proper social behavior" or "socially inappropriate" and "inconsistent with moral or proper social behavior".

If you give the same response as that most frequently given by other people, then you will receive an additional 200 cents. This means that you will earn most money if you select the response given most frequently by other participants.

If you have any questions, please raise your hand and one of us will come to your desk.

Part 5

Individual 2 also participates in an experiment. In this experiment, 2 is a seller in a market. 2 negotiates with other participants, and can earn money for every trade. However, **whenever a trade is made in this market, a donation with four doses of measles vaccines for UNICEF is cancelled.** A donation of four doses costs approximately 1.5 Euro.

2 can now make different offers or accept others' offers in this market, see the options below.

| Individual 2's choice | Very socially inappropriate | Somewhat socially inappropriate | Somewhat socially appropriate | Very socially appropriate |
|---|-----------------------------|---------------------------------|-------------------------------|---------------------------|
| 2 decides to accept an offer which allows him to earn 1 Euro. | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| 2 decides to accept an offer which allows him to earn 2 Euro. | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| 2 makes an offer in the market. If a trade is concluded based on this offer, 2 would earn 1 Euro. | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |

Submit

As a reminder, you are asked to evaluate the different possible choices available to Individual 2 and to decide, for each of the possible actions, whether taking that action would be "socially appropriate" and "consistent with moral or proper social behavior" or "socially inappropriate" and "inconsistent with moral or proper social behavior".
If you give the same response as that most frequently given by other people, then you will receive an additional 200 cents. This means that you will earn most money if you select the response given most frequently by other participants.

If you have any questions, please raise your hand and one of us will come to your desk.

F.3 Instructions part 6 (risk aversion)

For this part, you choose one gamble you would like to play from among six different gambles. The six different gambles are listed below. You must select one and only one of these gambles.

Each gamble has two possible outcomes (Roll Low or Roll High). For every gamble, each Roll has a 50% chance of occurring. At the end of the study, it will be randomly determined which event will occur.

For example, if you select Gamble 4 and Roll High occurs, you will be paid 260 cents. If Roll Low occurs, you will be paid 80 cents.

Your decision:

| Gamble | Choice | Roll | Payoff | Chances |
|--------|-----------------------|-------------|------------------------|------------|
| 1 | <input type="radio"/> | High Low | 140 cents 140 cents | 50% 50% |
| 2 | <input type="radio"/> | High Low | 180 cents 120 cents | 50% 50% |
| 3 | <input type="radio"/> | High Low | 220 cents 100 cents | 50% 50% |
| 4 | <input type="radio"/> | High Low | 260 cents 80 cents | 50% 50% |
| 5 | <input type="radio"/> | High Low | 300 cents 60 cents | 50% 50% |
| 6 | <input type="radio"/> | High Low | 350 cents 0 cents | 50% 50% |

[Confirm your choice](#)

If you have any questions, please raise your hand and one of us will come to your desk.