

Morals in multi-unit markets*

Andreas G. B. Ziegler¹, Giorgia Romagnoli², Theo Offerman²

¹University of Essex, ²CREED, University of Amsterdam

October 25, 2023

Abstract

We examine how the erosion of morals, norms, and norm compliance in markets depends on the market power of individual traders. Previously studied markets allow traders to exchange at most one unit and provide market power to individual traders by de-activating 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 aggregate quantities. In an experiment, we compare single-unit to (more common) multi-unit markets, which may activate these forces. Multi-unit markets, in contrast to single-unit markets, lead to a complete erosion of morals. This is associated primarily with a deterioration in norm compliance: the observed level of immoral trade is in contrast with the prevailing social norm. The replacement logic is the main mechanism driving this finding.

Keywords replacement logic, selection, markets, negative externalities

JEL Codes C91, C92, D62

*Acknowledgements: We thank Björn Bartling, Sourav Bhattacharya, Guillaume Fréchette, Jana Friedrichsen, Hartmut Kliemt, Boris van Leeuwen, Klaus M. Schmidt, 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, Virtual Market Design Seminar, M-BEES 2021, Spring Meeting of Young Economists 2021, TIBER 2021, ESEM 2021, New Trends in Decision-Making: Judgement and Reasoning, ECORES Summer School 2022, 6th Solomon Lew Conference on Behavioral Economics, ACBC workshop on morals and repugnance in markets, MPI-CREED workshop 2023, ESA World 2023, CERGE-EI Prague, CeDEx Nottingham, GATE Lyon and Tilburg University 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 and RCT ID AEARCTR-0008306.

Email: mail@andreasziegler.eu (Ziegler, corresponding author), g.romagnoli@uva.nl (Romagnoli), t.j.s.offerma@uva.nl (Offerman).

1 Introduction

People’s morals may easily take a back seat in markets. Consider the market for air travel. Passengers may think that someone else could take their place if they refrain from buying a cheap ticket, leaving total emissions unchanged. This reasoning, the so-called replacement logic, may explain why frequent flying also occurs among environmentally conscious consumers (Barr, Shaw & Coles 2011). Airlines themselves may justify 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 selfish trading behavior as well as the disproportional activity of a few irresponsible actors are features common to several morally questionable or highly polluting markets, such as the opioids market, 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 participants are willing to kill a mouse for €10 in individual decision-making, 75.9% do so in single-unit markets. In the multi-lateral bargaining setting, they also find a decline in prices as a result of competition, which the authors interpret as further evidence for moral erosion. However, key results of Falk & Szech (2013) are contested. Market prices can also decline 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 (2023) show that the partial erosion in markets disappears

¹In the opioids market, a spokeswoman 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 contrast, also in this market, single firms can be responsible for a significant share of overall harm: Purdue Pharma’s marketing campaign for OxyContin increased sales and the associated overdose deaths (Alpert, Evans, Lieber & Powell 2022). 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, Vuillemeys (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 from operating in their country, but that then someone else would step in to supply the arms that they supplied (Bartling & Özdemir 2023, Falk, Neuber & Szech 2020).

under repetition of both environments. So far, the evidence that people’s morals are eroded in markets is inconclusive.²

In our view, many real-world markets are poorly approximated by the single-unit market paradigms studied so far. In addition, these markets inhibit forces that may contribute to a strong erosion of morals. This may have led to an underestimation of the effect of markets on morals so far. In this paper, we focus on more realistic multi-unit markets.³

In our experiment, people’s (im)morality is reflected in their willingness to increase a negative externality in exchange for personal monetary gains. The externality is implemented as canceled donations for measles vaccines to UNICEF—consistently with our interpretation, canceling these donations for relatively small returns is considered immoral by our participants.⁴ As our main measure of morality in markets, we study aggregate quantities traded, as an increase in trading volumes is associated with an increase in the generated externalities.⁵ As a secondary measure of morality, we compare the willingnesses to pay to avoid negative externalities between markets and individual decision-making. This replicates the methodology used by Falk & Szech (2013) and ties us to the existing literature.

We investigate the moral erosion of markets that vary in the power that individuals have in preventing trade. The main contribution of our paper is threefold: (i) To measure whether and to what extent multi-unit markets erode morals, and in particular whether a change in individuals’ market power affects moral erosion; (ii) To determine the extent to which norms and norm compliance change in these markets; (iii) To distinguish between two forces that may drive moral erosion in multi-unit markets: *market selection* and *replacement logic*.

For market selection, we assume market participants trade whenever the material profits exceed their moral costs associated with causing negative externalities. Multi-unit markets remove individuals’ constraint to trade at most once. This allows the less-moral participants to capture a larger share of the market, as they can also trade units associated with low profits. Trade stops only when even the participants least concerned about the externalities are no longer willing to trade. Market selection then

²In Section 2, we position our paper more precisely in the literature.

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).

⁴Using the elicitation method by Krupka & Weber (2013), we find that 841 out of 1022 participants rate taking €1 as a payment to one-self instead of donating €1.50 to UNICEF in an individual decision-making task as “socially inappropriate” and “inconsistent with moral or proper social behavior”.

⁵In Appendix Section A.1 we provide a framework to predict how trade would unfold in markets if traders’ morals are assumed unchanged between individual decision-making and markets, an exercise we call moral competitive equilibrium. This approach leads to very similar results as the simpler one presented in the main text.

implies that the least moral traders predominantly determine aggregate outcomes in multi-unit markets, as the abstention of the more-moral traders no longer restricts the exchange of additional units. This effect is further enhanced when preferences are characterized by diminishing marginal moral costs for the negative externality, as trading repeatedly generates an additional competitive advantage for the least moral traders.

According to the replacement logic, market participants may decide to trade selfishly, as they realize that their individual actions do not affect adverse aggregate outcomes. This can justify their own trading and reaping the profits for themselves (Sobel 2007). A necessary ingredient for the replacement logic to be active is that no trader is pivotal. Crucially, in the single-unit markets studied so far, for each pair of active traders, at least one of them is pivotal: The total quantity traded would be reduced if this trader refrains from trading. This reduces the scope for the replacement logic. In the multi-unit markets we study, no trader is pivotal. Thus, traders on both sides of the market can excuse their trading with the argument that someone else would have taken advantage of the opportunity if they had not traded.

So far, the literature could not identify whether any erosion of morality in markets was due to a change in norms (i.e., in what people consider socially appropriate behavior), or a change in the degree of compliance to norms (i.e., the extent to which people’s behavior agrees with their norms). For policy applications, this distinction is essential. If norms are unchanged while norm compliance deteriorates, and thereby even the traders themselves regard their behavior in the market as inappropriate, this could lead to a stronger case for market regulation. We disentangle these two hypotheses by separately eliciting norms.

Our experiment is based on four main between-subject treatments: Three multi-lateral market treatments and, as in the previous literature, an individual decision-making control treatment, MPL. We repeat individual decision-making in MPL as often as we repeat all markets. This allows us to control for a potential erosive effect of repetition. In addition to the separate MPL treatment, we also use the individual decision-making task to elicit individual preferences at the start of all market treatments.

Across our market treatments, we vary how many units each market participant can trade. Our first market, treatment SINGLE, is a single-unit market. This treatment is comparable to the markets studied in the current literature and acts as a benchmark for the main market treatments of interest, MULTI and FULL. MULTI is a scaled-up version of SINGLE, where instead of one unit, three units per participant can be traded in each market period. In MULTI, each trader is similarly pivotal as

in SINGLE. 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.

In all market treatments, we use a common supply and demand schedule. With this schedule, costs and values are equalized across all traders, i.e., they only change in the aggregate quantity exchanged by all traders. A common schedule allows us to study the behavior of the traders holding constant monetary gains from all trades. It models features that are typical of markets with negative externalities, such as the ones for weapons and flights, where the common components of costs and values are very salient.

We first shed light on the debate of Falk & Szech (2013) and Bartling et al. (2023), that is, do single-unit markets erode people’s morals? We find evidence for a partial erosion of morals in single-unit markets. We then move to our key focus, trade in multi-unit markets. We find that the erosion in treatment MULTI is partial and comparable to SINGLE. Thus, moral erosion is not tied to the number of units a trader can trade *per se*. Instead, we detect a complete erosion of morals in FULL, which activates both the replacement logic and market selection. Trading in this unrestricted market is statistically indistinguishable from selfish competitive equilibrium, consistent with participants completely disregarding that their trading causes negative externalities.

Next, we disentangle the role of norms and norm compliance. While we document a small erosion of norms in all markets compared to individual decision-making tasks, remarkably, norms do not differ after experiencing all market treatments: On average, trading is considered approximately equally socially inappropriate in all markets. Since norms are unchanged, we infer that morals in multi-unit markets are eroded by a deterioration in norm compliance.

We further show that the replacement logic largely drives the complete moral erosion in FULL. In FULL, trade is carried forward by the vast majority of traders and not only the less moral fraction of participants. 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 at these monetary terms.

Additional treatments provide direct evidence for the two mechanisms of market selection and replacement logic. To shed light on the selection argument, we include a treatment similar to FULL in which we divide participants based on their individual decision-making preferences in either a homogeneous group or a heterogeneous group. In the homogeneous group, participants know that they are matched with

traders who, just like them, are close to the median moral preference, which should substantially reduce the scope for market selection. Even under these circumstances, the market exhibits the same degree of erosion documented in the FULL treatment. We infer that market selection does not contribute to an erosion of morals when the replacement logic is available. To shed direct light on the replacement logic, we include treatments similar to FULL and MULTI in which we elicit participants' beliefs about whether they are pivotal. In agreement with the replacement logic, we observe that participants believe to be more likely to be replaced in FULL than in MULTI and are more active when they think that they are more replaceable. In addition, we exogenously manipulate beliefs about their pivotality by selectively revealing information from earlier sessions. Again, consistent with the replacement logic, traders who are exogenously induced to believe that more others are active are more active themselves.

In the following, we start by positioning our paper in the related literature. We then describe the experimental design and 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 Related literature

In this section, we discuss how our paper contributes to the literature on erosion of morals in markets and the literature on erosion in other interactions. Following 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. In a market, traders affect each others' outcomes when they compete to buy and sell valuable products or services. According to this definition, the decisions people individually make when trading off money and a negative external effect in individual decision-making are not considered *market* decisions. There is no competition for a scarce good in individual decision-making, and people's decisions do not affect other traders' outcomes.⁶

We start with the related literature on moral erosion in markets. The paper by Falk & Szech (2013) inspired a follow-up literature that investigates how different market forces affect traders' morals. Bartling, Weber & Yao (2015) show that fair and unfair products can co-exist in a market and that it is not necessarily the case that unfair products crowd out fair products. They find only a modest role of erosion.

⁶Our finding that participants find trading less socially inappropriate in markets than in individual decision-making reveals that markets and individual decision-making do not only differ technically, but also in the minds of our participants.

In their Swiss sample, consumers make the fair choice on average 14 percentage points more often in the individual decision-making task than in the market, and the difference is not consistently significant across all specifications (in their Chinese sample they find slightly more erosion).⁷ Other papers have investigated the role played by other factors on moral erosion, such as anonymity, market framing, joint decision-making or relative share of buyers versus sellers affect traders' morals in markets (Kirchler, Huber, Stefan & Sutter 2016, Irlenbusch & Saxler 2019, Sutter et al. 2020). Engelmann, 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. All these papers exclusively focus on single-unit markets that deactivate the selection effect and the replacement logic. Instead, the forces they focus on are active across all our market treatments, so are held constant in the comparison between market treatments we are focusing on. All these studies also do not independently elicit participants' perceptions of norms, so they cannot distinguish between norm erosion and the erosion of norm compliance.

Besides Bartling et al. (2015), there are also some other papers that study specific market structures that allow markets to partially sustain pro-social behavior. Schneider, Brun & Weber (2020) document an endogenously arising wage premium, and associated sorting, for morally questionable tasks. Dufwenberg, Johansson-Stenman, Kirchler, Lindner & Schwaiger (2022) argue that reciprocity concerns can increase pro-sociality in market structures that allow for cycles among all traders. Other examples in which competition and pro-social behavior can be mutually reinforcing are provided by Byambadalai, Ma & Wiesen (2023) and van Leeuwen, Offerman & Schram (2020). In a large non-student sample, Riehm, Fugger, Gillen, Gretscho & Werner (2022) highlight the importance of norms in these types of markets: Traders prefer to condition their decisions on others' entry and punishment opportunities for immoral trading are frequently used. Kaufmann & Koszegi (2023) theoretically study equilibria and regulation of markets with socially responsible consumers. Ockenfels, Werner & Edenhofer (2020) and Herweg & Schmidt (2022) compare (experimentally the former, theoretically the latter) taxes and cap-and-trade schemes to regulate moral markets with negative externalities.

Our conjecture that market selection can be an important force is based on a literature that shows that there is substantial heterogeneity in people's social preferences (Offerman, Sonnemans & Schram 1996, Fischbacher, Gächter & Fehr 2001, Burlando & Guala 2005). Falk, Becker, Dohmen, Enke, Huffman & Sunde (2018) document heterogeneity in social preferences within and across many countries. Given that

⁷Bartling et al. (2015)'s findings are robust to different specifications of the externalities (Bartling, Valero & Weber 2019).

the most immoral traders are the ones who may determine how much is traded in a market, heterogeneity can furnish selfish aggregate outcomes.

Our paper also contributes to a literature that investigates how the replacement logic and diffusion of pivotality affect behavior in non-market games. Dana, Weber & Kuang (2007) show that a diffused responsibility for moral outcomes erodes moral behavior in dictator games. Grossman (2014) demonstrates that this effect survives when participants have to actively seek to remain ignorant. In an individual decision-making context, Falk & Szech (2014) find that almost a third of their participants pay for a diffused notion of being pivotal for a questionable moral outcome. Serra-Garcia & Szech (2022) study how the demand for moral ignorance depends on monetary incentives. They find that the demand for ignorance does not respond to social norm messages. Exley (2016) demonstrates that uncertainty about the impact of a charity may serve as an excuse not to give. Falk et al. (2020) find support for the replacement logic in committee decisions. A string of papers study diffusion of pivotality in ultimatum games with proposer or responder competition. Roth, Prasnikar, Okuno-Fujiwara & Zamir (1991), Prasnikar & Roth (1992) and Fischbacher, Fong & Fehr (2009) find that the side with competition receives almost nothing of the endowment.⁸

There are also studies that find only limited support for the replacement logic. Bartling & Özdemir (2023) find that participants do not employ the replacement excuse if a social norm exists that classifies the selfish action as immoral. In a voting context, Brütt, Schram & Sonnemans (2020) find mixed evidence for the effect of decreased pivotality.

An important contribution of Behavioral Economics is to study how findings from stylized, simple settings generalize to market settings (e.g. List 2003, Enke & Zimmermann 2019, Enke, Graeber & Oprea 2023). In this light, our paper studies the generalizability of the replacement logic to markets. Compared to the previous stylized settings, we can study the importance of the replacement logic in a market environment where competing forces are active. These can be previously studied forces that erode morals already in single-unit markets, as well as the market selection effect we introduce in multi-unit markets. Our findings show that the replacement logic substantially increases the erosion of morals in markets, beyond the erosion in single-unit markets. Lastly, our paper sheds novel light on how norms and norm compliance are shaped by the availability of the replacement logic argument. In particular, we find a full erosion of morals driven by the replacement logic, against the prevailing norm.

⁸There is also theoretical work on the replacement logic. Besides Sobel (2007), the papers of Huck & Konrad (2005), Grossman & Van Der Weele (2017), and Rothenhäusler, Schweizer & Szech (2018) have theoretically studied diffused notions of pivotality.

3 Experimental design

The experiment consisted of three main parts.

Parts 1 and 3 were identical to each other and the same in all treatments. In these parts, participants faced an individual decision-making task which elicited their willingness to accept (WTA) to cancel donations towards UNICEF for varying stakes. In Section 3.3, we give more details on the donation opportunity. We employed multiple price lists where participants 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. Each participant 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 participant's moral costs equal to the payment at which the participant switched.⁹ We set the moral costs of participants who never choose to cancel a donation equal to the upper bound of the multiple price list.

Part 2 varied in the between-subject treatments. In our control treatment (employing only individual decision-making tasks, we call the treatment MPL), part 2 presented a repetition of the task of part 1 for four times. In the three main market treatments, four market periods were implemented.

3.1 Markets

We implemented two-sided posted offer markets characterized by common supply and demand schedules. We here explain these features and the rationale behind them.

3.1.1 Two-sided posted offer markets

We implemented the market as a two-sided posted offer market with induced values and costs. Each market consisted of five buyers and five sellers interacting repeatedly and anonymously. Buyers posted bids, sellers asks, and all traders could accept an offer of the other market side. If accepted, a trade was implemented at the price of the accepted offer. The buyer received a payment corresponding to the induced value minus the price and the seller received a payment equal to the price minus the induced costs. For every unit traded, a donation to UNICEF which costs approximately €1.50 was cancelled.

Buyers and sellers moved in turns, trading unit by unit. In each market period, one side of the market – i.e., the buyers or the sellers – was randomly determined

⁹We do this to match behavior in the markets, where we can only infer that a participant's moral costs is at most equal to the profit margin of a submitted or accepted offer.

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 pre-existing 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) participants have 14 seconds to decide, which gives participants sufficient time to think and simultaneously generates observations on the willingness to trade for all active traders (and not only the fastest to react). This goes beyond what is normally observed in a traditional double auction where trade is implemented immediately after agreement. Notice further that the posted offer element fits the product markets that we target, whereas standard double auction rules are more representative of financial markets.

To ensure that the negative externalities were salient, each time when participants traded a unit and at the conclusion of a market period, traders were reminded about the consequences of their trading for the charity.

3.1.2 The common schedule

In our markets, we use a common schedule. In a common schedule, a seller's cost for supplying a unit and a buyer's value for buying a unit depend on the total quantity already traded in the market, while they are held constant between traders. As a consequence, costs and values depend on the timing of when the trade happens,

compared to the other trades in the market. In the common schedule of our paper, for any trader, profit margins of early trades are larger than profit margins of later trades. In contrast, in a private schedule, each trader's costs and values depend only on the quantity traded by themselves, and they differ across traders.

Our motivation for choosing a common schedule is threefold. First, a common schedule captures essential features of the markets that we target. While real-world market schedules have both private and common elements, we think that in markets with negative external effects common elements are often particularly salient. Consider for instance the market for weapons. In a war, the buyers of weapons benefit much more from guns that they are able to secure early in the conflict than guns that they obtain later, while at every moment the strategic advantage the weapons afford are first-order similar across potential buyers. Likewise, in the short run, there is only a limited number of factories in the world that produce for instance AK-47 guns, and a trader who acquires these guns early may do so at lower costs than a trader who does it later when the factories are closer to their capacity constraints. Thus, in the market for weapons, the willingness to pay for the products and the costs of the products depend to a large extent on the timing of the trade. Similar common schedule features characterize other important markets with negative external effects. In the aviation market, airlines lease a substantial part of the aircrafts. This feature represents a strong common cost element for airlines in this market. Consumers may prefer to fly to interesting places before they become less attractive for everyone due to overtourism. In the market for illegal construction permits, constructors will prefer to acquire early permits which allow them to choose the best spots to build their resorts. Corrupt officials will find it easier to hand out early permits before public opposition becomes organized.¹⁰

Second, such a schedule has the advantage of providing a clean interpretation of trading data: For each unit traded, all buyers (sellers) face the same values (costs). As the gains from trade for all participants are equated, differences in willingness to trade cannot be ascribed to differences in costs or values.

Third, equalizing the monetary terms across participants after each trade ensures that traders remain fully replaceable with each other. This means that both the replacement logic argument and market selection have the same opportunity to arise, irrespective of traders' earlier behavior. In contrast, with a private schedule, participants who had refrained from trading gain a competitive advantage, which inhibits

¹⁰For some background on these markets, see <https://www.theguardian.com/world/2001/jul/09/armstrade.iantraynor>; "Mid-life aircraft trading patterns and the impact of lessors". Flightglobal, 7 March 2017; <https://www.theguardian.com/world/2020/jan/25/overtourism-in-europe-historic-cities-sparks-backlash>; <https://www.phnompenhpost.com/national/apsara-raises-concerns-over-illegal-construction-angkor>.

both forces.

Opportunities to replace other traders can also occur in markets with private schedules. Here, the shape and slope of the private schedules affect the size of the maximal potential impact for moral erosion that can be produced by the replacement logic argument and market selection. In Appendix Section A.13, we provide a few examples of private schedules that may trigger replacement thinking.

3.1.3 Main market treatments

We ran three main market treatments: SINGLE, MULTI and FULL. In the single-unit market treatment, SINGLE, each trader is restricted to trade at most one unit, so up to five units could be traded in the entire market. This treatment allows for most market forces of erosion considered in the current literature.

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. Therefore, MULTI represents a scaled-up version of SINGLE, where each trader remains pivotal for the units assigned to her.

We allowed each trader to cater to the entire market in the unrestricted market, FULL. Treatment FULL has the same demand and supply schedules as MULTI. It differs from it in one key aspect: FULL removes the capacity constraints of each trader. This means that each participant is able to trade up to 15 units and thus serve the entire market.

In all treatments, costs and values each trader faces are identical (as a consequence of the common schedule) 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 in subsequent units where market participants could split €0.60, €0.40 and €0.20.

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.

Participants' trading in the practice market without externalities allows us to see if our design features lead to different market outcomes than previously established in the literature. Across all groups, all units were traded in the practice market period.




Figure 1: Induced common costs and values

Therefore, our trading institution produces standard results for experimental markets in the absence of externalities. Lower trading volumes can be cleanly attributed to the introduction of negative externalities.

3.1.4 Other treatments

We included some follow-up treatments that allow us to further investigate the mechanisms behind our main results. To provide direct evidence on the selection effect, we ran two additional FULL markets differentially activating market selection. On the basis of participants' moral costs elicited in individual decision-making in part 1, we formed groups either consisting of the middle two quartiles (so, participants close to the median preference) or of the first and fourth quartiles. The latter, HET, fully activates market selection as participants are very heterogeneous in their preference for the external effect. The former, HOM, generates homogeneous market groups, where market selection has less scope to affect outcomes. To ensure that participants hold correct beliefs about their fellow traders' morals, we informed participants both at the start of part 1 and part 2 of the group formation procedure, in part 2 they also learned which type of group they belonged to.¹¹

To shed direct light on the replacement logic, we included three treatments in which we directly elicited beliefs about other traders' activity in markets. Treatments B-MULTI and B-FULL replicate MULTI and FULL with additional belief elicitations about the trading of unit 10, 12, 13 and 15. Just before trading of these units started, traders reported their non-incentivized beliefs about the probability that the next unit will be traded, either with or without their participation. In addition,

¹¹This information was processed well, as beliefs about the median participants' morals are more precise in HOM (average absolute error of 38.8) than in HET (average absolute error of 69.8), the difference is statistically significant (MWU, 8 observations per treatment, p -value=.003).

we elicited the (cognitively less demanding) prediction of how many of the other traders will attempt to trade the next unit. This last prediction was incentivized: If and only if participants predicted this number correctly, they would earn a bonus of €1.50. Next to the treatments with direct belief elicitation in the markets, we ran a treatment with spectators, SPEC. The spectators were not directly involved in any market transaction. Instead, they followed the series of screens and received the identical information of a randomly matched participant from B-FULL and were asked to report their own beliefs in the same fashion. Comparing B-FULL and SPEC allows us to test for self-serving belief reports in B-FULL.

We include two additional market treatments to identify a causal effect of beliefs on trading behavior, as would be expected on the basis of the replacement logic. Treatments B-LOW and B-HIGH are similar to treatment B-FULL in that we elicit (incentivized) beliefs about the number of other active traders in the trading of units 4, 13 and 15. To identify a causal effect of beliefs, we differentially manipulate beliefs by providing information about the trading of that unit in an earlier session just before eliciting traders' beliefs. In B-HIGH, we provide information from a session in treatment B-FULL, in which many other traders had been active. For example, an average of 8.5 traders were active for unit 13 in that session. In B-LOW, we provide information from a session in treatment B-MULTI, in which few other traders had been active. For example, an average of 0 traders were active for unit 13. To avoid deception, traders were informed in the instructions that the data was either obtained from a market treatment which followed their own market rules, "or slightly different in that each participant could only trade a maximum number of units that was different to yours".

As preregistered, our main interest in this treatment was to study beliefs and individuals' trading behavior the first time that their beliefs were shocked differentially, which was at unit 13 in period 1. In addition, our main interest was in the first movers at this point, so the side of the market that first made offers. We did so, as we expected the effects of revealing information to be strongest early on.¹² As the trading at this unit would be the main focus, we also shortened the second part of the experiment to only include two market periods where externalities are present.

3.1.5 Treatment overview

In Table 1, we provide an overview of all experimental treatments discussed in the main text. The first four treatments are our main treatments, the remaining seven

¹²At all later points, traders learn more about their own markets. Based on the full trading in all other market treatments with FULL rules, we expected this to be intensive trading, which limits our ability to induce beliefs that others abstain from trading.

treatments are follow-up treatments that shed light on the mechanisms.

Table 1: Treatment overview

Treatment	Part 1	Part 2	Part 3	Market rules	Manipulation
MPL	IDM	4 periods IDM	IDM	-	-
SINGLE	IDM	1+4 market periods	IDM	SINGLE	-
MULTI	IDM	1+4 market periods	IDM	MULTI	-
FULL	IDM	1+4 market periods	IDM	FULL	-
HOM	IDM	1+4 market periods	IDM	FULL	Homogeneous moral costs
HET	IDM	1+4 market periods	IDM	FULL	Heterogeneous moral costs
B-MULTI	IDM	1+4 market periods	IDM	MULTI	Belief elicitation
B-FULL	IDM	1+4 market periods	IDM	FULL	Belief elicitation
SPEC	Beliefs	-	-	FULL	Spectators in B-FULL
B-HIGH	IDM	1+2 market periods	IDM	FULL	Beliefs with manipulation
B-LOW	IDM	1+2 market periods	IDM	FULL	Beliefs with manipulation

Note: Overview of the first three parts across all experimental treatments. IDM refers to the individual decision-making task, where we elicit participants’ willingness to cancel UNICEF donations using multiple price lists. 1+4 and 1+2 market periods, respectively, refer to one practice market period without externalities followed by four, or two, market periods where trading cancels UNICEF donations.

3.2 Additional elicitations

In all treatments, we included additional measurements of participants’ views and attitudes after part 3. We elicited: (i) beliefs about the median participant’s WTA to cancel donations within the individual decision-making task; (ii) norms about behavior in individual decision-making and markets; (iii) risk preferences. For the beliefs, participants 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 participants’ perception of the norms for canceling donations in either individual decision-making or the market, we followed the procedure by Krupka & Weber (2013) and asked participants 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, participants 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.6. We also elicited risk attitudes using the method introduced by Eckel & Grossman (2002).

3.3 Experimental procedures

For the treatments MPL, SINGLE, MULTI and FULL, 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 women, with an average age of 21. We had 100 participants per market treatment and 81 participants in MPL. Sessions lasted on average 1.5 hours, with average payments of €19 per participant, besides payments to UNICEF.

We conducted the follow-up treatments from October 2021 to January 2022. These were preregistered separately (Offerman, Romagnoli & Ziegler 2021). In total, 441 subjects participated in this second set of sessions. Out of those, 208 participants were recruited from the pool at the CREED laboratory at the University of Amsterdam. The remaining 233 participants were recruited from the pool at the CentERlab at Tilburg University. Treatments were balanced in the composition of participants from Amsterdam and Tilburg (between 63% and 69% of participants were from Tilburg). All treatments consisted of 80 participants, apart from SPEC with 41 participants. 55% of participants were women, with an average age of 21. Sessions lasted on average 1.7 hours, with average payments of €20.4.

We conducted additional follow-up treatments in March and April 2023. These were again preregistered separately (Offerman, Romagnoli & Ziegler 2023). In total, 200 subjects from the CREED laboratory at the University of Amsterdam participated in this third set of sessions. 53% of participants were women, with an average age of 21. Sessions lasted on average 1.2 hours, with average payments of €24.8. To account for higher minimum average earnings at the CREED laboratory in 2023, we included a bonus payment of €3 for completing the questionnaire.

In Appendix Section A.2, we show that participant characteristics are balanced across all treatments.

Participants 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 participant. 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, participants received a show-up fee of €7, all earnings from the three additional elicitations 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.

Participants 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. Participants were required to complete a set of test questions before they could proceed. Participants were paid in cash and in private at the end of the experiment.

In the experiment, several choices affected donations to UNICEF. As in Kirchler et al. (2016) and Sutter et al. (2020), donations were intended for measles vaccine. We used a text of UNICEF to inform participants about the consequences of measles.¹³ 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 participants in the instructions and the handout.¹⁴ In the instructions, participants were presented with sample receipts of such a donation to UNICEF.¹⁵ At the end of each experimental session, the donation was immediately implemented by the experimenter. Participants 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 participants so desired. These emails were collected on separate handouts and thus could not be linked to specific participants or choices in the experiment. Participants were made aware of this procedure at the start of the experiment. In total, approximately €2634 (€889 in 2019, €1222 in 2021/22 and €523 in 2023) was donated to UNICEF as a result of participants' choices.

4 Hypotheses

In this section, we elaborate on the hypotheses behind the main contributions of this paper, namely (i) the role played by market power in eroding morals in markets; (ii) the distinction between norm erosion and the erosion of norm compliance; and (iii) the separation of the role played by the replacement logic vis-à-vis market selection. These hypotheses, preregistered in (Offerman et al. 2019), (Offerman et al. 2021),

¹³“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.”

¹⁴This particular donation was only available in packs of 40 doses, excess donations were made over to UNICEF as a generic donation, which participants were aware of and could verify as well.

¹⁵At the time of the sessions in 2019, this donation is available at <https://market.unicef.org.uk/inspired-gifts/measles-vaccines-to-protect-20-children/S359163X/>, which we also communicated to participants. In 2021/22, we instead donated to UNICEF in Austria, <https://unicef.at/shop/produkte/>. Costs per dose were approximately constant and all procedures were kept identical otherwise.

and (Offerman et al. 2023), are summarized and motivated below.

The erosion of morals in single-unit markets

We start by exploring the erosion of morals in single-unit markets by comparing our treatment SINGLE to individual decision-making data. In doing so, we replicate the treatment effects from prior literature in our setting. Falk & Szech (2013) report limited erosion of morals in single-unit markets. Bartling et al. (2015) find limited erosion in most specifications. Bartling et al. (2023) fails to reject this hypothesis. Our first hypothesis is thus:

H1. There is no erosion of morals in single-unit markets.

The erosion of morals in multi-unit markets with market power (MULTI)

The following hypothesis bridges our multi-unit markets to the current literature, which studied single-unit markets. Treatment MULTI is a scaled-up version of SINGLE. In both treatments, a single trader can trade up to 1/5th of the maximal market size and retains full pivotality in that she can unilaterally decide to reduce the maximum aggregate quantity by not trading her units.

H2A. Compared to single-unit markets (SINGLE), there is no additional erosion in restricted multi-unit markets (MULTI).

The erosion of morals in multi-unit markets without market power (FULL)

While MULTI serves as a benchmark treatment for the introduction of multi-unit trading, the next hypothesis is the key hypothesis in our paper. Here, we focus on unrestricted multi-unit markets with treatment FULL. Between MULTI and FULL, the market structure remains identical, apart from removing individual traders' capacity constraints, so each trader can serve the entire market.

H2B. Unrestricted multi-unit markets (FULL) do not lead to more moral erosion than restricted multi-unit markets (MULTI).

Norm erosion and erosion of norm compliance

Our next hypothesis is concerned with the question of whether differences in the degree of moral erosion across treatments are due to changes in norms or in the degree of norm compliance.

H3. Norms are (A) not eroded in markets in comparison to individual decision-making and (B) not differentially affected by the specific market institution.

H3 is also a key hypothesis of our paper. Our independent measures for subjects' norms allow us to distinguish between norm erosion and the erosion of norm compliance. Previous literature highlighted the importance of norms for the availability of the replacement logic (Bartling & Özdemir 2023).

The mechanisms behind moral erosion in unrestricted markets: Market selection versus replacement logic

Our remaining hypotheses are concerned with investigating the relative role played by the two mechanisms of market selection and replacement logic that we expect to detect in treatment FULL. We here provide a definition of both forces.

Market selection. According to this mechanism, traders compare the material profit from trading to the moral costs that they incur from imposing the associated externality. Each trader continues to trade until their own moral costs no longer justify the monetary returns. As trade progresses, the profit margins get smaller, justifying trade for an ever smaller number of traders, i.e., those for whom moral costs are lowest. The final units will be traded by the traders with the lowest moral costs within their market. Additionally, a potential decrease in the least moral traders' marginal moral costs further increases the quantity traded.

The replacement logic. The replacement logic is a mechanism based on the following strategic thinking: Traders ask themselves whether their trading will affect the aggregate quantity traded in the market, assuming that other traders behave as if they are selfish (thus willing to trade all units available to them). If under this assumption their own behavior would not impact the aggregate volume traded, then this motive convinces them to trade irrespective of their own moral costs.

Notice that the belief of other traders behaving selfishly will be correct not only when other traders *are* actually selfish (i.e., genuinely unconcerned with the negative externality), but also when other moral traders *act* selfishly because they themselves apply replacement logic thinking, in a self-fulfilling cycle. Because traders can always replace each other in the unrestricted FULL market, the application of the replacement logic could lead to full trade and thus a full erosion of morals in this treatment. In the case of SINGLE or MULTI, traders' unilateral withdrawal from trade diminishes the aggregate quantity. This remains to hold even when all other traders act selfishly. Therefore, traders conclude that their behavior will matter for

the aggregate outcome and not trade units where moral costs exceed their profits. Notice that this view of the replacement logic is similar in spirit to Falk et al. (2020).

Our hypotheses regarding the mechanisms of moral erosion are thus:

- H4.** Any erosion of morals in FULL compared to MULTI is not driven by market selection.
- H5.** Any erosion of morals in FULL compared to MULTI is not driven by the replacement logic.

5 Results

In this section, we present the results of the experiment. For all market outcomes, we perform tests on the basis of averages of matching-group data, which yields 10 observations for each market treatment SINGLE, MULTI, FULL, B-LOW, and B-HIGH (10 groups with 10 participants each per treatment), as well as 8 observations for HOM, HET, B-MULTI, and B-FULL. MPL and SPEC feature no interaction, with 81 and 41 observations, respectively. For all tests on the individual level, for which participants do not interact, we study individual level data. To construct the confidence intervals in the graphs, we used a bootstrap procedure. We do this to correct for floor and ceiling effects of proportions close to 0% or 100%.¹⁶

5.1 Morals in individual decision-making

Across all treatments, we elicited participants' moral costs for canceling donations to UNICEF at the start of the experiment. In this section, we use this data to show that (i) our participants value this donation opportunity, as they require substantial payments to cancel donations, (ii) the moral cost distribution allows market selection to be active in the markets, (iii) the measure of moral costs is stable.

First, we study per-unit moral costs of all subjects in part 1 of the experiment. We show the payment that a subject requires to be willing to cancel a donation to UNICEF, averaged across all units at the subject level. That is, from the choice data for units 1, 2, 3, 5, 7, 10 and 15, we calculate the average per-unit valuation of a €1.50 donation to UNICEF.


In the individual decision-making task, the moral costs connected to causing the negative externality are quite substantial, with an average evaluation of €1.42 for a €1.50 donation to UNICEF.

¹⁶In the bias-corrected confidence intervals that we plot, we introduce clustering at the matching group level (the market group for market treatments and the participant for MPL or SPEC) and use 10,000 replications.

Next, we turn to the potential for market selection to matter. Two factors contribute to a potential effect of market selection in multi-unit markets: (i) initial heterogeneity in how traders value donations, and (ii) decreasing marginal moral costs in traders' preferences for causing the negative externality.

In Figure 2, we provide a histogram of the moral costs. Evidently, there is substantial heterogeneity in how participants value the opportunity to donate to UNICEF. A minority of participants hardly cares about donating to UNICEF. There is also a remarkable share of participants whose moral costs are estimated to be above €1.50, implying that they value donating *more* than the corresponding monetary value.¹⁷

Figure 2: Heterogeneity in valuations of donations



Notes: Histogram of participants' average moral costs for canceling a donation with a value of €1.5. For each participant, we use the switching points from all multiple price lists for canceling donations in part 1. Kernel density is displayed in green, the mean in red.

We also detect decreasing marginal moral costs and provide an analysis in Appendix Section A.3. Given these data, there is a clear possibility for market selection to play an important role.

Last, we can use the fact that this elicitation was repeated across all treatments in part 3, and especially an additionally four times in part 2 for the individual decision-making treatment MPL in part 2, to study whether we obtain a reliable measure of moral costs.

To do so, we calculate correlation coefficient of part 1 and part 3 moral costs

¹⁷Bénabou, 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, since image motives are affected differently by these methods. In our experiment, we keep the elicitation method constant across treatments. In our data, we find only few “observationally deontological” participants, those who never cancel a donation across all price lists, as only 32 out of 981 participants do so across part 1, compared to 26% of participants who do not cancel the donation for any monetary amount in Bénabou et al. (2020).

pooling across all treatments, as well as across-period correlations for treatment MPL. We find a very high correlation of this measure: The correlation between part 1 and 3-measures is 0.90. Calculating the across-period correlations in MPL generates a minimum correlation of 0.91, with most correlations being even higher. We report the full correlation matrix in Appendix Section A.4. There, we also show the robustness of this measure with Spearman rank-order correlations.


5.2 Moral erosion in markets

In this section, we investigate whether market behavior and outcomes display moral erosion. We start with measuring erosion in single-unit markets, as in Falk & Szech (2013). We compare decisions to cancel donations across individual decision-making tasks (in both part 1 of the market treatment and across all parts in the separate treatment MPL) and single-unit markets.¹⁸ In Figure 3, we plot the share of participants who cancel a donation in exchange for €1.50 (i.e., its value) or less in different treatments and periods. In the first two bars, we plot this variable as elicited from the individual decision-making task in part 1, identical to both treatments. We notice that the treatments are balanced in this dimension. The following two pairs of bars compare the same variable in the subsequent parts, differentiated by treatment: repeated individual decision-making in MPL or market behavior in SINGLE. For the markets, we study whether a trader concluded a trade for which she was paid at most €1.50. This is the comparison that speaks to the literature on erosion in single-unit markets. In the middle panel, we compare behavior in the first period in part 2. We observe that there appears to be an erosion of morals in markets. In the right panel, we use the entire four periods of the experiment and plot the share of participants who at least once cancelled a donation for at most €1.50 in part 2.

Table 2 reproduces estimation results of the corresponding effect. The dependent variable is a dummy variable equal to one if a participant cancelled a donation for at most its value either (i) in period 1 of part 2 or (ii) at least once in periods 1-4 of part 2. Models (1) and (2) suggest that there is erosion through repetition, as in Bartling et al. (2023): more participants cancel a donation in the entire part 2 compared to the first period only. In our setup, we do find evidence for an erosion in markets: models (3) and (4) suggest that more participants cancel a donation in SINGLE than in the corresponding time interval in MPL. Model (5) confirms that this is particularly strong when testing for erosion in the pooled data of part 2, compared to only the

¹⁸This analysis was not preregistered. It follows the analysis comparing single-unit markets to individual decision-making as in Falk & Szech (2013).

Figure 3: Cancellation of donations between environments and treatments



Notes: Share of participants who cancelled a donation for at most its value (€1.50) in individual decision-making tasks and in trades in the market. The left panel shows cancellation rates in part 1 of the experiment and the middle panel plots cancellation rates in the first period of part 2. The right panel displays the share of participants who, in the four periods of part 2, at least once cancelled a donation.

first period.¹⁹

¹⁹A more conservative approach would be to halve the moral costs in the market as a result of shared responsibility. The effect of erosion in SINGLE in models (4) and (5) is robust to defining erosion within markets as the decision to cancel a donation for a payment of €0.75 or less. For example, the estimate on SINGLE corresponding to (4) is .247 (p -value < .001). In Section 5.3, we also present direct evidence on norm erosion between individual decision-making and markets. Summarizing, we find evidence for both a partial erosion of morals in markets as well as erosion when measured by a participant canceling a donation at least once in a repeated task, compared to a non-repeated measurement.

Table 2: Erosion in single-unit markets and through repetition

	(1) MPL	(2) SINGLE	(3)	(4) MPL & SINGLE	(5)
			Period 1	Period 1-4	Pooled data
Period 1-4	0.099*** (0.033)	0.270*** (0.052)			0.099*** (0.033)
SINGLE			0.126* (0.074)	0.297*** (0.059)	0.126* (0.075)
SINGLE \times Period 1-4					0.171*** (0.060)
Constant	0.494*** (0.056)	0.620*** (0.051)	0.494*** (0.056)	0.593*** (0.055)	0.494*** (0.056)
Observations	162	200	181	181	362

Note: Dependent variable is a dummy equal to one if a participant cancelled a donation for a payment of at most its value (€1.50) either in markets (SINGLE) or in individual decision-making (MPL). Period 1-4 is a dummy variable equal to one if the choice is measured as occurring at least once in period 1 to 4 in part 2 of the experiment, the omitted category is cancellation in period 1 of part 2. SINGLE is a dummy equal to one if the choice occurred in treatment SINGLE, with the omitted category MPL. Standard errors, clustered on participant level for MPL and matching group level for SINGLE, are presented in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Result 1 *We reject hypothesis H1, and find partial erosion of morals in single-unit markets.*


Our key hypotheses are on behavior in multi-unit markets. Is there an erosion in these markets, *in excess* of the erosion we find in single-unit markets? To measure erosion, we will focus on aggregate quantities traded. Higher quantities imply larger negative externalities, so they are a natural measure of the overall effect of the market structure on the morality of trading outcomes. In addition, we can exploit that our markets featured decreasing gains from trade, while damages to UNICEF are kept constant at €1.50 per unit traded. Thus, the trading of larger volumes also implies that traders are willing to accept lower trading margins, which directly ties to the measure of moral erosion commonly used in the literature.

In Figure 4, we plot the observed market quantities. All quantities are relative to the selfish competitive equilibrium outcome, according to which 5 units are traded in SINGLE, and 15 units in MULTI and FULL.

The bars show traded quantities relative to the competitive equilibrium across the three treatments. SINGLE and MULTI show similar traded quantities, consistent with a comparable amount of erosion in these markets. In contrast, we observe that market outcomes in FULL are fully selfish. Traded quantities exceed quantities in other market treatments, indicating substantially stronger erosion in FULL.

The degree of erosion in FULL is striking if we take into account that gains from

Figure 4: Market outcomes



Notes: Average quantities relative to selfish competitive equilibrium. Trading units below 40% is efficient (gains from trade exceed the externality). Compared to the negative externality of €1.50 per unit, each unit between 40% and 60% yields gains from trade of €0.60, each unit between 60% and 80% yields €0.40 and each unit between 80% and 100% yields €0.20.

trade are rapidly shrinking, 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 of trade from 40% to 60% leads to additional negative externalities of €4.50, while traders receive €1.80. A further increase from 80% to 100% again yields damages of €4.50, however, traders only receive the meager total payments of €0.60.²⁰

In Table 3, 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.^{21 22}

Result 2 *We detect full erosion of morals in unrestricted multi-unit markets (FULL).*

²⁰This result is also supported by using part 1 data to predict market outcomes under the assumption that moral costs are not changing in a market environment. When we compute the moral competitive equilibrium, we find ample scope for market selection and erosion of morals in FULL. We provide details in Appendix Section A.1.

²¹These 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.5 for results.

²²We also included an additional control treatment in which we implemented a standard double auction with a private schedule, with a multi-unit design and a scope for replacement similar to MULTI. In this treatment, we assigned values and costs in such a way that the aggregate supply and demand coincides with MULTI. We report on these results in Appendix Section A.13. The main takeaway from this treatment is that morals are eroded to an approximately similar extent as in MULTI.

Table 3: Treatment effects

		SINGLE	MULTI	FULL
Quantity in %		75.5	78.3	99
<i>p</i> -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.

Erosion in MULTI is similar to erosion in SINGLE.


5.3 Norms and norm compliance

The preceding section presented evidence for a complete erosion of morals only in FULL markets. An important question is whether this change can be attributed to a change in norms or whether it is the result of an erosion of norm compliance. Did traders feel that canceling donations in exchange for minuscule profits in FULL was “consistent with moral or proper social behavior”?

To this end, we elicited participants’ norms about the appropriateness of canceling donations at the end of the experiment, after participants took part in the experimental markets or repeated individual decision-making tasks. We used the method proposed by Krupka & Weber (2013), so that participants were incentivized to report what they believed was their session’s modal answer on a 4-point scale from “very socially inappropriate” (indexed 1), to “very socially appropriate” (indexed 4) in response to scenarios in which a participant in an experiment chose to cancel donations of €1.50 when paid €1 either in an individual decision-making task or in an experimental market.

In Figure 5, we display the mean answers to two (otherwise identical) questions regarding the social appropriateness of canceling a €1.5 donation in exchange for €1 in an individual decision-making task (left panel), and in a market (right panel). We observe that across all market treatments and both environments, canceling such donations is rated on average at best as “somewhat socially inappropriate”. Thus, there does exist a clear norm that canceling donations and trading is not appropriate. This norm particularly contradicts the rather frenzied trading behavior observed in FULL.

Figure 5: Norms in individual decision-making and in markets



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

In accordance with even single-unit markets eroding morals, causing an externality in a market is perceived as less inappropriate as the same choice in individual decision-making (Wilcoxon signed-rank, 300 observations, p -value < .001).

Somewhat surprisingly, differences in elicited norms do not map one-to-one to differences in behavior between market treatments. In particular, the more selfish behavior in FULL is not supported by a further erosion of the norm compared to the other market treatments.²³ We cannot reject equality of norms in markets comparing SINGLE and MULTI (MWU, 100 observations per treatment, p -value=.238) and between MULTI and FULL (MWU, 100 observations per treatment, p -value=.705).²⁴ We report additional descriptive statistics for other scenarios in the Appendix Section A.6, which yield similar conclusions.

Even though norms do not further erode in FULL compared to the other treatments, we see a complete break-down of norm compliance. When traders can take advantage of trading opportunities foregone by other traders, norms take a back seat

²³We find no evidence of excuse-driven norm reports, see Appendix Section A.12.

²⁴Results are similar when regressing participants’ norms (2 elicitations for 781 participants, so 1562 observations) on treatment fixed effects, a dummy for the market scenario and interactions of this dummy with the treatment fixed effects, clustering standard errors on the matching group. Significant is the dummy for the market scenario (p -value = 0.02), but none of the interactions is significant (all p -values > .1). This confirms that there is not a specific treatment effect on norms in markets.

in participants' decision making. In the next section, we shed light on the question of whether the complete breakdown of norm compliance is caused by market selection or the replacement logic.

Result 3 *We reject hypothesis 3A. Traders find canceling a donation less inappropriate in markets than in individual decision-making. We do not reject hypothesis 3B. Norms are not differentially affected by market treatments. The finding that market outcomes are most selfish in FULL is caused by a breakdown of norm compliance.*

5.4 Mechanisms: Market selection versus replacement logic

A crucial question is the mechanism behind the full erosion of morals in FULL. In this section, we aim at providing direct evidence for market selection and the replacement logic separately.

Our first approach is to study which traders are active in the market. Under market selection, only the least moral participants trade the last units, while all other participants abstain. In contrast, the replacement logic can be used by any trader and is most powerful if many traders become active. We thus study which traders are active in submitting or accepting offers for the final units, the least profitable units which yield gains from trade of €0.20. To evaluate which type of trader is active we split the sample into those with below- and above-median moral costs in part 1. If market selection drives erosion in FULL, we would expect that few very immoral traders are active. If, in turn, the replacement logic is active, we expect many active traders, and there need not be a correlation between individual activity and the valuations in individual decision-making.

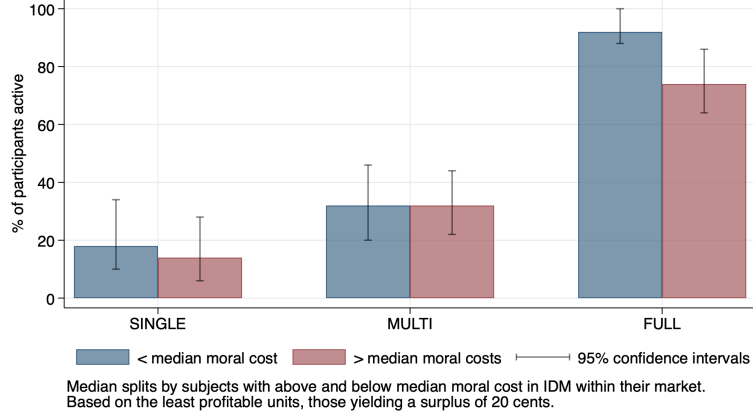
In Figure 6, we plot the share of traders who are active at least once at these least profitable units. We see that in both SINGLE and MULTI, both groups of market participants are similarly active. However, the share of active participants is much higher in FULL, where 94% of traders with below-median moral costs are active, but even 72% of traders with above-median moral costs are active. The difference between the above- and below-median group is significant only in FULL (MWU, 10 observations per above- or below-median group per treatment, p -value=.023).²⁵ This is however not robust to using a regression, see Appendix Section A.7.

This points to only a minor role for market selection. Traders with above-median moral costs are only slightly less active than traders with below-median moral costs.

²⁵Differences in other treatments move in the expected direction for earlier units with larger associated gains from trade, e.g. in MULTI 78% of above-median participants are active for units 10 to 12, while 92% of below-median participants are active.

This evidence hints at a major role for replacement thinking. A large share of participants are actively trading when the replacement logic is available, providing justification for the trading of others.²⁶

Figure 6: 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 are based on moral costs within the matching group.

A set of follow-up treatments distinguish between market selection and replacement logic more directly.

Direct evidence for market selection. In two new treatments with the same market structure as FULL, we divide participants into homogeneous groups that consist of traders close to the median preference for canceling donations (HOM), and heterogeneous groups that include traders on both extremes (HET). The main interest is in comparing outcomes in the HOM groups to the original FULL treatments as well as to HET. If market selection drives the erosion of morals in FULL, limiting its scope in HOM would lead to less erosion compared to the erosion in HET and FULL.

In Appendix Section A.2, we show that the participants in these two groups are balanced across other characteristics we observe. Yet, crucially, participants in HOM are more homogeneous than those in HET. Therefore, this treatment successfully manipulates the potential for market selection to drive outcomes, while other characteristics are not affected.

²⁶In the Appendix Section A.7, we provide further evidence in line with this analysis. While traders in SINGLE and MULTI submit or accept less than 1.4 offers on average, traders in FULL engage in 8.2 actions per trader. In addition, we show that a similar picture emerges for the traders who revealed to not use consequentialist reasoning in individual decision-making tasks, since they declined to cancel donations even when paid more than the monetary value of these donations.

Strikingly, market outcomes are similarly selfish in HOM, HET and FULL. Whereas in FULL 99% of units were traded, this is almost unchanged in HET with 100% of units traded, and, importantly, in HOM, where 99.8% of units were traded. Average quantities in HOM are not statistically distinguishable between HOM and FULL (MWU, 8 observations in HOM and 10 in FULL, p -value=.632) as well as between HOM and HET (MWU, 8 observations per treatment, p -value=.317). This indicates that even when limiting the scope of market selection, the replacement logic is sufficient to produce fully selfish market outcomes.

Result 4 *We do not reject Hypothesis 4. Both more and less moral traders are active. Market selection does not contribute to the complete erosion of morals in FULL.*

Direct evidence for the replacement logic: Analysis of beliefs. Treatments B-MULTI, B-FULL and SPEC allow us to shed direct light on the replacement logic. In these treatments, we directly elicited beliefs about others’ propensity to trade units 10, 12, 13 and 15 just before trading of these units started. We measured these beliefs in both an incentivized way (by asking ‘How many participants other than you will attempt to trade this unit?’) and an unincentivized way (by asking about the chances of pivotality: “What is the probability that whether or not the next unit is traded depends on your behavior?”). In the preregistration, we announced that we would focus on the non-incentivized measure if the two measures correlated. Unfortunately, the two measures do not correlate. Within B-MULTI, the Spearman correlation coefficient is -0.016 (p -value=.718, 500 observations). The same correlation in B-FULL is -0.003 (p -value=.910, 1280 observations).²⁷ In the main text, we therefore focus on the simpler incentivized measure.²⁸ Results for the non-incentivized measures are presented in the Appendix Section A.8 and are in line with these results unless otherwise noted in the main text.

We analyze belief data in two ways. First, we test whether our treatments induced differences in beliefs on others’ activity. If the replacement logic drives the enhanced trading in FULL, we would expect that participants believe that more traders are

²⁷This analysis assumes independence of observations, even though e.g. the same participant reports multiple beliefs. The conclusions are robust to using participant-level averages or regressions with standard errors clustered on a matching group level.


²⁸Other reasons to focus on the incentivized measure are that: (i) it correlates more strongly with the underlying true values; and (ii) while we do not find a correlation between incentivized and unincentivized measures for traders in B-MULTI and in B-FULL, we do find the expected correlation in the treatment using spectators (SPEC). The latter suggests that we may have been asking too much of our traders, and that they may have decided to focus on the incentivized questions. See Appendix Section A.8 for details.

active in FULL than in MULTI. Second, we check whether within-participant correlations between actions and beliefs are in line with replacement logic thinking, which would imply that participants who believe to be more replaceable are those who are more active.

Beliefs are only observed insofar as trade continues, which could bring selection issues for the between-treatment analysis.²⁹ To avoid selection issues, while still ensuring we analyze beliefs at the cusp of behavioral change, we preregistered that we would focus on beliefs that are observed for at least 13 out of 16 groups – i.e., where at least 13 groups have continued to trade up to the preceding unit. This turned out to be unit 10. For subsequent units, beliefs in B-MULTI are only available for a self-selected sample, as already at unit 12 only 40.6% of groups had continued to trade.³⁰

In Figure 7, we report the average number of other traders believed to be active in the trading of unit 10, including the corresponding target in the data. Traders in B-FULL believe that more other traders will be active than traders in B-MULTI do, consistent with replacement logic thinking. The difference between these two treatments is significant, with a p -value of .002 (MWU, 8 observations per treatment).

Figure 7: Beliefs about other traders' activity



Notes: Number of other traders believed to be active (grey bar), actual number of others active (the target; green diamond) and belief of spectators (orange circle).

Figure 7 also presents the target for these reports, based on the actual trading behavior of the other traders. Consistent with the beliefs, we observe more activity

²⁹Selection is a concern as groups that engage in more active trading are different than groups that stop earlier, along arguably relevant characteristics.

³⁰Treatments B-FULL and B-MULTI also allow us to investigate the robustness of the original results. In Appendix Section A.10 we reproduce the other analysis presented in the main text including the new treatments. Results are qualitatively in line with the original treatments.

in B-FULL than in B-MULTI already at unit 10. Lastly, we show the corresponding reports for the spectators, in SPEC. Directionally, this data is in line with self-serving reports, but differences between spectators' beliefs and traders' beliefs are minor and not significant (MWU, 8 observations in B-FULL and 41 in SPEC, p -value=.393).

We now test whether traders who believe to be more replaceable are those traders who trade most frequently. In Table 4, we regress the decision to be active in the market on participants' beliefs about others' activity. As we do not compare data across treatments, selection is not an issue, and we use the full data set. In columns (1) and (2), we observe that both in B-MULTI and B-FULL, participants who expect others to be more active are more inclined to trade themselves, again consistent with the replacement logic.³¹

Direct evidence for the replacement logic: Beliefs manipulation. To substantiate that beliefs are a causal driver of trading behavior (and not mere correlates), we ran two additional treatments similar to B-FULL in which we exogenously manipulated beliefs by selectively revealing trading behavior observed in an earlier session. In treatment B-HIGH, we provided information from groups where many participants attempted to trade units 13 and 15, in order to encourage participants to update their beliefs to be more replaceable. In contrast, in B-LOW, we provided information from groups where no participants attempted to trade units 13 and 15. After revealing this information, we elicited the incentivized belief we used in B-FULL. In the left panel of Figure 8, we show that we were successful in differentially moving beliefs between treatments.

Importantly, in the right panel of Figure 8, we show that the changes in beliefs also changed participants' trading behavior, in agreement with the replacement logic driving behavior in FULL markets. When participants beliefs were manipulated to expect to be less likely to be replaceable, they themselves decided to trade less. This allows to conclude that beliefs are not mere correlates of trading behavior, but their exogenous manipulation does move the propensity to trade in the direction predicted by the replacement logic.


The tests corresponding to these findings are reported in Table 4. In column (3), we see that beliefs about the number of other active traders decreased at the first time beliefs were shocked differentially, in period 1 for unit 13 (p -value< .001). Correspondingly, column (4) shows that less first movers are active at this unit after

³¹This is the only beliefs analysis that does not generalize when we use the unincentivized belief report (see Table A10 in Appendix Section A.8). Falk et al. (2020) show in a similar analysis that beliefs correlate with choice in line with the replacement logic. In their experiment, their incentivized belief did correlate with their unincentivized measure.

their beliefs have been shocked (p -value=.005).³²

This treatment also allows us to run an analysis of the causal effect of beliefs on behavior, using the treatment as an instrument.³³ For validity, we need to assume that the treatment affects trading behavior only through beliefs, which is plausible given our experimental setup. We report the first stage in Appendix Section A.9; the results indicate that the first stage is strong (Kleibergen-Paap F -statistic=87.1). The estimates reported in column (6) in Table 4 show that, as expected, a belief in more active others leads to higher activity. The point estimate is very similar to the correlation between action and beliefs reported in column (5).

Figure 8: Beliefs and actions with exogenously manipulated beliefs



Notes: Left panel: Number of other traders believed to be active (blue bars), right panel: actual number of active traders (red bars).

³²These outcomes were the preregistered main hypotheses for the new treatments. In Appendix Section A.9, we report on the robustness of this outcome at later periods and units. As expected, results are similar but are somewhat less pronounced, consistent with the participants learning over time about their particular market continuing to trade in B-LOW.

³³The IV analysis was not preregistered.

Table 4: Beliefs and activity

	(1) B-MULTI	(2) B-FULL	(3)	(4) B-LOW & B-HIGH	(5)	(6)
	Own activity		Belief activity		Own activity	
Belief activity	0.053*** (0.011)	0.069*** (0.013)			0.054*** (0.012)	
Average moral costs	-0.019 (0.023)	-0.111*** (0.025)			-0.147*** (0.033)	-0.147*** (0.035)
Period	0.001 (0.039)	-0.028* (0.013)				
B-LOW			-3.690*** (0.384)	-0.260*** (0.083)	-0.000 (0.080)	
Belief activity (IV)						0.055*** (0.016)
Constant	0.076 (0.119)	0.391** (0.121)	7.270*** (0.303)	0.700*** (0.060)	0.454*** (0.107)	0.453*** (0.121)
Observations	500	1280	200	100	200	200
Period	1-4	1-4	1	1	1	1
Units	10, 12, 13, 15	10, 12, 13, 15	13	13	13	13
Unit FE	Yes	Yes	No	No	No	No
Market-sides	both	both	both	first movers	both	both

Note: In (1), (2), and (4) to (6), the dependent variable is a dummy equal to one if a participant submitted or accepted an offer. In (3), the dependent variable is the belief about the number of other traders that are active, denoted belief activity. Average moral costs are the average moral costs for a participant, based on averaging per-unit moral costs based on the part 1 individual decision-making task. Standard errors clustered on matching group level in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Result 5 *We reject Hypothesis 5. Fully selfish market outcomes in unrestricted multi-unit markets are driven by the replacement logic.*

5.5 Additional findings

Our experimental design also allows us to test whether morals are eroded within an identical decision environment. All participants faced identical individual decision-making tasks in parts 1 and 3. In Appendix Section A.11, we present evidence showing that the average per-unit moral costs decrease slightly in MPL and SINGLE. In the multi-unit markets MULTI and FULL, erosion after market exposure is most drastic, with decreases of moral costs of 19.8 cents and 20.5 cents (relative to a donation of €1.50), respectively. Surprisingly, there seems to be some persistence of the erosion of morals outside of markets, especially so in multi-unit markets.

In Appendix Section A.11, we also report evidence for biased social learning. We elicited participants' beliefs about the median moral costs of canceling a donation in individual decision-making tasks. These incentivized beliefs were elicited at the end of the experiment, after participants experienced the markets or repeated individual decision-making task. Participants of multi-unit markets believe that their peers are more selfish than they truly are, in contrast to more accurate beliefs in MPL and SINGLE. This suggests that the participants do not sufficiently account for the fact that observed trading may not reflect the preferences of an average participant outside of the market.

6 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 traders to take advantage of trading opportunities foregone by other traders, 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 more participants canceling donations in markets than in individual decision-making. These results support Falk & Szech (2013)'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 these markets. At the same time, they are willing to forgo substantial amounts of money to prevent the externality, both before and after markets, in an individual decision-making task.

We further investigate the relative role played by market selection and the replacement logic in deteriorating market outcomes. We show that there is substantial heterogeneity in our traders' preferences for canceling donations, which leaves substantial scope for the selection effect to play a role. However, in our markets we find that less moral traders are hardly more active than more moral traders. Moreover, when we create homogeneous groups of traders who know that their preferences for the negative external effect are close to the median preference, we continue to see that all units in the market are traded. We conclude that the selection effect plays at most a minor role in our data. In contrast, and in agreement with the replacement logic, we find that (i) participants become more active in trading when they are more convinced that their behavior does not have an impact on the aggregate outcome and (ii) participants expect that their own behavior has less consequences for outcomes in FULL than in MULTI. Furthermore, our participants' beliefs are hardly biased in a self-serving direction, instead they correctly predict that many participants are trading.

It is particularly interesting and worrisome to see the extent to which replacement thinking can deteriorate market outcomes. Absent pivotality, large shares of participants engage in frenzied trade of units which cause large damages compared to the available gains from trade: 83% of participants are willing to trade when they can share gains from trade of €0.2, whereas only 9% of these same participants are willing to cancel the first donation when each is paid €0.2 in individual decision-making tasks, averaged on part 1 and 3-data.

Strikingly, this frenzied trading contrasts with the prevailing norm. Even though we observe some deterioration in participants' norms in markets compared to individual decision-making tasks, we do not see that norms are further eroded when pivotality of trading is removed. Still, norm compliance is completely eroded when participants can be replaced when others refrain from trading. This led to widespread frustration among participants, some of whom spontaneously wrote down their thoughts after the experiment. One participant commented: "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".

Our findings suggest implications for policy. Because selection effects hardly play a role, efforts to restrain the more immoral players in a market may not affect market

outcomes as long as these immoral players can be replaced by others. For instance, we think that it is doubtful that the recent dissolution of Purdue Pharma will solve the crisis in the opioids market. Instead, it may be more promising to pursue measures that restore or create pivotality in the market. One way to accomplish this would be to individually constrain traders in the quantities that they can trade. Treatment MULTI, which implements this constraint, shows much less moral erosion. Further, because even the traders themselves normatively disapprove of the outcomes in the unrestrained markets, we expect that there may be support for measures that restore pivotality. As an alternative to individual capacity constraints, externalities can be mitigated by introducing taxes on the relevant behavior (Plott 1983). On the other hand, aggregate quotas (i.e., *cap-and-trade* systems) can crowd-out moral behavior as they remove pivotality and make traders replaceable in the acquisition of the permits (Herweg & Schmidt 2022).

The large erosion of morals we detect has also implications for our understanding of markets as aggregators of preferences. Using existing market outcomes to infer individuals' preferences regarding damages to third parties seems not easily generalize to other (market) institutions. 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, as we discuss in Appendix Section A.11: Our participants 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 how much fellow members of our society would actually want to prevent the externalities we collectively cause.

References

- Alpert, A., Evans, W. N., Lieber, E. M. & Powell, D. (2022), ‘Origins of the opioid crisis and its enduring impacts’, *The Quarterly Journal of Economics* **137**(2), 1139–1179.
- 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. (2023), ‘Does market interaction erode moral values?’, *Review of Economics and Statistics* **105**(1), 226–235.
- Bartling, B. & Özdemir, Y. (2023), ‘The limits to moral erosion in markets: Social norms and the replacement excuse’, *Games and Economic Behavior* **138**, 143–160.
- 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.
- Burlando, R. M. & Guala, F. (2005), ‘Heterogeneous agents in public goods experiments’, *Experimental Economics* **8**(1), 35–54.
- 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. (2023), ‘Changing preferences: An experiment and estimation of market-incentive effects on altruism’, *Journal of Health Economics* **92**, 102808.
- Dana, J., Weber, R. A. & Kuang, J. X. (2007), ‘Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness’, *Economic Theory* **33**(1), 67–80.

- Dufwenberg, M., Johansson-Stenman, O., Kirchler, M., Lindner, F. & Schwaiger, R. (2022), ‘Mean markets or kind commerce?’, *Journal of Public Economics* **209**, 104648.
- 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’.
- Enke, B., Graeber, T. & Oprea, R. (2023), ‘Confidence, self-selection, and bias in the aggregate’, *American Economic Review* **113**(7), 1933–1966.
- Enke, B. & Zimmermann, F. (2019), ‘Correlation neglect in belief formation’, *The Review of Economic Studies* **86**(1), 313–332.
- Exley, C. L. (2016), ‘Excusing selfishness in charitable giving: The role of risk’, *The Review of Economic Studies* **83**(2), 587–628.
- Falk, A., Becker, A., Dohmen, T., Enke, B., Huffman, D. & Sunde, U. (2018), ‘Global evidence on economic preferences’, *The Quarterly Journal of Economics* **133**(4), 1645–1692.
- 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.
- Falk, A. & Szech, N. (2014), ‘Diffusion of being pivotal and immoral outcomes’.
- Fischbacher, U., Fong, C. M. & Fehr, E. (2009), ‘Fairness, errors and the power of competition’, *Journal of Economic Behavior & Organization* **72**(1), 527–545.
- Fischbacher, U., Gächter, S. & Fehr, E. (2001), ‘Are people conditionally cooperative? evidence from a public goods experiment’, *Economics letters* **71**(3), 397–404.
- Gössling, S. & Humpe, A. (2020), ‘The global scale, distribution and growth of aviation: Implications for climate change’, *Global Environmental Change* **65**, 102194.
- Grossman, Z. (2014), ‘Strategic ignorance and the robustness of social preferences’, *Management Science* **60**(11), 2659–2665.
- Grossman, Z. & Van Der Weele, J. J. (2017), ‘Self-image and willful ignorance in social decisions’, *Journal of the European Economic Association* **15**(1), 173–217.

- Herweg, F. & Schmidt, K. M. (2022), ‘How to regulate carbon emissions with climate-conscious consumers’, *The Economic Journal* **132**(648), 2992–3019.
- Huck, S. & Konrad, K. A. (2005), ‘Moral cost, commitment, and committee size’, *Journal of Institutional and Theoretical Economics* pp. 575–588.
- 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.
- Kaufmann, M. & Koszegi, B. (2023), ‘Understanding markets with socially responsible consumers’.
- 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.
- List, J. A. (2003), ‘Does market experience eliminate market anomalies?’, *The Quarterly Journal of Economics* **118**(1), 41–71.
- Ockenfels, A., Werner, P. & Edenhofer, O. (2020), ‘Pricing externalities and moral behaviour’, *Nature Sustainability* **3**(10), 872–877.
- 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>
- Offerman, T., Romagnoli, G. & Ziegler, A. (2021), ‘Morals in multi-unit markets: Study II’, *AEA RCT Registry* .
URL: <https://doi.org/10.1257/rct.8306-1.0>
- Offerman, T., Romagnoli, G. & Ziegler, A. (2023), ‘Morals in multi-unit markets: Study III’, *AEA RCT Registry* .
URL: <https://doi.org/10.1257/rct.11080-1.0>

- Offerman, T., Sonnemans, J. & Schram, A. (1996), ‘Value orientations, expectations and voluntary contributions in public goods’, *The Economic Journal* **106**(437), 817–845.
- 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.
- Riehm, T., Fugger, N., Gillen, P., Gretschno, V. & Werner, P. (2022), ‘Social norms, sanctions, and conditional entry in markets with externalities: Evidence from an artefactual field experiment’, *Journal of Public Economics* **212**, 104701.
- Roth, A. E., Prasnikar, V., Okuno-Fujiwara, M. & Zamir, S. (1991), ‘Bargaining and market behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An experimental study’, *American Economic Review* **81**(5), 1068–1095.
- Rothenhäusler, D., Schweizer, N. & Szech, N. (2018), ‘Guilt in voting and public good games’, *European Economic Review* **101**, 664–681.
- 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.
- Serra-Garcia, M. & Szech, N. (2022), ‘The (in) elasticity of moral ignorance’, *Management Science* **68**(7), 4815–4834.
- 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.

A.1 Predicting moral costs and moral competitive equilibria

In addition to the analysis presented in the main text, we can use individual decision-making data to predict outcomes in markets. For this, we proceed in two steps. First, we explain how we fit a moral cost curve to individual decision-making data. Second, we can use predicted moral costs to simulate market outcomes under the assumption that moral costs are not affected by moving to markets, to predict a moral competitive equilibrium.

A.1.1 Moral cost curves

We begin by fitting a moral cost curve to individual decision-making data. Denote $\Theta_i(q)$ the total moral costs participant i incurs for canceling q units of donation. 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, where $\epsilon_{j,q}$ is an individual-unit error:

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

After estimating the above equation, we can use $\hat{\alpha}_i, \hat{\beta}_i$ to predict moral costs for any quantity $q \in \{1, 2, \dots, 15\}$ and each individual i . This predicts total moral costs $\hat{\Theta}_i(q)$, so the total moral costs for canceling q units of donation. However, we often are interested in per-unit marginal moral costs $\theta_i(q)$, for unit q . These are the moral costs for canceling an additional q -th unit of donation, after having cancelled $q - 1$ units earlier. So, we want to decompose predicted total moral costs $\hat{\Theta}_i(q)$ into a sum of q per-unit, marginal moral costs $\theta_i(j)$: $\hat{\Theta}_i(q) = \sum_{j=1}^q \hat{\theta}_i(j)$. To obtain per-unit moral costs $\hat{\theta}_i(q)$ for unit q , we use the predicted total moral costs for unit q , $\hat{\Theta}_i(q)$, and unit $q - 1$, $\hat{\Theta}_i(q - 1)$, and take their difference. By repeating this exercise for all $q \in \{1, 2, \dots, 15\}$, we obtain all per-unit moral costs for all units for all individuals.

A.1.2 Moral competitive equilibria

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 participants are participating in. Since the literature finds repetition to be

a force behind erosion (Bartling et al. 2023), we correct moral costs in each period by the average erosion we find in the corresponding repetition of the treatment MPL. We estimate moral costs each period, and rescale estimated moral costs for each market period with the average erosion found in MPL.

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.³⁴

To be precise, 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 estimated 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 the 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

³⁴In 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.

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 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 trader’s 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 decreasing marginal moral costs from estimated moral costs. We then perform the above procedure, which again yields a predicted quantity to be traded.

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. This exercise is meaningful, as we observe full trade in the first practice market period, which is incentivized but features no externalities, across all treatments featuring our market institution. This is consistent with the standard competitive equilibrium arising in the absence of negative externalities. Any decrease in trading volume can thus cleanly be interpreted as traders’ concern for preventing the negative externality.

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. 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. By comparing outcomes in the heterogeneous to the homogeneous moral competitive equilibrium, we measure of how severe market selection is in this benchmark, 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.

In Figure A1, we present the results for this exercise. 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 participants, and are not related to underlying market features.³⁵ As it turns out, our participants valued donations to UNICEF somewhat higher in FULL and MULTI than in SINGLE.

The second bar, in red, shows predicted quantities given the 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

³⁵Note 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 A1: 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.

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.

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 A1. 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 A1. 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.

Table A1: The size of erosion in markets

	SINGLE	MULTI	FULL	B-MULTI	B-FULL	HOM	HET
Normalized units	4.8	5.3	5.8	3.7	9.0	9.1	3.5
Damage per €1 gain per unit	3.2	3.2	4.9	2.9	4.6	4.5	6.0


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.

The results of the moral competitive equilibrium exercise hinge on assumptions on the moral cost curve we use to fit individual decision-making data. In the following, we provide results on two exercises to test the robustness of the above conclusion. First, we use a linear moral cost curve to fit data from individual decision-making. This assumes marginal moral costs to be constant. Figure A2 presents the results of this exercise. The results are in line with the findings when allowing for non-linear moral cost curves.

Second, we repeat the procedure assuming that moral costs are halved when moving from individual decision-making to markets. This can account for the fact that decisions in markets always involve two participants and these trades generate payoff for two participants. Figure A3 presents the results. We continue to observe erosion compared to this benchmark. In particular, 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 falls substantially short of the observed traded quantities of 14.9 units.³⁶


³⁶We pre-registered a second method of evaluating moral erosion, that was based on the information conveyed in traders' offers. The results of this analysis are qualitatively similar to the results reported in this section. Details will be sent on request.

Figure A2: Market outcomes and competitive equilibria: Linear moral costs



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, using linear cost curves to estimate moral costs. 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.

Figure A3: Market outcomes and competitive equilibria: Halved moral costs



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, which are divided by two. 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.

A.2 Balancing

In Table A2, we show covariate balance across treatments. Of particular interest is the comparison between HOM and HET. Apart from the intended manipulation of heterogeneity, these treatments are balanced. Note that the data for the treatments MPL, SINGLE, MULTI and FULL was collected first at the CREED laboratory in Amsterdam, in September and October 2019. The data for the treatments B-MULTI, B-FULL, HOM, HET and SPEC was collected October 2021 to January 2022. Sessions were ran both at the CREED laboratory in Amsterdam as well as at the Centerlab of Tilburg University. The data for the treatments B-LOW and B-HIGH was collected at the CREED laboratory in Amsterdam, in March and April 2023.

Table A2: Balancing table


	Age	% women	% international	Switching point part 1	Risk
MPL	21.6	42	86	9.9	3.5
SINGLE	20.6	46	80	9.5	3.6
MULTI	20.7	52	81	10.3	3.5
FULL	21.7	45	76	9.8	3.6
B-MULTI	21.1	48	74	10.4	3.6
B-FULL	21.3	51	78	11.1	3.4
HOM	21.5	56	75	10.9	3.6
HET	21.5	63	76	11.5	3.4
SPEC	21.6	59	80	-	-
B-HIGH	21.4	53	80	10.4	3.7
B-LOW	21.3	53	80	10.4	3.7
B-HIGH vs. B-LOW (p -values)	.852	1.00	1.00	.936	.819
HOM vs. HET (p -values)	.973	.424	.856	.520	.455
Kruskal-Wallis (p -values)	.236	.617	.989	.402	.938

Notes: Average characteristic by treatment. Switching point part 1 and Risk were not elicited for the SPEC treatment. In the third-last row we report p -values of a t -test comparing B-HIGH with B-LOW, 100 observations per treatment. In the second-last row we report p -values of a t -test comparing HOM with HET, 80 observations per treatment. In the last row we report p -values of a Kruskal-Wallis test, comparing equality across all treatments.

In Figure A4, we show histograms of the average marginal moral costs for HOM and HET, using part 1 data. We do observe that generating homogeneous groups in HOM and heterogeneous groups in HET was successful.

In Table A3, we report additional descriptive statistics comparing the distribution of average moral costs for HOM and HET. Across all measures, the distribution is more compressed in HOM than in HET, differentially allowing for stronger market selection in HET. Additionally, if calculating the minimum moral costs within each matching group, the median minimum is 1.10 in HOM and .27 in HET. This indicates

Figure A4: Histograms of average marginal moral costs



Notes: Histograms and kernel densities of average moral costs based on the elicited valuations, with a value of €1.5 each.

that the participant with the lowest moral costs is considerably more moral in HOM than in HET.

Table A3: Summary statistics for HOM and HET

	Mean	Median	SD	Min	P_{10}	P_{25}	P_{75}	P_{90}	Max
HOM	1.49	1.54	.296	.58	1.11	1.34	1.65	1.83	2.20
HET	1.57	1.43	1.045	.14	.29	.53	2.66	2.99	3


Notes: Average moral costs based on the elicited valuations, with a value of €1.5 each, by treatment. P_X refers to the X -th percentile.

A.3 Are marginal moral costs decreasing?

In Figure A5, we provide evidence of decreasing marginal moral costs. We plot the average valuation implied by choice data in individual decision-making, averaged on the unit level. At larger stakes, subjects need to be paid less, averaged per unit, such that they are willing to cancel a donation. The effect is quite strong: For the first unit, subjects on average reported moral costs of €1.68, this decreases to €1.27 for the fifteenth unit.

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 A4, show that this variable is empirically important. The estimate on Unit is negative and significant.

Figure A5: Decreasing marginal moral costs



Notes: Average moral costs based on the elicited valuations, with a value of €1.5 each.

Table A4: Evidence for decreasing marginal moral costs

Marginal moral costs	
Unit	-.0248*** (.00166)
Constant	1.575*** (.0101)
Observations	5,467
# of subjects	781
Subject FE	Yes
Adjusted R^2	0.814

Note: Dependent variable is average per-unit valuation elicited in individual decision-making, in Euros. Unit captures the unit number, from 1, 2, 3, 5, 7, 10, 15. Subject fixed effects control for level differences in valuations across subjects. Standard errors clustered on matching group level in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A.4 Reliability of the individual-decision making measure of morality

In the main text, we discuss the correlation between the repeated measurements of moral costs. In Table A5, we report the full matrix of correlation coefficients between each repeated measurement in treatment MPL.

Table A5: Correlation between individual-decision making-measurements in MPL

	Part 1	Part 2: Per. 1	Part 2: Per. 2	Part 2: Per. 3	Part 2: Per. 4	Part 3
Part 1	1					
Part 2: Period 1	.933	1				
Part 2: Period 2	.910	.971	1			
Part 2: Period 3	.958	.934	.941	1		
Part 2: Period 4	.932	.924	.944	.969	1	
Part 3	.938	.942	.956	.967	.969	1

Notes: Correlation coefficient of the moral costs, averaged on a subject level, across repeated measurements in MPL.

We repeat this exercise using Spearman rank-order correlation coefficients. Results are in line with the results reported in the main text: The correlation across part 1 and 3 across all treatments is .888, and the lowest correlation between the reported measurements in treatment MPL is .906.

A.5 Robustness of treatment effects

In Table A6, 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 A6: 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.6 Norms

We elicited 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 A6: Norms in individual decision-making


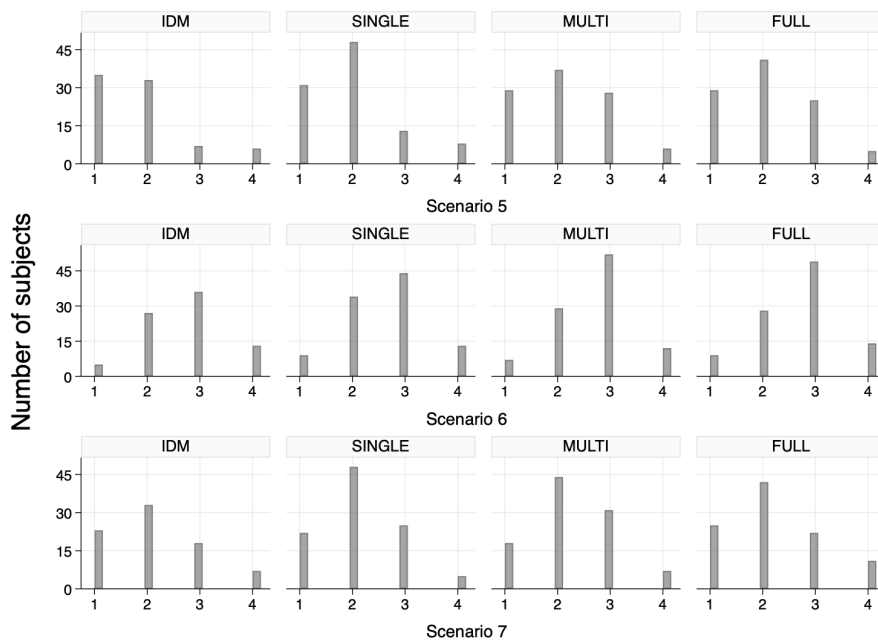


Figure A7: Norms in markets




A.7 The replacement logic: Intensive margins and deontological 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 Figure A8, 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.

Figure A8: 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.

The difference between the above and below median group is only significant for FULL (MWU, 10 observations per group, p -value=.003). Table A7 repeats the analysis using a regression. Model (1) repeats the analysis from the main text, regressing a dummy equal one if a subject was active for the last units in the markets on treatment dummies, a dummy equal one if a subject had above median moral costs and their interactions. We confirm that participants are more active in FULL. However, the interaction for above median participants in FULL is not significant. Model (2) uses the dependent variable from the appendix analysis, counting the number of acceptances and offers per participant. This analysis is robust to this specification, where above median participants are significantly less frequently trading.

Table A7: Replacement logic or market selection?

	(1) Active	(2) Nr. actions
MULTI	0.140 (0.090)	-0.160 (0.660)
FULL	0.740*** (0.069)	9.160*** (0.789)
Above median	-0.040 (0.049)	-0.300 (0.572)
MULTI \times Above median	0.040 (0.076)	0.680 (0.897)
FULL \times Above median	-0.140 (0.097)	-4.340*** (1.411)
Constant	0.180*** (0.061)	1.380** (0.530)
Observations	300	300


Note: Dependent variable is a dummy equal to one if a subject made or accepted an offer at least once for the units with gains from trade of €0.20 in (1), or how many offers or acceptances a subject made for these units in (2). Above median is a dummy equal one if a subject has above median moral costs. MULTI (FULL) is a dummy equal to one if the choice occurred in treatment MULTI (FULL), with the omitted category SINGLE. Standard errors, clustered on matching group level, are presented in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

In this analysis, 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.

Interestingly, this activity also carries over to subjects 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.³⁷

In Figure A9, 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.

Figure A9: 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.

In Table A8, 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 1 to 7 whether subjects agreed with a given statement. What appears to matter

³⁷Note 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.

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 A8: 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)
Norm in ind. dec.-making	0.219 (0.306)	0.029 (0.041)	0.024 (0.044)
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 Non-incentivized belief measures

In the main text, we analyze the incentivized belief measure. In this subsection, we provide some additional analysis and replicate the main analysis with our second belief measure, which is non-incentivized.

The non-incentivized belief measure consisted of three questions, each eliciting the three potential scenarios for the trade of the upcoming unit. The questions were:

1. What is the probability that whatever you do, the next unit will be traded?
2. What is the probability that whether or not the next unit is traded depends on your behavior?
3. What is the probability that whatever you do, the next unit will not be traded?

Participants received a payment of 300 cents irrespective of the correctness of their reports. Traders are pivotal only in the second scenario, thus the probability of being pivotal is measured by the likelihood ascribed to scenario 2. Sometimes the analysis requires the probability of being replaced. This is the probability of not being pivotal conditional on trade happening (i.e., the chance to have at least one other trader active on the own side of the market, conditional on the other side being active), and formally calculated by the probability of scenario 1 divided by the sum of the probabilities of scenarios 1 and 2.

For the incentivized question, we asked participants to report the following: “How many participants other than you will attempt to trade this unit?”. When correctly reporting the number of active traders, they received a payment of 150 cents.

In the main text, we focus on the incentivized measure as this correlates more strongly with the underlying true values. The Spearman correlation coefficients between the predicted and actual number of active traders averaged on a subject-level is 0.422. The same correlation between traders belief to be pivotal and the realized event to have actually been pivotal is 0.181. When bootstrapping the difference in test statistics this difference is significant with a p -value of .021 (160 observations, 1000 repetitions). The same pattern arises when calculating correlations treating each report as an independent observation. The correlation coefficient for the incentivized measure is 0.239, for the non-incentivized measure it is 0.120. The difference is significant with a p -value of .001 (1780 observations, 1000 repetitions).

First, Table A9 presents Spearman correlation coefficients of the two measurements of the beliefs. All data are based on the four market periods with externalities. The first row uses individual report-level data, the second row presents correlations between averages on a participant level. Both in B-MULTI and B-FULL there are

no detectable correlations. We do find the expected correlation for SPEC, which suggests that eliciting beliefs while simultaneously trading in markets inhibited this correlation. We think that we asked too much of our subjects in B-MULTI and B-FULL, which made them focus on the incentivized questions and pay less attention to the unincentivized ones.


Table A9: Correlation between belief measures

	SPEC	B-MULTI	B-FULL
All data	-.278 (.000)	-.016 (.718)	-.003 (.910)
Participant averages	-.447 (.003)	-.131 (.247)	.011 (.923)

Notes: Spearman correlation coefficients between the incentivized and non-incentivized belief measure, p -values in parentheses.

In Figure A10, we report the non-incentivized belief of the probability to be replaced across different treatments. These are calculated as the belief to have at least one other trader active on the own side of the market, conditional on the other side being active. Conclusions are in line with the analysis in the main text. The beliefs are significantly different between B-MULTI and B-FULL (MWU, 8 observations per treatment, p -value=.0209), while they do not differ between B-FULL and SPEC (MWU, 8 observations for B-FULL and 41 for SPEC, p -value=.704).

Figure A10: Unincentivized beliefs about own probability to be replaced



Notes: Probability to be replaced (grey bar), actual probability to be replaced (green diamond) and belief of spectators (orange circle).

Table A10 presents the results of an analysis in which we regress a dummy indicator of being active in the market on the unincentivized belief to be replaced. Surprisingly, participants who believe to be more replaceable are less likely to trade.

Possibly, some of our participants may have become confused about the questions that we were asking and may have thought that if they planned not to trade the subsequent unit, it will not be traded even if they were allowed to trade.

Table A10: Beliefs and activity		
	(1) B-MULTI	(2) B-FULL
Prob. to be replaced	-0.141* (0.066)	-0.687*** (0.092)
Average moral cost	-0.001 (0.000)	-0.001*** (0.000)
Period	-0.031 (0.030)	-0.022 (0.015)
Constant	0.589*** (0.073)	1.485*** (0.097)
Unit FE	yes	yes
Observations	466	1279
Adjusted R^2	0.025	0.124

Note: Dependent variable is a dummy equal to one if a subject submitted or accepted an offer at least once for units with gains from trade of €0.20. Average moral costs are the average moral costs for a participant, based on average estimated per-unit moral costs based on part 1 individual decision-making. Standard errors clustered on matching group level are presented in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A.9 Beliefs and activity in B-LOW and B-HIGH

In the main text in Table 4, we report on the pre-registered main hypotheses that compare behavior between B-LOW and B-HIGH. In this section, we report on the robustness of these results in Table A11. In column (1), we report the first stage of the 2SLS estimation reported in the main text. In columns (2), (3) and (4), we report robustness checks on the beliefs of participants. In (2), we use data from unit 4 in period 1, at which point participants in both B-HIGH and B-LOW receive very similar information: in both treatments, 9 or 10 participants were active for these early units. Reassuringly, beliefs do also not differ significantly between the two treatments. In (3), we use belief data of the first movers at unit 13 in period 1. Effects are similar to the results in the main text, where we pool both first- and second-movers. In (4), we pool data across all units and all periods in which beliefs were shocked for units for which trading is not efficient: units 13 and 15 in period 1 and 2. In columns (5), (6) and (7) we study the robustness of the change in own actions. In (5), we show that there is no difference in behavior at unit 4 in period 1

for the first movers, consistent with there not being a difference in induced beliefs. In (6), we study the effect of the first differential belief shock on both the first- and second-movers, at unit 13 in period 1. Results are as expected: less traders are active in B-LOW, similar to the results for the first-movers; however, the effect is already moderated, consistent with the second-movers learning that some first-movers continue to be active. In (7), we report on the activity of first- and second-movers for units 13 and 15 in periods 1 and 2. Traders continue to be less active in B-LOW, but that effect is considerably weakened and only just significant; consistent with the participants learning that still many traders are active, which allows participant to employ the replacement logic to justify trading themselves.

Table A11: Robustness: Beliefs and activity

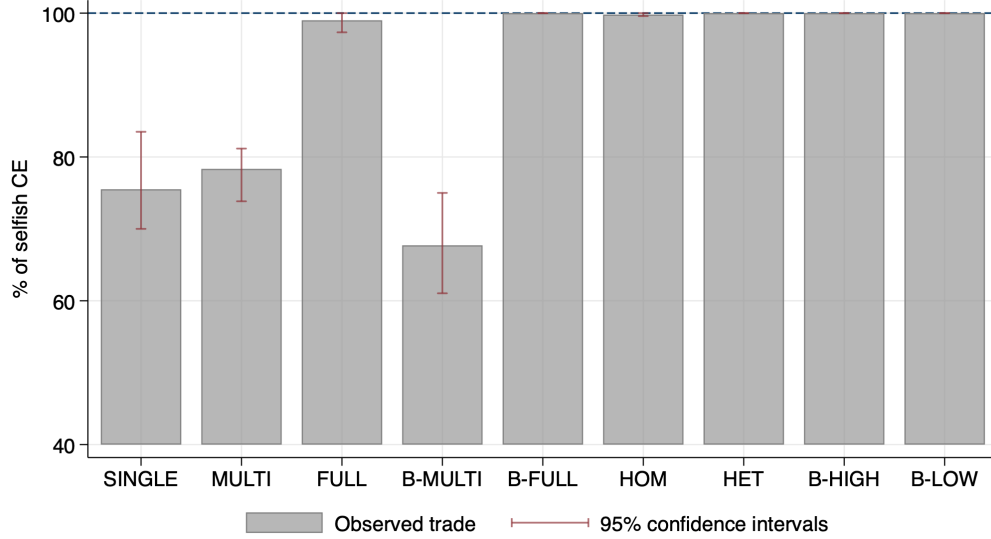
	(1) IV first stage	(2)	(3) # active others	(4)	(5)	(6) Own activity	(7)
B-LOW	-3.693*** (0.396)	0.260 (0.237)	-3.620*** (0.506)	-3.288*** (0.239)	0.140 (0.083)	-0.200*** (0.061)	-0.092* (0.050)
Average moral costs	-0.399 (0.295)						
Period				0.052 (0.173)			0.007 (0.025)
Constant	7.837*** (0.504)	8.210*** (0.212)	7.600*** (0.355)	6.878*** (0.331)	0.760*** (0.076)	0.640*** (0.046)	0.578*** (0.059)
Observations	200	200	100	800	100	200	800
Period	1	1	1	1, 2	1	1	1, 2
Units	13	4	13	13, 15	4	13	13, 15
Unit FE	No	No	No	Yes	No	No	Yes
Market-sides	both	both	first movers	both	first movers	both	both

Note: In (1) to (4), the dependent variable is the belief about the number of other traders that are active for units with gains from trade of €0.20. (1) is the first stage associated with the IV regression from the main text. In (5) to (7), the dependent variable is a dummy equal to one if a subject submitted or accepted an offer. Average moral costs are the average moral costs for a participant, based on averaging per-unit moral costs based on part 1 individual decision-making. Standard errors clustered on matching group level in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A.10 Main analysis in HOM, HET, B-MULTI and B-FULL

In Figure A11 and Table A12, we report market outcomes across all treatments. Quantities across all treatments using FULL market rules (FULL, B-FULL, HOM and HET) are all fully selfish and statistically indistinguishable. Quantities in B-MULTI are slightly below those in MULTI, suggesting that additionally eliciting beliefs in this treatment leads to slightly more moral behavior.

Figure A11: Market outcomes



Notes: Average quantities relative to selfish competitive equilibrium. Trading units below 40% is efficient (gains from trade exceeds the externality). Compared to the negative externality of €1.50 per unit, each unit between 40% and 60% yields gains from trade of €0.60, each unit between 60% and 80% yields €0.40 and each unit between 80% and 100% yields €0.20.

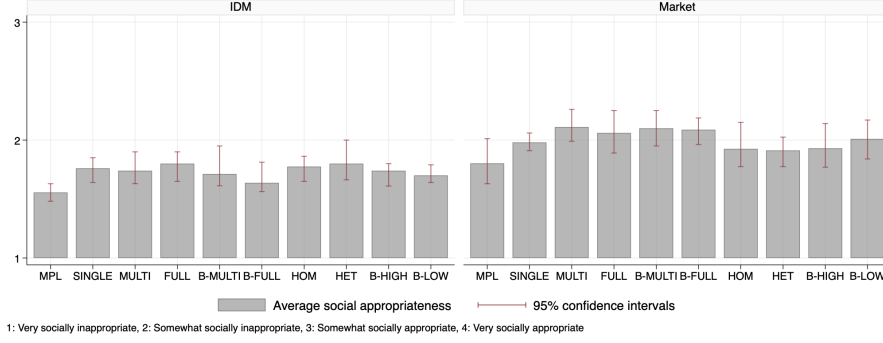
Table A12: Treatment effects

		SINGLE	MULTI	FULL	B-MULTI	B-FULL	HOM	HET	B-HIGH	B-LOW
Quantity in %		75.5	78.3	99	67.7	100	99.8	100	100	100
<i>p</i> -values	vs. SINGLE	-	.378	.0005	.0899	.0006	.0009	.0006	.0002	.0002
	vs. MULTI	-	-	.0001	.0308	.0002	.0002	.0002	.0000	.0000
	vs. FULL	-	-	-	.0002	.1931	.6318	.1931	.1468	.1468
	vs. B-MULTI	-	-	-	-	.0003	.0004	.0003	.0001	.0001
	vs. B-FULL	-	-	-	-	. -	.3173	1.000	1.000	1.000
	vs. HOM	-	-	-	-	. -	-	.3173	.2636	.2636
	vs. HET	-	-	-	-	. -	-	-	1.000	1.000
	vs. B-LOW	-	-	-	-	. -	-	-	-	1.000

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

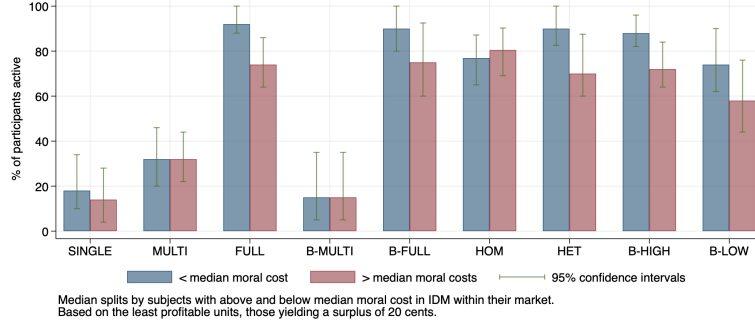
Figure A12 reports data on norms, complementing Figure 5. Figure A13 reports data on the share of participants active at the units with gains from trade of €0.20, complementing Figure 6.

Figure A12: Norms in individual decision-making and in markets



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

Figure A13: 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. Based on the least profitable units, those yielding a surplus of 20 cents.


A.11 Effects of market participation

Our experimental design also allows us to test whether morals are eroded within an identical decision environment, as participants faced identical individual decision-making tasks in parts 1 and 3. Treatment MPL allows 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. We hypothesized that any erosion of moral costs we find within markets is temporary.

In Figure A14, we plot the average predicted moral costs per treatment, by parts. In treatment MPL, 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.

Figure A14: Persistence of erosion



Notes: Average per-unit valuations in individual decision-making, for €1.50 donations, by part. In part 2, only MPL employs individual decision-making.

We do observe that moral costs are decreasing over time. In MPL, average per-unit moral costs in part 3 decrease by 6.5 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 9.5 cents. In the multi-unit markets MULTI and FULL, erosion is most drastic, with decreases of moral costs of 19.8 cents and 20.5 cents, respectively, after market exposure. This decrease is significant across all market treatments (Wilcoxon signed-rank tests, 100 observations per market treatment, 81 in MPL, p -values of MPL=.108, SINGLE=.002, MULTI=.000, FULL=.000). Comparing the decrease between treatments, we do not find significant differences between MPL and SINGLE (Mann-Whitney U-test, 81/100 observations, p -value=.289). 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=.008), while the decrease between MULTI and FULL is similar (MWU, 100 observations per treatment, p -value=.799). 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.

We further investigate how trading experience in our experimental markets affects traders' perceptions about the morality of other traders. For this, we elicited subjects' beliefs about the median moral costs of canceling a donation in individual decision-making. These beliefs were elicited at the end of the experiment, after participants experienced the markets or repeated individual decision-making. Participants were

paid a bonus of €1 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 A15, we report for each main treatment the mean difference between predicted and actual moral cost of the median participant in the left panel. In the right panel, we report the absolute prediction error, that is the absolute difference between the predicted and actual moral cost of the median participant.

Observing their fellow peers does not help participants to improve their estimate, measured by the absolute prediction error: the absolute error is not decreasing in the markets compared to MPL. Also, there do not appear to be strong differences between the market treatments.

Interestingly, the direction of the error changes systematically between treatments.³⁸ If anything, participants in MPL 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.

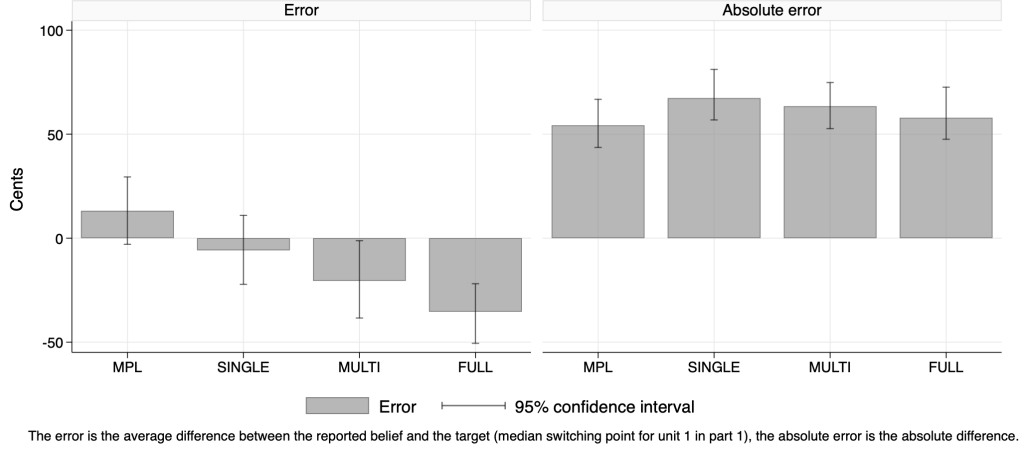
Summarizing, there is biased social learning in the sense that participants believe that their peers are more selfish than they truly are. Participants do not sufficiently take into account that other traders’ behavior in the market is to a large extent shaped by market forces. This is also consistent 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 may not reflect the preferences of an average participant outside of the market.

One potential concern may be that participants 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 participants are most 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 the most: those traders who we observe to be active at the least profitable units, those yielding profits of €0.20. In Table A13, 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 to one if a trader was active at the least profitable units, with

³⁸Regressing participants’ absolute errors on treatment fixed effect shows insignificant dummies (781 observations, clustering standard errors on matching group level; all p -values $> .1$ for SINGLE, MULTI and FULL). Regressing the error on treatment fixed effects shows differences in fixed effects, compared to the MPL baseline, for SINGLE (estimate of -19.0, p -value=.117), MULTI (estimate of -33.7, p -value=.009) and FULL (estimate of -48.5, p -value $< .001$).

Figure A15: Errors in beliefs about median participant's moral cost



Notes: Average error in estimating the session's median participant's moral cost for canceling one unit of donation of €1.50 in part 1 of the experiment. The left panel displays the average difference between prediction and target, the right panel the absolute distance between prediction and target.

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 A13: 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 participant was active at the final units, those with gains from trade of €0.20.

A.12 Norms as excuses

As for the beliefs, it may be the case that participants report perceived norms to excuse their behavior in part 2 of the experiment. We report the same analysis for the elicited norms in Table A14.

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 to one if a trader was active at the least profitable units, with treatment fixed effects,

Table A14: 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 canceling donations of €1.50 in return for a payment of €1 in individual decision-making (MPL) or with in an experimental market (Market). 2 corresponds to “somewhat socially inappropriate”. Split by whether the participant was active at the final units, those with gains from trade of €0.20.

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).

A.13 Morals in a double auction with a private schedule

Our experimental markets use a two-sided posted offer institution and a common schedule. In this section, we compare outcomes of our treatment MULTI to an additional control treatment PRIV, in which we implement a standard double auction with a private schedule.

In the double auction, buyers and sellers submitted offers simultaneously and units were traded if participants agreed on a price (if a bid exceeded an ask). Each market period ran for eight minutes. In this market, traders face a private cost/value schedule. For comparability, we mapped the schedule we use in our multi-unit markets to a private schedule. For this, every trader received three randomly drawn values or costs from the common cost or value schedule, redrawn every market period. This way aggregate costs and values were kept identical, and each participant was restricted to trade at most three units. In agreement with standard double auction procedures, participants did not know other traders’ costs or values. This is a further difference with our two-sided posted offer market.

All other elements of our experiment were kept identical. In a first market period, participants could trade in a market without externalities. In this market, quantities were slightly below the competitive equilibrium quantity, at 13.89 units on average. Before the start of the last four periods, participants learned that per trade donations of €1.50 to UNICEF were cancelled.


The sessions for treatment PRIV were ran in Amsterdam in December 2021 and January 2022, with 80 participants in total. We did so as data from PRIV was directly compared to the original treatments, which were also only ran in Amsterdam.

In Figure A16, we plot the traded amounts in PRIV, compared to MULTI and

FULL. Quantities in PRIV are the first two bars. The first bar normalizes quantities relative to the selfish competitive equilibrium, the second bar relative to the quantity traded in the market period without externalities. For MULTI and FULL, presented in the in the third and fourth bar, these normalizations are identical.

Quantities in PRIV are below quantities in MULTI and FULL. Comparing the normalization relative to selfish competitive equilibrium, quantities in PRIV are significantly different from MULTI (MWU, 8 observations in PRIV and 10 in MULTI, $p\text{-value}=.003$) and from FULL (MWU, 8 observations in PRIV and 10 in MULTI, $p\text{-value}<.001$). However, they are still consistent with a partial erosion of morals, as trade continued beyond 40%, which implies that units where the damage to UNICEF exceeds the associated gains from trade have been traded.


Figure A16: Market outcomes in PRIV



Notes: Average quantities traded relative to selfish competitive equilibrium (first, third and fourth bar) and relative to quantities traded in markets without externalities (second bar).

An important question is to which extent morals are eroded in the standard double auction, and how this compares to the two-sided posted offer market. To investigate this question, we repeat the analysis of how many traders trade a unit for a payment of less than €1.50. In Figure A17, we present the results, including the treatment MULTI for comparison. In this comparison we observe erosion in PRIV, at least as large as the erosion in SINGLE, and slightly below the erosion detected in MULTI.

Figure A17: Cancellation of donations between environments and treatments



Notes: Share of participants who cancelled a donation for at most its value (€1.50) in individual decision-making and in implemented trades in the market. The left panel shows shows cancellation rates in part 1 of the experiment and the middle panel plots cancellation rates in the first period of part 2. The right panel displays the share of participants who, in the four periods of part 2, at least once cancelled a donation.

In Table A15, we repeat the analysis from the main text comparing the three market treatments SINGLE, MULTI and PRIV. When we pool the data, we see that at the start there is similar erosion in PRIV as in MULTI, while there is more erosion in PRIV than in SINGLE. Then over time the difference with SINGLE disappears.

Table A15: Erosion in markets and through repetition

	(1) SINGLE	(2) MULTI	(3) PRIV	(4) SINGLE, MULTI & PRIV	(5)	(6)
				Period 1	Period 1-4	Pooled data
Period 1-4	0.270*** (0.052)	0.090*** (0.023)	0.100* (0.046)			0.100** (0.044)
SINGLE				-0.155** (0.071)	0.015 (0.033)	-0.155** (0.071)
MULTI				0.085 (0.053)	0.075** (0.029)	0.085 (0.053)
SINGLE \times Period 1-4						0.170** (0.067)
MULTI \times Period 1-4						-0.010 (0.050)
Constant	0.620*** (0.051)	0.860*** (0.016)	0.775*** (0.053)	0.775*** (0.050)	0.875*** (0.024)	0.775*** (0.050)
Observations	200	200	160	280	280	560

Note: Dependent variable is a dummy equal to one if a participant cancelled a donation for a payment of at most its value (€1.50). Period 1-4 is a dummy variable equal to one if the choice is measured as occurring at least once in period 1 to 4 in part 2 of the experiment, the omitted category is cancellation in period 1. SINGLE (MULTI) is a dummy equal to one if the choice occurred in treatment SINGLE (MULTI), with the omitted category PRIV. Standard errors, clustered on matching group level, are presented in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The possibility of replacement thinking is not limited to markets with a common schedule. Here, we discuss two ways in which replacement excuses can be introduced in markets with private schedules. The first possibility arises when aggregate supply and demand are horizontally extended beyond the competitive equilibrium, as is illustrated in Figure A18. This schedule results when, compared to the PRIV schedule, we allow traders on both sides of the market to trade additional units at a price close to the competitive equilibrium price. If the buyers' values for these additional units are slightly below the competitive equilibrium price, while the sellers' costs are slightly above, then the competitive equilibrium is not affected. This schedule allows traders to take full advantage of trading opportunities foregone by others. Assuming that other traders are selfish, traders would anticipate that their trading decisions will not matter for the aggregate outcome, and replacement thinking will excuse their trading. Note that this type of schedule does not feature elements of markets we want to capture, for example traders' costs and values do no longer depend on others' trading.

The second possibility is to add traders on both sides, with similar cost and demand schedules. Combined with a restriction for aggregate trade not to exceed the original competitive quantity, traders will again recognize that the replacement




Figure A18: Cost and value schedule with replacement for PRIV

excuse applies, and feel free to trade as much as they can.

B Experimental interface

Below is an example screenshot from the experimental markets.

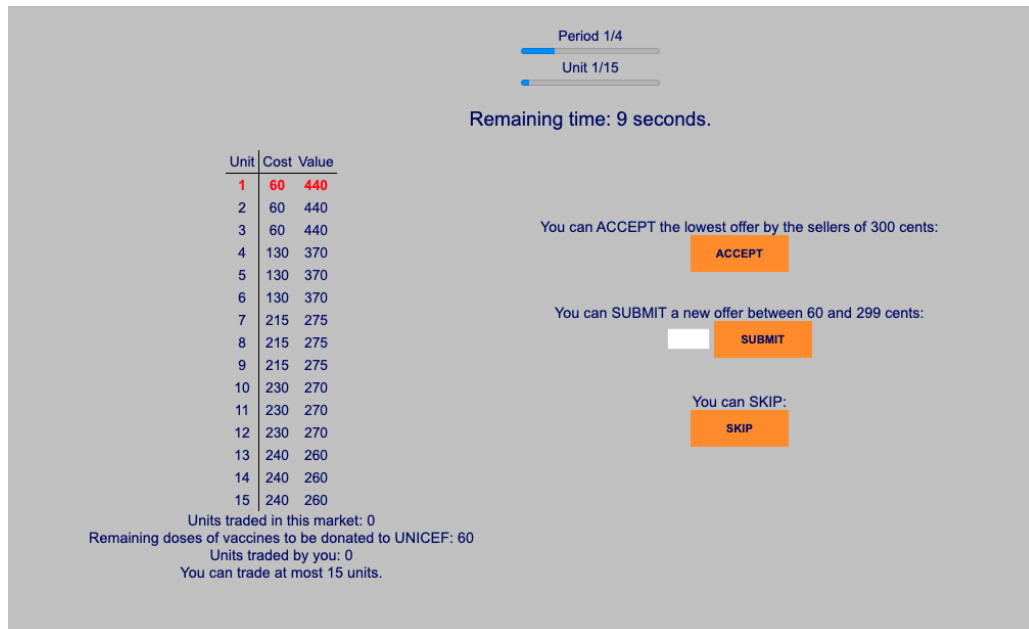


Figure A19: Experimental market interface

After each unit, traders received the following feedback:

You accepted an offer. You earn 140 cents.

The price agreed upon was 300. The buyer earns $440-300=140$ cents, the seller $300-60=240$ cents.

Four doses of measles vaccines will not be donated to UNICEF because of this trade. If no more units are traded in this market, 56 doses of measles vaccines will be donated to UNICEF.

In this market period, you have traded 1 unit and earned 140 cents so far.

Figure A20: 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 A21: 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 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. {NOT in HOM/HET 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).

{HOM/HET Your decisions in part 1 will affect with whom you will interact in a later part of this experiment. Like the other participants, you will either be assigned to a group of participants that made quite similar choices as the average participant did in part 1, or to a group in which participants behaved quite differently from the average participant.}

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.

Choice list 3/7

Your Decision

You see 21 choices on this choice list, for each choice you will have to decide between Option A and Option B. Amounts for option A are in cents.
If this choice list will be randomly determined for payment, the computer will determine randomly which of the twenty-one choices will be used.
Now, your decisions are between payments to you and **12 doses** of measles vaccine.

• Option A: 0 cents for you	or	• Option B: 12 doses of measles vaccine for UNICEF
• Option A: 45 cents for you	or	• Option B: 12 doses of measles vaccine for UNICEF
• Option A: 90 cents for you	or	• Option B: 12 doses of measles vaccine for UNICEF
• Option A: 135 cents for you	or	• Option B: 12 doses of measles vaccine for UNICEF
• Option A: 180 cents for you	or	• Option B: 12 doses of measles vaccine for UNICEF
• Option A: 225 cents for you	or	• Option B: 12 doses of measles vaccine for UNICEF
• Option A: 270 cents for you	or	• Option B: 12 doses of measles vaccine for UNICEF
• Option A: 315 cents for you	or	• Option B: 12 doses of measles vaccine for UNICEF
• Option A: 360 cents for you	or	• Option B: 12 doses of measles vaccine for UNICEF
• Option A: 405 cents for you	or	• Option B: 12 doses of measles vaccine for UNICEF
• Option A: 450 cents for you	or	• Option B: 12 doses of measles vaccine for UNICEF
• Option A: 495 cents for you	or	• Option B: 12 doses of measles vaccine for UNICEF
• Option A: 540 cents for you	or	• Option B: 12 doses of measles vaccine for UNICEF
• Option A: 585 cents for you	or	• Option B: 12 doses of measles vaccine for UNICEF
• Option A: 630 cents for you	or	• Option B: 12 doses of measles vaccine for UNICEF
• Option A: 675 cents for you	or	• Option B: 12 doses of measles vaccine for UNICEF
• Option A: 720 cents for you	or	• Option B: 12 doses of measles vaccine for UNICEF
• Option A: 765 cents for you	or	• Option B: 12 doses of measles vaccine for UNICEF
• Option A: 810 cents for you	or	• Option B: 12 doses of measles vaccine for UNICEF
• Option A: 855 cents for you	or	• Option B: 12 doses of measles vaccine for UNICEF
• Option A: 900 cents for you	or	• Option B: 12 doses of measles vaccine for UNICEF

OK

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.

Choice list 5/7

Your Decision

You see 21 choices on this choice list, for each choice you will have to decide between Option A and Option B. Amounts for option A are in cents.

If this choice list will be randomly determined for payment, the computer will determine randomly which of the twenty-one choices will be used.

Now, your decisions are between payments to you and **28 doses** of measles vaccine.

<input type="radio"/> Option A: 0 cents for you	or	<input type="radio"/> Option B: 28 doses of measles vaccine for UNICEF
<input type="radio"/> Option A: 105 cents for you	or	<input type="radio"/> Option B: 28 doses of measles vaccine for UNICEF
<input type="radio"/> Option A: 210 cents for you	or	<input type="radio"/> Option B: 28 doses of measles vaccine for UNICEF
<input type="radio"/> Option A: 315 cents for you	or	<input type="radio"/> Option B: 28 doses of measles vaccine for UNICEF
<input type="radio"/> Option A: 420 cents for you	or	<input type="radio"/> Option B: 28 doses of measles vaccine for UNICEF
<input type="radio"/> Option A: 525 cents for you	or	<input type="radio"/> Option B: 28 doses of measles vaccine for UNICEF
<input type="radio"/> Option A: 630 cents for you	or	<input type="radio"/> Option B: 28 doses of measles vaccine for UNICEF
<input type="radio"/> Option A: 735 cents for you	or	<input type="radio"/> Option B: 28 doses of measles vaccine for UNICEF
<input type="radio"/> Option A: 840 cents for you	or	<input type="radio"/> Option B: 28 doses of measles vaccine for UNICEF
<input type="radio"/> Option A: 945 cents for you	or	<input type="radio"/> Option B: 28 doses of measles vaccine for UNICEF
<input type="radio"/> Option A: 1050 cents for you	or	<input type="radio"/> Option B: 28 doses of measles vaccine for UNICEF
<input type="radio"/> Option A: 1155 cents for you	or	<input type="radio"/> Option B: 28 doses of measles vaccine for UNICEF
<input type="radio"/> Option A: 1260 cents for you	or	<input type="radio"/> Option B: 28 doses of measles vaccine for UNICEF
<input type="radio"/> Option A: 1365 cents for you	or	<input type="radio"/> Option B: 28 doses of measles vaccine for UNICEF
<input type="radio"/> Option A: 1470 cents for you	or	<input type="radio"/> Option B: 28 doses of measles vaccine for UNICEF
<input type="radio"/> Option A: 1575 cents for you	or	<input type="radio"/> Option B: 28 doses of measles vaccine for UNICEF
<input type="radio"/> Option A: 1680 cents for you	or	<input type="radio"/> Option B: 28 doses of measles vaccine for UNICEF
<input type="radio"/> Option A: 1785 cents for you	or	<input type="radio"/> Option B: 28 doses of measles vaccine for UNICEF
<input type="radio"/> Option A: 1890 cents for you	or	<input type="radio"/> Option B: 28 doses of measles vaccine for UNICEF
<input type="radio"/> Option A: 1995 cents for you	or	<input type="radio"/> Option B: 28 doses of measles vaccine for UNICEF
<input type="radio"/> Option A: 2100 cents for you	or	<input type="radio"/> Option B: 28 doses of measles vaccine for UNICEF

OK

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. {NOT in HOM/HET 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 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

{NOT in HOM/HET: 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.} {HOM/HET: In this part of the experiment you will repeatedly trade in a market. In the market, 5 sellers can trade with 5 buyers. As explained in part 1, you will either be assigned to a group of participants that made quite similar choices as the average participant did in part 1, or to a group in which participants behaved quite differently from the average participant. You will participate in a group in which participants' choices are similar/different from the average participant.

You will be a BUYER 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

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

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: $\text{PRICE} - \text{COST} = 180 - 60 = 120$ cents.

- The buyers get: $\text{VALUE} - \text{PRICE} = 440 - 180 = 260$ 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 Predictions {ONLY IN B-FULL/B-MULTI}

{ONLY IN B-FULL/B-MULTI: While you are trading, you will occasionally be asked to predict how future trading will proceed. At these moments, you will be asked what you think will happen in the trading of the next unit.

Each time, we will ask you to predict four things.

The first three predictions concern probabilities of whether trading will occur for the next unit. For each next unit, there are three possible events:

1. Whatever you do, the unit will be traded. This means that even if you do not participate in trading, the unit will be traded by the others.
2. Your behavior will determine whether the unit is traded or not. This means that if you do not participate, the unit will not be traded, while if you do participate, the unit will be traded.
3. Whatever you do, the next unit will not be traded. This means that even if you do try to trade the next unit, this will not happen because the buyers are not participating.

We will ask you for the probabilities that each of these events occurs. These probabilities are your predictions of how likely it is that each possible event will happen. A higher probability means that an event is more likely to happen. As a probability, your predictions can be between 0% (will not occur) and 100% (will certainly occur). As the three events above include all possible scenarios in which this experiment progresses, the probabilities you report across 1. to 3. need to add up to 100%.

The last prediction concerns the number of sellers and buyers who will be active for the next unit. We will ask you to predict how many buyers and sellers other than you within your market will attempt to trade the next unit. By this we mean the total number of participants other than you who will either submit an offer to and/or accept an offer. Your predictions can be between 0 participants (no other participant will be active) and 9 participants (all other participants will be active).

At the end of the experiment, if this part is selected for payment, you will be paid for a set of predictions for one unit in one period of the markets. This will be another period than the period for which your trading determines your earnings. For the first three predictions, you will receive 300 cents. For the fourth prediction, you will receive an additional payment of 150 cents if you correctly predict how many participants other than you attempt to trade the next unit.}

D.5 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. {ONLY IN B-FULL/B-MULTI How much will you earn if you correctly predict how many other participants will attempt to trade the next unit? [FREE FORM: 150] cents}
4. 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.6 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]

- 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.7 Page 2

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

D.8 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 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 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 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

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	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
1 chooses to receive 2 Euro instead of making a donation of 4 doses of measles vaccine to UNICEF.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
1 chooses to receive 3 Euro instead of making a donation of 12 doses of measles vaccine to UNICEF.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
1 chooses to receive 6 Euro instead of making a donation of 12 doses of measles vaccine to UNICEF.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

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.

G Additional instructions and screenshots for B-LOW and B-HIGH

On instructions page 2 of part 2, participants received these additional instructions. Below, we also reproduce screenshots of how the information was revealed to participants and how they reported their beliefs.

Information from earlier session While trading proceeds, we will also provide information about trading of the next unit from an earlier session of this experiment. All participants within your market will see this information.

We will show how many participants in that session attempted to trade the next unit, by submitting and/or accepting an offer. In addition, we will show whether that unit was traded.

The information is based on an earlier session of this experiment. In that session, the market rules were either identical to your rules, or slightly different in that each participant could only trade a maximum number of units that was different to yours. However, all other market rules will be identical: In total, 15 units could be traded, these markets consisted of 10 participants, featured the same cost and values, and the trading of additional units led to a reduction of a donation to UNICEF, as in this experiment.

Information from the earlier session

In the earlier session of this experiment, **10 participants** attempted to trade the next unit.

The unit **was traded**.

Remember that all participants within your market receive this information.

CONTINUE

Your predictions

Please make a prediction on the trade of the next unit, unit 15.

We want you to predict how many of the other buyers and sellers in your market will attempt to trade the next unit (so, will either submit or accept an offer). Remember that there are 9 other traders besides you, so your prediction can be between 0 and 9 other participants.

If you correctly predict this question, you will earn 150 cents if this prediction is randomly selected for payment.

How many of the other participants in your market will attempt to trade this unit?

participants

OK

For your reference, these are the cost and values:

Unit	Cost	Value
1	60	440
2	60	440
3	60	440
4	130	370
5	130	370
6	130	370
7	215	275
8	215	275
9	215	275
10	230	270
11	230	270
12	230	270
13	240	260
14	240	260
15	240	260

H Instructions for SPEC

Participants read instructions as their counterpart in B-FULL did. They did not see tasks related individual decision-making or any of the tasks at the end of the experiment (beliefs, risk aversion, norms).

The instructions started with:

H.1 Page 1: Your task

In this experiment, you will first read instructions similar to the instructions that participants received in earlier sessions of this experiment. However, you will *not* participate in that experiment, and will not be able to make the choices described in these instructions.

It is important that you read and understand these instructions well, as your earnings depend on decisions in these earlier sessions.

After you have read the instructions, you will make predictions about choices these earlier participants made. Your earnings in this experiment will depend on the accuracy of your predictions about choices these earlier participants made.

H.2 Instructions and quiz of their matched participant (see B-FULL)

H.3 Page 2: Your prediction task

We now describe the choices that you will make in this experiment. You will follow one particular participant in an earlier session of this experiment. We will call this person your matched participant. We will show you the market interaction that your matched participant experienced in the experiment. Within a dark grey box, you will see a screen identical to the screen your matched participant observed. There is also some basic explanation on how to interpret this screen in the markets.

We will ask you to make predictions about the choices of *other* participants in the session of your matched participant.

Each time, we will ask you to predict four things.

The first three predictions concern probabilities of whether trading will occur for the next unit. For each next unit, there are three possible events:

1. Whatever your matched participant does, the unit will be traded. This means that even if your matched participant does not participate in trading, the unit will be traded by the others.

2. Your matched participant's behavior will determine whether the unit is traded or not. This means that if your matched participant does not participate, the unit will not be traded, while if your matched participant does participate, the unit will be traded.
3. Whatever your matched participant does, the next unit will not be traded. This means that even if your matched participant does try to trade the next unit, this will not happen because the buyers/sellers are not participating.

We will ask you for the probabilities that each of these events occurs. These probabilities are your predictions of how likely it is that each possible event will happen. A higher probability means that an event is more likely to happen. As a probability, your predictions can be between 0% (will not occur) and 100% (will certainly occur). As the three events above include all possible scenarios in which this experiment progresses, the probabilities you report across 1. to 3. need to add up to 100%.

The last prediction concerns the number of sellers and buyers who will be active for the next unit. We will ask you to predict how many buyers and sellers other than your matched participant within your matched participant's market will attempt to trade the next unit. By this we mean the total number of participants other than your matched participant who will either submit an offer to and/or accept an offer. Your predictions can be between 0 participants (no other participant will be active) and 9 participants (all other participants will be active).

At the end of the experiment, you will be paid for three randomly selected sets of predictions. Within each set of predictions, for the first three predictions, you will receive 300 cents. For the fourth prediction, you will receive an additional payment of 150 cents if you correctly predict how many participants other than you attempt to trade the next unit.

H.4 Page 3: Practice questions

Please answer the following questions:

1. In this experiment, you will only make predictions about the choices of other participants in an earlier experiment. [TRUE/false]
2. You will observe the market interaction your matched participant observed in this earlier experiment. [TRUE/false]
3. You will also trade in a market similar to the market your matched participant participated in. [true/FALSE]

4. How much will you earn if you correctly predict how many participants other than your matched participant will attempt to trade the next unit? [FREEFORM: 150] cents

I Double auction

All participants received instructions on the individual decision-making task as in the other market treatments. Below we reproduce all instructions for part 2.

I.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.

I.1.1 Market 1

Trading profits

In market 1, up to 15 units can be traded. To trade one unit, one buyer and one seller can conclude a trade by agreeing on a price for that unit.

If sold, each unit has a certain cost. This will be denoted in cents. Earnings for concluding a trade are:

- You earn the difference between the price and the cost for this unit: $\text{PRICE} - \text{COST}$

You will only know your costs in the market. You will not know other sellers' costs. In general, all participants have different costs. It is randomly determined which participants has what costs.

In the example in the screenshot below, the seller can trade their first unit, which is highlighted by the red first line in the table.

Remaining time: **6m 54s**

Orderbook
Buyers Sellers

Your current offer:
New offer: **Submit**

	Unit	Cost	Price	Your profit
Prices of the last four concluded trades	1	60		
	2	215		
	3	240		

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

Example in the screenshot: Imagine you have a cost given by 60. You, as the seller, and one of the buyers agree on a price of 260. Then,

- You get: $\text{PRICE} - \text{COST} = 260 - 60 = 200$ cents.

In the screenshot, notice that the cost of the seller change with the unit transacted (e.g. for the first unit the cost for the seller is 60 cents, for the third unit the cost is 240 cents).

I.2 Page 2: Market instructions

Trading protocol

To agree on a price, all participants submit offers simultaneously. At any time during the trading, you can submit an offer to the buyers.

You can see how you can submit an offer in the screenshot below:

Remaining time: **6m 54s**

Orderbook
Buyers Sellers

Your current offer:
New offer: **Submit**

	Unit	Cost	Price	Your profit
Prices of the last four concluded trades	1	60		
	2	215		
	3	240		

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

If you submit an offer, one of two things will happen:

1. A trade happens immediately:

- If your new offer is at a price *below or equal to* the best current offer of the buyers, you will trade.
- You will trade with the buyer who made this offer. The price will be the price offered by this buyer.
- Your offer is removed and new ones can be made.

2. Your offer enters the order book:

- If your new offer is at a price above the best current offer of the buyers, you will not trade yet.
- Instead, your offer will enter the order book. In the order book, all current offers of all buyers and sellers are collected and shown to all buyers and sellers.
- If your offer is the lowest current offer among all sellers, and a buyer makes an offer *above* your price, you will trade with this buyer at a price equal to your offer.
- You can always decide to adjust your offer.

After you have traded a unit, you can submit new offers to trade additional units. Note that your cost for trading the next unit may be different.

End of trading

Trading ends automatically after 8 minutes.

I.3 Page 3: Market instructions

Additional details

- On the top of the trading screen you always see the remaining trading time.
- We will also inform you about how many units you have traded, as well as the price and profits for these units.
- You will see all offers currently in the order book. The most favorable offers are ranked highest in the order book, which are the highest offered price by the buyers and the lowest offered price by the sellers.
- You also see the prices of the last four concluded trades.

- Note that offers are executed at the time they arrive.
- No trader knows with whom in the room he or she has traded. That means that your anonymity is ensured.

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 earlier. That is, for each unit, you will be paid the difference between the price and the cost for this unit. In this market, the earnings for each participant are then given by the sum of earnings for all units traded by the participant.

In part 2, there will be a total of 5 market periods (1 period for market 1 and 4 periods for market 2). 2 out of the 5 market periods will be randomly selected to be paid.

I.4 Page 4: Practice questions

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 each unit.
true/FALSE
2. Trading automatically ends after 8 minutes. 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.
 - The market begins, and the two sellers as well as one buyer submit an offer.
 - First, the two sellers submit an offer. Seller S1 has a cost of 60 cents and submits an offer with a price of 200 cents. Seller S2 has a cost of 130 cents and submits an offer with a price of 190 cents.
 - All market participants see the two offers in the order book. As seller S2'S offer is more favorable, it will be shown first in the order book.
 - Next, one buyer submits an offer: Buyer B1 submits an offer with a price of 210 cents.

- As B1's offer is higher than S2's offer, B1 immediately trades with S2. They will trade at the price offered by S2.
- Both B1's and S2's offers are removed from the order book and trading can continue.

Please calculate the earnings of S1 and S2 at this point in the market:

How many cents does seller S1 earn? [FREEFORM: 0] cents

How many cents does seller S2 earn? [FREEFORM: 60] cents

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

I.5 Page 1: Part 2

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 60. 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 60 doses are donated to UNICEF for this market. If at the end of the market 3 units have been traded then in total 48 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: [TABLE WITH NUMBER OF UNITS AND CORRESPONDING DONATION AMOUNTS, from 15 units traded (0 donations) to 0 units traded (60 units donation).]

- 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.

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.

I.6 Page 2

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

I.7 Page 3: Practice questions

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.
2. 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? [FREEFORM: 4] doses
3. We will ask you a question about the scenario below. Note that the behavior in this scenario is randomly determined, only for the purpose of asking the question.
 - The market begins, and one seller as well as one buyer submit an offer.
 - First, one seller submits an offer. Seller S1 has a cost of 130 cents and submits an offer with a price of 300 cents.
 - All market participants see the offer in the order book.

- Next, one buyer submits an offer: Buyer B1 submits an offer of 350 cents.
- As B1's offer is higher than S1's offer, B1 immediately trades with S1. They will trade at the price offered by S1.
- As one unit was traded, four doses of measles vaccine are *not* donated to UNICEF.
- Both B1's and S1's offers are removed from the order book. Trading can continue.

Please calculate the earnings of S1 at this point in the market:

How many cents does seller S1 earn? [FREEFORM: 170] cents