Morals in Multi-Unit Markets

Ziegler, A.G.B.; Romagnoli, G.; Offerman, T.

DOI
10.1093/jeea/jvae001

Publication date
2024

Document Version
Final published version

Published in
Journal of the European Economic Association

MORALS IN MULTI-UNIT MARKETS

Andreas G. B. Ziegler
University of Essex, UK

Giorgia Romagnoli
CREED, University of Amsterdam, The Netherlands

Theo Offerman
CREED, University of Amsterdam, The Netherlands

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 and (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. (JEL: C91, C92, D62)

Teaching Slides
A set of Teaching Slides to accompany this article is available online as Supplementary Data.

The editor in charge of this paper was Imran Rasul.

Acknowledgments: We thank the editor Imran Rasul, the anonymous referees, and Björn Bartling, Sourav Bhattacharya, Katharina Brütt, Guillaume Fréchette, Jana Friedrichsen, Hartmut Kliemt, Boris van Leeuwen, Klaus M. Schmidt, Florian H. Schneider, Nora Szech, 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, and 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, RCT ID AEARCTR-0008306, and RCT ID AEARCTR-0011080.

E-mail: mail@andreasziegler.eu (Ziegler); g.romagnoli@uva.nl (Romagnoli); t.j.s.offerman@uva.nl (Offerman)

Journal of the European Economic Association 2024 22(5):2225–2260
https://doi.org/10.1093/jeea/jvae001

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, and 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 and Humpe (2020) find that in the USA, 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 and 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 and Szech (2013) are contested. Market prices can also decline without moral erosion (Sutter et al. 2020). Moreover, while Falk and Szech (2013) compare a single decision in individual decision-making with repeated decisions in a market, Bartling, Fehr, and Özdemir (2023) show that the partial erosion in markets disappears 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

1. In the opioids market, a spokeswoman for McKesson, which was the largest distributor in the USA 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 et al. 2022). In 2017, the number of Americans dying from an overdose of opioids (47,600) surpassed the number dying from car accidents (Scholl et al. 2019). Empirically, Vuillemey (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 (Falk, Neuber, and Szech 2020; Bartling and Özdemir 2023).

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

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 and 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; and (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 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

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

4. Using the elicitation method by Krupka and Weber (2013), we find that 841 out of 1,022 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".

5. In Online 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.

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 multilateral market treatments and, as in the previous literature, an individual decision-making control treatment, multiple price-lists (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 and Szech (2013) and Bartling, Fehr, and Özdemir (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. Out of the total, 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 and 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 and Szech (2013) inspired a follow-up literature that investigates how different market forces affect traders' morals. Bartling, Weber, and 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. 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 et al. 2016; Irlenbusch and Saxler 2019; Sutter et al. 2020). Engelmann, Friedrichsen, and 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, Weber, and Yao (2015), there are also some other papers that study specific market structures that allow markets to partially sustain pro-social behavior. Schneider, Brun, and Weber (2020) document an endogenously arising wage premium, and associated sorting, for morally questionable tasks. Dufwenberg et al. (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, and Wiesen (2023) and van Leeuwen, Offerman, and Schram (2020). In a large non-student sample, Riehm et al. (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. Theoretically, Kaufmann and Köszegi (2023) study equilibria and regulation of markets with socially responsible consumers, whereas Dewatripont and Tirole (2023) show that competition may not

6. Our 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.

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

crowd out consequentialist ethics. Ockenfels, Werner, and Edenhofer (2020) and Herweg and 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, and Schram 1996; Fischbacher, Gächter, and Fehr 2001; Burlando and Guala 2005). Falk et al. (2018) document heterogeneity in social preferences within and across many countries. Given that 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, and 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 and Szech (2016) find that almost a third of their participants pay for a diffused notion of being pivotal for a questionable moral outcome. Serra-Garcia and 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, Neuber, and Szech (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 et al. (1991), Prasnikar and Roth (1992), and Fischbacher, Fong, and 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 and 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, and 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 and Zimmermann 2019; Enke, Graeber, and 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

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

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.

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 toward UNICEF for varying stakes. In Section 3.3, we give more details on the donation opportunity. We MPL 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 MPL.

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 canceled.

9. We 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.

Buyers and sellers moved in turns, trading unit by unit. In each market period, one side of the market—that is, the buyers or the sellers—was randomly determined to move first. The starting side had the opportunity to submit offers to the second movers within a time constraint of 14 seconds. We restricted all offer submissions to yield non-negative profits for both market sides. Afterward, 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 and (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 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 Online 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.

10. For 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, March 7, 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.

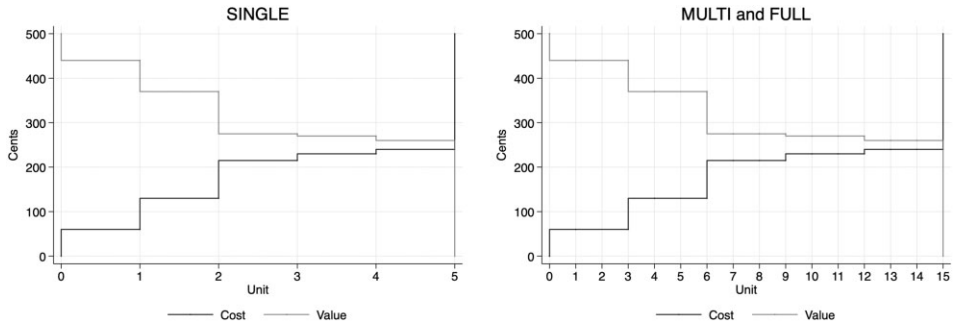

FIGURE 1. Induced common costs and values.

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 canceled 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. 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, heterogeneous participants (HET), fully activates market selection as participants are very heterogeneous in their preference for the external effect. The former, homogeneous participants (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, 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 were 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".

11. This 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).

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 MPL. 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.

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

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; and (iii) risk preferences. For the beliefs, participants were asked to fill in a MPL reporting what they "think the average participant did" in the first list of part 1. If their belief matched the choice

12. At 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.

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 and 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 Online Appendix Section A.6. We also elicited risk attitudes using the method introduced by Eckel and 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, and 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, and 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, and 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 Online 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 MPL 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 Online 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; and (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

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."

14. This 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.

15. At 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.

separation of the role played by the replacement logic vis-à-vis market selection. These hypotheses, preregistered in Offerman, Romagnoli, and Ziegler (2019), Offerman, Romagnoli, and Ziegler (2021), and Offerman, Romagnoli, and Ziegler (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 and Szech (2013) report limited erosion of morals in single-unit markets. Bartling, Weber, and Yao (2015) find limited erosion in most specifications. Bartling, Fehr, and Özdemir (2023) fail to reject this hypothesis. Our first hypothesis is thus.

HYPOTHESIS 1. 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.

HYPOTHESIS 2A. 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.

HYPOTHESIS 2B. 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.

HYPOTHESIS 3. 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 and 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, that is, 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, Neuber, and Szech (2020).

Our hypotheses regarding the mechanisms of moral erosion are thus:

HYPOTHESIS 4. Any erosion of morals in FULL compared to MULTI is not driven by market selection.

HYPOTHESIS 5. 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, and (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.

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]

16. In 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.

17. Bénabou et al. (2020) show that elicited moral costs can be affected by the method of elicitation, when using either direct elicitation or MPL, 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

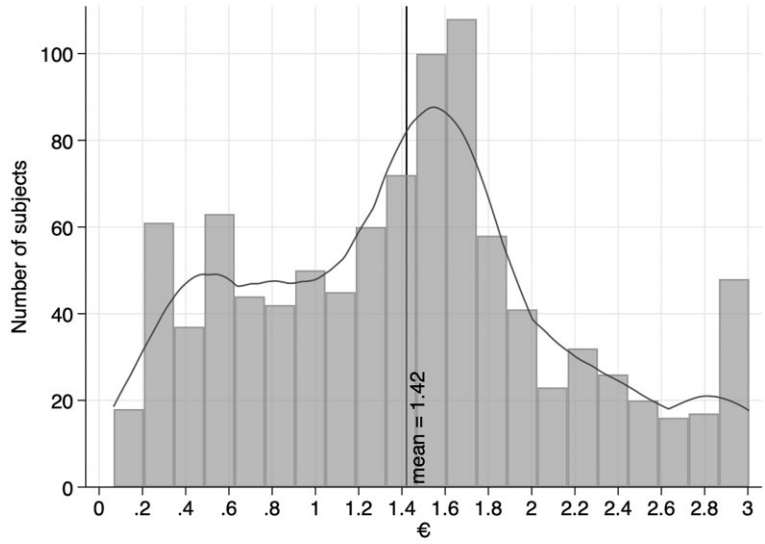

FIGURE 2. Heterogeneity in valuations of donations. Histogram and kernel density of participants' average moral costs for canceling a donation with a value of €1.5. For each participant, we use the switching points from all MPL for canceling donations in part 1.

We also detect decreasing marginal moral costs and provide an analysis in Online 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 additional 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 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 Online 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 and

"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).

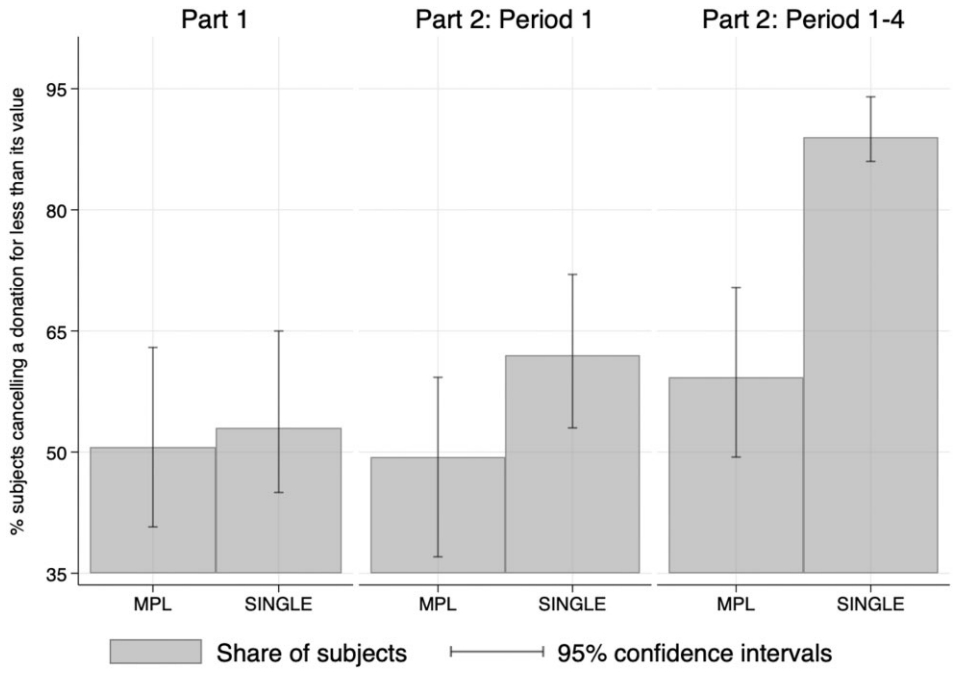

FIGURE 3. Cancellation of donations between environments and treatments. Share of participants who canceled a donation for at most its value (€1.50) in individual decision-making tasks and in trades in the market. The left panel shows cancellation rates in part 1 of the experiment and the middle panel 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 canceled a donation.

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

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

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***	0.270***			0.099***
	(0.033)	(0.052)			(0.033)
SINGLE			0.126*	0.297***	0.126*
			(0.074)	(0.059)	(0.075)
SINGLE × Period 1–4					0.171***
					(0.060)
Constant	0.494***	0.620***	0.494***	0.593***	0.494***
	(0.056)	(0.051)	(0.056)	(0.055)	(0.056)
Observations	162	200	181	181	362

Note: Dependent variable is a dummy equal to one if a participant canceled 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 1–4 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, $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

use the entire four periods of the experiment and plot the share of participants who at least once canceled 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 canceled 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, Fehr, and Özdemir (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 first period.[19]

RESULT 1. We reject Hypothesis H1, and find partial erosion of morals in single-unit markets.

19. A 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.

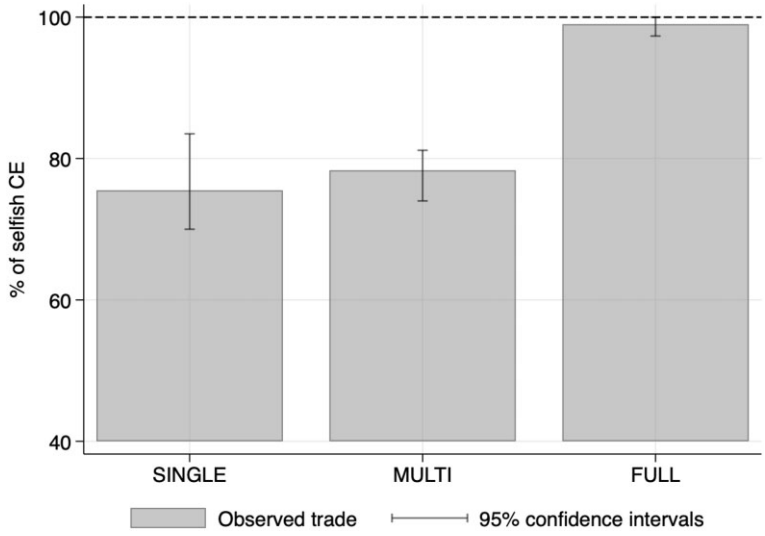

FIGURE 4. Market outcomes. Average quantities relative to selfish competitive equilibrium. Trading units below 40% is efficient (gains from trade exceed the externality). Compared to the negative externality of €1.50 per unit, each unit between 40% and 60% yields gains from trade of €0.60, each unit between 60% and 80% yields €0.40, and each unit between 80% and 100% yields €0.20.

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 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

TABLE 3. Treatment effects.

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

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

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). 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 and 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

20. This 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 Online Appendix Section A.1.

21. These 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 Online Appendix Section A.15 for results.

22. We 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 Online Appendix Section A.13. The main takeaway from this treatment is that morals are eroded to an approximately similar extent as in MULTI.

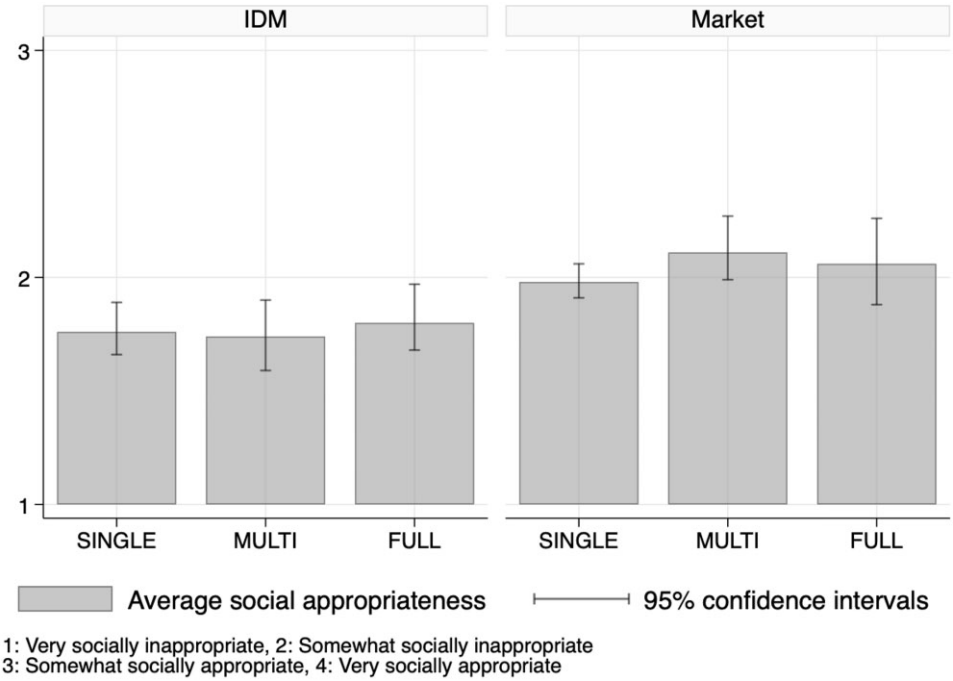

FIGURE 5. Norms in individual decision-making and in markets. Average norm about canceling one donation of €1.50 when paid €1 in an individual decision-making task (left panel) and in an experimental market (right panel). A rating of 2 corresponds to "somewhat socially inappropriate".

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.

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

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

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 Online 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 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,

24. Results are similar when regressing participants' norms (2 elicitations for 981 participants, so 1962 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 =.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.

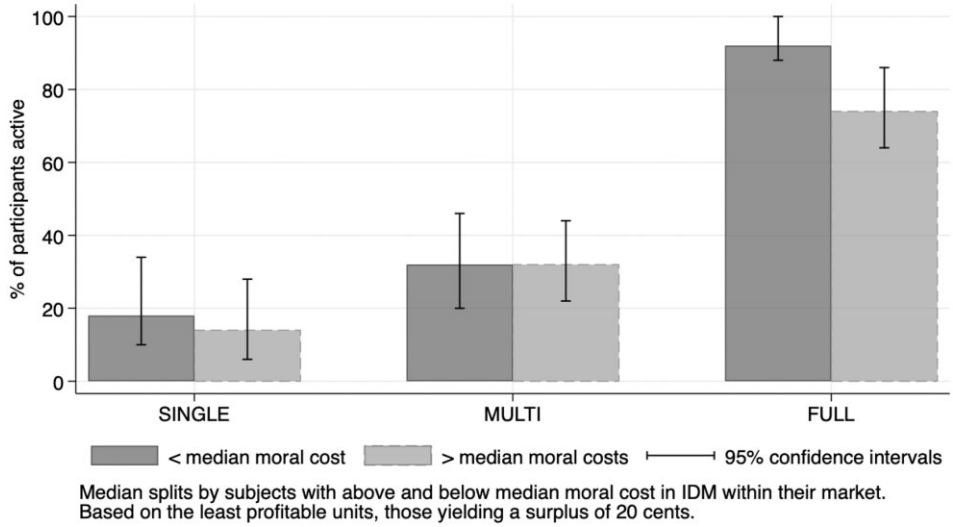

FIGURE 6. Share of traders active at the least profitable units. Share of traders who submit or accept an offer at the final units, which yield gains from trade of €0.20 in exchange for an externality of €1.50. Median splits are based on moral costs within the matching group.

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 Online 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. 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]

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

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

26. In the Online 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.

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 Online 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.

Strikingly, market outcomes are similarly selfish in HOM, HET, and FULL. Whereas in FULL 99.0% 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 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

27. This analysis assumes independence of observations, even though, for example, 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.

28. Other 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 Online Appendix Section A.8 for details.

presented in the Online 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 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—that is, 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 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 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]

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

30. Treatments B-FULL and B-MULTI also allow us to investigate the robustness of the original results. In Online 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.

31. This is the only beliefs analysis that does not generalize when we use the unincentivized belief report (see Table A.10 in Online Appendix Section A.8). Falk, Neuber, and Szech (2020) show in a similar analysis

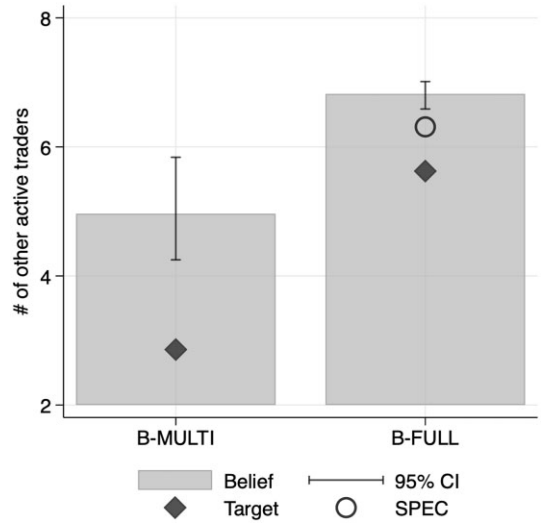

FIGURE 7. Beliefs about other traders' activity. Number of other traders believed to be active (bars), actual number of others active (the target, diamond) and belief of spectators (circle).

TABLE 4. Beliefs and activity.

	(1) B-MULTI	(2) B-FULL	(3)	(4)	(5) B-LOW & B-HIGH	(6)
	Own activity		Belief activity		Own activity	
Belief activity	0.053***	0.069***			0.054***	
	(0.011)	(0.013)			(0.012)	
Average moral costs	−0.019	−0.111***			−0.147***	−0.147***
	(0.023)	(0.025)			(0.033)	(0.035)
Period	0.001	−0.028*				
	(0.039)	(0.013)				
B-LOW			−3.690***	−0.260***	−0.000	
			(0.384)	(0.083)	(0.080)	
Belief activity (IV)						0.055***
						(0.016)
Constant	0.076	0.391**	7.270***	0.700***	0.454***	0.453***
	(0.119)	(0.121)	(0.303)	(0.060)	(0.107)	(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)–(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, $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

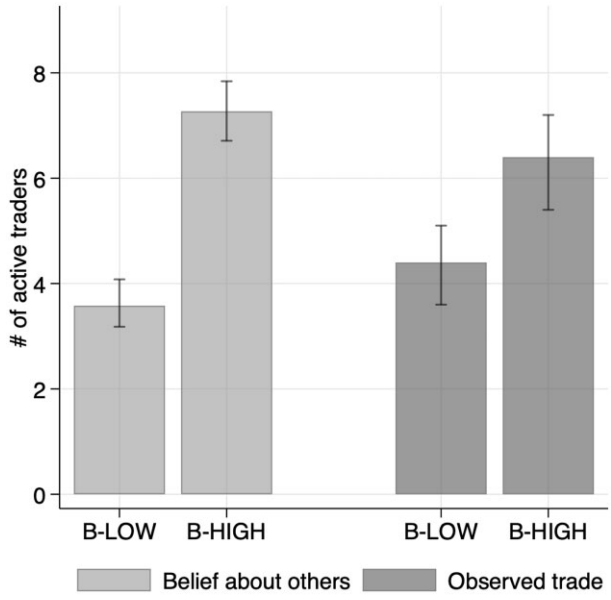

FIGURE 8. Beliefs and actions with exogenously manipulated beliefs. Left panel: number of other traders believed to be active, right panel: actual number of active traders.

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.

that beliefs correlate with choice in line with the replacement logic. In their experiment, their incentivized belief did correlate with their unincentivized measure.

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 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 Online 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).

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 Online 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 Online 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.

32. These outcomes were the preregistered main hypotheses for the new treatments. In Online 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.

33. The IV analysis was not preregistered.

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 and 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 and 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 Online 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, Abby, William N. Evans, Ethan M. J. Lieber, and David Powell (2022). "Origins of the Opioid Crisis and its Enduring Impacts." *The Quarterly Journal of Economics*, 137, 1139–1179.

Barr, Stewart, Gareth Shaw, and Tim Coles (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, 1234–1244.

Bartling, Björn, Ernst Fehr, and Yagiz Özdemir (2023). "Does Market Interaction Erode Moral Values?" *Review of Economics and Statistics*, 105, 226–235.

Bartling, Björn and Yagiz Özdemir (2023). "The Limits to Moral Erosion in Markets: Social Norms and the Replacement Excuse." *Games and Economic Behavior*, 138, 143–160.
Bartling, Björn, Vanessa Valero, and Roberto Weber (2019). "On the Scope of Externalities in Experimental Markets." *Experimental Economics*, 22, 610–624.
Bartling, Björn, Roberto A. Weber, and Lan Yao (2015). "Do Markets Erode Social Responsibility?" *The Quarterly Journal of Economics*, 130, 219–266.
Bénabou, Roland, Armin Falk, Luca Henkel, and Jean Tirole (2020). "Eliciting Moral Preferences: Theory and Experiment." Princeton University Working Paper. https://rbenabou.scholar.princeton.edu/sites/g/files/toruqf3331/files/rbenabou/files/moral_preference_29_05_22.pdf
Brütt, Katharina, Arthur Schram, and Joep Sonnemans (2020). "Endogenous Group Formation and Responsibility Diffusion: An Experimental Study." *Games and Economic Behavior*, 121, 1–31.
Burlando, Roberto M. and Francesco Guala (2005). "Heterogeneous Agents in Public Goods Experiments." *Experimental Economics*, 8, 35–54.
Byambadalai, Undral, Ching-to Albert Ma, and Daniel Wiesen (2023). "Changing Preferences: An Experiment and Estimation of Market-Incentive Effects on Altruism." *Journal of Health Economics*, 92, 102808.
Dana, Jason, Roberto A. Weber, and Jason Xi Kuang (2007). "Exploiting Moral Wiggle Room: Experiments Demonstrating an Illusory Preference for Fairness." *Economic Theory*, 33, 67–80.
Dewatripont, Mathias and Jean Tirole (2024). "The Morality of Markets." *Journal of Political Economy*. https://doi.org/10.1086/729445
Dufwenberg, Martin, Olof Johansson-Stenman, Michael Kirchler, Florian Lindner, and Rene Schwaiger (2022). "Mean Markets or Kind Commerce?" *Journal of Public Economics*, 209, 104648.
Eckel, Catherine C. and Philip J. Grossman (2002). "Sex Differences and Statistical Stereotyping in Attitudes Toward Financial Risk." *Evolution and Human Behavior*, 23, 281–295.
Engelmann, Dirk, Jana Friedrichsen, and Dorothea Kübler (2018). "Fairness in Markets and Market Experiments." WZB Discussion Paper No. SP II 2018–203.
Enke, Benjamin, Thomas Graeber, and Ryan Oprea (2023). "Confidence, Self-Selection, and Bias in the Aggregate." *American Economic Review*, 113(7), 1933–1966.
Enke, Benjamin and Florian Zimmermann (2019). "Correlation Neglect in Belief Formation." *The Review of Economic Studies*, 86, 313–332.
Exley, Christine L. (2016). "Excusing Selfishness in Charitable Giving: The Role of Risk." *The Review of Economic Studies*, 83, 587–628.
Falk, A., T. Neuber, and N. Szech (2020). "Diffusion of Being Pivotal and Immoral Outcomes." *The Review of Economic Studies*, 87, 2205–2229.
Falk, Armin, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman, and Uwe Sunde (2018). "Global Evidence on Economic Preferences." *The Quarterly Journal of Economics*, 133, 1645–1692.
Falk, Armin and Nora Szech (2013). "Morals and Markets." *Science*, 340, 707–711.
Falk, Armin and Nora Szech (2016). "Diffusion of Being Pivotal and Immoral Outcomes." Working Paper No. 2016-013, Human Capital and Economic Opportunity Global Working Group.
Fischbacher, Urs, Christina M. Fong, and Ernst Fehr (2009). "Fairness, Errors and the Power of Competition." *Journal of Economic Behavior & Organization*, 72, 527–545.
Fischbacher, Urs, Simon Gächter, and Ernst Fehr (2001). "Are People Conditionally Cooperative? Evidence from a Public Goods Experiment." *Economics Letters*, 71, 397–404.
Gössling, Stefan and Andreas Humpe (2020). "The Global Scale, Distribution and Growth of Aviation: Implications for Climate Change." *Global Environmental Change*, 65, 102194.
Grossman, Zachary (2014). "Strategic Ignorance and the Robustness of Social Preferences." *Management Science*, 60, 2659–2665.
Grossman, Zachary and Joel J. Van Der Weele (2017). "Self-Image and Willful Ignorance in Social Decisions." *Journal of the European Economic Association*, 15, 173–217.
Herweg, Fabian and Klaus M. Schmidt (2022). "How to Regulate Carbon Emissions with Climate-Conscious Consumers." *The Economic Journal*, 132, 2992–3019.
Huck, Steffen and Kai A. Konrad (2005). "Moral Cost, Commitment, and Committee Size." *Journal of Institutional and Theoretical Economics*, 161, 575–588.

Irlenbusch, Bernd and David J. Saxler (2019). "The Role of Social Information, Market Framing, and Diffusion of Responsibility as Determinants of Socially Responsible Behavior." *Journal of Behavioral and Experimental Economics*, 80, 141–161.
Kaufmann, Marc and Botond Köszegi (2023). "Understanding Markets with Socially Responsible Consumers." Working paper, briq.
Ketcham, Jon, Vernon L. Smith, and Arlington W. Williams (1984). "A Comparison of Posted-Offer and Double-Auction Pricing Institutions." *The Review of Economic Studies*, 51, 595–614.
Kirchler, Michael, Jürgen Huber, Matthias Stefan, and Matthias Sutter (2016). "Market Design and Moral Behavior." *Management Science*, 62, 2615–2625.
Krupka, Erin L. and Roberto A. Weber (2013). "Identifying Social Norms using Coordination Games: Why Does Dictator Game Sharing Vary?" *Journal of the European Economic Association*, 11, 495–524.
List, John A. (2003). "Does Market Experience Eliminate Market Anomalies?" *The Quarterly Journal of Economics*, 118, 41–71.
Ockenfels, Axel, Peter Werner, and Ottmar Edenhofer (2020). "Pricing Externalities and Moral Behaviour." *Nature Sustainability*, 3, 872–877.
Offerman, Theo, Giorgia Romagnoli, and Andreas Ziegler (2019). "Morals in Multi-Unit Markets." *AEA RCT Registry*. https://doi.org/10.1257/rct.4603-1.3. Accessed January 18, 2024.
Offerman, Theo, Giorgia Romagnoli, and Andreas Ziegler (2021). "Morals in multi-unit markets: Study II." *AEA RCT Registry*. https://doi.org/10.1257/rct.8306-1.0. Accessed January 18, 2024.
Offerman, Theo, Giorgia Romagnoli, and Andreas Ziegler (2023). "Morals in Multi-Unit Markets: Study III." *AEA RCT Registry*. https://doi.org/10.1257/rct.11080-1.0. Accessed January 18, 2024.
Offerman, Theo, Joep Sonnemans, and Arthur Schram (1996). "Value Orientations, Expectations and Voluntary Contributions in Public Goods." *The Economic Journal*, 106, 817–845.
Plott, Charles R. (1983). "Externalities and Corrective Policies in Experimental Markets." *The Economic Journal*, 93, 106–127.
Prasnikar, Vesna and Alvin E. Roth (1992). "Considerations of Fairness and Strategy: Experimental Data from Sequential Games." *The Quarterly Journal of Economics*, 107, 865–888.
Riehm, Tobias, Nicolas Fugger, Philippe Gillen, Vitali Gretschko, and Peter Werner (2022). "Social Norms, Sanctions, and Conditional Entry in Markets with Externalities: Evidence from an Artefactual Field Experiment." *Journal of Public Economics*, 212, 104701.
Roth, Alvin E, Vesna Prasnikar, Masahiro Okuno-Fujiwara, and Shmuel Zamir (1991). "Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study." *American Economic Review*, 81(5), 1068–1095.
Rothenhäusler, Dominik, Nikolaus Schweizer, and Nora Szech (2018). "Guilt in Voting and Public Good Games." *European Economic Review*, 101, 664–681.
Samuelson, Paul A. and William D. Nordhaus (2005). *Economics*, 18th ed. Columbus, OH, McGraw-Hill Irwin.
Schneider, Florian, Fanny Brun, and Roberto A. Weber (2020). "Sorting and Wage Premiums in Immoral Work." University of Zurich, Department of Economics, Working Paper No. 353.
Scholl, Lawrence, Puja Seth, Mbabazi Kariisa, Nana Wilson, and Grant Baldwin (2019). "Drug and Opioid-Involved Overdose Deaths'United States, 2013–2017." *Morbidity and Mortality Weekly Report*, 67, 1419.
Serra-Garcia, Marta and Nora Szech (2022). "The (in) Elasticity of Moral Ignorance." *Management Science*, 68, 4815–4834.
Smith, Vernon L. (1962). "An Experimental Study of Competitive Market Behavior." *Journal of Political Economy*, 70, 111–137.
Sobel, Joel (2007). "Do Markets Make People Look Selfish?" *UCSD*. https://econweb.ucsd.edu/~jsobel/Papers/20070225_markets.pdf. Accessed January 18, 2024.
Sutter, Matthias, Jürgen Huber, Michael Kirchler, Matthias Stefan, and Markus Walzl (2020). "Where to Look for the Morals in Markets?" *Experimental Economics*, 23, 30–52.

van Leeuwen, Boris, Theo Offerman, and Arthur Schram (2020). "Competition for Status Creates Superstars: An Experiment on Public Good Provision and Network Formation." *Journal of the European Economic Association*, 18, 666–707.
Vuillemey, Guillaume (2020). "Evading Corporate Responsibilities: Evidence from the Shipping Industry." CEPR Discussion Paper No. 15291.

Supplementary Data

Supplementary data are available at *JEEA* online.

Journal of the European Economic Association 2024 22(5):2225–2260
https://doi.org/10.1093/jeea/jvae001