

May 2024

“Doing the right thing (or not) in a lemons-like situation:
on the role of social preferences and Kantian moral
concerns”

Ingela Alger and José Ignacio Rivero-Wildemauwe

Doing the right thing (or not) in a lemons-like situation: on the role of social preferences and Kantian moral concerns ^{*}

Ingela Alger[†]

José Ignacio Rivero-Wildemauwe[‡]

6th May 2024

Abstract: We conduct a laboratory experiment using framing to assess the willingness to “sell a lemon”, i.e., to undertake an action that benefits self but hurts the other (the “buyer”). We seek to disentangle the role of other-regarding preferences and (Kantian) moral concerns, and to test if it matters whether the decision is described in neutral terms or as a market situation. When evaluating an action, morally motivated individuals consider what their own payoff would be if—hypothetically—the roles were reversed and the other subject chose the same action (universalization). We vary the salience of role uncertainty, thus varying the ease for participants to envisage the role-reversal scenario. We find that subjects are (1) more likely to “sell a lemon” in the market frame, and (2) less likely to do so when the role uncertainty is salient. We also structurally estimate other-regarding and Kantian moral concern parameters.

JEL codes: C91, D01, D91.

Keywords: market framing, lemons, social preferences, Kantian morality, experiment.

^{*}I.A. acknowledges funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 789111 - ERC EvolvingEconomics), as well as IAST funding from the French National Research Agency (ANR) under grant ANR-17-EURE-0010 (Investissements d’Avenir program). We are grateful to Lony Bessagnet, Jean-François Bonnefon, Eric Danan, Pierre Fleckinger, Peter Hammond, Astrid Hopfensitz, Nicolas Jacquemet, Pau Juan-Bartroli, Fabrice Le Lec, Thierry Magnac, Pauline Morault, Esteban Muñoz-Sobrado, Sébastien Pouget, Régis Renault, Angelo Secchi, Sigrid Suetens, Ao Wang and Boris van Leeuwen for helpful discussions, and to Stéphane Cézera for invaluable assistance in the lab. We thank audiences at CY Cergy Paris University, IAST, the Paris School of Economics Summer School, Toulouse School of Economics, Catholic University of Uruguay, Max Planck Institute, Economic Science Association, Uruguayan Central Bank, Universidad de la República and University of Gothenburg for feedback.

[†]Toulouse School of Economics, CNRS, University of Toulouse Capitole, and Institute for Advanced Study in Toulouse, France, and CEPR. ingela.alger@tse-fr.eu

[‡]Department of Business and Administration, Catholic University of Uruguay, Montevideo, Uruguay. joseignacio.rivero@ucu.edu.uy

1 Introduction

There is by now robust evidence from experimental games played by subjects in laboratory settings in a host of countries, that a large share of individuals give up own material payoff to increase that of others (e.g., [Zelmer \(2003\)](#); [C. Engel \(2011\)](#)), and exhibit lying aversion (e.g., [Gneezy, Rockenbach, & Serra-Garcia \(2013\)](#); [Abeler, Nosenzo, & Raymond \(2019\)](#)). What are the implications of this evidence for behaviour in real market settings? In particular, is there reason to believe that sellers are less willing to sell goods at a price above their value than many theoretical models assume ([Akerlof, 1970](#))? We aim at providing some insights by conducting a laboratory experiment that we designed with two objectives in mind: (1) to examine behavior of subjects in a situation reminiscent of that of a seller of a “lemon”, using both a neutral and a market frame; and (2) to disentangle the role of other-regarding, or distributional, preferences from Kantian moral concerns in this context.

Subjects in our experiment face a series of anonymous binary choices, each of which affects the payoff of a randomly drawn other subject, hence avoiding that subjects’ decisions be influenced by strategic or repeated interaction considerations, ([Roth & Schoumaker, 1983](#); [Andreoni & Miller, 1993](#)), and social image or reciprocity concerns ([Bénabou & Tirole, 2006](#); [D. K. Levine, 1998](#); [Charness & Rabin, 2002](#)). In each binary choice one of the options—the *selfish option*—entails a gain for self and a loss for the other, compared to the other option—the *status quo* option (in other words, each decision situation is a Dictator game, with taking rather than giving ([Dreber, Ellingsen, Johannesson, & Rand, 2013](#))). The payoff consequences of this choice thus capture a lemons-like situation, in which the selfish option amounts to selling an object of low quality at such a price that the seller gets better off but the buyer gets worse off, while the *status quo* option amounts to keeping the object.

To achieve our first objective, we let half of the decisions be taken in a standard *neutral frame* and the other half in a *market frame*, keeping the payoff consequences identical under both frames. Only the wording used to describe the situation differs. Specifically, in the market frame subjects get the following information: “As you can see on the decision screen above, the Buyer would be better off if you chose Not Sell, while you are better off if you Sell. Think of this as representing a situation in which

the good that you sell has a defect which makes the Buyer enjoy owning the good less than you do.” The wording was chosen to mimic a situation in which a seller can sell a “lemon” to an uninformed buyer (who in the experiment is passive). In the neutral frame the options are simply referred to as X and Y, and no context is given.

Our second objective is to disentangle two distinct motivations for refraining from the selfish option. The first motivation is consequentialist: since the selfish option increases the payoff gap between the decision-maker and the passive subject, the *status quo* may be selected due to an other-regarding concern in the form of altruism or inequity aversion, formalized as a weight β on the payoff gap (Becker, 1974; Fehr & Schmidt, 1999). The second motivation relies on a (possibly partial) universalization argument: it makes the subject evaluate each decision in the light of what his payoff would be, if—hypothetically—the other were to take the same decision as the subject himself with some probability κ , if the roles were reversed. This motivation, which is reminiscent of a Kantian moral concern, can be expressed with *Homo moralis* preferences, where κ is the individual’s degree of morality (Alger & Weibull, 2013). We posit that subjects have preferences that combine other-regard with a Kantian moral concern. Such a combination is indeed predicted by recent theoretical work on the evolutionary foundations of preferences in a model where interactions yield payoff consequences with negligible effects on overall reproductive success (Alger, Weibull, & Lehmann, 2020), and it has also received empirical support in another experimental study (Van Leeuwen & Alger, in press). Moreover, recent theoretical studies on the effects of Kantian moral concerns on market outcomes that feature an unequal distribution of information (including cases where product quality is the relevant information asymmetry) document important effects of such moral concerns on equilibrium outcomes (Rivero-Wildemaue, 2023a,b). In equilibrium, Kantian motivations reduce the likelihood of “lemons” being exchanged, or may even eliminate it altogether. Importantly, these references also underline that other-regard may have similar effects. In light of this, it is of particular relevance to be able to distinguish between both behavioral drivers.

Because the Kantian concern is triggered by consideration of a possible role reversal, the experimental design consists in varying the salience of role reversal. Specifically, half of the decisions we observe are taken in a *VOI frame* and the other

half in a *non-VOI frame*, where VOI stands for *veil of ignorance*. In the VOI frame, a subject is told that (s)he stands an equal chance of being the decision-maker and the passive subject. In the non-VOI frame (s)he is simply told that (s)he is the decision-maker. In reality, however, the two situations are the same, since upon entering the experiment any subject effectively stands an equal chance of being the decision-maker and the passive subject in any given match. Hence, the only difference is that we make this role uncertainty explicit in the VOI frame. Our theory thus predicts that if subjects are driven by a Kantian moral concern and if they are not fully aware of the arbitrariness in the role distribution in the non-VOI decisions, they are more likely to sell in non-VOI decisions than in VOI decisions. We also estimate the contribution of the other-regarding and the Kantian motivation to the subjects' decisions, and examine whether the estimated motivations differ in the market and the neutral frames.

We find that, on average, subjects are significantly more prone to selecting the selfish option in the market frame compared to the neutral frame. In line with our theoretical prediction, they also select the selfish option significantly more often in non-VOI than in VOI decision situations. Comparing the effect of the VOI frame on the propensity to select the selfish decision between the market frame and the neutral frame, the magnitude in absolute terms is larger in the market frame, but the effect of VOI relative to non-VOI decisions is the same in both frames. The effects of the market frame and of the VOI frame are larger in magnitude in the within-subjects regressions, suggesting significant heterogeneity across subjects.

We also structurally estimate the preference parameters β and κ , in line with several earlier experimental studies (e.g., [Fisman, Kariv, & Markovits \(2007\)](#); [Bruhin, Fehr, & Schunk \(2018\)](#); [Van Leeuwen & Alger \(in press\)](#)). A novelty of our study compared to existing ones, is that we obtain estimates of the preference parameters in two different frames: the market frame and the neutral frame. Overall, we find differences in both the degree of aheadness aversion and the Kantian morality estimates across the market and the neutral frame, some of the differences being large.

Our study contributes to three strands of the literature. First, it adds insights into an old question: do market interactions render people more selfish ([Bowles, 1998](#))? Some authors have relied on lab-in-the-field experiments to compare the be-

havior of individuals with varying degrees of exposure to market interactions (e.g., [Henrich et al. \(2005, 2016\)](#) and [Agneman & Chevrot-Bianco \(2022\)](#)), while others have used laboratory to examine how variation in the competition between subjects affects behaviors (e.g., [Cabrales, Miniaci, Piovesan, & Ponti \(2010\)](#); [Sutter, Huber, Kirchler, Stefan, & Walzl \(2020\)](#); [J. Engel & Szech \(2020\)](#); [Dufwenberg, Johansson-Stenman, Kirchler, Lindner, & Schwaiger \(2022\)](#); [Bartling & Özdemir \(2023\)](#)). Our experiment explores a different possibility, by adopting a design where we examine the effect of framing the task as the sale of an lemon (recall the wording above), compared to a standard neutral framing, the objective being to quantify the effect of the market frame (for a recent discussion of the use of frames in economic experiments, see [Alekseev, Charness, & Gneezy \(2017\)](#)). Our hypothesis is close to that put forward elsewhere (e.g., [Bowles \(1998\)](#); [Kirman & Teschl \(2010\)](#)): preferences may be endogenous and thus shaped differently by market interactions than by, say, interactions with friends and family. Some experiments lend support to the importance of frames, for example the famous study that showed that a Prisoner’s dilemma generates more cooperation if it is labeled as “the community game” than if it is called “the Wall street game” ([Lieberman, Samuels, & Ross, 2004](#)) (see also [Kay & Ross, 2003](#)). Other studies, however, found small or no effects ([Dreber, Ellingsen, Johannesson, & Rand, 2013](#); [Dufwenberg, Gächter, & Hennig-Schmidt, 2011](#)). By contrast to other studies that use some kind of market-related framing (e.g., [Dufwenberg, Gächter, & Hennig-Schmidt \(2011\)](#); [Thöni & Gächter \(2015\)](#)), our instructions to the participants explicitly describe an action as the sale of a lemon, and we compare behavior under this wording to behavior under neutral wording. Furthermore, our design allows us to compare the estimated preference parameters between the neutral and the market frames.

Second, our study adds to a recent set of experiments that have sought to detect moral concerns as drivers of behavior, where moral concerns are distinct from distributional preferences ([Capraro & Rand, 2018](#); [Bursztyn, Fiorin, Gottlieb, & Kanz, 2019](#); [Miettinen, Kosfeld, Fehr, & Weibull, 2020](#); [Chen & Schonger, 2022](#); [Feess, Kerzenmacher, & Timofeyev, 2022](#); [Bénabou, Falk, Henkel, & Tirole, 2023](#); [Bénabou, Falk, & Henkel, 2024](#); [Van Leeuwen & Alger, in press](#)). The most closely related experiments are those by [Miettinen, Kosfeld, Fehr, & Weibull \(2020\)](#) and [Van Leeuwen & Alger \(in press\)](#), who also posit a utility function with a Kantian moral concern à

la *Homo moralis* (Alger & Weibull, 2013), and who also seek to disentangle this concern from distributional preferences. Our study differs from both of them in two important ways. First, while they only use neutral wording, our experimental design allows us to test whether preferences differ between the neutral and the market frame. Second, they rely exclusively on what we call VOI decisions; our comparison of decisions in VOI and non-VOI decisions gives an indication about the extent to which the explicit mention of the role uncertainty helps trigger the Kantian moral concern.

This last point brings us to the third literature to which our study contributes: that which examines if decisions depend on whether they are elicited using the direct-response or the strategy method — i.e., whether they are taken after or before subjects learn the role distribution. Several studies find that decisions differ significantly, while other studies find mixed or no effects (see the surveys by Brandts & Charness (2011) and Grech, Nax, & Soos (2022)). In the studies where an effect appears, decisions taken behind the veil of ignorance tend to be more pro-social on average (Sutter & Weck-Hannemann, 2003; Iriberry & Rey-Biel, 2011; Huang, Greene, & Bazerman, 2019; S. Levine, Kleiman-Weiner, Schulz, Tenenbaum, & Cushman, 2020; García-Pola, Iriberry, & Kovářik, 2020; Ortiz-Riomalo, Koessler, & Engel, 2021; Herne, Hietanen, Lappalainen, & Palosaari, 2022). Our findings are in line with this empirical regularity. But while previous studies have sought to explain the greater pro-sociality in decisions taken under role uncertainty by referring mainly to aspects of the decision-making process (e.g., number of decisions taken, complexity of the information presented, likelihood that emotions are triggered by the task), our experiment relies on a preference-based theory which predicts and thus explains it. According to our theory, the explicit mention of role uncertainty awakens or reinforces a Kantian moral concern, which triggers pro-social behavior in our dictator game design.

The next section describes the experimental design and procedures, and in Section 3 we test whether behaviors are different under the market and the neutral frames, and under the VOI and the non-VOI frames. In Section 4 we report the results from the structural estimations, and Section 5 concludes.




Figure 1: The game protocol used in the experiment. The payoff vectors state subject i 's and j 's payoff as the first and second component, respectively.

2 The experiment

2.1 Game protocols and preferences

In the experiment subjects are matched into pairs to play anonymous one-shot interactions. In each interaction one subject is assigned to the *Player 1* role and the other to the *Player 2* role, each role assignment being equally likely. Player 1 has to choose between two actions, call them X and Y . If she chooses X , both individuals obtain their initial endowment, denoted e_1 and e_2 , respectively, while if she chooses Y , she gets $e_1 + G > e_1$ while Player 2 gets $e_2 - L < e_2$, where $e_1 > e_2$, $G > 0$, and $L > 0$. That is, Player 1 makes a net gain G from choosing Y rather than X , while this choice entails a net loss L for Player 2. Note that whether she chooses X or Y , Player 1 receives a higher payoff than Player 2. The game tree for any matched pair of subjects—here called i and j —is depicted in Figure 1.

We use two frames in the experiment: the **Neutral frame** and the **Market frame**. In the latter, we replace *Player 1* by *Seller* and *Player 2* by *Buyer*, and we describe actions X and Y by *Not Sell* and *Sell*, respectively. Indeed, the payoff structure captures in a stylised manner the case where a seller must decide whether to sell a low-quality item to a buyer in exchange for the going (fixed) price in the market—a classic “lemons” situation, in which Y amounts to selling and X to keeping the good. The Buyer has no say here: we adopt this simplification to concentrate on the Seller’s willingness to sell a lemon. To facilitate the exposition, we will henceforth




Figure 2: The matched subjects (i and j) are informed about the whole game tree in a VOI decision, but only about the realized role distribution in a non-VOI decision (if the Player 1 role is assigned to subject i , she only receives information about the part highlighted in red).

use the term *Selfish option* to refer to *Sell* in the Market frame and action Y in the Neutral frame, and the term *Status quo option* to refer to *Not Sell* in the Market frame and action X in the Neutral frame. We will also refer to the two roles as the active and the passive roles.

In addition to the aforementioned frames, we vary the information that subjects receive about the decision situation. In the **VOI decision situation** subjects are asked to state their choice in the active role *before* learning the role distribution. Participants are told that they have an equal chance of being cast in either role. Being thus explicitly informed about the whole game tree in Figure 2, they are led to reason behind the “Veil of Ignorance” (VOI) with respect to the role distribution. In the **non-VOI decision situation**, a subject is asked to state her decision *after* being informed that she is in the active role and is therefore not behind the “Veil of Ignorance”. Hence, even though the game actually being played between any given matched subject pair in the experiment is the one depicted in Figure 2 (since each subject stands an equal chance of being handed the active role), in a non-VOI decision situation a subject cast in the active role is given explicit information from us only about part of the game tree (the part highlighted in red in Figure 2, should the active player be if i is the active subject).

The decisions being anonymous and one-shot, the experimental design removes motivations such as social image concerns, repeated interaction, and reciprocity ef-

fects. Hence, a subject’s decision should be driven by his or her intrinsic preferences and beliefs. We posit a utility function that combines material self-interest, other-regard, and a Kantian moral concern. As we will now show the key point is that the Kantian moral concern is expected to be fully triggered in a VOI decision situation, but only partially so or not at all in a non-VOI decision situation. The reason is that in the former, subjects explicitly take into account the possibility that they may end up in either role, while in the latter individuals are either fully unaware or only partially aware of the arbitrariness of the role allocation.

Formally, consider subjects i and j and denote their decisions by $x \in \{0, 1\}$ and $y \in \{0, 1\}$ respectively, where $x = 1$ if i chooses the selfish option and $x = 0$ if she selects the *status quo* option when active, while $y = 1$ if j selects the selfish option and $y = 0$ if she chooses the *status quo* option when active. We posit the following (expected) utility obtained by i in the VOI decision situation:

$$\begin{aligned}
 U_i(x, y) &= (1 - \kappa_i) \cdot \left[\frac{1}{2} \cdot (e_1 + xG) + \frac{1}{2} \cdot (e_2 - yL) \right] & (1) \\
 &+ \kappa_i \cdot \left[\frac{1}{2} \cdot (e_1 + xG) + \frac{1}{2} \cdot (e_2 - xL) \right] \\
 &- \frac{1}{2} \cdot \beta_i \cdot [e_1 - e_2 + x(G + L)] \\
 &- \frac{1}{2} \cdot \alpha_i \cdot [e_1 - e_2 + y(G + L)]
 \end{aligned}$$

The term inside the square brackets in the first line is i ’s expected material payoff: with probability $1/2$ i gets the active role, and i ’s decision x determines his material payoff; with probability $1/2$ j gets the active role, and j ’s decision y determines i ’s material payoff. The term inside the square brackets in the second line captures i ’s Kantian moral concern: it is the expected material payoff that i (in fact, any subject in this interaction) would get if—hypothetically—the strategy x was universalized. The parameter $\kappa_i \in [0, 1]$ is i ’s *degree of morality*.¹

¹Mathematically, the Kantian moral concern in *Homo moralis* preferences has a similar effect as the false belief that one’s action affects the action of the opponent, a belief known as magical (or quasi-magical) thinking (Shafir & Tversky, 1992; Daley & Sadowski, 2017). Moreover, under the special case of *Homo moralis* preferences where the degree of morality κ is equal to one, predictions based on Nash equilibrium play sometimes coincide with predictions based on Kantian equilibrium, an equilibrium concept introduced by Roemer (2010). The preference-based approach adopted here presents several advantages: it avoids reliance on false beliefs, it allows for intermediate degrees

The third line measures the effects on i 's utility of being materially ahead of the other when in the active role, while the last line measures the effects on i 's utility of being materially behind the other when in the passive role. The parameter β_i represents i 's *aheadness aversion* (if $\beta_i > 0$) or *love for aheadness* (if $\beta_i < 0$), where aheadness means that i 's monetary payoff is larger than j 's.

Similarly, the parameter α_i represents i 's *behindness aversion* (if $\alpha_i > 0$) or *love for behindness* (if $\alpha_i < 0$), where behindness means that i 's monetary payoff is smaller than j 's. While it is natural to include attitudes towards being both ahead and behind materially, in our experimental design any active decision-maker is always ahead of the other, passive subject; hence we henceforth omit the last term of the utility function.²

To see why the Kantian moral concern is expected to generate different decisions in the VOI and the non-VOI decision situations, consider now a subject in a non-VOI decision situation, cast in the active role. Assuming that this subject is completely unaware of the fact that she could have been allocated to the passive role instead, her utility reduces to:

$$V_i(x, y) = e_1 + xG - \beta_i \cdot [e_1 - e_2 + x(G + L)]. \quad (2)$$

Comparing this with (1), we see that as long as $L > 0$, a subject may select a different x in VOI and non-VOI decisions if $\kappa_i > 0$. Specifically, Kantian morality reduces the subject's willingness to select the selfish decision: from (1) and (2), we see that she is willing to select the selfish decision if

$$G \geq \beta_i(G + L) \quad (3)$$

in the non-VOI decision situation, but only if

$$G - \kappa_i L \geq \beta_i(G + L) \quad (4)$$

of Kantian concern, it allows for the co-existence of other-regarding and Kantian concerns, and it allows for the adoption of the standard Nash equilibrium concept.

²The utility function in (1) is the same as that posited in the main analysis of (Van Leeuwen & Alger, in press). Moreover, such a utility function, which describes the individual's preferences over (potentially trivial) material payoffs, has been shown to be favored by evolution by natural selection (Alger, Weibull, and Lehmann, 2020).

in the VOI decision situation. The VOI decision situation renders explicit the fact that the selfish decision would hurt her if she were cast in the passive rather than the active role and her strategy was universalized to the other individual.

Importantly for our experiment, these two conditions imply that for a given payoff vector $(e_1, e_2, e_1 + G, e_2 - L)$ a subject *makes a switch* from the selfish action in the non-VOI decision to the *status quo* act in the VOI decision situation if her aheadness aversion β_i is not too pronounced,

$$\beta_i \leq \frac{G}{G+L} \equiv z, \quad (5)$$

and her Kantian moral concern κ_i is sufficiently pronounced,

$$\kappa_i \geq \frac{G - \beta_i(G+L)}{L} = \frac{z - \beta_i}{1 - z}. \quad (6)$$

Figure 3 shows these threshold values for two payoff configurations used in the experiment: the solid vertical and oblique lines show, respectively, the threshold values z and $(z - \beta_i)/(1 - z)$ for payoff configuration $(e_1, e_2, e_1 + G, e_2 - L) = (150, 100, 165, 90)$; the dashed lines for payoff configuration $(e_1, e_2, e_1 + G, e_2 - L) = (200, 190, 210, 100)$. To see how variation of the payoffs enables to distinguish between different preference types, note that the two solid lines and the two dashed lines divide the (β, κ) -space into six regions. An individual with a preference type (β_i, κ_i) :

- in region A always selects the selfish option;
- in region B selects the selfish option under payoff (150, 100, 165, 90), but makes a switch under payoff (200, 190, 210, 100);
- in region C makes a switch under both payoffs;
- in region D selects the selfish option under payoff (150, 100, 165, 90), but always the *status quo* option under payoff (200, 190, 210, 100);
- in region E makes a switch under payoff (150, 100, 165, 90), but always selects the *status quo* option under payoff (200, 190, 210, 100);
- in region F always selects the *status quo* option.

In the experiment we use 20 different payoffs $(e_1, e_2, e_1 + G, e_2 - L)$. Table 1 shows, for each of these 20 payoffs, the threshold values z and $(z - \beta)/(1 - z)$ (see (5) and (6)). Note that the 20 payoff configurations used in the experiment entail 20 different values of z . Let Z denote the set of these 20 values. Figure 4 shows the lines




Figure 3: Threshold values z and $(z - \beta_i)/(1 - z)$ for payoffs $(e_1, e_2, e_1 + G, e_2 - L) = (150, 100, 165, 90)$ (solid) and $(e_1, e_2, e_1 + G, e_2 - L) = (200, 190, 210, 100)$ (dashed)

$\kappa = (z - \beta)/(1 - z)$ for the 20 values of z . Except for individuals with $\beta_i \geq 0.6$ (who always select the *status quo* option for all payoffs), and those with $\beta_i \leq 0.03$ and $\kappa_i \leq 0.03 - 1.03\beta_i$ (who always select the selfish option for all payoffs), these payoffs are expected to generate behavioral variation across payoffs and/or across the two conditions (VOI vs. non-VOI).

The payoffs were chosen so as to maximize statistical power in the between-subject estimation of the VOI treatment's effects (conditional on the fact that we decided to have no more than 20 decisions in each sequence), where power computations were based on the results presented by [Van Leeuwen & Alger \(in press\)](#). More precisely, we carried out 1000 simulations where in each one, the 109 subjects for which [Van Leeuwen & Alger \(in press\)](#) estimate individual preference parameters were randomly assigned to either VOI (55 participants) or non-VOI (54 people). Then, we computed their decisions based on the (β, κ) -estimates obtained by [Van Leeuwen & Alger \(in press\)](#), assuming that subjects in the non-VOI treatment were completely unaware of the possibility of role reversal. Finally, we estimated the effect of the VOI treatment on the decision to sell, through the specifications

presented in Section 3. The results of this exercise indicate that the chosen payoffs generate behavioral variations that our main specification is able to detect a significant effect (at a 5% confidence level) 75% of the time.




Figure 4: Each line represents the threshold value $\bar{\kappa}(\beta)$ for one of the 20 payoffs.

In the analysis above we assumed that in a non-VOI decision situation a subject is completely unaware of the complete game tree. However, some subjects may in fact understand that they could have been in their counterpart's place with some non-null probability. The following equation shows the utility of a subject i who assigns a probability \hat{p}_i that she was cast in the Seller role.

$$\begin{aligned}
 \hat{U}_i(x, y) &= (1 - \kappa_i) \cdot \left[\hat{p}_i \cdot (e_1 + xG) + (1 - \hat{p}_i) \cdot (e_2 - yL) \right] & (7) \\
 &+ \kappa_i \cdot \left[\hat{p}_i \cdot (e_1 + xG) + (1 - \hat{p}_i) \cdot (e_2 - xL) \right] \\
 &- \hat{p}_i \cdot \beta_i \cdot [e_1 - e_2 + x(G + L)] \\
 &- (1 - \hat{p}_i) \cdot \alpha_i \cdot [e_1 - e_2 + y(G + L)]
 \end{aligned}$$

The expressions in (1) and (2) correspond to the special cases $\hat{p}_i = 1/2$ and $\hat{p}_i = 1$, respectively. We will say that a subject with $\hat{p}_i = 1$ is *fully unaware*, that one with $\hat{p}_i \in (1/2, 1)$ is *partially aware*, and that one with $\hat{p}_i = 1/2$ is *fully aware* of the full game tree.

Table 1: Payoff combinations (e_1, e_2, e_1+G, e_2-L) (first four columns) and (approximate) threshold values for β and κ as per equations (5) and (6) (last two columns)

Payoff number (s)	e_1	e_2	$e_1 + G$	$e_2 - L$	$z = G/(G + L)$	$(z - \beta_i)/(1 - z)$
1	150	100	165	90	0.6	$1.5 - \beta_i \cdot 2.5$
2	150	100	160	90	0.5	$1 - \beta_i \cdot 2$
3	150	100	165	80	0.43	$0.75 - \beta_i \cdot 1.75$
4	150	100	165	70	0.33	$0.5 - \beta_i \cdot 1.5$
5	250	240	300	100	0.26	$0.36 - \beta_i \cdot 1.36$
6	250	240	300	90	0.25	$0.33 - \beta_i \cdot 1.33$
7	250	240	300	80	0.24	$0.31 - \beta_i \cdot 1.31$
8	250	240	300	70	0.23	$0.29 - \beta_i \cdot 1.29$
9	250	240	300	60	0.22	$0.28 - \beta_i \cdot 1.28$
10	150	120	170	20	0.17	$0.2 - \beta_i \cdot 1.2$
11	200	190	220	60	0.13	$0.15 - \beta_i \cdot 1.15$
12	200	190	220	50	0.125	$0.14 - \beta_i \cdot 1.14$
13	200	190	210	100	0.1	$0.11 - \beta_i \cdot 1.11$
14	200	190	210	90	0.09	$0.1 - \beta_i \cdot 1.1$
15	200	190	210	80	0.08	$0.09 - \beta_i \cdot 1.09$
16	250	220	265	20	0.07	$0.08 - \beta_i \cdot 1.08$
17	250	220	260	50	0.06	$0.06 - \beta_i \cdot 1.06$
18	250	220	260	10	0.05	$0.05 - \beta_i \cdot 1.05$
19	250	220	260	5	0.04	$0.05 - \beta_i \cdot 1.05$
20	250	220	255	60	0.03	$0.03 - \beta_i \cdot 1.03$

A partially aware subject switches from the selfish option in the non-VOI decision situation to the *status quo* option in the VOI decision situation if

$$\beta_i < \frac{G + \frac{1-\hat{p}_i}{\hat{p}_i} \kappa_i L}{G + L} \quad (8)$$

and

$$\kappa_i > \frac{G - \beta_i(G + L)}{L}. \quad (9)$$

Our experiment is based on the premise that in the non-VOI decisions at least some subjects are partially aware or fully unaware, i.e., they hold some belief $\hat{p} > 1/2$, since a fully aware subject would never switch (i.e., conditions (8) and (9) are incompatible). We deliberately chose not to elicit the subjects' beliefs \hat{p} : eliciting them prior to the presentation of the decision situations might have affected the choices, and eliciting them afterwards would make no sense since the subjects would then have been exposed already to VOI decisions, in which they are told explicitly that both role distributions are equally likely.

2.2 Treatments

Each subject is asked to state his/her choice in two sequences, both of which consist of the same set of 20 payoff configurations, listed in Table 1. Each sequence of 20 decisions is either entirely non-VOI or entirely VOI.

Clearly, to avoid raising awareness of the arbitrariness of the role distribution in the non-VOI decision situations, subjects should be presented a non-VOI sequence before being exposed to a VOI sequence. However, always letting a VOI sequence be preceded by a non-VOI one may lead to anchoring and/or fatigue effects. To deal with these issues, and to enable meaningful comparisons between the Neutral and the Market frames, we adopt the following four treatments:

1. **Neutral (N)**: Non-VOI + VOI — neutral frame
2. **Market (M)**: Non-VOI + VOI — market frame
3. **Mixed Frame A (A)**: VOI neutral frame + VOI market frame
4. **Mixed Frame B (B)**: VOI market frame + VOI neutral frame

These treatments allow us to carry out a large number of comparisons: (1) a within-subject comparison of decisions in the non-VOI and VOI sequences in the **Neutral** treatment; (2) a within-subject comparison of decisions in the non-VOI and VOI sequences in the **Market** treatment; (3) a between-subject comparison of decisions in the non-VOI sequence in the **Neutral** treatment and the first sequence in the **Mixed Frame A** treatment; (4) a between-subject comparison of decisions in the non-VOI sequence in the **Market** treatment and the first sequence in the **Mixed Frame B** treatment; and (5) within-subject comparisons of decisions in differently framed VOI sequences, using treatments **Mixed Frame A** and **Mixed Frame B**.

For further use below, let D denote the set of eight decision sequences in our four treatments: $D = \{N1, N2, M1, M2, A1, A2, B1, B2\}$, where a 1 refers to the first sequence and a 2 to the second sequence of 20 decisions in the treatment.

2.3 Procedures and data

We conducted the experiment between November 2021 and March 2022 at the Toulouse School of Economics TSE Lab for Experimental Social Sciences. The software oTree was used to program the experiment, and participants were recruited via email using the Laboratory’s participant pool (people who had signed up to be informed of laboratory experiments). Overall, we recruited 453 participants, all of whom participated in only one session. Tables [A.1](#) and [A.2](#) in Appendix [A1](#) provide information on the total number of subjects in each treatment and associated summary statistics. The experiment (the design and the empirical analysis reported below) was pre-registered on [aspredicted.org](#) on November 21st 2021, the first session taking place on November 22nd.

Each session was allocated to one treatment, and lasted between 25 and 45 minutes. After being randomized to the lab booths, participants read the instructions on their desktop computer screens, were allowed to ask questions privately, and completed a comprehension test. Then the two sequences of 20 decisions were presented, upon which the participants filled out a post-experiment questionnaire (with questions on sex, age, nationality, and field of study). The English version of the instructions are included in Appendix [A3](#) (the experiment was conducted in French). In each decision situation the participants had to answer questions about

the payoffs correctly before being allowed to state their decision (see an example of the two screens shown for one decision in Appendix A3.4).

The participants' payoffs were determined based on two randomly drawn matches, one from each sequence. Each unit of the payoffs used in the instructions (see Table 1) was converted to 2.5 eurocents. For VOI decisions, for the randomly drawn match the role distribution was determined randomly, and the payoffs were calculated based on the active subject's decision. For non-VOI decisions, for the randomly drawn match each subject received the payoff corresponding to their decision, and also the payoff corresponding to the other's decision (of which they were not aware during the experiment). Participants who answered all the questions received a show-up fee of 4 euros. Together with the payoffs obtained from the tasks, participants earned on average 15.5 euros. Participants privately received their payoffs in cash, upon which they left the premises.

Table 2 below and Figures A.1-A.3 in Appendix A1 give a first glimpse of the data, here pooled for all the $n = 453$ participants' decisions across all the treatments and sessions. In the figures we define one observation as the number of selfish options selected by one participant in one sequence of 20 decisions (there are thus two observations per participant). Figure A.1 shows the observations from the Neutral frame on the left and those from the Market frame on the right. The dashed (resp. solid) horizontal lines show the average (resp. median) number of selfish choices. It appears that, overall, the participants were more prone to select the selfish option in the Market than in the Neutral frame. Figure A.2 compares the cumulative distributions of the observations in the VOI decision sequences (solid line, $n = 686$) and the observations in the non-VOI decision sequences (dashed line, $n = 220$). Participants were more prone to select the selfish option in the non-VOI than in the VOI decision situations. Figure A.3 confirms these two tendencies by splitting the observations into the four different frame combinations: Neutral VOI (black line), Market VOI (red line), Neutral non-VOI (blue line), and Market non-VOI (green line). Finally, Table 2 provides the summary statistics on the number of selfish decisions: in the first two rows for the Neutral and the Market frames, followed by two rows for the VOI and non-VOI frames, and finally the four different frame combinations. We see that the differences in the total number of selfish decisions between the Neutral and the Market frame, and between the VOI and non-VOI frame, are significantly

different from zero.

Table 2: Total number of selfish decisions per sequence of 20 decisions, by frame

Frame	N	Mean	Median	Min	Max	Q1	Q3
Neutral	449	7.08	6	0	20	2	10
Market	457	9.03	9	0	20	4	13
t-test and Wilcoxon test p-values		0.00	0.00				
VOI	686	7.35	7	0	20	3	10
non-VOI	220	10.30	10	0	20	6	15
t-test and Wilcoxon test p-values		0.00	0.00				
VOI and Neutral	341	6.52	6	0	20	2	9
VOI and Market	345	8.17	8	0	20	4	11
non-VOI and Neutral	108	8.86	8	0	20	4	13
non-VOI and Market	112	11.7	11	0	20	8	17

3 Hypothesis testing

In this section we perform regressions to test whether the Market and VOI frames have significant effects on behavior in the experiment.

3.1 Does the Market frame significantly affect behavior?

To assess the effects of the Market frame we estimate linear models of the form:

$$x_{izd} = \lambda_0 + \lambda_1 z + \lambda_2 M_{izd} + v_{izd} \quad (10)$$

where $x_{izd} \in \{0, 1\}$ is the decision of subject i in payoff configuration $z = G/(G + L)$ in sequence $d \in \{N1, N2, M1, M2, A1, A2, B1, B2\}$, where $x_{izd} = 1$ if she chooses the selfish option and $x_{izd} = 0$ otherwise; M_{izd} is a dummy variable that takes value 1 if the frame is Market and 0 if the frame is Neutral; and v_{izd} is a mean-zero random variable assumed to be uncorrelated with z and M_{izd} . Our null hypothesis is then that $\lambda_2 = 0$. Recalling also condition (5) for a (fully unaware subject i to choose the selfish option in a non-VOI decision, our theory predicts that *ceteris paribus* an increase

in $z = G/(G + L)$ should lead to a higher propensity to choose the selfish option, and hence we conjecture that $\lambda_1 > 0$. Because the values of z are not uniformly spaced, we also run regressions where instead of including z as an independent variable we include payoff fixed effects.

Table 3: Effect of Frame on Selling under non-VOI conditions - between subjects (N1-M1)

Dependent Variable:	Sell (Yes or No)					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
<i>Variables</i>						
Frame	0.1413*** (0.0149)	0.1413*** (0.0144)	0.1413*** (0.0396)	0.1413*** (0.0149)	0.1413*** (0.0396)	0.1425*** (0.0403)
z		0.8027*** (0.0461)	0.8027*** (0.0655)			
Socio-dem. controls	No	No	No	No	No	Yes
Clustering of SE	No	No	Subject	No	Subject	Subject and Payoff
<i>Fixed-effects</i>						
Payoff				Yes	Yes	Yes
<i>Fit statistics</i>						
Observations	4,400	4,400	4,400	4,400	4,400	4,400
R^2	0.01998	0.08323	0.08323	0.09530	0.09530	0.10272
Within R^2				0.02161	0.02161	0.02964

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: Column (3) includes clustering of SE's at the Subject level, column (4) includes Payoff fixed effects, column (5) includes Payoff fixed effects with clustering of SE's at the Subject level, column (6) includes Payoff fixed effects with clustering of the SE's at the Subject and Payoff level. Socio-demographic controls are: nationality, gender, couple status, previous experience in experiments and attendance to a private school.

Overall, the results reported in Tables 3-5 show that subjects are more prone to selecting the selfish option under the Market than under the Neutral frame, that

this effect is more pronounced in non-VOI than in VOI decisions, and that within-subject effects are stronger than between-subject effects in VOI decisions.

Specifically, Table 3 shows between-subject effects of the Market frame in non-VOI decisions, based on the decisions in the first decision sequences of the Neutral and the Market treatments (N1 and M1). The Market frame reduces the propensity to select the selfish option, an effect which is consistent in magnitude at 14.13 percentage points, and consistently significant across specifications.

Table 4: Effect of Frame on Selling under VOI conditions - between subjects (A1-B1)

Dependent Variable:	Sell (Yes or No)					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
<i>Variables</i>						
Frame	0.0319** (0.0142)	0.0319** (0.0138)	0.0319 (0.0367)	0.0319 (0.0185)	0.0319 (0.0368)	0.0488 (0.0362)
z		0.7354*** (0.0442)	0.7354*** (0.0546)			
Socio-dem. controls	No	No	No	No	No	Yes
Clustering of SE	No	No	Subject	No	Subject	Subject and Payoff
<i>Fixed-effects</i>						
Payoff				Yes	Yes	Yes
<i>Fit statistics</i>						
Observations	4,660	4,660	4,660	4,660	4,660	4,660
R ²	0.00107	0.05714	0.05714	0.06439	0.06439	0.08209
Within R ²				0.00114	0.00114	0.02004

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: Column (3) includes clustering of SE's at the Subject level, column (4) includes Payoff fixed effects, column (5) includes Payoff fixed effects with clustering of SE's at the Subject level, column (6) includes Payoff fixed effects with clustering of the SE's at the Subject and Payoff level. Socio-demographic controls are: nationality, gender, couple status, previous experience in experiments and attendance to a private school.

Table 4 also shows between-subject effects of the Market frame, but now in VOI decisions, based on regressions using only the first decision sequences of Mixed Frame A and Mixed Frame B (A1 and B1): these subjects would thus not have been exposed to the other frame. The effect of the Market frame is much smaller in magnitude than for the non-VOI decisions, and also insignificant as soon as subject-specific and/or payoff-specific effects are accounted for.

Finally, Table 5 shows within-subject effects of the Market frame in VOI decisions, based on decisions in the mixed treatments *A* and *B*. Framing the decision as a market interaction on average increases the likelihood of selecting the selfish option by 7 percentage points, and the significance level is consistently high across specifications. Although large, this effect is about half as large as for the non-VOI decisions.

Note that, as expected, the estimate of λ_1 is positive and highly significant (see columns (2) and (3) in the tables), and that its inclusion substantially increases R^2 . The specifications which instead use “payoff fixed effects” yield similar results.

3.2 Does the VOI frame significantly affect behavior?

To evaluate the effects of the VOI frame, we estimate linear models of the form:

$$x_{izd} = \gamma_0 + \gamma_1 z + \gamma_2 V_{izd} + \epsilon_{izd}. \quad (11)$$

where V_{izd} is a dummy variable that equals 1 if the decision situation is VOI and 0 if the decision is non-VOI, and ϵ_{izd} is a mean-zero random variable, assumed to be uncorrelated with c and V_{izd} . Our null hypothesis is then that $\gamma_2 = 0$. The results, reported in Tables 6-9, show that this hypothesis is clearly rejected, for both within- and between-subject analyses, and that the effect of VOI is large.

Beginning with the neutral frame, Table 6 shows the results based on a within-subjects comparison ($d \in \{N1, N2\}$), while Table 7 reports results from the between-subjects comparison ($d \in \{N1, A1\}$). Independent of the specification, the effect is highly significant, and on average the VOI frame reduces the likelihood of selecting the selfish option by 15.28 percentage points.³ The between-subjects effect is

³This indicates that the within-subjects treatment does not suffer too much from subjects’ desire to act consistently in the VOI and the non-VOI decision situations with the same payoff configuration.

Table 5: Effect of Frame on Selling under VOI conditions - within subjects (A1-A2-B1-B2)

Dependent Variable:	Sell (Yes or No)					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
<i>Variables</i>						
Frame	0.0700*** (0.0100)	0.0700*** (0.0096)	0.0700*** (0.0181)	0.0700*** (0.0096)	0.0700*** (0.0181)	0.0700*** (0.0082)
z		0.8806*** (0.0308)	0.8806*** (0.0513)			
Socio-dem. controls	No	No	No	No	No	Yes
Clustering of SE	No	No	Subject	No	Subject	Subject and Payoff
<i>Fixed-effects</i>						
Sub.			Yes			
Payoff				Yes	Yes	
Sub.-Payoff						Yes
<i>Fit statistics</i>						
Observations	9,320	9,320	9,320	9,320	9,320	9,320
R ²	0.00517	0.08557	0.33028	0.09207	0.09207	0.66812
Within R ²			0.11329	0.00567	0.00567	0.01535

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: Column (3) includes clustering of SE's at the Subject level, column (4) includes Payoff fixed effects, column (5) includes Payoff fixed effects with clustering of SE's at the Subject level, column (6) includes Payoff fixed effects with clustering of the SE's at the Subject and Payoff level. Socio-demographic controls are: nationality, gender, couple status, previous experience in experiments and attendance to a private school.

smaller, at 7.45 percentage points for most specifications, and also less robust in terms of significance across the specifications.

One explanation for the difference between the within- and the between-subjects effect of VOI appears to be heterogeneity across subjects, as indicated by the fact that

Table 6: Effect of VOI on Selling under Neutral frame - within subjects (N1-N2)

DependentVariable:	Sell (Yes or No)					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
<i>Variables</i>						
VOI	-0.1528*** (0.0145)	-0.1528*** (0.0140)	-0.1528*** (0.0277)	-0.1528*** (0.0140)	-0.1528*** (0.0278)	-0.1528*** (0.0122)
z		0.7685*** (0.0447)	0.7685*** (0.0828)			
Clustering of SE	No	No	Subject	No	Subject	Subject*Payoff
<i>Fixed-effects</i>						
Sub.			Yes		Yes	
Payoff				Yes	Yes	
Sub.-Payoff						Yes
<i>Fit statistics</i>						
Observations	4,320	4,320	4,320	4,320	4,320	4,320
R ²	0.02513	0.08750	0.32862	0.09302	0.33415	0.65531
Within R ²			0.11530	0.02696	0.03637	0.06795

Signif. Codes: ***: 0.01, **: 0.05, *: 0.1

Note: Column (3) includes Subject fixed effects with clustering of SE's at the Subject level, column (4) includes Subject and Payoff fixed effects with clustering of SE's at the Subject level, column (5) includes Subject and Payoff fixed effects with clustering of SE's at the Subject level, column (6) includes Subject by Payoff fixed effects with clustering of the SE's at the Subject by Payoff level.

the between-subjects effect is only weakly significant for the specifications which cluster the standard errors at the subject level.

Turning now to the market frame, Tables 8 and 9 report, respectively, how VOI affects the propensity to choose Sell within subjects ($d \in \{M1, M2\}$) and between subjects ($d \in \{M1, B1\}$). Qualitatively, the within-subjects results are similar to those in the neutral frame (see Table 6): the effect is highly significant and identical in

Table 7: Effect of VOI on Selling under Neutral frame - between subjects (N1-A1)

Dependent Variable:	Sell (Yes or No)					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
<i>Variables</i>						
VOI	-0.0745*** (0.0145)	-0.0745*** (0.0141)	-0.0745* (0.0395)	-0.0745*** (0.0143)	-0.0745* (0.0396)	-0.0728* (0.0403)
z		0.6780*** (0.0451)	0.6780*** (0.0586)			
Socio-dem. controls	No	No	No	No	No	Yes
Clustering of SE	No	No	Subject	No	Subject	Subject and Payoff
<i>Fixed-effects</i>						
Payoff				Yes	Yes	Yes
<i>Fit statistics</i>						
Observations	4,580	4,580	4,580	4,580	4,580	4,580
R ²	0.00574	0.05257	0.05257	0.05625	0.05625	0.06045
Within R ²				0.00604	0.00604	0.01047

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: Column (3) includes clustering of SE's at the Subject level, column (4) includes Payoff fixed effects, column (5) includes Payoff fixed effects with clustering of SE's at the Subject level, column (6) includes Payoff fixed effects with clustering of the SE's at the Subject and Payoff level. Socio-demographic controls are: nationality, gender, couple status, previous experience in experiments and attendance to a private school.

size across all specifications. However, the effect is even stronger than in the neutral frame: the VOI wording reduces the propensity to sell by 18.66 percentage points. Contrary to the neutral frame, here the between-subjects effect is similar in size to the within-subjects effect, and it is also robustly significant across the specifications. Overall there is thus a large and robust effect of the VOI wording on the propensity to choose Sell in the market frame.

Table 8: Effect of VOI on Selling under Market frame - within subjects (M1-M2)

Dependent Variable:	Sell (Yes or No)					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
<i>Variables</i>						
VOI	-0.1866*** (0.0147)	-0.1866*** (0.0138)	-0.1866*** (0.0265)	-0.1866*** (0.0138)	-0.1866*** (0.0265)	-0.1866*** (0.0112)
z		1.042*** (0.0442)	1.042*** (0.0913)			
Clustering of SE	No	No	Subject	No	Subject	Subject*Payoff
<i>Fixed-effects</i>						
Sub.			Yes		Yes	
Payoff				Yes	Yes	
Sub.-Payoff						Yes
<i>Fit statistics</i>						
Observations	4,480	4,480	4,480	4,480	4,480	4,480
R ²	0.03483	0.14142	0.36707	0.15022	0.37586	0.72045
Within R ²			0.18263	0.03938	0.05286	0.11080

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: Column (3) includes Subject fixed effects with clustering of SE's at the Subject level, column (4) includes Subject and Payoff fixed effects with clustering of SE's at the Subject level, column (5) includes Subject and Payoff fixed effects with clustering of SE's at the Subject level, column (6) includes Subject by Payoff fixed effects with clustering of the SE's at the Subject by Payoff level.

As was the case for the regressions measuring the effects of the Market frame, the estimate of the coefficient for z is positive and highly significant (see columns (2) and (3) in the tables), and its inclusion substantially increases R². Moreover, the results are similar in the specifications which instead use “payoff fixed effects”.

Table 9: Effect of VOI on Selling under Market frame - between subjects (M1-B1)

Dependent Variable:	Sell (Yes or No)					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
<i>Variables</i>						
VOI	-0.1839*** (0.0147)	-0.1839*** (0.0141)	-0.1839*** (0.0367)	-0.1839*** (0.0220)	-0.1839*** (0.0368)	-0.1879*** (0.0402)
z		0.8602*** (0.0451)	0.8602*** (0.0609)			
Socio-dem. controls	No	No	No	No	No	Yes
Clustering of SE	No	No	Subject	No	Subject	Subject and Payoff
<i>Fixed-effects</i>						
Payoff				Yes	Yes	Yes
<i>Fit statistics</i>						
Observations	4,480	4,480	4,480	4,480	4,480	4,480
R ²	0.03384	0.10642	0.10642	0.11804	0.11804	0.14127
Within R ²				0.03695	0.03695	0.06231

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: Column (3) includes clustering of SE's at the Subject level, column (4) includes Payoff fixed effects, column (5) includes Payoff fixed effects with clustering of SE's at the Subject level, column (6) includes Payoff fixed effects with clustering of the SE's at the Subject and Payoff level. Socio-demographic controls are: nationality, gender, couple status, previous experience in experiments and attendance to a private school.

3.3 Interaction effects

The results reported above show that the VOI frame has a stronger absolute effect on the propensity to select the selfish option in the Market than in the Neutral frame. However, from the descriptive statistics we also know that subjects are on average more willing to choose the selfish option in the Market than in the Neutral frame. Here we test whether the relative effect of VOI is different under the two frames, by running regressions with a VOI-Market interaction, as follows:

$$x_{izd} = \gamma_0 + \gamma_1 z + \gamma_2 V_{izd} + \gamma_3 V_{izd} \times M_{izd} + \epsilon_{izd}. \quad (12)$$

The results are reported in Table A.5 in Appendix A1. The coefficient on the interaction $\text{VOI} \times \text{Market}$ is not significant, implying that, on average, the effect of VOI relative to non-VOI is the same in the Market as in the Neutral frame.

4 Structural estimation of β and κ

Besides the reduced-form estimates of treatment effects specified above, we structurally estimate the aheadness aversion and Kantian morality parameters β and κ .

4.1 Method

We run the estimations using the data obtained from subjects exposed to both non-VOI and VOI sequences. We will first report the estimates of the preference parameters, $(\hat{\beta}, \hat{\kappa})$, and the choice sensitivity, $\hat{\sigma}$, under the hypothesis that the observed data pooled over all the subjects (in a given treatment) had emanated from a representative agent. For non-VOI decisions we assume that any subject is fully unaware of the possible role reversal, i.e., we posit $\hat{p} = 1$ (see (7)). By contrast, for VOI decisions we assume that any subject is fully aware of the arbitrariness of the role distribution (by positing $\hat{p} = 1/2$), since we explicitly inform the subjects about it. These assumptions are extreme, and we do not claim that they are in line with the subjects' subjective beliefs about the role distribution. However, the advantage is that they imply a conservative estimate of $\hat{\kappa}$. Indeed, under the assumption that $\hat{p} = 1$ in non-VOI decisions, only a positive $\hat{\beta}$ can explain why a subject would refrain from selecting the selfish option. Hence, for any subject who is in fact fully or

partly aware of the role reversal in the non-VOI decisions (i.e., with $\hat{p} = 1 \in [1/2, 1)$) and who refrains from the selfish option in such a decision due to a combination of aheadness attitude and Kantian morality (see (8)), our assumption that $\hat{\kappa}$ cannot be at work leads to over-estimation of $\hat{\beta}$ and under-estimation of $\hat{\kappa}$.

We use a standard random utility model, by assuming that the representative agent's true utility from using pure strategy x is a random variable of the form

$$\tilde{U}(x; \beta, \kappa, \sigma, z) = 2U(x; \beta, \kappa, z) + \varepsilon_x \quad (13)$$

in a VOI decision situation with payoff configuration $z \in Z$ and

$$\tilde{V}(x; \beta, \sigma, z) = V(x; \beta, z) + \varepsilon_x \quad (14)$$

in a non-VOI decision situation with payoff configuration $z \in Z$, where the function $U(x; \beta, \kappa, z)$ is that specified in (1), the function $V(x; \beta, z)$ is that specified in (2), and ε_x represents randomness in the utility evaluation, assumed to be independent across all the decisions. This random variable is taken to follow a type 1 extreme value distribution with scale parameter $1/\sigma$. The factor 2 in equation (13) avoids that the magnitude of noise relative to z , β , and κ be twice as large for VOI than for non-VOI decisions (recall the factor 1/2 in (1)).

In a VOI decision situation, the selfish option is selected if

$$\tilde{U}(1; \beta, \kappa, \sigma, z) \geq \tilde{U}(0; \beta, \kappa, \sigma, z), \quad (15)$$

or

$$2U(1; \beta, \kappa, z) - 2U(0; \beta, \kappa, z) \geq \varepsilon_0 - \varepsilon_1, \quad (16)$$

which reduces to

$$G - \kappa L - \beta(G + L) \geq \varepsilon_0 - \varepsilon_1. \quad (17)$$

Likewise, in a non-VOI decision situation, the selfish option is selected if

$$\tilde{V}(1; \beta, \sigma, z) \geq \tilde{V}(0; \beta, \sigma, z), \quad (18)$$

or

$$V(1; \beta, z) - V(0; \beta, z) \geq \varepsilon_0 - \varepsilon_1, \quad (19)$$

which reduces to

$$G - \beta(G + L) \geq \varepsilon_0 - \varepsilon_1. \quad (20)$$

Letting ν be a dummy variable that takes the value 1 in a VOI decision situation and the value 0 in a non-VOI decision situation, it follows that the probability that the selfish option is selected in payoff configuration z equals (McFadden, 1974):

$$\frac{1}{1 + \exp \left[-\sigma \left[\nu [(1 - \beta)G - (\beta + \kappa)L] + (1 - \nu) [(1 - \beta)G - \beta L] \right] \right]}, \quad (21)$$

where σ is the representative agent's choice sensitivity with respect to differences in deterministic utility. If σ is close to 0, the agent chooses either option with a probability close to 1/2, independent of the difference in deterministic utilities. In contrast, arbitrarily large values of σ indicate that the probability of choosing the option that results in the higher deterministic utility approaches 1. Noting that the expression in (21) can be written as:

$$H \left(\sigma \left[\nu [(1 - \beta)G - (\beta + \kappa)L] + (1 - \nu) [(1 - \beta)G - \beta L] \right] \right),$$

where $H : \mathbb{R} \rightarrow (0, 1)$ is the logistic function, the probability density function can be written

$$\begin{aligned} f(\mathbf{x}, \beta, \kappa, \sigma) &= \prod_{i \in I} \prod_{z_i \in Z} H \left(\sigma \left(U(x^{(z_i)} = 1; \beta, \kappa, z_i) - U(x^{(z_i)} = 0; \beta, \kappa, z_i) \right) \right)^{\mathbb{1}_{\{x^{(z_i)} = 1\}}} \\ &\times \left[1 - H \left(\sigma \left(U(x^{(z_i)} = 1; \beta, \kappa, z_i) - U(x^{(z_i)} = 0; \beta, \kappa, z_i) \right) \right) \right]^{1 - \mathbb{1}_{\{x^{(z_i)} = 1\}}}, \end{aligned} \quad (22)$$

where I is the set of subjects, and z_i refers to the situation in which subject i faced payoff z : this notation addresses the fact that here all the subjects' decisions are pooled and treated as if taken by one single individual. In the expression, \mathbf{x} is the vector of observed choices, and $\mathbb{1}_{\{x^{(z_i)} = 1\}}$ is an indicator function that takes the value 1 if $x^{(z_i)} = 1$ and 0 if $x^{(z_i)} = 0$. Using (22), we obtain the log-likelihood

$$\ln(L(\beta, \kappa, \sigma; \mathbf{x})) = \log(f(\mathbf{x}, \beta, \kappa, \sigma)), \quad (23)$$

and we will report the vector of estimates $(\hat{\beta}, \hat{\kappa}, \hat{\sigma})$ that maximizes this log-likelihood.

Second, we use the finite-mixture approach used by [Bruhin, Fehr, & Schunk \(2018\)](#) and [Van Leeuwen & Alger \(in press\)](#) to evaluate whether the data is reasonably described by a finite set of estimated preference types, and if so, whether the Kantian moral concern appears relevant for most of them. This approach consists in

estimating preference parameters for a given number K of “preference types”. For each type $k = \{1, \dots, K\}$, we estimate the preference parameters (β_k, κ_k) and the noise parameter σ_k . The log-likelihood is then given by:

$$\ln(L(\boldsymbol{\beta}, \boldsymbol{\kappa}, \boldsymbol{\sigma}; \mathbf{x})) = \ln \left(\sum_{k=1}^K \phi_k \cdot f(\beta_k, \kappa_k, \sigma_k; \mathbf{x}) \right), \quad (24)$$

where ϕ_k is the share of subjects of type k , estimated together with the preference and noise parameters.

All structural estimations of the preference parameters β and κ are conducted both for the Neutral frame and for the Market frame, to examine whether any differences appear. Furthermore, we carry out all our estimation procedures for the full samples and for two subsets of the latter, which we deem “Core 1” and “Core 2”. Our Core 1 samples consist of the 98 subjects in the Neutral treatment and 100 subjects in the Market treatment who made at least one decision switch between non-VOI and VOI. Recalling that our theoretical model predicts that a subject should be less inclined to select the selfish decision in VOI situations, we say that a subject switches in the expected direction if (s)he selects the selfish option under non-VOI and the *status quo* option under VOI, and in the unexpected direction if (s)he selects the *status quo* option under non-VOI and the selfish option under VOI. Thus, for our Core 2 samples we remove subjects who made at least as many unexpected as expected switches from the Core 1 samples. This leaves us with 68 subjects in the Neutral treatment and 79 subjects in the Market treatment. In the following section, all reported estimates are obtained using the Core 1 samples, while results with the full and Core 2 samples can be found in Appendix A1.

4.2 Results

4.2.1 Representative agent estimates

Table 10 shows the representative agent estimates of the preference parameters β and κ , as well as that of the choice sensitivity parameter σ for our Core 1 sample for the Neutral frame (first column) and the Market frame (second column).

Both the aheadness aversion and the Kantian morality estimates significantly differ from zero and are positive. They are also both higher in the Neutral than in the

Table 10: Estimated preferences for the representative agent in Neutral and Market frames - Core 1 sample

	Neutral	Market	H_0 : Neutral = Market ⁺
β : Aheadness aversion	0.194*** (0.030) - [0.012]	0.099*** (0.020) - [0.008]	0.009
κ : Degree of morality	0.258*** (0.075) - [0.032]	0.228*** (0.037) - [0.018]	0.721
σ : Choice sensitivity	0.295*** (0.009) - [0.003]	0.040*** (0.010) - [0.003]	0.295
Number of Observations	3,600	3,720	
Number of subjects	90	93	
Log likelihood	-2,196.136	-2,255.419	

Notes: Standard errors clustered at the individual level in parentheses. Non-clustered standard errors in brackets. *** Significant at 1% using clustered standard errors.

⁺ p -value of z-test, using clustered standard errors.

Market frame: $\beta = 0.194$ and $\kappa = 0.258$ in the Neutral frame while $\beta = 0.099$ and $\kappa = 0.228$ in the Market frame. The absolute difference is thus larger for the aheadness aversion estimates than for the Kantian morality estimates, and the difference is statistically different only for the aheadness aversion estimates (β) as shown in the third column.

The estimates based on the full sample and the Core 2 samples, reported in Tables A.6-A.7 in Appendix A1, are qualitatively similar, except that for the Core 2 sample, for which the estimated aheadness aversion parameter is not significantly different under the Market than under the Neutral frame.⁴

⁴This result is partly driven by the clustering of the standard errors at the individual level, as they generate wider confidence intervals compared to non-clustered standard errors.

4.2.2 Finite-mixture estimates

Table 11 presents the finite mixture estimates using a two-type model ($K = 2$ in (24)), as well as the shares of subjects who are classified either as Type 1 or as Type 2, depending on the type that is more aligned with their choices.⁵ It is worth underscoring that both types in either frame show some degree of Kantian morality, which is in line with the significant VOI treatment effects presented in Section 3.

Table 11: Two-type Finite Mixture Model estimates, Core 1 sample

	Neutral		Market	
	Type 1	Type 2	Type 1	Type 2
β : Aheadness aversion	0.327*** (0.082) [0.023]	-0.065 (0.258) [0.049]	0.203*** (0.015) [0.007]	-0.143 (0.117) [0.046]
κ : Degree of morality	0.116** (0.049) [0.029]	0.342* (0.199) [0.070]	0.153*** (0.023) [0.017]	0.325*** (0.072) [0.056]
σ : Choice sensitivity	0.046*** (0.009) [0.005]	0.025** (0.012) [0.004]	0.068*** (0.013) [0.005]	0.030*** (0.010) [0.004]
Share	0.554*** (0.173) [0.058]	0.446*** (0.173) [0.058]	0.560*** (0.056) [0.052]	0.440*** (0.056) [0.052]
Observations	3,920		4,000	
Number of subjects	98		100	
Log likelihood	-2,155.558		-2,193.214	

Notes: Standard errors clustered at the individual level in parenthesis, non-clustered standard errors in brackets. *** Significant at 1%, ** Significant at 5%, * Significant at 10% using clustered standard errors.

⁵The estimates based on the full sample and the Core 2 samples are reported in Tables A.8-A.9 in Appendix A1.

Starting with the Neutral frame, Type 1, which accounts for 55.4% of the sample, displays aheadness aversion that is almost twice that of the representative agent (0.327 compared to 0.194), but a Kantian morality that is only about half as large as the representative agent’s (0.116 compared to 0.238). By contrast, Type 2 (with a share of 44.6%) seems to be indifferent to advantageous inequality (notice that the -0.065 estimate is not significant at any of the usual confidence levels) and exhibits a much larger Kantian morality estimate than the representative agent (0.342 versus 0.258).

Turning now to the Market frame, a similar pattern emerges. We again see a type (Type 1) who is about twice as aheadness averse as the representative agent -0.203 compared with 0.099– but less morally concerned, with a Kantian morality estimate of 0.153 (versus 0.228 for the representative agent). In addition, we estimates that 56% of subjects in the Market frame belong to Type 1. Type 2 appears to not display aheadness liking nor aversion (the negative β estimate of -0.143 is once more not statistically significant at the usual levels). Meanwhile, the Kantian morality estimate of this second type is larger than that of the representative agent (0.325 compared to 0.228). Here, 44% of the subjects are classified as this latter type.

The estimated parameter values described above allow us to endogenously classify each subject into the preference type that best fits her observed behaviour. More precisely, using Bayes’ rule, a subject i ’s posterior probability of belonging to Type $k = 1, 2$ is given by:

$$\tau_{ik} = \frac{\hat{\pi}_k f(\hat{\beta}_k, \hat{\kappa}_k, \hat{\sigma}_k; \mathbf{x}_i)}{\sum_{m=1}^2 \hat{\pi}_m f(\hat{\beta}_m, \hat{\kappa}_m, \hat{\sigma}_m; \mathbf{x}_i)}. \quad (25)$$

These probabilities indicate the type with which the participant’s behavior is the most compatible. A model that manages to capture the preferences and share of the underlying types should give way to most posterior probabilities being very close to 0 or to 1. Moreover, the share of subjects classified as belonging to a given type according to these probabilities should be similar to that same type’s share estimated by the model. Figure 5 shows that this is indeed the case here.

According to our theory, subjects who are classified as Type 1 should display behavior that is more consistent across the non-VOI and VOI-treatments, as indicated by the relatively low value of κ . By contrast, those who are classified as Type 2 ought to exhibit less consistent behavior across the non-VOI and VOI-treatments, as sug-

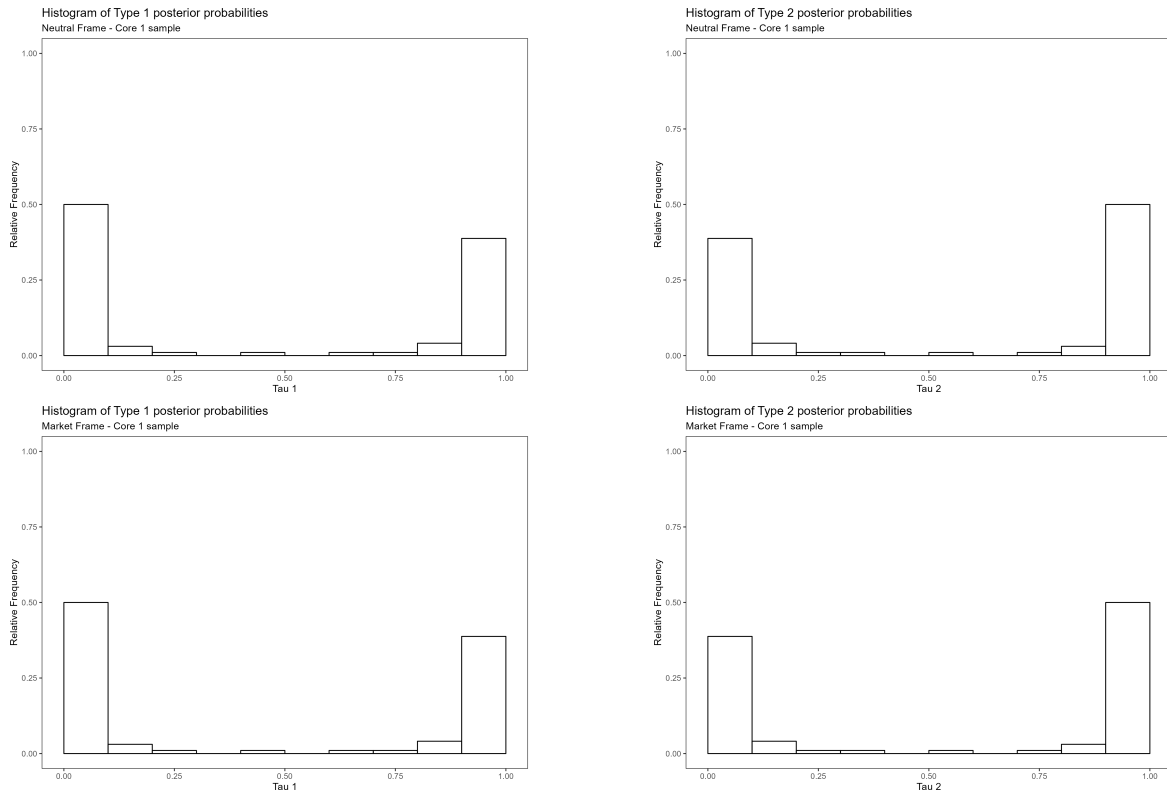



Figure 5: Distribution of posterior probabilities of individual type-membership in Neutral (first row) and Market frames (second row). Core 1 sample estimates.


gested by the high value of κ . This is clearly seen in Figure 6, where we show that the difference in the number of selfish actions in the non-VOI and VOI treatments is more concentrated around zero for participants classified as Type 1.⁶

We do not report estimates based on more than two types, since our experimental design is better suited for estimation at the aggregate level (i.e., treating the observations as if they emanated from one or a small number of preference types) than for estimation of each individual’s preference type (β_i, κ_i) (we refer to Appendix A2 for a detailed explanation).

⁶We refer to Figures A.4-A.7 in Appendix A1 for the corresponding analysis for the full and Core 2 samples.



(a) Neutral frame



(b) Market frame

Figure 6: Difference in selfish actions non-VOI versus VOI. Core 1 sample. Subjects classified as Type 1 if $\tau_1 > 0.95$ or Type 2 if $\tau_1 < 0.05$

5 Concluding remarks

In order to issue effective and desirable policy recommendations, economists need to achieve a realistic understanding of the motivations of human behavior. Indeed, insights about the functioning of markets and institutions derived in theoretical models where humans are driven solely by material self-interest may be misleading. We designed an experiment aimed at evaluating the willingness of subjects to sell a lemon – i.e., a good that makes the buyer worse off – and to disentangle the role of other-regarding concerns and Kantian moral concerns in this decision. Our design further allows us to structurally estimate preference parameters, and to test whether the obtained estimates differ between decisions taken in a frame that describes the situation as the sale of a lemon – the market frame – and those taken in a standard neutral frame.

The subjects in our experiment were on average much more prone to “sell lemons” in the market frame than they were to select the payoff-equivalent decisions in the neutral frame. The propensity to sell lemons was, however, strongly mitigated when the instructions made explicit reference to the arbitrariness in the role distribution. This result is in line with our theoretical model, according to which such instructions are expected to awaken or strengthen subjects’ Kantian moral concerns. As per our structural estimations of this Kantian moral concern and the aversion towards being ahead materially compared to the other subject in a match, the difference in behavior between the market and the neutral frame was driven mainly by a weaker aheadness aversion in the market frame.

Our results are in line with those of other studies that have also shown that preferences that combine a Kantian moral concern with other-regard significantly enhances the explanatory power compared to preferences without the Kantian moral concern (Miettinen, Kosfeld, Fehr, & Weibull, 2020; Van Leeuwen & Alger, *in press*). Like Van Leeuwen & Alger (*in press*), who also estimate aheadness aversion and Kantian moral preference parameters, we find estimates which indicate that some individuals appear to be driven by a combination of altruism (a positive β) and a Kantian moral concern, and all the estimates of the Kantian moral concern parameter are positive and significantly so. Interestingly, under the Market frame some individuals appear to be driven by a combination of spite (a negative β) and a Kan-

tian moral concern. This suggests that in our experiment the Market frame may have triggered more competitive preferences.

Our study is the first to explore – and find support for – the hypothesis that the Kantian moral concern is (at least somewhat) muted in situations where role uncertainty is not made explicit. If externally valid, this result would mean that the mere mention of role uncertainty in situations outside of the experimental laboratory could be used to awaken Kantian moral concerns, and thus trigger behavioral changes. In this line, [Bowles \(2023\)](#) suggests that indeed, in order to improve the functioning of markets, policies ought to combine both incentives and *moral messages*, leveraging synergies between the two ([Kranton, 2019](#)). Field experiments might be appropriate to examine this question in the future.

Our results further strongly suggest that both frames in our experiment matter. On average decisions are more selfish under the Market than under the Neutral frame, and the difference is large. However, the VOI wording also has a stronger dampening effect on selfish actions under the Market frame. Accordingly the estimates of the preferences based on subjects' decisions in market interactions differ from those based on decisions in interactions with similar payoff consequences but presented with a neutral frame.

Such disparities in preferences across different social situations would be in line with the theory of preference evolution, which predicts that as long as humans in our evolutionary past could discriminate between classes of interactions (say, market vs non-market), preferences would be expected to depend on the specifics of each class of interaction (for surveys, see [Alger & Weibull \(2019\)](#) and [Alger \(2023\)](#)). As such, they strengthen the more general case made elsewhere for further research on how the language used to describe a task matters for subjects' decisions ([Alekseev, Charness, & Gneezy, 2017](#); [Capraro, Halpern, & Perc, 2024](#)).

As a final note, it is worth mentioning that recent theoretical studies highlight the role that Kantian moral concerns can play in mitigating inefficiencies in equilibrium outcomes in bilateral trade situations plagued by information asymmetries ([Rivero-Wildemauwe, 2023a,b](#)). While our study focuses solely on the behavior of potential “lemon” sellers, by shutting down the buyers' decisions, these references find that morality decreases or even eliminates the exchange of “lemons” *in equilibrium*. Moreover, they document that altruism has effects that go in the same di-

rection, although they are generally weaker. It would thus be worth extending our present work to capture interactions between both sides of the market. With both sellers and buyers being active, beliefs should also play a role; the evidence that preferences may be affected by framing can help explain why beliefs would also be affected by framing ([Dufwenberg, Gächter, & Hennig-Schmidt, 2011](#)).

References

- Abeler, J., Nosenzo, D., & Raymond, C. (2019). Preferences for truth-telling. *Econometrica*, 87(4), 1115–1153.
- Agneman, G., & Chevrot-Bianco, E. (2022). Market participation and moral decision-making: Experimental evidence from Greenland. *Economic Journal*, 133(650), 537-581.
- Akerlof, G. A. (1970). The market for “lemons”: Quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3), 488–500.
- Alekseev, A., Charness, G., & Gneezy, U. (2017). Experimental methods: When and why contextual instructions are important. *Journal of Economic Behavior Organization*, 134, 48-59.
- Alger, I. (2023). Evolutionarily stable preferences. *Philosophical Transactions B*, 378, 20210505.
- Alger, I., & Weibull, J. W. (2013). Homo moralis—preference evolution under incomplete information and assortative matching. *Econometrica*, 81(6), 2269–2302.
- Alger, I., & Weibull, J. W. (2019). Evolutionary models of preference formation. *Annual Review of Economics*, 11, 329–354.
- Alger, I., Weibull, J. W., & Lehmann, L. (2020). Evolution of preferences in structured populations: Genes, guns, and culture. *Journal of Economic Theory*, 185, 104951.
- Andreoni, J., & Miller, J. H. (1993). Rational cooperation in the finitely repeated prisoner’s dilemma: Experimental evidence. *Economic Journal*, 103(418), 570–585.
- Bartling, B., & Özdemir, Y. (2023). The limits to moral erosion in markets: Social norms and the replacement excuse. *Games and Economic Behavior*, 138, 143–160.
- Becker, G. S. (1974). A theory of social interactions. *Journal of Political Economy*, 82(6), 1063–1093.

- Bénabou, R., Falk, A., & Henkel, L. (2024). *Ends versus means: Kantians, utilitarians, and moral decisions* (NBER Working Paper 32073).
- Bénabou, R., Falk, A., Henkel, L., & Tirole, J. (2023). *Eliciting moral preferences: Theory and experiment* (CRC TR 224 Discussion Paper Series).
- Bénabou, R., & Tirole, J. (2006). Incentives and prosocial behavior. *American Economic Review*, 96(5), 1652-1678.
- Bowles, S. (1998). Endogenous preferences: The cultural consequences of markets and other economic institutions. *Journal of Economic Literature*, 36(1), 75–111.
- Bowles, S. (2023). Moral economics. *Fiscal Studies*, 44(2), 151-160.
- Brandts, J., & Charness, G. (2011). The strategy versus the direct-response method: A first survey of experimental comparisons. *Experimental Economics*, 14(3), 375–398.
- Bruhin, A., Fehr, E., & Schunk, D. (2018). The many faces of human sociality: Uncovering the distribution and stability of social preferences. *Journal of the European Economic Association*, 17(4), 1025–1069.
- Bursztyn, L., Fiorin, S., Gottlieb, D., & Kanz, M. (2019). Moral incentives in credit card debt repayment: Evidence from a field experiment. *Journal of Political Economy*, 127(4), 1641–1683.
- Cabrales, A., Miniaci, R., Piovesan, M., & Ponti, G. (2010). Social preferences and strategic uncertainty: An experiment on markets and contracts. *American Economic Review*, 100(5), 2261–2278.
- Capraro, V., Halpern, J. Y., & Perc, M. (2024, March). From outcome-based to language-based preferences. *Journal of Economic Literature*, 62(1), 115-54.
- Capraro, V., & Rand, D. G. (2018). Do the right thing: Experimental evidence that preferences for moral behavior, rather than equity or efficiency per se, drive human prosociality. *Judgment and Decision Making*.
- Charness, G., & Rabin, M. (2002). Understanding social preferences with simple tests. *Quarterly Journal of Economics*, 117(3), 817–869.

- Chen, D. L., & Schonger, M. (2022). Social preferences or sacred values? Theory and evidence of deontological motivations. *Science Advances*, 8(19).
- Daley, B., & Sadowski, P. (2017). Magical thinking: A representation result. *Theoretical Economics*, 12(2), 909–956.
- Dreber, A., Ellingsen, T., Johannesson, M., & Rand, D. G. (2013). Do people care about social context? Framing effects in dictator games. *Experimental Economics*, 16(3), 349–371.
- Dufwenberg, M., Gächter, S., & Hennig-Schmidt, H. (2011). The framing of games and the psychology of play. *Games and Economic Behavior*, 73(2), 459–478.
- Dufwenberg, M., Johansson-Stenman, O., Kirchler, M., Lindner, F., & Schwaiger, R. (2022). Mean markets or kind commerce? *Journal of Public Economics*, 209, 104648.
- Engel, C. (2011). Dictator games: A meta study. *Experimental Economics*, 14(4), 583–610.
- Engel, J., & Szech, N. (2020). A little good is good enough: Ethical consumption, cheap excuses, and moral self-licensing. *PLoS ONE*, 15(1), 1–19.
- Feess, E., Kerzenmacher, F., & Timofeyev, Y. (2022). Utilitarian or deontological models of moral behavior—What predicts morally questionable decisions? *European Economic Review*, 149(September 2021), 104264.
- Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics*, 114(3), 817–868.
- Fisman, R., Kariv, S., & Markovits, D. (2007). Individual preferences for giving. *American Economic Review*, 97(5), 1858–1876.
- García-Pola, B., Iriberry, N., & Kovářík, J. (2020). Hot versus cold behavior in centipede games. *Journal of the Economic Science Association*, 6(2), 226–238.
- Gneezy, U., Rockenbach, B., & Serra-Garcia, M. (2013). Measuring lying aversion. *Journal of Economic Behavior and Organization*, 93, 293–300.

- Grech, P. D., Nax, H. H., & Soos, A. (2022). Incentivization matters: A meta-perspective on dictator games. *Journal of the Economic Science Association*, 8(1), 34–44.
- Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., Gintis, H., ... Tracer, D. (2005). "Economic man" in cross-cultural perspective: Behavioral experiments in 15 small-scale societies. *Behavioral and Brain Sciences*, 28(06), 795–855.
- Henrich, J., Ensminger, J., McElreath, R., Barr, A., Bolyanatz, A., Cardenas, J. C., ... Ziker, J. (2016). Markets, religion, community size, and the evolution of fairness and punishment. *Science*, 327(5972), 1480–1484.
- Herne, K., Hietanen, J. K., Lappalainen, O., & Palosaari, E. (2022). The influence of role awareness, empathy induction and trait empathy on dictator game giving. *PLoS ONE*, 17(3), 1–19.
- Huang, K., Greene, J. D., & Bazerman, M. (2019). Veil-of-ignorance reasoning favors the greater good. *Proceedings of the National Academy of Sciences of the United States of America*, 116(48), 23989–23995.
- Iriberri, N., & Rey-Biel, P. (2011). The role of role uncertainty in modified dictator games. *Experimental Economics*, 14(2), 160–180.
- Kay, A. C., & Ross, L. (2003). The perceptual push: The interplay of implicit cues and explicit situational construals on behavioral intentions in the prisoner's dilemma. *Journal of Experimental Social Psychology*, 39(6), 634–643.
- Kirman, A., & Teschl, M. (2010). Do markets foster selfishness? *Revue de philosophie économique*, Vol. 11(1), 113–140.
- Kranton, R. (2019). The devil is in the details: Implications of Samuel Bowles's *The Moral Economy* for economics and policy research. *Journal of Economic Literature*, 57(1), 147–160.
- Levine, D. K. (1998). Modeling altruism and spitefulness in experiments. *Review of Economic Dynamics*, 1(3), 593–622.

- Levine, S., Kleiman-Weiner, M., Schulz, L., Tenenbaum, J., & Cushman, F. (2020). The logic of universalization guides moral judgment. *Proceedings of the National Academy of Sciences of the United States of America*, 117(42), 26158–26169.
- Lieberman, V., Samuels, S. M., & Ross, L. (2004). The name of the game: Predictive power of reputations versus situational labels in determining Prisoner's Dilemma game moves. *Personality and Social Psychology Bulletin*, 30(9), 1175–1185.
- Miettinen, T., Kosfeld, M., Fehr, E., & Weibull, J. W. (2020). Revealed preferences in a sequential prisoners' dilemma: A horse-race between six utility functions. *Journal of Economic Behavior and Organization*, 173, 1-25.
- Ortiz-Riomalo, J. F., Koessler, A. K., & Engel, S. (2021). Inducing perspective-taking for prosocial behaviour in natural resource management. *Journal of Environmental Economics and Management*, 110(August), 102513.
- Rivero-Wildemauwe, J. I. (2023a). *Moral motivations in sequential buyer-seller interactions with adverse selection* (Tech. Rep.). THEMA (THéorie Economique, Modélisation et Applications), CY Cergy Paris Université.
- Rivero-Wildemauwe, J. I. (2023b). *Trade among moral agents with information asymmetries* (Tech. Rep.). THEMA (THéorie Economique, Modélisation et Applications), CY Cergy Paris Université.
- Roemer, J. E. (2010). Kantian equilibrium. *Scandinavian Journal of Economics*, 112(1), 1–24.
- Roth, B. A. E., & Schoumaker, F. (1983). Expectations and reputations in bargaining: An experimental study. *American Economic Review*, 73(3), 362–372.
- Shafir, E., & Tversky, A. (1992). Thinking through uncertainty: Nonconsequential reasoning and choice. *Cognitive Psychology*, 24(4), 449–474.
- Sutter, M., Huber, J., Kirchler, M., Stefan, M., & Walzl, M. (2020). Where to look for the morals in markets? *Experimental Economics*, 23(1), 30–52.
- Sutter, M., & Weck-Hannemann, H. (2003). Taxation and the veil of ignorance - A real effort experiment on the Laffer curve. *Public Choice*, 115(1-2), 217–240.

Thöni, C., & Gächter, S. (2015). Peer effects and social preferences in voluntary cooperation: A theoretical and experimental analysis. *Journal of Economic Psychology*, 48, 72-88.

Van Leeuwen, B., & Alger, I. (in press). Estimating social preferences and Kantian morality in strategic interactions. *Journal of Political Economy Microeconomics*.

Zelmer, J. (2003). Linear public goods experiments: A meta-analysis. *Experimental Economics*, 6(3), 299–310.

Appendix A1 Additional tables and figures

Table A.1: Participants' socio-demographic characteristics

Variable	N	Mean	Variable	N	Mean
Age	453	21	Education	413	
Gender	453		... Highschool	37	8%
... Female	193	43%	... Undergrad.	272	61%
... Male	260	57%	... Engineer	5	1%
Background	418		... Master 1	68	15%
... Economics	339	81%	... Master 2	44	10%
... Sciences, Eng.	10	2%	... PhD	3	1%
... Political Science	47	11%	... Other	19	4%
... Languages	8	2%	Experience	453	
... Literature, Philo.	11	3%	... Yes	140	31%
... History, Geog.	3	1%	... No	313	69%
Nationality	453		Couple	453	
... French	362	80%	... Yes	66	15%
... Other	91	20%	... No	387	85%

For Age, Std. Dev. = 3.9, Min = 18, Pctl. 25 = 19, Pctl. 75 = 22, Max = 54. Five participants did not report their education level, while 35 did not report their background.

Table A.2: Participants' socio-demographic characteristics

Treatment	N		M		A		B		Test
Variable	N	Mean	N	Mean	N	Mean	N	Mean	
Age	108	21.2	112	22	121	20.5	112	20.3	F: 0.003***
Gender	108		112		121		112		χ^2 : 0.29
... Female	53	49%	44	39%	54	45%	42	38%	
... Male	55	51%	68	61%	67	55%	70	62%	
Education	108		112		118		110		χ^2 : 0.12
... Highschool	10	9%	7	6%	11	9%	9	8%	
... Undergrad.	60	56%	67	60%	73	62%	72	65%	
... Engineer	1	1%	2	2%	1	1%	1	1%	
... Master 1	25	23%	18	16%	17	14%	8	7%	
... Master 2	10	9%	13	12%	11	9%	10	9%	
... PhD	1	1%	0	0%	2	2%	0	0%	
... Other	1	1%	5	4%	3	3%	10	9%	
Background	103		99		111		105		χ^2 : 0.034**
... Economics	84	82%	78	79%	98	88%	79	75%	
... Sciences, Eng.	3	3%	2	2%	4	4%	1	1%	
... Political Science	9	9%	16	16%	7	6%	15	14%	
... Languages	1	1%	0	0%	1	1%	6	6%	
... Literature, Philo.	5	5%	2	2%	0	0%	4	4%	
... History, Geog.	1	1%	1	1%	1	1%	0	0%	
Experience	108		112		121		112		χ^2 : 0.023**
... Yes	28	26%	44	39%	28	23%	40	36%	
... No	80	74%	68	61%	93	77%	72	64%	
Couple	108		112		121		112		χ^2 : 0.551
... Yes	15	14%	13	12%	22	18%	16	14%	
... No	93	86%	99	88%	99	82%	96	86%	
Nationality	108		112		121		112		χ^2 : 0.818
... French	83	77%	90	80%	99	82%	90	80%	
... Other	25	23%	22	20%	22	18%	22	20%	

Standard deviations for Age are respectively: 3.6, 5.6, 2.9 and 2.4. The last column reports the p-value of an independence test between each variable and the treatment arm. An F test is used for the "Age" variable and " χ^2 " for the remaining ones. Statistical significance markers: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. Five participants did not report their education level, while 35 did not report their background.




Figure A.1: Each dot represents one observation, defined as the total number of times that one subject selected the selfish option in one sequence of 20 decisions. Each subject was presented with two such sequences, hence each subject is represented by two observations. The left part shows the $n = 453$ observations collected under the Neutral frame in sequences $N1$, $N2$, $A1$, and $B2$, while the right part shows the $n = 453$ observations collected under the Market frame in sequences $M1$, $M2$, $A2$, and $B1$. The two-sided t-test for the difference in means results in a p -value of 0.0000004. The Wilcoxon rank-sum test for the difference in medians results in a p -value of 0.0000004.




Figure A.2: Cumulative distributions of observations—defined as the total number of times that one subject selected the selfish option in one sequence of 20 decisions—for decisions collected in the VOI sequences (solid line, for the $n = 686$ observations in sequences $N2$, $M2$, $A1$, $A2$, $B1$, and $B2$) and those collected in the non-VOI sequences (dashed line, for the $n = 220$ observations collected in sequences $N1$ and $M1$). The Kolmogorov–Smirnov test’s D statistic was 0.21206, while the p -value was .0000006231.




Figure A.3: Cumulative distributions of observations—defined as the total number of times that one subject selected the selfish option in one sequence of 20 decisions—for decisions collected in the Neutral VOI sequences *N2*, *A1*, and *B2* (black line, $n = 341$), the Market VOI sequences *M2*, *A2*, and *B1* (red line, $n = 345$), the Neutral non-VOI sequence *N1* (blue line, $n = 108$), and the Market non-VOI sequences *M1* (green line, $n = 112$). All pairwise Kolmogorov-Smirnoff tests result in p -values at or below 0.01, with the exception of Market VOI versus Neutral non-VOI, which gives a p -value of 0.6.

Table A.3: Share of subjects choosing selfish action - By Frame

Payoff	n-VOI - between (N1-M1)		VOI - between (A1-B1)		VOI - within (A1-A2-B1-B2)	
	Neutral	Market	Neutral	Market	Neutral	Market
1	0.69	0.84	0.59	0.80	0.65	0.77
2	0.69	0.82	0.55	0.58	0.60	0.65
3	0.64	0.74	0.54	0.56	0.62	0.66
4	0.49	0.72	0.48	0.49	0.50	0.50
5	0.53	0.72	0.45	0.61	0.45	0.57
6	0.53	0.67	0.42	0.44	0.40	0.48
7	0.45	0.66	0.37	0.47	0.34	0.48
8	0.51	0.60	0.36	0.34	0.31	0.39
9	0.52	0.73	0.38	0.53	0.35	0.52
10	0.39	0.60	0.31	0.25	0.23	0.32
11	0.38	0.58	0.29	0.30	0.27	0.33
12	0.36	0.50	0.32	0.32	0.26	0.34
13	0.42	0.46	0.32	0.42	0.36	0.40
14	0.34	0.39	0.24	0.29	0.25	0.31
15	0.31	0.54	0.36	0.30	0.28	0.30
16	0.36	0.46	0.30	0.21	0.23	0.24
17	0.24	0.39	0.24	0.34	0.19	0.33
18	0.44	0.60	0.25	0.24	0.19	0.26
19	0.29	0.35	0.35	0.27	0.26	0.29
20	0.29	0.30	0.26	0.24	0.21	0.25
Mean	0.44	0.58	0.37	0.40	0.35	0.42

Table A.4: Share of subjects choosing selfish action - By VOI

Payoff	Neutral, bet. (N1-A1)		Neutral, within (N1-N2)		Market, within (M1-M2)		Market, bet. (M1-B1)	
	non-VOI	VOI	non-VOI	VOI	non-VOI	VOI	non-VOI	VOI
1	0.69	0.59	0.69	0.60	0.84	0.83	0.84	0.80
2	0.69	0.55	0.69	0.42	0.82	0.74	0.82	0.58
3	0.64	0.54	0.64	0.60	0.74	0.83	0.74	0.56
4	0.49	0.48	0.49	0.40	0.72	0.61	0.72	0.49
5	0.53	0.45	0.53	0.40	0.72	0.53	0.72	0.61
6	0.53	0.42	0.53	0.31	0.67	0.46	0.67	0.44
7	0.45	0.37	0.45	0.30	0.66	0.41	0.66	0.47
8	0.51	0.36	0.51	0.28	0.60	0.34	0.60	0.34
9	0.52	0.38	0.52	0.30	0.73	0.39	0.73	0.53
10	0.39	0.31	0.39	0.23	0.60	0.29	0.60	0.25
11	0.38	0.29	0.38	0.26	0.58	0.31	0.58	0.30
12	0.36	0.32	0.36	0.19	0.50	0.27	0.50	0.32
13	0.42	0.32	0.42	0.29	0.46	0.32	0.46	0.42
14	0.34	0.24	0.34	0.23	0.39	0.26	0.39	0.29
15	0.31	0.36	0.31	0.23	0.54	0.26	0.54	0.30
16	0.36	0.30	0.36	0.19	0.46	0.21	0.46	0.21
17	0.24	0.24	0.24	0.17	0.39	0.27	0.39	0.34
18	0.44	0.25	0.44	0.15	0.60	0.25	0.60	0.24
19	0.29	0.35	0.29	0.15	0.35	0.18	0.35	0.27
20	0.29	0.26	0.29	0.12	0.30	0.20	0.30	0.24
Mean	0.44	0.37	0.44	0.29	0.58	0.40	0.58	0.40

Table A.5: Effects of VOI (within) and Market frame (between) on Selling (N1-N2-M1-M2)

Dependent Variable:	Sell (Yes or No)				
Model:	(1)	(2)	(3)	(4)	(5)
<i>Variables</i>					
VOI	-0.1528*** (0.0147)	-0.1528*** (0.0141)	-0.1528*** (0.0276)	-0.1528*** (0.0277)	-0.1528*** (0.0294)
Market	0.1413*** (0.0146)	0.1413*** (0.0140)	0.1413*** (0.0396)	0.1413*** (0.0396)	0.1411*** (0.0403)
VOI x Market	-0.0338 (0.0206)	-0.0338* (0.0197)	-0.0338 (0.0382)	-0.0338 (0.0383)	-0.0338 (0.0412)
z		0.9079*** (0.0315)	0.9079*** (0.0623)		
Socio-dem. controls	No	No	No	No	Yes
Clustering of SE	No	No	Subject	Subject	Subject and Payoff
<i>Fixed-effects</i>					
Payoff				Yes	Yes
<i>Fit statistics</i>					
Observations	8,800	8,800	8,800	8,800	8,800
R ²	0.04555	0.12801	0.12801	0.13370	0.13965
Within R ²				0.04995	0.05648

Signif. Codes: ***: 0.01, **: 0.05, *: 0.1

Note: Column (3) includes clustering of SE's at the Subject level, column (4) includes Payoff fixed effects with clustering of SE's at the Subject level, column (5) includes Payoff fixed effects with clustering of SE's at the Subject and Payoff level. Socio-demographic controls are: nationality, gender, couple status, previous experience in experiments and attendance to a private school.

Table A.6: Estimated preferences for the rep. agent in Neutral and Market frames
- Full sample

	Neutral	Market	H_0 : Neutral = Market ⁺
β : Aheadness aversion	0.230*** (0.034) - [0.013]	0.106*** (0.023) - [0.008]	0.003
κ : Degree of morality	0.223*** (0.071) - [0.030]	0.196*** (0.035) - [0.019]	0.738
σ : Choice sensitivity	0.024*** (0.007) - [0.002]	0.034*** (0.007) - [0.002]	0.331
Number of Observations	4320	4480	
Number of subjects	108	112	
Log likelihood	-2665.795	-2845.855	

Notes: Standard errors clustered at the individual level in parentheses. Non-clustered standard errors in brackets. *** Significant at 1% using clustered standard errors.

⁺ p -value of z-test, using clustered standard errors.

Table A.7: Estimated preferences for the representative agent in Neutral and Market frames - Core 2 sample

	Neutral	Market	H_0 : Neutral = Market ⁺
β : Aheadness aversion	0.147*** (0.035) - [0.013]	0.089*** (0.022) - [0.008]	0.159
κ : Degree of morality	0.441*** (0.112) - [0.058]	0.255*** (0.041) - [0.019]	0.120
σ : Choice sensitivity	0.025*** (0.011) - [0.003]	0.045*** (0.013) - [0.003]	0.254
Number of Observations	2,720	3,160	
Number of subjects	68	79	
Log likelihood	-1.576.849	-1,846.436	

Notes: Standard errors clustered at the individual level in parentheses. Non-clustered standard errors in brackets. *** Significant at 1% using clustered standard errors.

⁺ p -value of z-test, using clustered standard errors.

Table A.8: Two-type Finite Mixture Model estimates, Full sample

	Neutral		Market	
	Type 1	Type 2	Type 1	Type 2
β : Aheadness aversion	0.373*** (0.086) [0.026]	-0.057 (0.159) [0.036]	0.188*** (0.012) [0.007]	-0.104 (0.072) [0.034]
κ : Degree of morality	0.125*** (0.064) [0.034]	0.309*** (0.054) [0.053]	0.229*** (0.020) [0.020]	0.136*** (0.040) [0.033]
σ : Choice sensitivity	0.043*** (0.009) [0.005]	0.026*** (0.003) [0.004]	0.055*** (0.011) [0.004]	0.037*** (0.011) [0.005]
Share	0.553*** (0.134) [0.053]	0.447*** (0.134) [0.053]	0.643*** (0.026) [0.046]	0.357*** (0.026) [0.046]
Number of Observations	4,320		4,480	
Number of subjects	108		112	
Log likelihood	-2297.761		-2403.156	

Notes: Standard errors clustered at the individual level in parenthesis, non-clustered standard errors in brackets. *** Significant at 1%, ** Significant at 5%, * Significant at 10% using clustered standard errors.

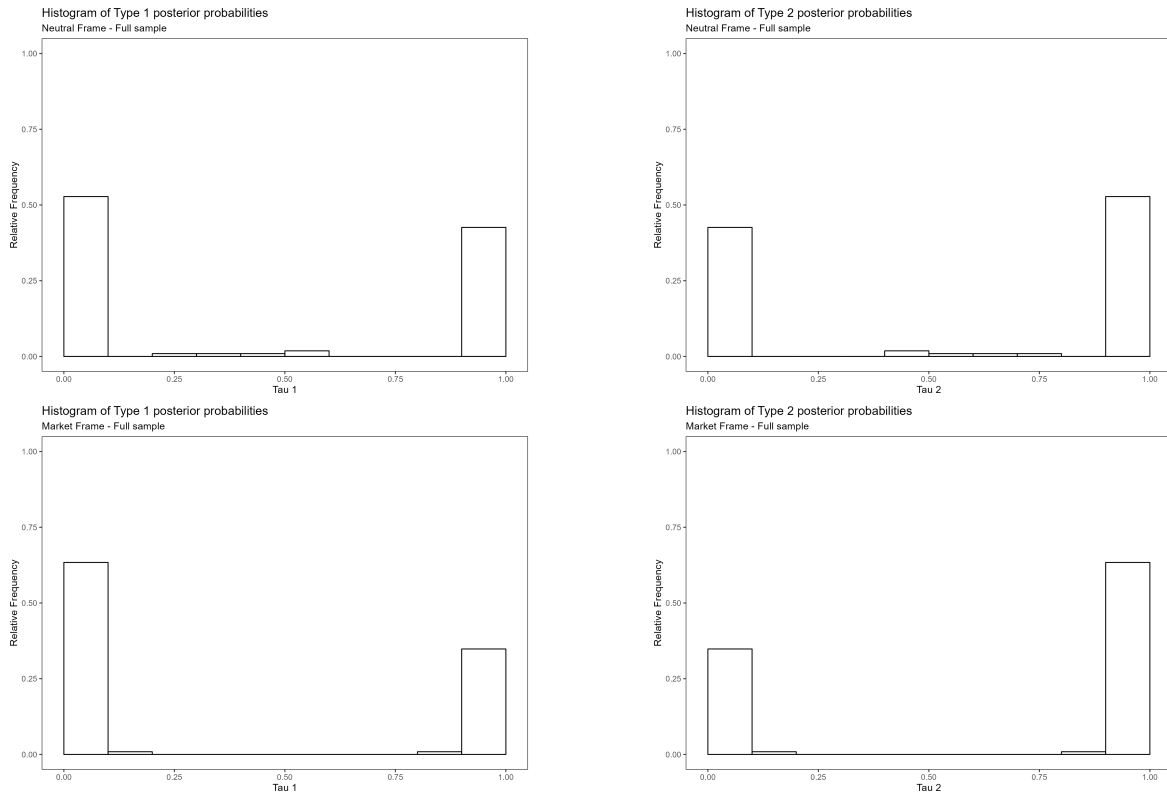




Figure A.4: Distribution of posterior probabilities of individual type-membership in Neutral (first row) and Market frames (second row). Full sample estimates.



(a) Neutral frame



(b) Market frame

Figure A.5: Difference in selfish actions non-VOI versus VOI. Full sample. Subjects classified as Type 1 if $\tau_1 > 0.95$ or Type 2 if $\tau_1 < 0.05$

Table A.9: Two-type Finite Mixture Model estimates, Core 2 sample

	Neutral		Market	
	Type 1	Type 2	Type 1	Type 2
β : Aheadness aversion	0.287*** (0.024) [0.020]	-0.145 (0.194) [0.066]	0.18*** (0.012) [0.007]	-0.305 (0.283) [0.109]
κ : Degree of morality	0.344*** (0.085) [0.061]	0.595*** (0.146) [0.124]	0.176*** (0.020) [0.016]	0.597** (0.297) [0.143]
σ : Choice sensitivity	0.047*** (0.018) [0.007]	0.023*** (0.006) [0.005]	0.073*** (0.012) [0.005]	0.029* (0.015) [0.006]
Share	0.522*** (0.084) [0.065]	0.478*** (0.084) [0.065]	0.623*** (0.028) [0.056]	0.377*** (0.028) [0.056]
Number of Observations	2,720		3,160	
Number of subjects	68		79	
Log likelihood	-1384.981		-1617.195	

Notes: Standard errors clustered at the individual level in parenthesis, non-clustered standard errors in brackets. *** Significant at 1%, ** Significant at 5%, * Significant at 10% using clustered standard errors.

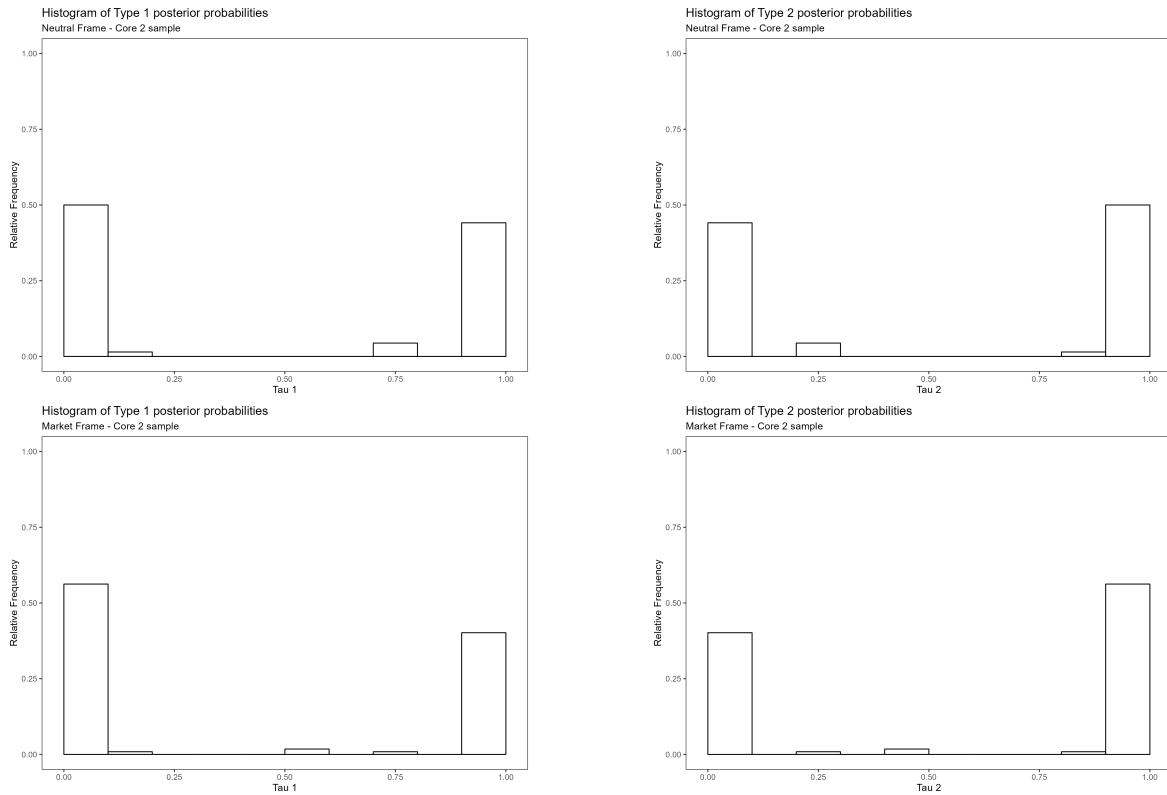




Figure A.6: Distribution of posterior probabilities of individual type-membership in Neutral (first row) and Market frames (second row). Core 2 sample estimates.



(a) Neutral frame



(b) Market frame

Figure A.7: Difference in selfish actions non-VOI versus VOI. Core 2 sample. Subjects classified as Type 1 if $\tau_1 > 0.95$ or Type 2 if $\tau_1 < 0.05$


Appendix A2 Preference parameter estimation at the individual level

Figures A.8 and A.9 show the distributions of the estimates of κ at the individual level, for the Market and the Neutral frames, respectively.


These figures show that for a large number of subjects the estimates are negative and/or exceed 2 in absolute value (this is especially true in the Neutral frame). This suggests that our experimental design is not ideal to structurally estimate individual preference parameters. To see why, consider Figure A.10, which shows the threshold values for β and κ (recall (5) and (6)) for three of the payoffs used in the experiment: those labeled 1, 4, and 20 in Table 1. As indicated by the numbers in the blue balls, the solid lines correspond to payoff 1, the dashed ones to payoff 4, and the dotted ones to payoff 20. We discuss the challenge of identifying (β_i, κ_i) for several different behavioral scenarios (a discussion which clearly generalizes to subjects taking decisions in the 20 payoff configurations of the experiment). For this discussion, recall that, for any given payoff, a subject i : (a) does not sell under either non-VOI or VOI if (β_i, κ_i) is to the right of the vertical line; (b) sells under both non-VOI or VOI if (β_i, κ_i) is below the downward-sloping curve; and (c) makes a switch (i.e., sells under non-VOI but not under VOI) if (β_i, κ_i) is between the vertical and the downward-sloping curve (recall Figure 3).

The first issue is the lack of bounds on β and/or κ inherent in individual estimations. Consider a subject who switches under payoff 1 and never sells under payoffs 4 and 20. Such behavior is consistent with any (β, κ) in zone I (where the red ball labeled I appears). While the value of β is thus bounded below by 0.33 and above by 0.6, the value of κ is bounded below only. In a similar vein, a subject who switches under payoff 20 and sells under payoffs 1 and 4 is consistent with any (β, κ) in zone II: for this subject, the value of β is bounded above only, while the value of κ is bounded below by the downward-sloping dashed and dotted lines. As a third example, a subject who switches under the three payoffs is consistent with any (β, κ) in zone III. Again, identification of this subject's (β, κ) is clearly impossible.

The second issue arises for subjects whose behavior is compatible with disjoint sets of values of (β, κ) . For example, consider a subject who switches under payoff 1 and 20, but not under payoff 4. Such a subject could have a (β, κ) in zone I




(a)




(b)

Figure A.8: Distributions of individual estimates of κ for the 100 subjects for whom the decision was different under VOI and non-VOI for at least one of the 20 payoffs, in the Market frame (treatment M). In panel (a), no restriction was imposed on the estimates. In panel (b) the estimates were restricted to the interval $[0, 1]$. In panel (a) the estimates whose value is above 2 in absolute value are lumped together at -2 and 2.

(consistent with the switch under payoff 1), in which case (s)he would have made a large mistake by switching under payoff 20; or (s)he could have a (β, κ) in zone II



(a)



(b)

Figure A.9: Distributions of individual estimates of κ for the 98 subjects for whom the decision was different under VOI and non-VOI for at least one of the 20 payoffs, in the neutral frame (treatment N). In panel (a), no restriction was imposed on the estimates. In panel (b) the estimates were restricted to the interval $[0,1]$. In panel (a) the estimates whose value is above 2 in absolute value are lumped together at -2 and 2.

(consistent with the switch under payoff 20), in which case (s)he would have made a large mistake by switching under payoff 1.

By contrast, estimation of (β, κ) at the aggregate level is possible. To see why, consider again Figure A.10. Suppose that 20% of the subjects switched under all the payoffs (zone III), 25% switched under payoffs 1 and 4 and did not sell under payoff 20 (zone IV), 35% switched under payoff 1 and did not sell under payoffs 4 and 20 (zone I), while the remaining 20% switched only under payoff 20 and sold under payoffs 1 and 4 (zone II). If these decisions are interpreted as emanating from one single individual who sometimes makes mistakes, the estimated (β, κ) of this hypothetical individual would likely be in zone V, or in zone IV close to zone V. The between-subject heterogeneity in behavior would thus help put boundaries on the estimates of β and κ .




Figure A.10: The numbers in the blue balls indicate the payoff to which the V-shaped lines correspond (see Table 1). The numbers in the red balls indicate zones with different switching behaviors between non-VOI and VOI.

Appendix A3 Experimental instructions

Screen 1 [common to all the treatments]:

Welcome to this experiment!

Please read the following instructions carefully and from now on, do not communicate with any of the other participants. If you have any questions during the experiment, please raise your hand and wait until we come to you to answer your questions in private.

All the participants here today will be asked to take decisions. These decisions will generate points. At the end of the experiment points will be converted into money. Every point is worth 0.025 euros. Each participant also receives 4 euros for attending until the end and answering all questions.

Your decisions during the experiment are anonymous. They will not be linked to your name in any way. Other participants can never trace your decisions back to you. Moreover, the amount of money you receive at the end will be handed over to you in an opaque envelope, and no other participant will see what is inside the envelope.

During the experiment, your cell phone should be switched off and out of reach. And remember not to talk with the other participants. We would need to exclude you from the experiment (and the payment) if you breach these rules.

The experiment consists of two parts, followed by a short questionnaire. At the beginning of each part, you will receive new instructions. Your decisions made in one part will never affect outcomes in the other part, so you can treat both parts as independent.

A3.1 Remaining screens for the Market treatment

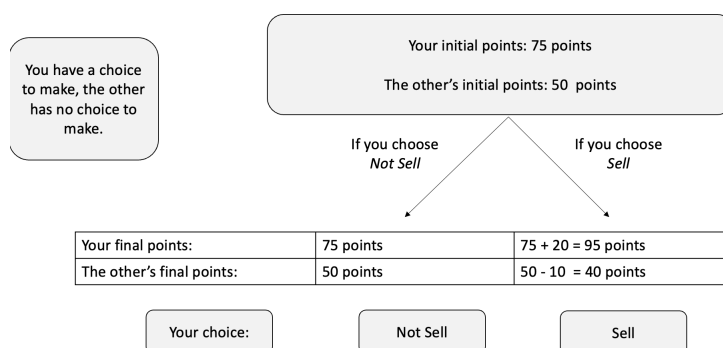
Screen 2:

Part I

In this part you will be asked to make choices in 20 different decision situations, which represent an interaction between a Seller and a Buyer.

In each decision situation you will be paired with one of the other participants here today, each time with a different participant.

In each decision situation you are the Seller and the person you are paired with is the Buyer. You will be asked to choose one of two options, Sell and Not Sell, while the Buyer has no choice to make. Think of this as representing a situation in which the Buyer is willing to pay the price for a good that you own. Here is an example of a decision screen:



The number of points that you get if you choose Not Sell can be interpreted as the value you attach to owning the good. The number of points that you get if you choose Sell can be interpreted as the amount of money you have if you Sell the good.

The number of points that the Buyer gets if you choose Not Sell can be interpreted as the amount of money he/she has initially. The number of points that the Buyer gets if you choose Sell can be interpreted as the value he/she attaches to owning the good.

As you can see on the decision screen above, the Buyer would be better off if you chose Not Sell, while you are better off if you Sell. Think of this as representing a situation in which the good that you sell has a defect which makes the Buyer enjoy owning the good less than you do. At the end of the experiment one of these 20

decision situations will be randomly drawn. Your decision in this situation (and only this situation) will have an effect on the number of points you and the person you were paired with will get. This other participant will make no decision that affects your number of points, however.

Remember that:

- all your decisions are anonymous, and no participant will ever learn with whom he/she was paired in any decision situation;
- all the decision situations are equally important, in the sense that they are all equally likely to count towards the amount of money you and the person you were paired with will receive at the end;
- each point will be converted to 0.025 euros at the end of the experiment.

Screen 3 [comprehension quiz]:

1. Among the decision situations I will be presented with in this part:
 - all will be taken into account to calculate my final number of points.
 - one randomly drawn decision will be taken into account to calculate my final number of points.
 - two randomly drawn decisions will be taken into account to calculate my final number of points.
2. In each decision situation:
 - I will be paired with the same participant.
 - I will be paired with a new, randomly determined participant.
3. At the end of the experiment:
 - I will learn with whom I was paired with.
 - I will not learn with whom I was paired with.
4. In this experiment, each point is worth:
 - 0.02 euros.
 - 0.025 euros.
 - 0.03 euros.

Screens 4-33: //the first 20 decision situations//


Screen 34:

Part II

In this part you will be asked to make choices in 20 different decision situations.

In each decision situation you will be paired with one of the other participants here today, each time with a different participant.

In each decision situation there are two roles: the Seller role and the Buyer role. Either you or the person you are paired with will be assigned to the Seller role, while the other will be assigned to the Buyer role. The two different role assignments are equally likely. The person assigned to the Seller role chooses between two options, Sell and Not Sell, while the other person has no choice to make. However, because the role assignment will only be made at the end of the experiment, both you and the person you were paired with will be asked to state what you would do in Role A in each decision situation. Here is an example of a decision screen:



The number of points that the Seller gets from choosing Not Sell can be interpreted as the value that the Seller attaches to owning the good. The number of points that the Seller gets from choosing Sell can be interpreted as the amount of money the Seller has if he/she Sells the good. The number of points that the Buyer gets if the Seller chooses Not Sell can be interpreted as the amount of money the Buyer has initially. The number of points that the Buyer gets if the Seller chooses Sell can be interpreted as the value the Buyer attaches to owning the good.

As you can see on the decision screen above, the Buyer would be better off if the Seller chose Not Sell, while the Seller is better off if he/she chooses Sell. Think of

this as representing a situation in which the good that the Seller sells has a defect which makes the Buyer enjoy owning the good less than the Seller does.

At the end of the experiment one of these 20 decision situations will be randomly drawn. The role assignment for this decision situation will be randomly drawn. Recall that the two roles are equally likely. If you are assigned to the Seller role (this happens with probability $1/2$), it is your decision in this situation that will have an effect on the number of points you and the person you were paired with will get. If you are assigned to the Buyer role (this happens with probability $1/2$), it is the decision of the person you were paired with in this situation that will have an effect on the number of points you and the person you were paired with will get.

Remember that:

- all your decisions are anonymous, and no participant will ever learn with whom he/she was paired in any decision situation;
- all the decision situations are equally important, in the sense that they are all equally likely to count towards the amount of money you and the person you were paired with will receive at the end;
- each point will be converted to 0.025 euros at the end of the experiment.

Screen 35 [comprehension quiz]:

1. Among the decision situations I will be presented with in this part:
 - all will be taken into account to calculate my final number of points.
 - one randomly drawn decision will be taken into account to calculate my final number of points.
 - two randomly drawn decisions will be taken into account to calculate my final number of points.
2. In each decision situation:
 - there are two roles and I will be paired with the same participant.
 - I will be paired with a new, randomly determined participant.
3. In each decision situation:
 - I will be paired with the same participant.
 - I will be paired with a new, randomly determined participant.
4. At the end of the experiment:
 - I will learn who you were paired with.
 - I will not learn who you were paired with.
5. In this experiment, each point is worth:
 - 0.02 euros.
 - 0.025 euros.
 - 0.03 euros.

Screens 36-65: the 20 decision situations for Part II

A3.2 Remaining screens for the Neutral treatment

Screen 2:

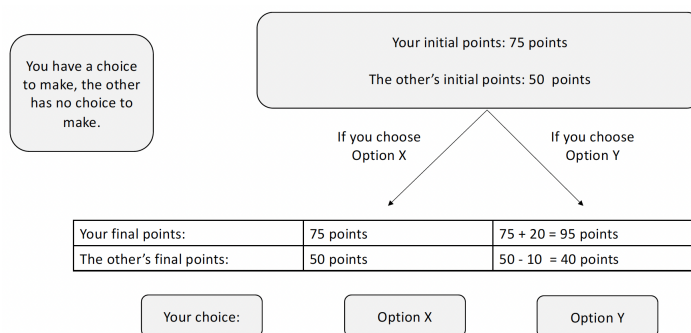
Part I

In this part you will be asked to make choices in 20 different decision situations.

In each decision situation you will be paired with one of the other participants here today, each time with a different participant.

In each decision situation you will be asked to choose one of two options, X and Y, while the other participant has no choice to make.

Here is an example of a decision screen:



At the end of the experiment one of these 20 decision situations will be randomly drawn. Your decision in this situation (and only this situation) will have an effect on the number of points you and the person you were paired with will get. This other participant will make no decision that affects your number of points, however.

Remember that:

- all your decisions are anonymous, and no participant will ever learn with whom he/she was paired in any decision situation;

- all the decision situations are equally important, in the sense that they are all equally likely to count towards the amount of money you and the person you were paired with will receive at the end;
- each point will be converted to 0.025 euros at the end of the experiment.

Screen 3 [comprehension quiz]:

1. Among the decision situations I will be presented with in this part:
 - all will be taken into account to calculate my final number of points.
 - one randomly drawn decision will be taken into account to calculate my final number of points.
 - two randomly drawn decisions will be taken into account to calculate my final number of points.
2. In each decision situation:
 - I will be paired with the same participant.
 - I will be paired with a new, randomly determined participant.
3. At the end of the experiment:
 - I will learn with whom I was paired with.
 - I will not learn with whom I was paired with.
4. In this experiment, each point is worth:
 - 0.02 euros.
 - 0.025 euros.
 - 0.03 euros.

Screens 4-33: the 20 decision situations for Part I

Screen 34:

Part II

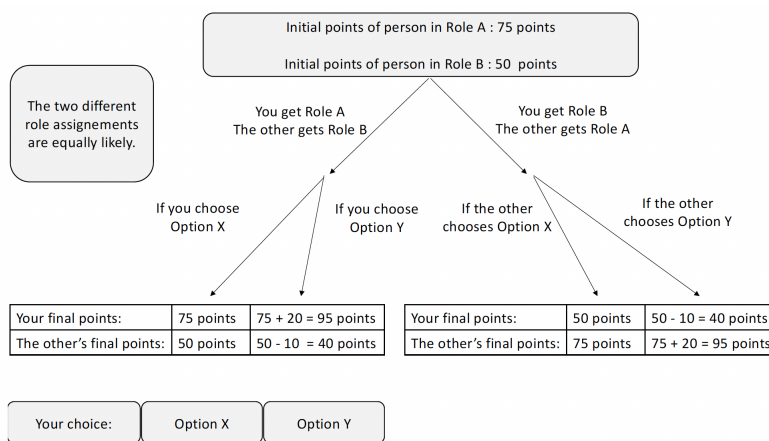
In this part you will be asked to make choices in 20 different decision situations.

In each decision situation you will be paired with one of the other participants here today, each time with a different participant.

In each decision situation there are two roles: Role A and Role B. Either you or the person you are paired with will be assigned to Role A, while the other will be assigned to Role B. The two different role assignments are equally likely. The person assigned to Role A chooses between two options, X and Y, while the other person

has no choice to make. However, because the role assignment will only be made at the end of the experiment, both you and the person you were paired with will be asked to state what you would do in Role A in each decision situation.

Here is an example of a decision screen:



At the end of the experiment one of these 20 decision situations will be randomly drawn. The role assignment for this decision situation will be randomly drawn. Recall that the two roles are equally likely. If you are assigned to Role A (this happens with probability 1/2), it is your decision in this situation that will have an effect on the number of points you and the person you were paired with will get. If you are assigned to Role B (this happens with probability 1/2), it is the decision of the person you were paired with in this situation that will have an effect on the number of points you and the person you were paired with will get.

Remember that:

- all your decisions are anonymous, and no participant will ever learn with whom he/she was paired in any decision situation;
- all the decision situations are equally important, in the sense that they are all equally likely to count towards the amount of money you and the person you were paired with will receive at the end;
- each point will be converted to 0.025 euros at the end of the experiment.

Screen 35 [comprehension quiz]:

1. Among the decision situations I will be presented with in this part:
 - all will be taken into account to calculate my final number of points.
 - one randomly drawn decision will be taken into account to calculate my final number of points.
 - two randomly drawn decisions will be taken into account to calculate my final number of points.

2. In each decision situation:
 - there are two roles and I will be paired with the same participant.
 - I will be paired with a new, randomly determined participant.

3. In each decision situation:
 - I will be paired with the same participant.
 - I will be paired with a new, randomly determined participant.

4. At the end of the experiment:
 - I will learn who you were paired with.
 - I will not learn who you were paired with.

5. In this experiment, each point is worth:
 - 0.02 euros.
 - 0.025 euros.
 - 0.03 euros.

Screens 36-65: the 20 decision situations for Part II

A3.3 Remaining screens for the two Mixed treatments

Each mixed treatment combines the Market VOI instructions (shown in Part II of subsection “Remaining screens for the Market treatment”) and the Neutral VOI instructions (shown in Part II of subsection “Remaining screens for the Neutral treatment”).

A3.4 Example of the two screens shown for one decision

