Strategy vs. Direct-Response Method: Evidence from a Large Online Experiment

Marcus Roel
mcs.roel@gmail.com
Beijing Normal University

Zhuoqiong Chen*
chenzq926@gmail.com
Harbin Institute of Technology, Shenzhen

April 3, 2023

Abstract

This paper studies the behavioral differences that arise from eliciting choices in sequential games by the direct-response method, where players observe the choices made by players who acted before them and make a single choice in response, or the strategy method, where they make their choice(s) in response to all possible choices by those who acted prior regardless of whether such actions were actually taken. We conduct a 2×2 between-subject design, large-scale online experiment with over 8000 participants on Amazon MTurk, where, in addition to varying the elicitation method, we also manipulate the ex-ante beliefs of participants about player 1's likely choices via an information-provision treatment. In two neutrally framed binary-choice sequential games, a sequential Prisoners' Dilemma and a mini-Ultimatum Game, we document that the elicitation method does not alter player 2's preferences (their tendency to reward cooperation or reject unfair offers) yet significantly reduces mistakes (rewarding defection or rejecting fair offers). Our results suggest that the more economical strategy method should be the go-to choice for eliciting behavior in sequential games, which may be particularly relevant for all experiments conducted online nowadays.

Keywords: Strategy method, Direct-response method, Online experiments, Information-provision experiment, Elicitation, Methodology

JEL Codes: C72, C91, C92, D83, D91

*We would like to thank participants at the Experimental Economics Workshop at Beijing Normal University, and the CeDEx China Workshop at Nottingham Ningbo for their helpful comments and feedback. Zhuoqiong Chen gratefully acknowledges support from the National Natural Science Foundation of China (No. 71903046) and the "Shenzhen Peacock Plan" start-up research grant (No. GA11409002).

1. Introduction

Social sciences, such as economics, management, and psychology, often rely on experiments and surveys to collect data. The validity and reliability of such data depends on whether it captures the participants' genuine preferences, feelings, or beliefs. A key methodological question is therefore how such behavior and beliefs can be best elicited – ideally in a way that is simple to understand for participants, as well as cheap to implement and easy to conduct for researchers.

For sequential games, the two most common approaches to elicit choices from players who do not act first are the direct-response method and the strategy method. In the former, players directly observe the choices made by players who acted before them and make a (single) choice in response. In the latter, they make their choice(s) without knowing anything about the choices made by those that acted before them. Instead, they are asked to take an action in response to *all possible choices* by those who acted prior, regardless of whether such actions were actually taken. In other words, the direct-response method elicits a single choice[1] whereas the strategy method elicits a full strategy profile. Nevertheless, the general methodological approach as such is not just relevant for sequential games, it extends to all settings with conditional elicitations (e.g., state-dependent preferences, etc.), highlighting the importance of understanding when best to use what.

If players' preferences, understanding of the game, and likelihood of making mistakes are not affected by which of the two elicitation methods is used, they will generate the same data. In this case, the strategy method has the key advantage of allowing the researcher to capture choices at all points in the game for every player, which greatly increases the effective sample size, lowering the cost of running experiments compared to the direct-response method. Moreover, running experiments becomes easier as they can be conducted as a de-facto single-user survey with ex-post matching of players to determine payoffs, which is especially useful for large online experiments that have become much more popular recently, partially due to Covid-19.[2]

The strategy method has, however, been criticized for being artificial and psychologically "cold", as forming a contingent plan of actions may not be a natural decision process for most people, and thus may result in different behaviors compared to the direct-response method (e.g., Roth (1995)).[3] So far, the evidence of whether the elicitation method affects the choices taken has been mixed; see Brandts and Charness (2011) for a review. In light of this, we conduct a large online experiment on Amazon MTurk with over 8000 participants to investigate the role of the strategy method and the direct-response method in shaping behaviors.

Participants play two sequential, binary-choice games (neutrally framed and in random order), corresponding to a sequential Prisoners' Dilemma (sPD) and a mini-Ultimatum Game (mUG; with an unfair 85-15 or a fair 50-50 split). The choice of games is kept intentionally simple in order to allow for a systematic analysis of the elicitation method in games where social preferences may give rise to player 2 rewarding a helpful action (i.e., cooperate after cooperation) or punishing an unfair choice (i.e., reject the 85-15 offer) of player 1. We adopt a 2×2 between-subject design, where, in addition to varying the elicitation method, we manipulate the ex-ante belief of all participants about player 1's choice by providing them with information about typical player 1 behavior in the same games from past published studies. In particular, such information highlighted that the majority of first movers either took the "selfish" (defect, 85-15 offer) or

[1] This may happen more than once if the player needs to respond to other's choices at multiple points in time.

[2] Figure A.1 in the Appendix documents the recent popularity of online experiments in more detail.

[3] It is well understood that beliefs and choices may (partially) be a result of how they are elicited (Gigerenzer and Hoffrage (1995)), or how the problem is presented or framed (Gigerenzer et al. (1988), Tversky and Kahneman (1981)). It is not always obvious, however, what the "true data" is or how behavior will be affected. Both, observing a particular fair/unfair choice that stirs the player's emotions (direct-response), or an increased fairness-concern due to considering all nodes of the game (strategy-method), may lead to more non-selfish choices. Indeed, Roth (1995) suggests that it's not clear whether forcing participants to think about all their information sets is an advantage or disadvantage compared to the direct response method.

"non-selfish" action (cooperate, 50-50 offer). We refer to the two respective information treatments as the *selfish* and *non-selfish belief treatment*.

The experimental design is motivated by our conjecture that previous mixed findings were driven by a combination of two factors: (1) ex-ante beliefs about player 1's behavior that varied across experiments; (2) incomplete conditional thinking[4]. The first factor is, in our view, the most likely candidate for an orthogonal dimension whose value may vary across experiments. However, in order to explain the experimental data, such ex-ante beliefs cannot exert a constant effect on behavior. Instead, they must affect player 2's behavior elicited with the strategy method in a different way than those with the direct response method, where the key behavior of interest is their likelihood of engaging in non-selfish choices (cooperate/defect) in response to cooperation or unfair offers. The second factor provides a possible mechanism for this differential effect, as ex-ante beliefs about player 1's choices can shape the preferences of player 2 particularly in the strategy method if they fail to fully condition their choice on the particular point in the game at which they are making it.[5]

Overall, we find that Mturk workers' behavior is broadly consistent with conditional cooperation in the *sPD*, with over 60% of player 2 choosing to cooperate after player 1 cooperates and over 75% opting to defect after player 1 defects in all treatments. Yet in light of this pro-social behavior in the *sPD*, only few participants are willing to punish unfair offers in the *UG*: rejection rates of unfair offers are below 15% in all treatments (and never higher than 4% for even splits).

Our belief manipulation strongly affects players' beliefs about player 1's choices and, to a slightly lesser extent, player 1's choices themselves – in line with the information provided. The direct effect on player 2's propensity to take a non-payoff maximizing choice is either very small or non-existent, however.

We find no evidence in support of a model with incomplete conditional thinking[6] as the belief treatment displays no significant differential effect (by elicitation methods) on player 2's response to cooperation or unfair offers. Moreover, for these choices, the elicitation method appears inconsequential: there is no significant difference in player 2's tendency to reward cooperation by player 1 and only a small, albeit significant, 4 percentage point difference of rejecting unfair offers between the two methods. That does not mean that the choice of elicitation method is immaterial, however. Our data suggests that the strategy method reduces "mistakes", in the sense that it reduces both the frequency of cooperation in response to player 1 defecting (rewarding the un-helpful, payoff minimizing choice) and rejecting fair offers (punishing the payoff maximizing choice). This effect is economically significant at around 10 percentage points and constant across belief treatments in the *sPD*, whereas it mainly affects the selfish-belief treatment in the *mUG*. We also provide further evidence that such choices indeed represent mistakes, e.g., they are made more frequently by inattentive participants.

We contribute to the literature in several ways. Our paper is the largest experiment to this date that analyzes the important question of whether eliciting choices via the direct-response method or the strategy method fundamentally influences behavior. In the survey paper by Brandts and Charness (2011), prior studies that adopt both methods often show that subjects are more likely to punish selfish behavior when elicited by the direct response method (Brandts and Charness (2003); Brosig et al. (2003); Oxoby and McLeish (2004); Falk et al. (2005)), while others study show there is no difference between the two methods

[4]See, for instance, Esponda and Vespa (2014), Esponda and Vespa (2019), Martínez-Marquina et al. (2019).

[5]Previous first-order explanations for behavioral differences across elicitation methods are generally independent of variation in beliefs (or other orthogonal dimensions) and thus cannot explain the existing experimental results. For example, initial beliefs play no role if the emotional response to a certain action is stronger in direct-response method (hot-versus-cold, e.g., Brandts and Charness (2000)). For further details, consult the theory section in the Online Appendix (part C). Here, the reader will find a formal model of social-preferences with incomplete contingent thinking and all respective hypotheses.

[6]Or indeed any other model that would generate such a differential effect.

(e.g., Brandts and Charness (2000); Naef and Schupp (2009)). Our results suggests that preferences seem to be unaffected by the elicitation method yet mistakes are significantly lowered by the strategy method.

This insight that may be particularly important for online experiments that are increasingly relied-upon nowadays.[7] For (at the very least) the typical simple experiments that are run online, the strategy methods appears to be the superior choice as it reduces mistakes, makes running experiments cheaper and, in many cases, easier to run, without distorting participants' preferences. We hope our large-scale experiment provides the methodological justification for using the strategy-method in these settings.

Moreover, our data further highlights a small, yet noteworthy detail, which suggests that the strategy method appears to facilitate the strategic thinking of participants. In particular, we observe that player 1's behavior is more responsive to the information, which indirectly provides information about player 2's likely response, in the strategy method treatment. This effect may prove particularly useful for online-experiments that focus on one-shot games.[8]

Finally, we contribute to the literature of information-provision experiments (Haaland et al. (2023)) by providing a very simple, yet powerful method to shape participants beliefs about the likely choices of others in games. Similarly to the literature of framing in games (Ellingsen et al. (2011), Dreber et al. (2013), Ockenfels and Werner (2014)), the belief treatment appears to have little direct effect on players' preferences (as documented by player 2 behavior) and thus seems to operate mostly through beliefs when shaping player 1s' choices.

The rest of the paper is structured as follows: section 2 details our experimental design. The results are presented in section 3. Our paper concludes with a discussion in section 4. Other supportive tables and figures (e.g., variable definitions, summary statistics, etc.) can be found in appendix A. The instructions for our experiments are provided in appendix B.

2. Experimental Design

The study was pre-registered and has been approved by the (School of Economics and Management) Research Ethics Committee at the Harbin Institute of Technology, Shenzhen.[9] The experimental instructions can be found in section B of the appendix. Screenshots of the experiment can be found in the Online Appendix.

Our participants played two sequential games, the sequential Prisoner's Dilemma (sPD) and the mini Ultimatum Game (mUG). The payoffs in the sPD were (\$1, \$1), (\$1.5, \$0), (\$0, \$1.5), and (\$0.5, \$0.5). In the mUG, the proposer could either split \$2 equally or according to 85-15 (\$1.7, \$0.3). Rejection resulted in a payoff of 0 for both parties. The games were presented in a neutral frame, referred to as task 1 or 2, with their order randomized. The only discernible difference between the games were their payoffs, which were presented in a matrix-like format using actual \$-values.[10] Player 1 (he), referred to as the first mover

[7]At this time, the social science literature that uses online experiment is simply becoming too large to provide an exhaustive list of references. To mention a few, see Horton et al. (2011) for economics, Aguinis et al. (2021) for management, Goodman and Paolacci (2017) for marketing, and Hunt and Scheetz (2019) for information system. The studies that are relevant to our particular experimental setting will be discussed in the next section.

[8]In an early survey of the bargaining literature, Güth and Tietz (1990) make a similar observation when discussing Güth et al. (1982), who included a treatment with role-reversal and strategy method as a "consistency check" in their experiment, pointing out that subjects' attention to strategic aspects can be shaped by the experimental setting. Surveying the impact of players' role-reversal more generally, Brandts and Charness (2011) conclude that it is difficult to draw any conclusion about its effect.

[9]Chen, Zhuoqiong and Marcus Roel. 2019. "Strategy vs. Direct Response Method." AEA RCT Registry. September 20. https://doi.org/10.1257/rct.4737. Our main analysis follows the pre-specified plan. In order to provide further insight into these results and our interpretation thereof, we explored mistakes in more detail, framed the analysis using a basic selfish and social theory, and classified player 2 types. This was not pre-specified. Clearly, any pre-Covid/Covid data comparisons were also not specified in advance.

[10]Many of our design choices were aimed towards making games as easy and as quickly to understand as possible in view of the online-nature of our experiment. One implication of this design principle was using real currency values in the games over

in the experiment, chose between actions $\{A, B\}$ and player 2 (she), referred to as the second mover, chose between actions $\{C, D\}$.

2.1. Treatments

We implemented a 2×2 between subject design, in which we manipulated both the beliefs of participants, by providing them with behavioral data from past experiments, and how the second mover' choices were elicited.

Belief-Manipulation. Participants were randomly assigned to the *selfish* or *non-selfish beliefs* treatment, which determined what information was presented to a player for a given game. In particular, a short sentence that *described player 1's behavior in a past study* was displayed in addition to their particular payoff matrix and their particular role. For the *sPD* in the non-selfish belief treatment, the sentence read: *"In a well-known study of this task by Watabe, Terai, Hayashi, and Yamagishi, published in the year 1996, 82.6% of the first movers chose **A**."* For the selfish belief treatment, participants were informed that *"In a well-known study of this task by Bolle and Ockenfels, published in the year 1990, 82.7% of the first movers chose **B**."* In order to keep the belief treatments as similar as possible, and to prevent potential confusion, we emphasized the more likely action in the sentence instead of highlighting the same action. Note that in the *sPD*, action **A** represents cooperation whereas **B** stands for defection. In other words, players in the *non-selfish* belief manipulation were informed that the action of player 1 that improves player 2's (set of) payoffs is the more commonly chosen one, and vice versa in the *selfish* belief treatment.[11] For the non-selfish belief treatment in the *mUG*, we cited Güth et al. (2001), where 70.6% offer an equal split, using the same sentence format as in the *sPD*.

We selected these papers based on two criteria: (1) a similar payoff structure as our experiment in order to provide accurate information to our participants, and (2) similar frequencies of the two opposing actions.[12] Given these restrictions, we unfortunately did not find a suitable datapoint for the *mUG* and the selfish belief treatment.[13] The closest study in this regard involved Chimpanzees (Jensen et al. (2007)), who documented an 75% offer rate of unequal splits. Such a study is unsuitable for an online experiment, however, where subjects can look up data provided to them - even if such data is predictive of human behavior. We had used this study in a small classroom pilot that tested how our belief manipulation affects elicited beliefs. Among first year economics and business majors, who were unfamiliar with the games, 80% opted for unequal splits.[14] Instead of citing any alternative study with more balanced offer rates, we truthfully told participants in the selfish belief treatment in the *mUG* *"In our previous experiment of this task, 80% of the first movers chose **B**."* We note that this belief manipulation satisfies our previous criteria and involves no deception. The only remaining concern is thus whether it is as powerful as citing an existing published study. In the next section, we will see that it is.

---

some fake experimental-currency with some specific exchange rate. This also guided our thinking in how to present payoffs. We opted for a payoff-table, where the subject was always the "row player", indicated with *You*, and the other player was the column player, indicated with *Other Participant*. The participant's payoffs were preceded with "*You earn:*", whereas the other player's payoff were preceded with "*Other earns:*", for each element in the table.

[11] Throughout this paper, we use the terms *selfish* and *non-selfish* beliefs treatments mainly for the sake of expositional convenience. Whether player 1's action is selfish or non-selfish from his perspective generally depends on player 2's response.

[12] The payoff used in Watabe et al. (1996) were (10, 10), (0, 15), (15, 0) and (5, 5), in Bolle and Ockenfels (1990) (50, 50), (0, 75), (75, 0), and (10, 10). In Güth et al. (2001) 20 units were either equally split or in (17, 3).

[13] One contributing factor for this was the use of the binary-choice *mUG* over the discrete version with offers of $0, 1, \ldots, 10$ or the approximately continuous game, which enabled us to keep our two game-tasks as similar as possible. From a practical perspective, this ensured that subjects could be shown a single, simple instruction before any game is played. More importantly, it allows us to make reasonable comparisons between games in our analysis.

[14] 80% offered unequal splits in the selfish-belief treatment and 79% did so in the non-selfish belief treatment.

Behavior-Elicitation. Participants were randomly assigned to the *direct-response* or *strategy method* treatment at a rate of 3-to-1. We use a larger sample for the direct response treatment, in which we only observe player 2's response to player 1's chosen action, in order to balance the observations of player 2 at each node.[15] While the choice of elicitation method makes no difference for how player 1's choices are elicited, all participants were informed in the common set of instruction how choices are elicited for player 1 and player 2. Indeed, we aimed to create common knowledge among players with regards to the game played, including the fact they they will receive the same information as the other person they are matched with.

2.2. Experimental Setting and other Procedures

In view of the sizable sample requirement of a 2×2 design, we conducted the experiment online, with participants recruited from Amazon's online job platform mechanical turk (MTurk).[16] In order to be eligible to participate in our experiment, MTurk workers had to have completed at least 100 jobs, possess an approval rate of at least 99%, and be located in the USA or Canada.[17] In the job description on MTurk, we invited subjects to take part in a large online experiment, informed potential participant that they will be paid a participation fee of $1 upon successful completion, with a possible additional payment of up to $3, and highlighted that the task can be finished within 10 minutes. On average, subjects earned a generous hourly wage of $37.14.[18]

After accepting the job on MTurk, interested participants were redirected to our experiment on an external website, created with oTree (Chen et al., 2016), that provided more general details about the job's nature, the experiment itself, and elicited consent. At this point in time, workers who were uninterested to participate in our experiment could quit the website and return the job on the MTurk platform.[19] Those who chose to continue were provided with a detailed instruction about the nature of the games played, including how choices are elicited. After completing a short test that checked whether they understood the instructions, they proceeded to the play stage. After both games were played, we elicited their beliefs about player 1's behavior while reminding them about the particular games.[20] The experiment concluded with a

[15] If player 1 chooses each action with probability 0.5, a randomization of 66.6% to 33.3% results in the same number of observations for player 2 at each node. Ex-ante, we did not expected player 1's behavior to be balanced, which led us to use a randomization of 75%/25% in order to increase the power in the less likely node.

[16] By now, many studies have evaluated the use of MTurk for laboratory research and found that it provides consistent, reliable, high-quality data that replicate traditional laboratory experiments and/or nationally representative studies (Paolacci et al., 2010; Buhrmester et al., 2011; Amir et al., 2012; Paolacci and Chandler, 2014; Johnson and Ryan, 2020). Most related to our study is Horton et al. (2011), who (i) find similar levels of cooperation in a prisoner's dilemma on MTurk and a traditional laboratory environment and (ii) show that MTurkers respond to framing (Tversky and Kahneman, 1981). Arechar et al. (2018) run a complex interactive repeated public goods game, with two rounds of 10 periods, replicating the typical behavioral patters of cooperation and punishment in the laboratory. More importantly, they document that even in such a complex environment, dropouts are exogenous to the experiment.

[17] Such quality requirements are typical to ensure attentive and high-quality responses. A good resource in this regard is Hauser et al. (2019), who outline common concerns regarding MTurk experiments, provide empirical evidence about them, and offer practical solutions on how to run experiments. Note, that ineligible MTurk workers cannot see or access the job-ad itself.

[18] We set payoffs and participation fees so that most people earn a good hourly-wage, expecting most participants to finish the experiment within 10 minutes or less. If anything, we expected our round payoffs to be too generous but preferred them over scaling payments down. Payments varied across subjects but were reasonable overall: the hourly wage was $14.79 and $21.91 at the 10th and 25th percentile, and $31.91 for the median earner. 90% of participants finished within 10 minutes.

[19] MTurk workers can only accept and work on a single task at any given time. Hence, if they are uninterested to participate in our particular experiment, e.g., because they don't want to interact with other MTurk workers, they can return the un-completed task without any repercussions. Instead of returning a task themselves, they may (choose to) let it time out, in which case the task is returned after a pre-determined time. From a practical perspective, note that MTurk only displays the job-ad to potential workers if there are outstanding jobs (the total number is set upon posting the ad). When a job is accepted, the number of outstanding jobs is reduced by one. It is increased by one if an un-completed task is returned.

[20] Belief elicitation was incentivized, with an additional $0.25 paid for beliefs within 5 percentage points of the correct answer. We opted to elicit beliefs after the play stage to keep the play-stage simple and comparable to the usual environment where beliefs are typically not elicited before play. The downside to this approach is that the second mover's elicited beliefs in the direct-response method can be (and are) influenced by their particular experience. As the belief measure is not a primary outcome measure of interest, we viewed this to be a sensible tradeoff to make.

short survey that asked participants about their gender, age, degree, household income, as well as whether they have participated in similar experiments before.

Our sessions were conducted in two stages, from October to November 2019, and again in October 2021, with up to 500 participants per session.[21] We will refer to data from the first timeframe as the *pre-Covid* sample and from the second as the *Covid* sample. In order to ensure an ideal user-experience with no wait-times, we ran player 1 and player 2 separately and matched them afterwards to determine their payoffs.[22]

When paying subjects after the experiment, we explicitly informed them of their opponents' choices for each game as well as whether their beliefs-guesses resulted in additional payments. We also reminded them about their role, choices, and provided them with a link to the payoff-tables for their particular sequence of tasks.

3. Results

In this section, we present our main findings. The definitions and descriptions of our variables can be found in Table A.1. Summary statistics are provided in Table A.2. In our regressions, we include the personal information elicited via the post-experiment survey as categorical control variables. We also control for whether the game was played second, and whether the data is from the Covid sample. Finally, we conducted a randomization analysis (Table A.3) that shows that the randomization was a success.[23]

Before presenting our findings, we outline two classical theories that will help us categorize player 2's behavior. According to the *selfish theory*, players choose those actions that maximize their material payoffs. For our games, the selfish theory predicts that player 2 always defects (D) in the *sPD* and always accepts (A) any offer in the *mUG*. The alternative explanation we consider is a *social theory*, according to which players are motivated by reciprocity, i.e., they want to reward nice and punish nasty behavior. In particular, player 2 cooperates (C) after cooperation but defects after defection in the *sPD*, accepts the fair offer (50-50) but rejects (R) unfair splits of 85-15. In line with these theories, we will refer to players who act according to the selfish-theory (social-theory) prediction as selfish/selfish-types (social/social-types).[24]

Table 1 summarizes these predictions and illustrates two key observations that will guide our data analysis in this section and beyond. First, in both games, player 2's behavior after one particular action of player 1 can be used to classify her as either selfish or socially motivated. In the *sPD*, in response to cooperation, cooperation is only consistent with the social theory whereas defection is only consistent with the selfish theory. In the *mUG*, player 2's acceptance of the unequal offer is predicted by the selfish theory while the

[21]In our Pre-Analysis Plan, we had pre-committed to a sample of 4000 in case of no order-effects and 8000 otherwise. An preliminary check of the data, however, revealed the existence of simple order-effects, and so we continued the data collection. Indeed, order-effects are still present in our full data. However, our main conclusions is robust. As a result, we report results from both tasks jointly, noting that the task-order is a control variable in all regressions. For a detailed discussion of the order effects, see Appendix D.2. By November 2019 we unfortunately ran into technical issues due to the fact that Amazon had ceased operations in China and hence no longer officially accepted payments from China. We resumed the experiment in the fall of 2021 (with the delay due to Covid-19), cooperating with the experimental lab of a large British University. Except for the change in the official MTurk account, all details of the procedures and implementation remained the same.

[22]For a given batch of players, we first ran the experiment for player 1. We then used their actual behavior in each treatment arm/task order to determine the action that player 2 was informed of. Workers were never able to retake the experiment. Lastly, note that the nature of running the experiment on MTurk instead of a tightly controlled laboratory environment typically leads to not perfectly balanced samples, e.g., there are a total 4009 first and 4020 second movers. Such small inconsistencies were in line with our expectations. In general, we did not encounter any issues when running our 28 sessions.

[23]In particular, we ran pairwise (across four treatments) Kolmogorov-Smirnov tests of equality of distributions for the end-of-experiment questionnaire responses. We find no significant differences at conventional levels for all but one out of 36 tests.

[24]We use those two theories to classify behavior in the 'as if' sense. In reality, it is often not one or the other, but rather, a question of how much players are motivated by these different notions. Note also that the selfish-theory can easily be rejected when non-selfish choices are observe. It is much more difficult, however, to reject social preferences when only selfish choices are taken as social-preferences may only induce non-selfish behavior if the social-motivation is sufficiently strong and/or the material cost of the non-payoff maximizing action is not too prohibitive, all of which depends on the game played.

Table 1: Predictions for Player 2 for Selfish and Social Theory

seq. Prisoners' Dilemma	after P1 cooperates	after P1 defects
Selfish theory	D	D
Social theory	C	D
mini Ultimatum Game	after P1 offers 50-50	after P1 offers 85-15
Selfish theory	A	A
Social theory	A	R

social theory predicts subsequent rejection. Second, after player 1's other action, there is an action for player 2 that can be explained by neither theories, namely cooperation after defection and rejecting the fair offer. For a selfish player, such behavior is clearly a *mistake* as it results in lower payoffs. It can similarly be viewed as a mistake for a socially minded player, since it would reward player 1 for choosing the action that results in a strictly lower set of feasible payoffs in the *sPD* and punish the strictly more generous offer in the *mUG*. In the rest of the paper, we refer to these two actions as *mistakes*. Before we continue, we do want to recognize however that there are theories that predict what we view as mistakes as normal behavior. For example, preferences for efficiency leader player 2 to always cooperates in the *sPD* (Engelmann and Strobel, 2004) while spiteful preferences (Levine, 1998) may result in the rejection of all offers in the *mUG*. Indeed, we would have a difficult time arguing that all such behavior represent genuine mistakes in our data. However, it is a useful notion to categorize this type of behavior and we will show how elicitation methods may help players to either avoid or facilitate them. Later, we will also provide both direct and indirect evidence that this behavior is, at least partially, driven by mistakes.

3.1. Overall behavior

We begin by describing the overall frequency of cooperation, fair offers, and rejections, both for our overall sample, as well as for the pre-Covid and Covid samples in Table 2. When differences between the two subsamples are statistically significant, we indicate this using the typical significance-stars in the Covid column.

Table 2: General Behavior across Games

	Full Sample	pre-Covid	Covid
seq. Prisoner's Dilemma			
Player 1 cooperates	0.57	0.56	0.58
Player 2 cooperates after C	0.65	0.65	0.66
Player 2 cooperates after D	0.18	0.16	0.20***
mini Ultimatum Game			
Player 1 offers 50-50	0.66	0.67	0.64**
Player 2 rejects 85-15	0.13	0.09	0.17***
Player 2 rejects 50-50	0.03	0.01	0.04***
Observations	8029	4647	3382

Notes: statistically significant differences between pre- and Covid behavior (based on t-tests) at significance levels of * $p<0.10$, ** $p<0.05$, *** $p<0.01$, are indicated by the respective stars in the Covid column.

The majority of first movers cooperate in the *sPD* (57%) and offer 50-50 in the *mUG* (66%), which indicates that a large fraction of player 1 expects player 2 to act in a non-selfish manner. Differences in player 1 behavior between the two samples are small, albeit significant in the *mUG*. Indeed, given our large sample, we will often find differences to be statically significant at conventional levels despite the fact that they are economically small.

On average, the behavior of player 2 in the *sPD* is more consistent with the reciprocal behavior suggested by the social theory: 65% cooperate in response to cooperation yet only 18% cooperate after defection. Player 2's behavior across the samples is fairly stable, albeit with significantly more mistakes being made in the Covid sample. For the *mUG*, player 2's behavior is largely consistent with the selfish theory as we observe very few rejections, although most of which occur in response to unequal offers as predicted by the social theory. The key difference between the two samples lies in the rejection rate of unequal offers, which almost double from the pre-Covid days. However, even at 17%, these rejection rates are fairly small, especially in view of the large social behavior in the Prisoner's dilemma.[25] As in the *sPD*, we again observe significantly more mistakes in the Covid sample.

As we mentioned previously, there are other preferences that predict behavior we call mistakes. In Table A.4 of the Appendix, we summarize the frequency distribution of player 2 types. There, we also consider the two most common forms of such preferences, namely preferences for efficiency or spitefulness. Using behavioral data from the strategy method, it is shown that efficiency preferences in the *sPD* and spitefulness in the *mUG* may explain at most 1/3 of the mistakes in each game (6% vs. 18% and 1% vs. 3%), and even less if we rely on both-games for the classification of such preferences. In other words, it is doubtful that mistakes are mainly reflections of preferences. Moreover, just because behavior is said to be consistent with such preferences, it may still represent honest mistake.

3.2. Belief manipulation and player 1's choices

We now turn to whether the belief manipulation influenced beliefs, whether it affected player 1's behavior, and how beliefs are related to player 1's choices. Part 1 of Table 3 tabulates the data from previous studies that was used for our belief manipulations, part 2 summarizes the elicited belief about player 1's behavior from our participants, while player 1's actual behavior is shown in part 3.

Table 3: Behavior of and Beliefs about Player 1

	Selfish	Non-Selfish
1. Provided data (from previous studies)		
Player 1 cooperates	0.173	0.826
Player 1 offers 50-50	0.20	0.706
2. Elicited Beliefs		
Belief Player 1 cooperates	0.34	0.74***
Belief Player 1 offers 50-50	0.33	0.74***
3. Player 1's Behavior		
Player 1 cooperates	0.47	0.67***
Player 1 offers 50-50	0.59	0.72***

Notes: statistically significant differences between the belief-treatments (based on t-tests) for part (2) and (3) at significance levels of * $p<0.10$, ** $p<0.05$, *** $p<0.01$, are indicated by the respective stars in column (2).

[25] Note also that rejecting the unequal offer only costs a dollar payoff of $0.3, while rewarding the first mover for cooperating costs $0.5. On top of that, the unequal offer was also very skewed at a division of 85 to 15. We find the low rate of rejecting the unequal offer in the *mUG*, compared to the high rate of cooperation after cooperation in the *sPD*, to be a surprising result.

The belief manipulation strongly influenced players' beliefs. In the selfish-belief treatment, on average, subjects believe that 34% (33%) of player 1 cooperate (offer 50-50), whereas they expect 74% (74%) to cooperate (offer 50-50) in the non-selfish belief treatment. These beliefs are significantly different from the respective data that subjects were provided with (t-tests: $p < 0.01$), and, with the exception of the non-selfish beliefs treatment in the mUG, tend to be less extreme than the provided data itself.[26]

The belief manipulation had a similar effect on player 1's actual behavior, leading to significantly more cooperation and equal offers in the non-selfish belief group compared to the selfish-belief treatment. This effect could, for example, be due to indirect learning about player 2's behavior, norms, or experimenter demand effects.[27] Relative to beliefs, player 1's behavior is less responsive to the belief manipulation. Finally, we observe significantly more cooperation (fair offers) in the selfish-belief treatment than what players expected (t-tests: $p < 0.01$) but significantly less than expected in the non-selfish belief group (sPD: $p < 0.01$, mUG: $p = 0.06$).

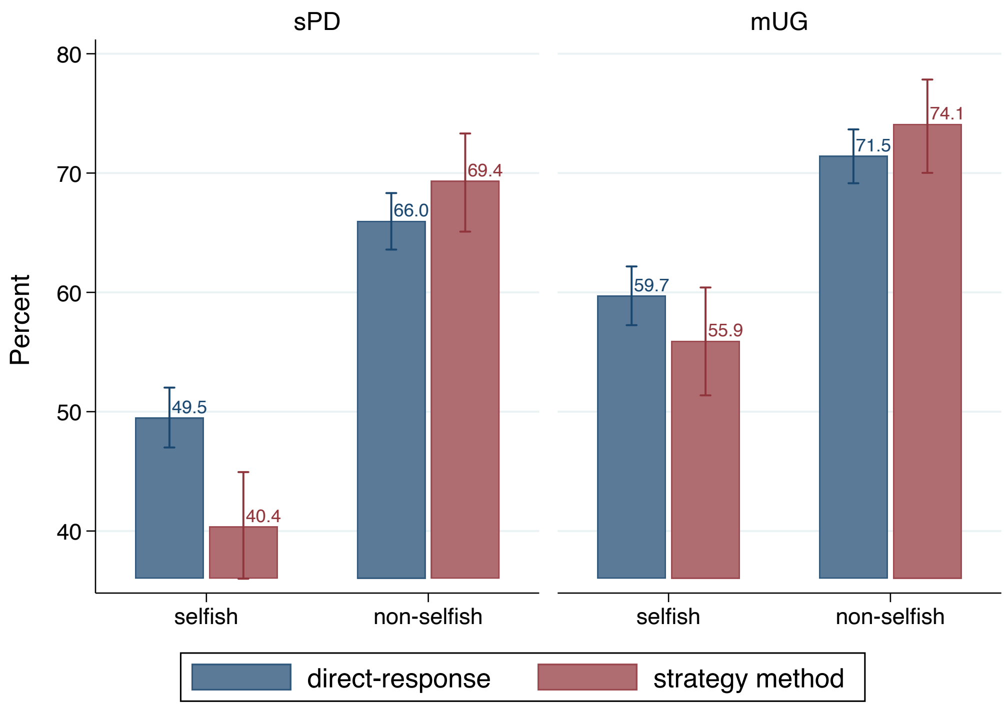

Figure 1: Player 1's behavior (frequency of cooperation or 50/50 offers) with 95% conf.-intervals (logit).

Figure 1 graphs player 1's behavior for all treatments showing in detail how the elicitation method impacts their behavior. For the selfish-beliefs group, the average rate of cooperation and fair offers tend to be higher when choices are elicited using the direct-response rather than the strategy method; for the non-selfish belief group, the pattern is reserved. Table 4, shows that this pattern is statistically significant. It reports the estimates from an OLS regression for player 1's behavior, i.e., whether he cooperates or makes the fair offer, on the non-selfish belief treatment dummy, the strategy method treatment dummy, as well as the interaction

[26]Belief histograms for both treatments can be found in Figure A.2 & A.3 in the Appendix. A regression analysis of beliefs can be found in table A.5. Regarding the regression estimates, it is not surprising that beliefs differ between the two elicitation methods for they are elicited at the end of the experiment.

[27]It was neither the aim of our study to differentiate between different causes for such behavioral change nor to explore player 1's preferences. Instead, the goal was to leverage the belief manipulation to understand how player 2's behavior depends on the elicitation method.

term of both dummies. As a result, the baseline is the direct-response, selfish beliefs treatment. In both games, we observe a relatively larger increase in cooperation and fair offers in the strategy-method as we move from selfish to non-selfish beliefs, indicated by the positive interaction term. For selfish beliefs, there is significantly less cooperation when the strategy method is employed. For the mUG, the direction is similar, but not statically significant.

Table 4: Player 1's behavior

	sPD		mUG	
Dep. Var: P1 cooperates; offers 50-50	(1)	(2)	(3)	(4)
Non-Selfish Belief	0.165***	0.166***	0.117***	0.115***
	(0.0176)	(0.0175)	(0.0170)	(0.0169)
Strategy Method	-0.0912***	-0.0907***	-0.0380	-0.0317
	(0.0262)	(0.0263)	(0.0263)	(0.0262)
Non-Selfish × Strategy Method	0.125***	0.123***	0.0646*	0.0663*
	(0.0357)	(0.0357)	(0.0349)	(0.0347)
Controls	No	Yes	No	Yes
Observations	4009	4009	4009	4009

Notes: this table reports estimates from OLS regressions. Control variables for individual characteristics include gender, age, income, highest-education, dummies for prior participation in experiments, task-order, and Covid sample. Estimates for control variables are not reported. Robust standard errors are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

The differential effect of the elicitation method on player 1's behavior is interesting in the sense that player 1's choice is elicited in the same way regardless of the elicitation method. However, player 1 is cognizant of how choices are elicited since experimental instructions were independent of role assignment, which only occurred thereafter. One reason for this could be that the elicitation method itself affects how player 1 approaches the game (fixing player 2's behavior). Qualitatively, player 1's behavior is consistent with a larger degree of strategic thinking in the strategy method. After all, a selfish-belief signal is not only indicative of player 1's behavior, but also suggestive of a more selfish response by player 2, and vice versa for the non-selfish belief signal. For the strategy method, player 1 appears more responsive to these signals.[28] An alternative explanation would be that player 1 believes that player 2's behavior is directly influenced by the elicitation method, in turn influencing player 1's preferred choice. Since we did not elicit player 1's belief about player 2's choices, we cannot test this directly with our data.

For the interested reader - and for completeness - we repeat the OLS-regressions using the 4 treatment dummies, i.e., direct-response *and* selfish beliefs, direct-response *and* non-selfish beliefs, etc., and report the difference (and the respective significance) between all treatments. These estimates can be found in table A.6 in the appendix.

3.3. Player 2

We now turn to the main focus of this paper: player 2's behavior. Player 2's choices in the sPD are depicted in Figure 2. Overall, player 2's behavior is consistent with the social theory: after player 1 cooperates,

[28] In terms of material (expected) payoffs payoffs, we find that they are often fairly similar (within \$0.1) in the sPD (see the by-treatment, by-sample data in Table E.24). Pre-Covid, cooperation (defection) is payoff maximizing for the strategy method (direct-response) regardless of belief-treatment. The same patter is true for the overall sample. In the Covid sample, defection is optimal except in the selfish-beliefs, strategy method treatment. For the mUG, the unfair offer does significantly better given the few rejections. As we do not elicit player 1's belief about player 2, we cannot test whether they act as if they maximize their own payoffs.

over 60% of player 2s in all combinations of elicitation methods and belief groups cooperate. In response to defection, more than 75% in all groups defect.

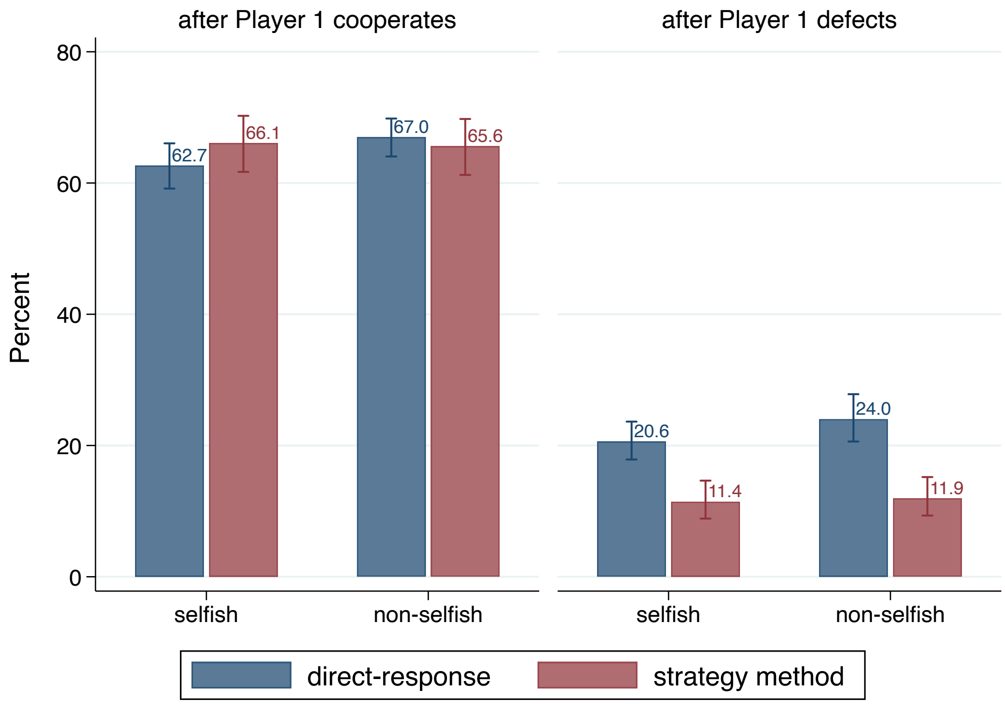

Figure 2: Player 2's behavior (frequency of cooperation) in sPD with 95% conf.-intervals (logit).

Overall, behavior in response to player 1 cooperating is remarkably similar across our four treatment groups. Indeed, the only significant difference between any of the groups is between the direct-response treatment with selfish and non-selfish beliefs ($p < 0.1$ without and $p < 0.05$ with controls), where we observe a small uptick in cooperation with non-selfish beliefs, which can be explained by social norms (among others).[29] In contrast to the economically large response of player 1 to the belief manipulation, the change is slight.[30] Table 5 reports our respective OLS-estimates for this setting. The elicitation method does not appear to have any effect on player 2's response to cooperation.

In contrast, the elicitation method strongly affects player 2's response to player 1 defecting: the strategy method cuts the rate of cooperation in half; subjects make significantly less mistakes if the strategy method is used. The best explanation for this effect is that the strategy method forces player 2 to pay more attention when they are required to make two choices, reducing (random) mistakes. For example, they are less likely to look at the incorrect node of the game/payoff column or think that player 1 took a different action.[31] It is important to emphasize that this reduction in mistakes cannot be explained by alternative theories that view cooperation in response to defection as true reflections of preferences. For example, while preferences for efficiency predict cooperation after defection, they cannot explain why these preferences vary

[29] This positive effect, in turn, also suggest that positive reciprocity due to surprise (Khalmetski et al. (2015)) may not be an important driver of behavior in our sample.

[30] The regression table that report treatment differences between groups can be found in the appendix, see Table A.7.

[31] The alternative hypothesis that player 2 simply understands the game better when she is forced to look at both nodes of the game cannot explain this behavior, however. While in isolation, this may explain why there is less cooperation in response to defection as player 2 realizes that player 1 took the worst action for her, she must also come to the opposite conclusion after cooperation. This in turn would predict more conditional cooperation in the strategy method, which is not true in the data.

Table 5: Player 2's behavior in sPD

	after P1 cooperates		after P1 defects	
Dep. Var: Player 2 cooperates	(1)	(2)	(3)	(4)
Non-Selfish Belief	0.0433*	0.0472**	0.0342	0.0302
	(0.0230)	(0.0226)	(0.0235)	(0.0232)
Strategy Method	0.0344	0.0401	-0.0917***	-0.110***
	(0.0280)	(0.0279)	(0.0208)	(0.0212)
Non-Selfish × Strategy Method	-0.0481	-0.0580	-0.0291	-0.0173
	(0.0384)	(0.0383)	(0.0315)	(0.0315)
Controls	No	Yes	No	Yes
Observations	2722	2722	2247	2247

Notes: this table reports estimates from OLS regressions, with control variables identical to those in Table 4. Estimates for control variables are not reported. Robust standard errors are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

with the elicitation method. Moreover, given that we do not see any influence of the elicitation method *after cooperation*, i.e., the part of the game where social-preferences play the largest role, it is difficult to conceive why there should be a differential effect on preferences after defection.[32] Finally, we do not observe any difference in behavior after defection due to different beliefs

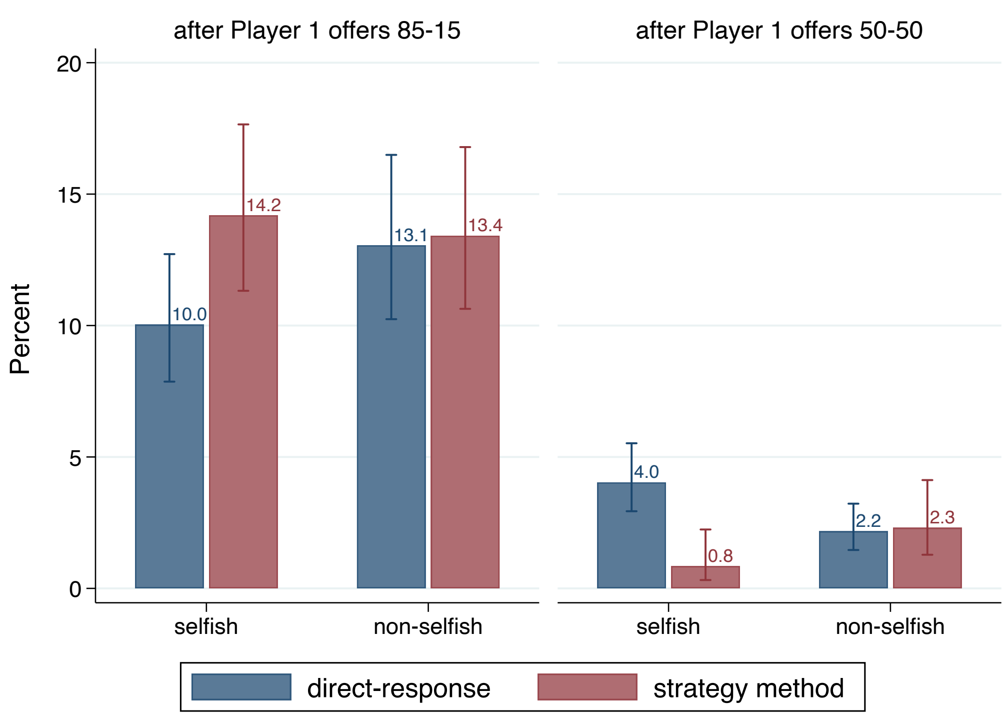

Figure 3: Player 2's behavior (frequency of rejection) in UG with 95% conf.-intervals (logit).

Figure 3 depicts Player 2's behavior in the *mUG*. Overall, rejection rates are low, below 15% in all treatment groups. Consequently, behavior is well predicted by the selfish theory. What the social-theory

[32]Previously, we had already highlighted that behavior consistent with preferences for efficiency (yet which could still be an honest mistake) is at most 6%, see Table A.4. This suggest that most of the reduction in cooperation is due to reduced mistakes.

gets correct, is the relative behavior across the two nodes: rejection rates of unequal offers are no less than 10% while they are no higher than 4% for equal splits.

Before we look at the treatment effects in details, it is worth emphasizing that in contrast to the *sPD*, the absolute differences for any treatment group are very small. Consequently, the economic significance of the elicitation method is minor when it comes to behavior. From the right-hand panel of Figure 3, we see that the selfish-belief, direct-response group features the least rejection.[33] As a result, our regression estimates in Table 6 indicate that the strategy method (dummy) results in more rejections. Like in the *sPD*, there is no evidence for a differential effect on behavior as the interaction term remains insignificant. Moreover, beliefs do affect rejections of unfair offers.

Table 6: Player 2's behavior in mUG

	after P1 offers 85-15		after P1 offers 50-50	
Dep. Var: Player 2 rejects	(1)	(2)	(3)	(4)
Non-Selfish Belief	0.0302	0.0277	-0.0186**	-0.0195**
	(0.0201)	(0.0202)	(0.00785)	(0.00773)
Strategy Method	0.0416**	0.0410**	-0.0319***	-0.0363***
	(0.0202)	(0.0201)	(0.00775)	(0.00810)
Non-Selfish × Strategy Method	-0.0380	-0.0399	0.0332***	0.0362***
	(0.0301)	(0.0300)	(0.0113)	(0.0116)
Controls	No	Yes	No	Yes
Observations	1999	1999	2970	2970

Notes: this table reports estimates from OLS regressions, with control variables identical to those in Table 4. Estimates for control variables are not reported. Robust standard errors are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

With regard to player 2's behavior after equal splits, the picture is more complicated than in the *sPD*. When choices are elicited using the strategy method, player 2 is significantly less likely to reject offers in the selfish-belief manipulation compared to the direct-response method. This observation is consistent with our previous argument that the strategy method results in less mistakes, as rejection of equal offers can neither be explained by selfish-preferences nor typical social preferences. Interestingly, we do not observe this pattern for non-selfish belief groups, where rejection rates are essentially identical for the two elicitation methods.

The behavior in the two games display a similar pattern: when the belief manipulation points towards the other action, we observe slightly more mistakes in the direct-response method. In the *sPD*, we observe more mistakes after defection in the direct-response method when the belief manipulation points towards cooperation (i.e., the non-selfish belief). In the *mUG*, player 2 rejects more equal offers in the direct-response method when the belief manipulation points towards the unequal offer (i.e., the selfish belief). Note that this effect occurs on top of the general increase in mistakes for the direct-response method relative to the strategy method. One reason for the higher rates of mistake could be that player 2 not solely response to the actual choice of player 1 is shown to her, but instead is influenced by her subjective (ex-ante) belief of about what player 1 is going to do, leading to confusion regarding which node she is at. Alternatively, it is not her belief that leads to such confusion but the information that is provided by the belief-treatment. Table 7 below lends credence to the first interpretation, which shows that the treatment itself is not correlated to the rate of mistakes but that the belief in the alternative node (same node) is associated with more (less)

[33]The difference is significant at $p < 0.05$ compared to the strategy method with selfish beliefs. For a comparison between all treatment groups, consult Table A.8 in the appendix.

mistakes for players in the direct-response method.

Table 7: Player 2's behavior: Mistakes and Beliefs

	sPD: after P1 defects		mUG: after P1 offers 50-50	
Dep. Var: Player 2 makes mistake	(1)	(2)	(3)	(4)
Non-Selfish Belief	0.00367	-0.00180	-0.00389	-0.00598
	(0.0288)	(0.0281)	(0.00755)	(0.00756)
Belief Player 1 cooperates	0.0798*	0.0757*		
	(0.0422)	(0.0413)		
Belief Player 1 offers 85-15			0.0399***	0.0372***
			(0.0131)	(0.0130)
Controls	No	Yes	No	Yes
Observations	1298	1298	2021	2021

Notes: this table reports estimates from OLS regressions, with control variables identical to those in Table 4, for player 2s in the direct-response method. Estimates for control variables are not reported. Robust standard errors are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

So far, we observed how the strategy method leads to a fewer cooperation after defection and rejections of fair offers. Guided by the theories, players' behavioral patterns, and implied distribution of player types, we classified such choices as mistakes and concluded that their frequency are reduced by the strategy method. We now provide further evidence for this idea.

In particular, if the strategy method reduces the likelihood of mistakes, we would expect players who are more prone to mistakes be more strongly affected by the strategy method than those who are not. To test this hypothesis, we leverage our experimental design. After the introduction of our experiment and all related rules, subjects were required to answer a series of control questions, which is typical for experiments of this type. We reckon that those who made mistakes in the questionnaire are more likely inattentive and/or generally care less about (their performance in) the experiment. Consequently, we expect them to make (i) more mistakes when playing the games and (ii) to be more affected by the strategy method with regards to the prevention of mistakes. We refer to this group as *inattentive* people below and test this idea by regressing player 2's choice, i.e., whether they made a mistake, on the strategy method dummy, whether the person is considered inattentive, and the interaction of the two. The results are reported in Table 8.

Table 8: Player 2's behavior: Mistakes and Inattention

	sPD: after P1 defects		mUG: after P1 offers 50-50	
Dep. Var: Player 2 makes mistake	(1)	(2)	(3)	(4)
Strategy Method	-0.116***	-0.0772***	-0.0172***	-0.00625
	(0.0157)	(0.0179)	(0.00568)	(0.00541)
Inattentive	0.0814***	0.135***	0.0342***	0.0451***
	(0.0180)	(0.0270)	(0.00728)	(0.0101)
Inattentive × Strategy Method		-0.118***		-0.0319**
		(0.0351)		(0.0140)
Controls	Yes	Yes	Yes	Yes
Observations	2247	2247	2970	2970

Notes: this table reports estimates from OLS regressions, with control variables identical to those in Table 4. Estimates for control variables are not reported. Robust standard errors are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

First, we note that the coefficients of the *inattentive* dummy is positive and significant, confirming the conjecture (i) that inattentive people make more mistakes in general. By itself, this also provides further support for the viewpoint that cooperation after defection and rejecting an equal offer are simple mistakes. After all, one would not expect preference for such behavior to be correlated with making more mistakes in a control questionnaire. Second, the interaction term in column (2) and (4) are negative and significant. In other words, the strategy method reduces the likelihood of making a mistake relatively more when the player is inattentive, confirming our prediction (ii).[34]

3.4. Robustness and Sample Comparisons

In the online appendix, we further investigate the robustness of our results as well as compare the pre-Covid (October to November 2019) and Covid sample (October 2021) in terms of participants' characteristic and behavior.

In particular, we first check for the robustness of our OLS estimates by re-estimating all tables from this section using Probit and Logit regressions (section D.1). The results are robust to such a change in estimator. We also revisit the topic of order effects in our data (section D.2), which had been highlighted previously in footnote 21. We show that, in general, order effects exist, yet that our analysis restricted to task 1 yields qualitatively similar results despite the reduction in sample size by 50%.

The sample comparison (section E) highlights that participants in the Covid sample tend to be slightly more educated, have higher incomes, and are slightly older. The most notable change in terms of behavior are the higher rejection rates of unfair offers in the later sample.

More detailed commentary for each aspect can be found in the respective sections of the appendix.

4. Conclusion

We conducted a large-scale online experiment, where, in addition to varying the elicitation method, we also manipulated the ex-ante beliefs of participants about player 1's likely choices via an information-provision treatment. In neutrally-framed sequential games, a sequential Prisoners' Dilemma and a mini-Ultimatum Game, we found that the elicitation method does not alter player 2's preferences (their tendency to reward cooperation or reject unfair offers) yet significantly reduces mistakes (rewarding defection or rejecting fair offers). The takeaway from our study is that unless researchers have a particular practical reason for using the direct-response method, they should opt for the more economical strategy method as their go-to method for eliciting behavior in sequential games.

This conclusion may be particularly relevant for many experiments conducted online nowadays. Here, the strategy methods appears to be the superior choice as it reduces mistakes, makes running experiments cheaper and, in many cases, easier to run. While we expect these insights to generally extend to classical on-site laboratory experiments (the strategy method is always more economical) – especially when the researcher is interested in true one-shot behavior – we recognize the inherent difference of such setting to unsupervised online-experiments. It is simply less feasible to run several rounds of practice games, with the opportunity for Q&A during the instructions and/or such practice rounds, in an online setting. Similarly, it is likely more challenging to maintain the attention of participants as well as limit dropouts over an extended amount of time online – although Arechar et al. (2018) suggests that it is possible.

[34]When presenting this table, we opted omit the belief-treatment dummy as it greatly improves the presentation of the interactions, and allows us to focus on the main question. The pattern remains true for more involved regressions. For completeness, estimates of Table 8 without controls can be found in the online appendix (Table D.15), with results being unaffected by the inclusion/exclusion of such controls.

One limitation of our study was the focus on simple one-shot binary-choice games, which may limit the external validity of our findings with regards to more complex settings. Those limitations were intentional however, as we regarded this setting not only to be the most suitable one for testing our initial conjecture (that ex-ante beliefs may affect behavior in a differential way across elicitation methods) but also the most relevant for informing future studies and relating to the literature at large. After all, Brandts and Charness (2011, p. 391) cautiously suggested in their literature review that differences between elicitation methods may be more likely for fewer contingent choices and less likely for games with more periods.

Finally, what really surprised us was that participants showed a strong tendency for positive reciprocity in the *sPD* yet seemed very unwilling to punish the low offer in the *mUG*, despite the fact that the personal cost at \$0.15 of doing so seemed low (to us). The work by Amir et al. (2012), who run a dictator, ultimatum, trust, and a public goods game, suggests that MTurk workers are not generally averse to rejecting low offers: the average minimum-acceptable offer in their study lies in the mid 30s for both a no stakes and a \$1 treatment.[35] If anything, our results are reminiscent of Charness and Rabin (2002), who test social preferences in a series of simple experimental games, which were framed and explained (both graphically and in words) in the most neutral of way possible and elicited with the strategy method. Similar to our results, they also document very weak negative reciprocity. One might thus conjecture that the usual, fairly harmless framing of an ultimatum bargaining game with "offers" and the opportunity to "reject" may increase or trigger the participants' willingness to forgo monetary payoffs in order to punish others compared to its fully neutrally framed counterpart.

[35] Their Ultimatum Game was the standard discrete version, with offers made from $\{0, 10, \ldots, 100\}$. Player 2 indicated whether they accept/reject each possible offer, i.e., their preferences were elicited by the strategy method.

A. Appendix: Supportive Tables

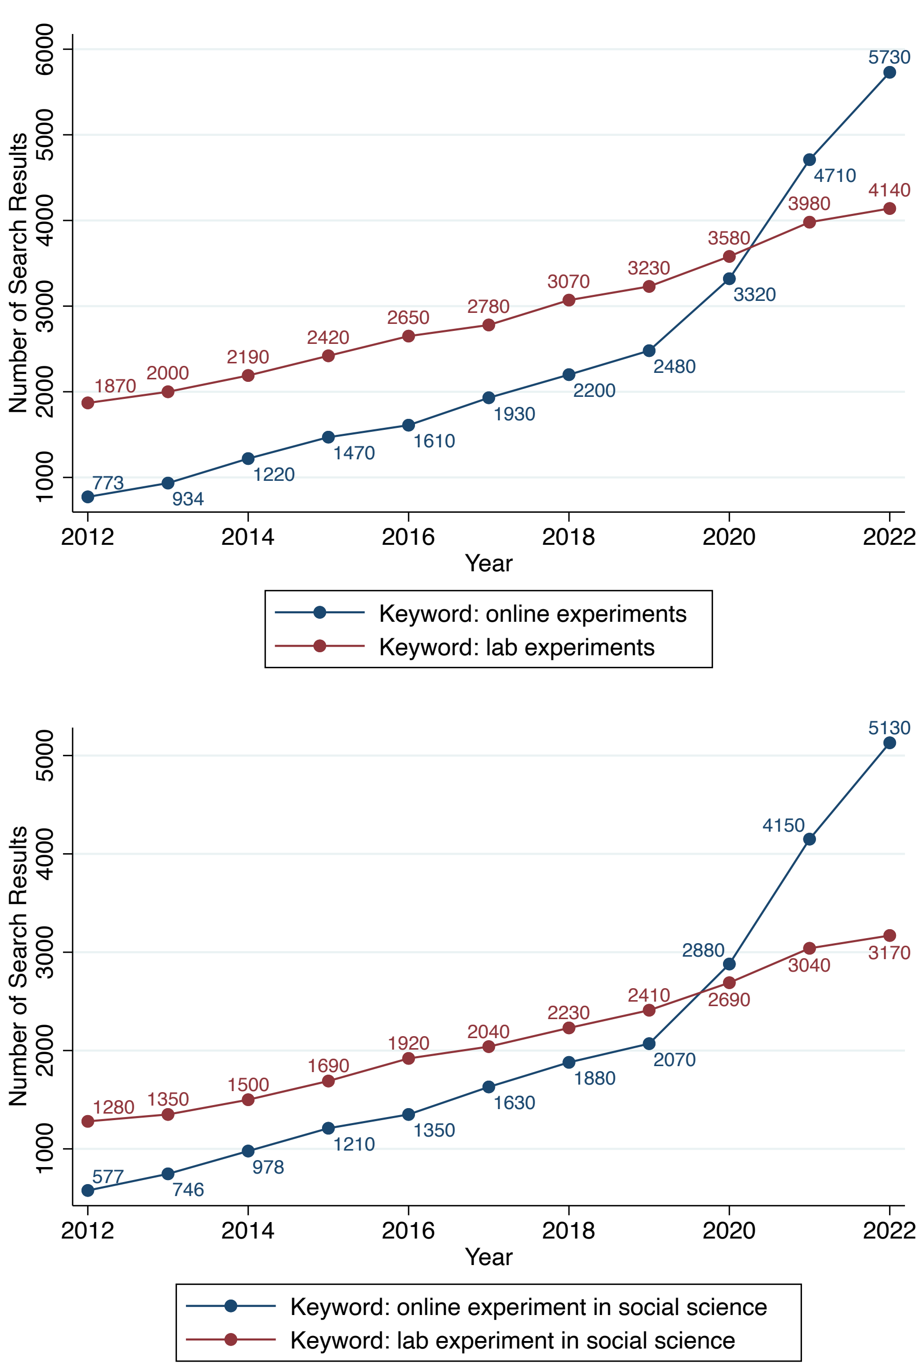

Figure A.1: Number of Search Results for Lab/Online Experiments on Google Scholar

Note: the figure displays the numbers of search results from a Google Scholar Search using the keywords online experiment" or "lab experiment", with or without being restricted to the social sciences, which relied on the keywords economics, psychology, management, business, or politics (*date accessed: 3. April 2023*).

Table A.1: Variable Defintions

Variable Name	Definitions
Behavior and Beliefs	
Player 1 cooperates	= 1 if player 1 chooses to cooperate in the sPD; 0 otherwise
Player 2 cooperates after C	= 1 if player 2 chooses to cooperate in response to player 1 cooperating in the sPD; 0 otherwise
Player 2 cooperates after D	= 1 if player 2 chooses to cooperate in response to player 1 defecting in the sPD; 0 otherwise
Belief Player 1 cooperates	A player's belief regarding the % of player 1 that cooperate in the sPD
Player 1 offers 50-50	= 1 if player 1 offers 50-50 in the mUG; 0 otherwise
Player 2 rejects 50-50	= 1 if player 2 rejects the 50-50 offer in the mUG; 0 otherwise
Player 2 rejects 85-15	= 1 if player 2 rejects the 85-15 offer in the mUG; 0 otherwise
Belief Player 1 offers 50-50	A player's belief regarding the % of player 1 that offer 50-50 in the mUG
Player 2 makes mistake	= 1 if player 2 cooperates after defection in the sPD (rejects the 50-50 offer in the mUG); 0 otherwise
Treatment Variables / Other Indicator Variables	
Strategy Method (SM)	= 1 if in strategy method treatment, i.e., behavior of player/opponent is elicited using the strategy method; 0 otherwise
Direct Response (DR)	= 1 if in direct response treatment, i.e., behavior of player/opponent is elicited using the direct response method; 0 otherwise
Selfish (Belief)	= 1 if in selfish belief treatment; 0 otherwise
Non-Selfish (Belief)	= 1 if in non-selfish belief treatment; 0 otherwise
Direct Response, Selfish	= 1 if in direct-response *and* selfish belief treatment; 0 otherwise
Direct Response, Non-Selfish	= 1 if in direct-response *and* non-selfish belief treatment; 0 otherwise
Strategy Method, Selfish	= 1 if in strategy method *and* selfish belief treatment; 0 otherwise
Strategy Method, Non-Selfish	= 1 if in strategy method *and* non-selfish belief treatment; 0 otherwise
Player 2	= 1 if participants plays in the role of the second mover; 0 otherwise
P1 (P2)	Short version for Player 1 (2)
Task 2	= 1 if a given game (sPD or mUG) was played second; 0 otherwise
pre-Covid	= 1 if observation is from the pre-Covid sample; 0 otherwise
Covid	= 1 if observation is from the Covid sample; 0 otherwise
Inattentive	=1 if participants makes a mistakes in the instruction test; 0 otherwise
x × y	Interaction term: = 1 if both x and y is true (=1); 0 otherwise
x × y × z	Interaction term: = 1 if x, y, and z are true (=1); 0 otherwise
Personal Information from Survey	
Gender	Categorical variable that indicates whether participant is female, male, other/prefer not to say
Age	Categorical variable that indicates whether participant is $<$ 12, 12-17, 18-24, 25-34, 35-44, 45-54, 55-64, 65-74, $\geq$ 75 years old; or prefers not to say
Income	Categorical variable that indicates whether the participant's annual household income (in USD) is $<$ 20 000, 20 000 - 34 999, 35 000 - 49 999, 50 000 - 74 999, 75 000 - 99 999, 100 000 - 140 999, or $\geq$ 150 000; or prefers not to say
Education	Categorical variable that indicates whether the highest degree that the participant is holding or currently pursuing is no degree, High School, Bachelor, Master, Doctorate, Other post-graduate degree; or prefers not to say
Participated in experiments before	=1 if participant has participated in similar experiments as this one before; 0 otherwise

Table A.2: Summary Statistics

	Median	Mean	Std. Dev.	Min.	Max.	Obs.
Sequential Prisoner's Dilemma						
Player 1 cooperates	1.00	0.572	0.49	0	1.0	4009
Player 2 cooperates after C	1.00	0.654	0.48	0	1.0	2722
Player 2 cooperates after D	0.00	0.177	0.38	0	1.0	2247
Belief Player 1 cooperates	0.64	0.541	0.32	0	1.0	8029
mini Ultimatum Game						
Player 1 offers 50-50	1.00	0.655	0.48	0	1.0	4009
Player 2 rejects 50-50	0.00	0.026	0.16	0	1.0	2970
Player 2 rejects 85-15	0.00	0.125	0.33	0	1.0	1999
Belief Player 1 offers 50-50	0.62	0.536	0.32	0	1.0	8029
Other Game Outcomes						
Total earnings from games (in USD)	2.00	1.752	0.60	0	3.6	8029
Total time (in sec.)	304.30	367.037	280.80	32	4846.3	8029
Number of mistakes in understanding test	0.00	0.311	0.46	0	1.0	8029
Treatments						
Direct Response, Selfish	0.00	0.379	0.49	0	1.0	8029
Direct Response, Non-selfish	0.00	0.385	0.49	0	1.0	8029
Strategy Method, Selfish	0.00	0.116	0.32	0	1.0	8029
Strategy Method, Non-selfish	0.00	0.120	0.32	0	1.0	8029
Players, Game Order, Sample						
Player 2	1.00	0.501	0.50	0	1.0	8029
Player 2 $\times$ Strategy Method	0.00	0.118	0.32	0	1.0	8029
Played sPD as 2nd task	1.00	0.501	0.50	0	1.0	8029
Covid	0.00	0.421	0.49	0	1.0	8029

Table A.2: Summary Statistics (continued)

	Median	Mean	Std. Dev.	Min.	Max.	Obs.
Personal Information from Survey						
Participated in experiments before	1.00	0.742	0.44	0	1.0	8029
Gender:						
Female	1.00	0.517	0.50	0	1.0	8029
Male	0.00	0.475	0.50	0	1.0	8029
Other / Prefer not to say	0.00	0.008	0.09	0	1.0	8029
Age:						
$<$ 12 years	0.00	0.000	0.00	0	0.0	8029
12-17 years old	0.00	0.000	0.01	0	1.0	8029
18-24 years old	0.00	0.073	0.26	0	1.0	8029
25-34 years old	0.00	0.370	0.48	0	1.0	8029
35-44 years old	0.00	0.277	0.45	0	1.0	8029
45-54 years old	0.00	0.153	0.36	0	1.0	8029
55-64 years old	0.00	0.091	0.29	0	1.0	8029
65-74 years old	0.00	0.031	0.17	0	1.0	8029
$\geq$ 75 years	0.00	0.003	0.05	0	1.0	8029
Prefer not to say	0.00	0.002	0.05	0	1.0	8029
Income:						
Less than 20 000	0.00	0.000	0.00	0	0.0	8029
20 000 to 34 999	0.00	0.152	0.36	0	1.0	8029
35 000 to 49 999	0.00	0.178	0.38	0	1.0	8029
50 000 to 74 999	0.00	0.241	0.43	0	1.0	8029
75 000 to 99 999	0.00	0.155	0.36	0	1.0	8029
100 000 to 140 999	0.00	0.108	0.31	0	1.0	8029
over 150 000	0.00	0.048	0.21	0	1.0	8029
Prefer not to say	0.00	0.022	0.15	0	1.0	8029
Education:						
No Degree	0.00	0.009	0.09	0	1.0	8029
High School Degree	0.00	0.262	0.44	0	1.0	8029
Bachelor Degree	1.00	0.503	0.50	0	1.0	8029
Master Degree	0.00	0.158	0.37	0	1.0	8029
Other Post-Grad Degree	0.00	0.031	0.17	0	1.0	8029
Doctorate Degree	0.00	0.025	0.16	0	1.0	8029
Prefer not to say	0.00	0.012	0.11	0	1.0	8029

Table A.3: Randomization Check

	DR & Selfish (1)			DR & Non-selfish (2)			SM & Selfish (3)			SM & Non-selfish (4)		
	Obs	Mean	Std. Dev.	Obs	Mean	Std. Dev.	Obs	Mean	Std. Dev.	Obs	Mean	Std. Dev.
Female	3,040	0.52	0.50	3,094	0.52	0.50	935	0.51	0.50	960	0.53	0.50
Age	3,036	39.41	12.08	3,087	39.14	12.20	935	39.35	12.12	954	39.29	11.80
Participated in experiments before	3,040	0.74	0.44	3,094	0.74	0.44	935	0.74	0.44	960	0.74	0.44
Household income	2,206	65,607.43	38,187.38	2,273	65,849.10	38,399.20	675	64,662.96	36,258.21	680	63,382.35	36,091.66
No Degree	3,040	0.01	0.10	3,094	0.01	0.10	935	0.01	0.10	960	0.01	0.07
High School Degree	3,040	0.27	0.44	3,094	0.27	0.44	935	0.24	0.42	960	0.26	0.44
Bachelor Degree	3,040	0.50	0.50	3,094	$0.49^{(3)**}$	0.50	935	0.54	0.50	960	0.52	0.50
Master or above	3,040	0.21	0.41	3,094	0.22	0.42	935	0.20	0.40	960	0.21	0.41

Notes: This table reports the Pairwise randomization tests between treatments. To keep the table succinct, we aggregate the age and household income categories. Moreover, small category values such as “Other” or “Prefer not to say” for gender and highest education are omitted, noting that none of these subcategories display significant differences. The superscript next to the mean of each treatment shows the column number to which treatment (column) is compared, and the asterisks mark the significance level of the difference following the conventional manner. If, for a given variable, two treatments are not significantly different at conventional levels, no superscript is added. This comparison is only conducted to the “right” to avoid double counting, i.e., DR & Selfish (1) is compared to DR & Non-selfish (2), SM & Selfish (3), SM & Non-selfish (4), DR & Non-selfish (2) is compared to SM & Selfish (3) and SM & Non-selfish (4), etc. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Table A.4: Frequency of Player 2 Types

Strategy Method	sPD only	mUG only	Both Games
Social	0.60	0.13	0.10
Selfish	0.28	0.85	0.27
Other	0.12	0.02	0.63
Direct Response	sPD only	mUG only	Both Games
Social	0.71	0.68	0.48
Selfish	0.53	0.94	0.51
Other	0.09	0.02	0.22
Alternative Preferences	sPD only	mUG only	Both Games
Efficiency	0.06	0.85	0.05
Spite	0.28	0.01	0.00

Notes: A player classified as social/selfish for *sPD* or *mUG* only if their behavior is consistent with the respective theoretical prediction for that particular game. If it is consistent with neither, s/he is classified as *other*. Similarly, if their behavior is consistent with the predictions of theory for *both games*, then she is classified into the respective category. Further note that for the strategy method, the three categories are distinct and exhaustive. However, given that only one action of player 2 is observed in the direct response method, a player may be categorized as more than one type, and as such, frequencies do not need to sum to 1.

The alternative preferences are computed for only based on the strategy-method. Player 2 has preferences for efficiency if she always cooperates and always accepts. Her preferences are classified as spiteful if she always defects and always rejects. Note that these preferences are only different from selfish preferences for one of the two games, i.e., efficiency (spiteful) preferences matches selfish-preferences for the *mUG* (*sPD*), or when we consider both games jointly.

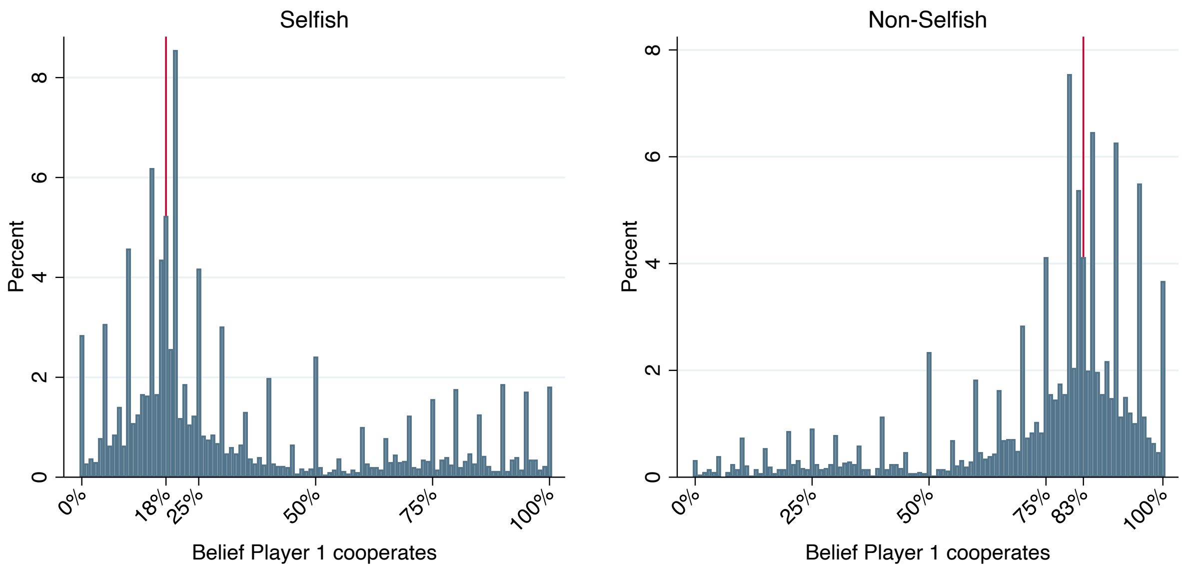

Figure A.2: Beliefs that Player 1 cooperates in the seq. Prisoner's Dilemma

Note: Red Vertical Lines indicates provided probabilities.

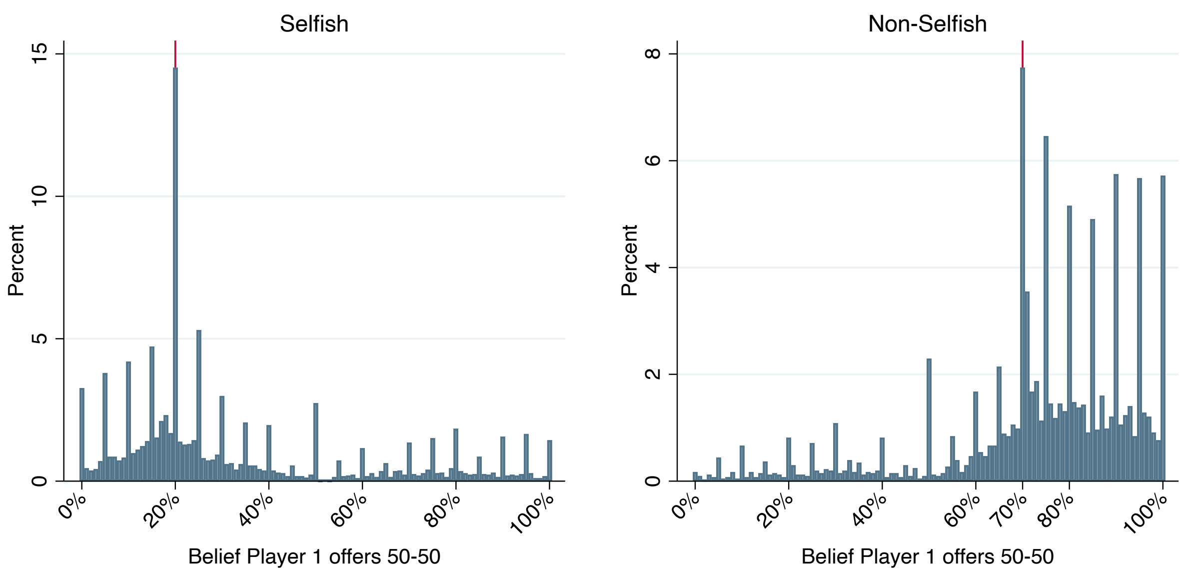

Figure A.3: Beliefs that Player 1 offers 50/50 in the mini Ultimatum Game

Note: Red Vertical Lines indicates provided probabilities.

Table A.5: Beliefs about Player 1's Behavior

	sPD				mUG			
Dep. Var: Belief P1 takes Non-Selfish Action	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Non-Selfish Belief	0.399***	0.399***	0.399***	0.434***	0.402***	0.402***	0.401***	0.407***
	(0.00649)	(0.00649)	(0.00641)	(0.00860)	(0.00623)	(0.00623)	(0.00619)	(0.00843)
Strategy Method	-0.0187*	-0.0196*	-0.0201**	-0.0142	-0.0205**	-0.0209**	-0.0213**	-0.00224
	(0.0101)	(0.0101)	(0.00989)	(0.0122)	(0.00976)	(0.00977)	(0.00976)	(0.0128)
Non-Selfish × Strategy Method	0.0149	0.0156	0.0166	0.0191	0.0140	0.0144	0.0150	0.0118
	(0.0130)	(0.0130)	(0.0128)	(0.0170)	(0.0125)	(0.0125)	(0.0125)	(0.0169)
Player 2			0.0785***	0.118***			0.0509***	0.0650***
			(0.00555)	(0.0100)			(0.00537)	(0.00974)
Player 2 × Strategy Method				-0.0121				-0.0378*
				(0.0197)				(0.0195)
Player 2 × Non-Selfish				-0.0707***				-0.0117
				(0.0128)				(0.0124)
Player 2 × Strategy Method × Non-Selfish				-0.00515				0.00610
				(0.0256)				(0.0249)
Controls	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Observations	8029	8029	8029	8029	8029	8029	8029	8029

Notes: this table reports estimates from OLS regressions with the dependent variable being a player's belief regarding the % of player 1 that either cooperate in the *sPD* or offer 50-50 in the mUG. The control variables are identical to those in Table 4. Robust standard errors are reported in parentheses * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

For column (4), the linear combinations indicate that player 2 is more optimistic for the Strategy Method and Selfish Beliefs (diff = 0.1056, p-value < 0.01) and Strategy Method and Non-Selfish Beliefs (diff = 0.0298, p-value = 0.036) than player 1 in the respective treatments. In column (8), while positive, neither of the two linear combinations is significant at 10% (with p-values of 0.107 and 0.109).

Table A.6: Player 1's behavior

	sPD		mUG	
Dep. Var: P1 cooperates; offers 50-50	(1)	(2)	(3)	(4)
Direct Response, Non-Selfish	0.165***	0.166***	0.117***	0.115***
	(0.0176)	(0.0175)	(0.0170)	(0.0169)
Strategy Method, Selfish	-0.0912***	-0.0907***	-0.0380	-0.0317
	(0.0262)	(0.0263)	(0.0263)	(0.0262)
Strategy Method, Non-Selfish	0.198***	0.198***	0.144***	0.150***
	(0.0246)	(0.0245)	(0.0236)	(0.0232)
Estimated Differences				
SM, Selfish − DR, Non-Selfish	-0.256***	-0.256***	-0.155***	-0.147***
	(0.0258)	(0.0260)	(0.0258)	(0.0258)
SM, Non-Selfish − DR, Non-Selfish	0.0336	0.0328	0.0266	0.0345
	(0.0242)	(0.0242)	(0.0230)	(0.0228)
SM, Non-Selfish − SM/Selfish	0.290***	0.289***	0.182***	0.181***
	(0.0310)	(0.0311)	(0.0305)	(0.0303)
Controls	No	Yes	No	Yes
Observations	4009	4009	4009	4009

Notes: this table reports estimates from OLS regressions in the top panel and estimated difference between the treatments, with control variables are identical to those in Table 4. Estimates for control variables are not reported. Moreover, note that the omitted category is the direct-response, selfish beliefs treatment, meaning that the regression estimates in the top panel represent the difference to this treatment group. Robust standard errors are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Table A.7: Player 2's behavior in the sPD

	after P1 cooperates		after P1 defects	
Dep. Var: Player 2 cooperates	(1)	(2)	(3)	(4)
Direct Response, Non-Selfish	0.0433*	0.0472**	0.0342	0.0302
	(0.0230)	(0.0226)	(0.0235)	(0.0232)
Strategy Method, Selfish	0.0344	0.0401	-0.0917***	-0.110***
	(0.0280)	(0.0279)	(0.0208)	(0.0212)
Strategy Method, Non-Selfish	0.0295	0.0292	-0.0866***	-0.0968***
	(0.0280)	(0.0279)	(0.0209)	(0.0210)
Estimated Differences				
SM, Selfish − DR, Non-Selfish	-0.00893	-0.00707	-0.126***	-0.140***
	(0.0263)	(0.0262)	(0.0235)	(0.0235)
SM, Non-Selfish − DR, Non-Selfish	-0.0138	-0.0179	-0.121***	-0.127***
	(0.0263)	(0.0263)	(0.0236)	(0.0234)
SM, Non-Selfish − SM/Selfish	-0.00483	-0.0108	0.00509	0.0129
	(0.0308)	(0.0309)	(0.0209)	(0.0212)
Controls	No	Yes	No	Yes
Observations	2722	2722	2247	2247

Notes: this table reports estimates from OLS regressions in the top panel and estimated difference between the treatments, with control variables are identical to those in Table 4. Estimates for control variables are not reported. Moreover, note that the omitted category is the direct-response, selfish beliefs treatment, meaning that the regression estimates in the top panel represent the difference to this treatment group. Robust standard errors are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Table A.8: Player 2's behavior in the mUG

	after P1 offers 85-15		after P1 offers 50-50	
Dep. Var: Player 2 rejects	(1)	(2)	(3)	(4)
Direct Response, Non-Selfish	0.0302	0.0277	-0.0186**	-0.0195**
	(0.0201)	(0.0202)	(0.00785)	(0.00773)
Strategy Method, Selfish	0.0416**	0.0410**	-0.0319***	-0.0363***
	(0.0202)	(0.0201)	(0.00775)	(0.00810)
Strategy Method, Non-Selfish	0.0338*	0.0289	-0.0173*	-0.0195**
	(0.0199)	(0.0198)	(0.00946)	(0.00943)
Estimated Differences				
SM, Selfish − DR, Non-Selfish	0.0114	0.0133	-0.0133**	-0.0168***
	(0.0226)	(0.0225)	(0.00609)	(0.00637)
SM, Non-Selfish − DR, Non-Selfish	0.00364	0.00118	0.00132	-0.0000752
	(0.0223)	(0.0222)	(0.00816)	(0.00824)
SM, Non-Selfish − SM/Selfish	-0.00778	-0.0122	0.0146*	0.0167**
	(0.0224)	(0.0223)	(0.00807)	(0.00844)
Controls	No	Yes	No	Yes
Observations	1999	1999	2970	2970

Notes: this table reports estimates from OLS regressions in the top panel and estimated difference between the treatments, with control variables are identical to those in Table 4. Estimates for control variables are not reported. Moreover, note that the omitted category is the direct-response, selfish beliefs treatment, meaning that the regression estimates in the top panel represent the difference to this treatment group. Robust standard errors are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

B. Appendix: Instructions of Experiment

Below, you will find the exact instructions of our experiment.[36]

Introduction

Thank you for accepting this HIT. If you choose to continue with this job, you will participate in an experiment on decision making. This experiment has three parts: two decision tasks and a short survey with 7 questions. The entire experiment will take 10 minutes to complete. You will earn $1 for completing the HIT and, depending on your choices and the choices of other MTurk workers, an additional amount of up to $3. If you wish to continue with this HIT, please ensure you have sufficient time to complete the whole study. In each decision task, you will be randomly matched with another participant. The interaction is completely anonymous. Neither you nor the other worker will know the other person's worker ID. This experiment follows a no-deception policy. All tasks are implemented exactly as outlined in the instructions. The instructions are the same for all participants that you may interact with. All participants are real MTurk workers. Finally, your earnings and decision in each part of the experiment do not depend on earnings and decisions in other parts. If you wish to continue with this HIT, please ensure you have sufficient time to complete the whole study. **Please do not close this page during the experiment.** If you leave the website during the experiment, you will **not** receive any earnings. Moreover, you will only be able to participate in this experiment once. By clicking the next button, you consent to taking part in this experiment and promise to do your best to complete the whole experiment. [Next Page Button]

Instructions

We now explain how the decision tasks work. Please read these instructions carefully as we will ask you some simple questions about it on the next page.

Who you interact with

In each of the two decision tasks, you will be randomly matched with another MTurk worker. The interaction is completely anonymous. Neither you nor the other worker will know the other person's worker ID. Moreover, you will not face the same worker twice, i.e. you will interact with one participant in task 1 and another in task 2.

The amount of money that you earn in these tasks will depend on your choice and the other participant's choice. For each task, you will be given a table, similar to the one below, that summarizes your potential earnings. The numbers in the table represent real dollars.

An example of your tasks (slightly different from the actual tasks)

We will now walk you through an example to illustrate the finer details. Note that you will not be paid for this particular example and that the earnings associated with the actual tasks will be quite different.

		Other Participant			
		C		**D**	
You	**A**	You earn:	$2.00	You earn:	$1.00
		Other earns:	$3.00	Other earns:	$2.00
	B	You earn:	$0.50	You earn:	$6.00
		Other earns:	$0.50	Other earns:	$5.00

[36] In order to save space, we have omitted many line-breaks as well as spacing between lines that were used to display the text unless it facilitates readability and/or improves understanding. All other formatting is exactly the same as in the experiment itself. Screenshots of our experiments can be found in the online appendix.

In this example, you can choose between option **A** and **B** (the rows) while the other participant decides between **C** and **D** (the columns). If, for example, you choose **B** and the other participant chooses **D**, you will earn $6 while the other participant will earn $5.

Who acts when

Either you or the other participant will move first. You will be randomly assigned to be the **first mover** or the **second mover**. Your role will be the same for both tasks, that is, you will be either a first mover for both tasks or a second mover for both tasks. The difference between these two roles is as follows:

`[Direct response instructions]` The first mover makes his or her decision first. Afterwards, the second mover will be informed about the first mover's choice and decides how to respond.

`[end of direct response instructions]`

`[Strategy method instructions]` The **first mover** makes his or her decision first. The **second mover** needs to make two choices, one in response to each of the first mover's possible decisions. For example, if you are the **second mover**, you will make the following choices:

If the first mover chooses **C**, I respond with [select **A** or **B**]

If the first mover chooses **D**, I respond with [select **A** or **B**]

The actual outcome will be determined by the first mover's choice and how the second mover responds to that particular choice. For instance, suppose the first mover chose **C** and you, as the second mover, chose **A** in response to **C** and **B** in response to **D**. In this case you earn $2 and the other participant earns $3.

`[end of strategy method instructions]`

Note: All information that you see as the first or second mover will also be available to the other participant

Your earnings Your total earnings from participating in this HIT will be sum of your earnings from the two decision tasks, money earned in the survey, and the participation fee. [Next Page Button]

Control Questions

Before we start with task 1, we want to ensure that you have understood the instructions. **In order to continue with this study, you will need to get at least 3 out of 4 questions correct.** If you aren't quite sure about your answers, have a look at the instructions at the bottom of this page again. Please answer the following questions:

1. Do you know the identity, i.e. their MTurk ID or any other personal information, of the participant you are matched with? [Yes/No (multiple choice list)]

2. Imagine you assume the role of the second mover in task 1. Will your role change in task 2? [Yes/No (multiple choice list)]

3. In the two decision tasks, will you interact with the same MTurk worker? [Yes/No (multiple choice list)]

4. Suppose you are the first mover and earnings are determined by the following table:

`[Same earnings table as in the introduction is displayed here]`

`[Direct response treatment]`

If you choose **A** and the second mover responds with **C**, how much do you and the other participant earn in this task?

`[end of direct response treatment]`

`[Strategy method treatment]`

Suppose you choose **A** and the second mover takes the following conditional choices:

In response to **A**, the second mover chooses **C**
In response to **B**, the second mover chooses **D**

How much do you and the other participant earn in this task? [`end of strategy method treatment`]
You earn [select among 1,2,3,4,5,6 (dropdown menu without default value)]
The other participant earns [select among 1,2,3,4,5,6 (dropdown menu without default value)]

[Next Page Button]

[`previous instructions are again displayed in a box at this position`]

Control Questions – Review
[`this page is displayed only if the participant answered at least one question incorrectly. The following are the statements for the respective questions if they were answered incorrectly. Subjects who answered at least two questions incorrectly are not able to continue.`][37]

Dear participant, you answered at least one control question incorrectly. Before you continue, please have a quick look at the correct answer(s) given below:
Question 1: Do you know the identity of the participant you are matched with? You answered this question with *Yes*. This is *incorrect*. You will never learn any information about the identity of the MTurk workers that you will interact with. Neither will they learn any information about your identity.
Question 2: Imagine you assume the role of the second mover in task 1. Will your role change in task 2? You answered this question with *Yes*. This is *incorrect*. Your role will never change. If you are the first mover in task 1, you will also be the first mover in task 2. Similarly, if you are the second mover in task 1, you will also be the second mover in task 2.
Question 3: In the two decision tasks, will you interact with the same MTurk worker? You answered this question with *Yes*. This is *incorrect*. In task 2, you will be randomly matched with a different worker.
Question 4: Suppose you are the first mover and earnings are determined by the following table:

[`Earnings table from control question 4 is displayed here`]

[`Direct response treatment`] If you choose **A** and the second mover responds with **C**, how much do you and the other participant earn in this task? You answered this question with: You earn [`their answer`], the other participant earns [`their answer`]. This is incorrect. As the first mover, which is you in this example, chooses **A** and the second mover responds with **C**, you will earn $2 and the other participant will earn $3. [`end of direct response treatment`]

[`Strategy method treatment`] Suppose you choose **A** and the second mover takes the following conditional choices:

- In response to **A**, the second mover chooses **C**
- In response to **B**, the second mover chooses **D**

You answered this question with: You earn [`their answer`], the other participant earns [`their answer`]. This is incorrect. As the first mover, which is you in this example, chooses **A** and the second mover responds

[37] Those who failed the control questions were redirected to a simple feedback page that informed them that the experiment has ended and offered them an opportunity to write down any feedbacks or complaints in a textfield.

to **A** with **C**, you will earn $2 and the other participant will earn $3.[end of strategy method treatment]

Decision Task 1 out of 2
When you are ready, please press the next button [Next Page Button]

[The task order is random. For the purpose of presentation, we use the sPD (mUG) for task 1(2). We present task 1(2) from the perspective of player 1(2).

Decision Task 1

		Other Participant: C		Other Participant: D	
You	**A**	You earn:	$1.00	You earn:	$0.00
		Other earns:	$1.00	Other earns:	$1.50
	B	You earn:	$1.50	You earn:	$0.50
		Other earns:	$0.00	Other earns:	$0.50

Your role: you are the **first mover.**

[Non-Selfish belief treatment] **Background Information:** In a well-known study of this task by Watabe, Terai, Hayashi, and Yamagishi, published in the year 1996, 82.6% of the first movers chose **A**. [end of non-selfish belief treatment]

[Selfish Belief treatment] **Background Information:** In a well-known study of this task by Bolle and Ockenfels, published in the year 1990, 82.7% of the first movers chose **B**. [end of selfish belief treatment]

As the first mover, I choose: [A/B (multiple choice list)]
[Next Page Button]

Decision Task 2 out of 2
When you are ready, please press the next button [Next Page Button]

Decision Task 2

		Other Participant: A		Other Participant: B	
You	**C**	You earn:	$1.00	You earn:	$0.30
		Other earns:	$1.00	Other earns:	$1.70
	D	You earn:	$0.00	You earn:	$0.00
		Other earns:	$0.00	Other earns:	$0.00

Your role: you are the **second mover.**

[Non-Selfish Belief treatment] **Background Information:** In a well-known study of this task by Güth, Huck, and Müller, published in the year 2001, 70.6% of the first movers chose **A**.

[end of non-selfish belief treatment]

[Selfish belief treatment] **Background Information:** In our previous experiment of this task, 80% of the first movers chose **B**. [end of selfish belief treatment]

[Direct response treatment]
The other participant chose: **A**

As the second mover, I respond with [C/D (multiple choice list)] [end of direct response treatment]

[Strategy method treatment] As the second mover,
if the first mover chooses **A**, I respond with: [C/D (multiple choice list)]
if the first mover chooses **B**, I respond with: [C/D (multiple choice list)] [end of strategy method treatment]

[Next Page Button]

Survey - page 1/3

The first decision task you completed today was the following interaction:
[The payoff matrix, role assignment and background information from task 1 is displayed inside a gray box, appearing exactly how they say it before]

Among the MTurk workers who participated in this experiment with you today, what percentage of first movers do you think will choose [**A** or **B** is shown depending on which action was highlighted in the belief-treatment]?
[slider from 0 to 100 is shown, with a default at 50]
Note: If you are within 5% of the correct answer you will receive an additional $0.25.
[Next Page Button]

Survey - page 2/3
[same as page 1 but for the second task]

Survey - page 2/3
Before finishing the experiment, we would like to know more about you. All answers will be processed anonymously and will not be connected to your MTurk worker ID.
What is your gender? [Male, Female, Other, Prefer not to say (dropdown menu)]
What is your age (in years)? [Under 12 years old, 12-17 years old, 18-24 years old, 25-34 years old, 34-44 years old, 45-54 years old, 55-64 years old, 65-74 years old, 75 years or older, Prefer not to say (dropdown menu)]
What is the highest degree you are holding or currently pursuing? [High School, Bachelor, Master, Doctorate, Other post-graduate degree, None, Prefer not to say (dropdown menu)]
What is the annual household income (in USD) you have at your disposal? [Less than $20,000, $20,000 to $34,999, $35,000 to $49,999, $50,000 to $74,999, $75,000 to $99,999, $100,000 to $140,999, Over $150,000 (dropdown menu)]

Have you participated in another similar experiment as this before? [Yes, No (dropdown menu)]
[Next Page Button]

End of Experiment[38]

Thank you for completing this HIT. Before you continue, please copy-paste the following survey completion-code into MTurk.
Completion Code: [`participant's completion code is shown`]
[check-box] I have copy-pasted the completion code.
Have a good day [Finish HIT button]

Feedback[39]

Thanks again for participating. If you have copy pasted the survey code to MTurk, you are done. We will calculate your earnings shortly and will provide you with a detailed summary of your choices, as well as the choices of the participants you were matched with, in the message that is sent alongside the bonus payment. If you encountered any technical or other difficulties today, it would be great if you would let us know that we can fix them. You can type in here: [large textfield]
Thank you and have a great day! [Exit button]

[38] At this point, the experiment had formally ended. No further button or check-box click was required.
[39] This page was fully optional. No button-click was required.

References

Herman Aguinis, Isabel Villamor, and Ravi S. Ramani. Mturk research: Review and recommendations. *Journal of Management*, 47(4):823–837, 2021.

Ofra Amir, David G. Rand, and Ya'akov Kobi Gal. Economic games on the internet: The effect of $1 stakes. *PLOS ONE*, 7(2):1–4, 2012.

Antonio A. Arechar, Simon Gächter, and Lucas Molleman. Conducting interactive experiments online. *Experimental Economics*, 21(1):99–31, 2018.

Friedel Bolle and Peter Ockenfels. Prisoners' dilemma as a game with incomplete information. *Journal of Economic Psychology*, 11:69–84, 1990.

Jordi Brandts and Gary Charness. Hot vs. cold: Sequential responses and preference stability in experimental games. *Experimental Economics*, 2:227–238, 2000.

Jordi Brandts and Gary Charness. Truth or consequences: An experiment. *Management Science*, 49(1): 116–130, 2003.

Jordi Brandts and Gary Charness. The strategy versus the direct-response method: a first survey of experimental comparisons. *Experimental Economics*, 14:375–398, 2011.

Jeannette Brosig, Joachim Weimann, and Chun-Lei Yang. The hot versus cold effect in a simple bargaining experiment. *Experimental Economics*, 6:75–90, 2003.

Michael Buhrmester, Tracy Kwang, and Samuel D. Gosling. Amazon's mechanical turk: A new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science*, 6(1):3–5, 2011.

Gary Charness and Matthew Rabin. Understanding social preferences with simple tests. *The Quarterly Journal of Economics*, 117(3):817–869, 2002.

Daniel L. Chen, Martin Schonger, and Chris Wickens. oTree - an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9:88–97, 2016.

Anna Dreber, Tore Ellingsen, Magnus Johannesson, and David G. Rand. Do people care about social context? Framing effects in dictator games. *Experimental Economics*, 16(3):349–371, 2013.

Tore Ellingsen, Magnus Johannesson, Johanna Mollerstrom, and Sara Munkhammar. Social framing effects: preferences or beliefs? *Games and Economic Behavior*, 76(1):117–130, 2011.

Dirk Engelmann and Martin Strobel. Inequality aversion, efficiency, and maximin preferences in simple distribution experiments. *American Economic Review*, 94(4):857–869, 2004.

Ignacio Esponda and Emanuel Vespa. Hypothetical thinking and information extraction in the laboratory. *American Economic Journal: Microeconomics*, 6(4):180–202, 2014.

Ignacio Esponda and Emanuel Vespa. Contingent preferences and the sure-thing principle: Revisiting classic anomalies in the laboratory. *working paper*, 2019.

Erik Eyster and Matthew Rabin. Cursed equilibrium. *Econometrica*, 73(5):1623–1672, 2005.

Armin Falk, Ernst Fehr, and Urs Fischbacher. Driving forces behind informal sanctions. *Econometrica*, 73 (6):2017–2030, 2005.

Gerd Gigerenzer and Ulrich Hoffrage. How to improve bayesian reasoning without instruction: frequency formats. *Psychological Review*, 102(4):684–702, 1995.

Gerd Gigerenzer, Wolfgang Hell, and Hartmut Blank. Presentation and content: The use of base rates as a continuous variable. *Journal of Experimental Psychology: Human Perception and Performance*, 14(3): 513–525, 1988.

Joseph K. Goodman and Gabriele Paolacci. Crowdsourcing consumer research. *Journal of Consumer Research*, 44(1):196–210, 2017.

Werne Güth and Reinhard Tietz. Ultimatum bargaining behavior: a survey and comparison of experimental results. *Journal of Economic Psychology*, 11(3):417–449, 1990.

Werner Güth, Rolf Schmittberger, and Bernd Schwarze. An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior & Organization*, 3(4):367–388, 1982.

Werner Güth, Steffen Huck, and Wieland Müller. The relevance of equal splits in ultimatum games. *Games and Economic Behavior*, 37(1):161–169, 2001.

Ingar Haaland, Christopher Roth, and Johannes Wohlfart. Designing information provision experiments. *Journal Of Economic Literature*, 2023. forthcoming.

David J. Hauser, Gabriele Paolacci, and Jesse Chandler. Common concerns with mturk as a participant pool: Evidence and solutions. In F. R. Kardes, P. M. Herr, and N. Schwarz, editors, *Handbook of research methods in consumer psychology*, pages 319–337. Routledge/Taylor & Francis Group, 2019.

John J. Horton, David G. Rand, and Richard J. Zeckhauser. The online laboratory: conducting experiments in a real labor market. *Experimental Economics*, 14(3):399–425, 2011.

Nicholas C. Hunt and Andrea M. Scheetz. Using mturk to distribute a survey or experiment: Methodological considerations. *Journal of Information Systems*, 33(1):43–65, 2019.

Keith Jensen, Josep Call, and Michael Tomasello. Chimpanzees are rational maximizers in an ultimatum game. *Science*, 318(5847):107–109, 2007.

David Johnson and John Barry Ryan. Amazon mechanical turk workers can provide consistent and economically meaningful data. *Southern Economic Journal*, 87(1):369–385, 2020.

Kiryl Khalmetski, Axel Ockenfels, and Peter Werner. Surprising gifts: Theory and laboratory evidence. *Journal of Economic Theory*, 159:163–208, 2015.

David K. Levine. Modeling altruism and spitefulness in experiments. *Review of Economic Dynamics*, 1: 429–431, 1998.

Alejandro Martínez-Marquina, Muriel Niederle, and Emanuel Vespa. Failures in contingent reasoning: The role of uncertainty. *American Economic Review*, 109(10):3437–74, October 2019.

Michael Naef and Jürgen Schupp. Measuring trust: Experiments and surveys in contrast and combination. *SOEPpaper No. 167*, 2009.

Axel Ockenfels and Peter Werner. Scale manipulation in dictator games. *Journal of Economic Behavior & Organization*, 97:138–142, 2014.

Robert J. Oxoby and Kendra N. McLeish. Sequential decision and strategy vector methods in ultimatum bargaining: evidence on the strength of other-regarding behavior. *Economics Letters*, 84(3):399–405, 2004.

Gabriele Paolacci and Jesse Chandler. Inside the turk: Understanding mechanical turk as a participant pool. *Current Directions in Psychological Science*, 23(3):184–188, 2014.

Gabriele Paolacci, Jesse Chandler, and Panagiotis G. Ipeirotis. Running experiments on amazon mechanical turk. *Judgment and Decision Making*, 5(5):411–419, 2010.

Alvin E. Roth. Bargaining experiments. In John H. Kagel and Alvin E. Roth, editors, *Handbook of experimental economics*, volume 1. Princeton University Press, Princeton, 1995.

Amos Tversky and Daniel Kahneman. The framing of decisions and the psychology of choice. *Science*, 211: 453–458, 1981.

Motoki Watabe, Shigeru Terai, Nahoko Hayashi, and Toshio Yamagishi. Cooperation in the one-shot prisoner's dilemma based on expectations of reciprocity. *Japanese Journal of Experimental Social Psychology Quarterly*, 51:265–271, 1996.

Strategy vs. Direct-Response Method: Evidence from a Large Online Experiment

Marcus Roel
mcs.roel@gmail.com
Beijing Normal University

Zhuoqiong Chen
chenzq926@gmail.com
Harbin Institute of Technology, Shenzhen

- *Online Appendix* -

April 3, 2023

C. Appendix: Theory

In this section, we outline a theory of incomplete conditional thinking with social preferences in which (a change in) the ex-ante beliefs of player 2 (she) about player 1's choices (he) shape her choices in a different way depending on how her response is elicited. We first provide a detailed explanation of the idea, using the sequential Prisoners' Dilemma as an example, before formalizing it mathematically.[40]

Suppose player 2 has reciprocity-like preferences, i.e., she wants to be nice (nasty) to those players who treat her nicely (nastily). In the sequential Prisoners' Dilemma, she considers cooperation to be the nice and defection to be the nasty action. Fundamentally, she thus prefers to cooperate in response to cooperation and to defect after defection. If choices are elicited using the direct-response method, player 2 will clearly observe player 1's action and thus has perfect knowledge of what happened. As a result, she cooperates conditionally. For the strategy method, she does not observe player 1's chosen action but must take an action for any possible choice of player 1. In other words, she must reason conditionally. Evidence suggests that this tends to be more challenging, and as such, individuals may make predictable "mistakes" (Esponda and Vespa, 2014, 2019; Martínez-Marquina et al., 2019). In the context of social preferences, we conjecture that in doing so, player 2's preferences at any node of the game are influenced by the action of player 1 she views as ex-ante most likely. In other words, she evaluates player 1's hypothetical action that leads to a particular node not in isolation, but as a combination of that correct action but also the action she originally thought (and since she doesn't learn anything about player 1's actual choice still thinks) to be most likely. In a sense, when she reasons through the game, she doesn't fully update her belief about player 1's action with respect to the action implied by a given node.[41]

To see how this affects behavior, suppose that player 2 believes that the majority of player 1 either (a) cooperate, or (b) defect. In the direct-response method, we already argued that player 2 does not make any mistake in terms of conditional thinking and responds to cooperation with cooperation and defection with defection regardless of her initial belief. In the strategy method, initial beliefs may influence her choices, however (assuming her conditional thinking is sufficiently lacking). If she initially believes that (a) player 1 usually cooperates, then at the node following cooperation, both her initial belief and (hypothetical) position in the game point to player 1 having cooperated, her updated belief is correct, and so she cooperates. However, if she initially believes player 1 defects (b), at the node following cooperation, her initial belief does not align with her hypothetical position in the game and so her updating with respect to player 1's choice will be incomplete. If she places too much weight on the ex-ante belief, her preferences are incorrectly swayed by her ex-ante view that player 1 defects, and so she doesn't think player 1's action warrants a nice response. It follows that player 2's response to player 1's cooperative behavior may be identical (or mostly similar) in the two elicitation methods (case (a)) or different (case (b)).

This shows how ex-ante beliefs may shape what researchers observe in experiments that test for difference in behavior induced by the elicitation method as initial beliefs can have a different effect on player 2's response depending on the elicitation method employed.[42] Of course, it may also be possible that initial beliefs directly affect player 2's preferences, e.g., due to norms, surprise, or kindness evaluations. However, such effect would typically apply to both elicitation methods equally, and thus would shift behavioral equally.

The Model. Player 1 (he) and player 2 (she) interact in a simple sequential game, where player 1 moves first and player 2 observers his action and responds. Assume that each player only has two choices and denote their respective actions by $a_1 \in A_1$ and $a_2 \in A_2$. Our focus will be on player 2, whose preferences are driven by her material payoffs $\pi_2(a_2, a_1)$ and other-regarding payoffs $g(a_2, a_1)$. The key idea behind $g(a_2, a_1)$ is that player 2 wants to treat player 1 well if he treated her well, and vice versa if he treated her badly. As a

[40] The purpose of this theory section is to formalize one potential model that generates differential effects - using a mechanism that we view as most likely. Needless to say, other potential behavioral explanations exists. Indeed, in the Pre-Analysis Plan, we also proposed a signaling model that also generate a differential effect, but with behavior being shifted in the "opposite" direction compared a model of incomplete contingent thinking. According to this model, player 2 prefers to signal their social type through behavior - either to herself or to other players - at the decision node that she views as unlikely.

[41] We express our model in terms of outcome based preferences. The model could also be written in terms of player 1's types, where player 2 may want to give up material payoffs to rewards (punish) the type of players she views as nice (nasty) (Levine, 1998). In such a model, a failure to properly conditionally update can be written similarly to Eyster and Rabin (2005).

[42] In this discussion, we picked the node that is typically associated with a non material-payoff maximizing social response, as most discussion surrounding the elicitation method has focussed on this choice, e.g., hot/cold, etc. Indeed, for the other node where player 1 has defected, a lack of conditional thinking may shift behavior - but do so the opposite direction: player 2 may actually respond with cooperation in case (a). However, under reasonable assumptions on payoffs and mistakes, one would expect such effect to be smaller even if it was present.

result, she may be willing to lower her own material payoff to increase (reduce) player 1's material payoff.[43] In general, player 2's preferences may also depend on (her belief about) the average behavior of player 1, denoted by σ_1, due to norms, kindness concerns, surprise, etc., and so we generalize the other-regarding payoffs to $g(a_2, a_1, \sigma_1)$. We are agnostic with regards to how σ_1 affects $g(\cdot)$, which depends on the exact nature of such preferences.[44] Let $\gamma \in \mathbb{R}^+$ capture the degree that player 2 is motivated by other regarding preferences. In general, it is possible that the elicitation method itself, $e \in \{dr, sm\}$, where dr stands for direct response and sm for strategy method, changes the way how player 2 experiences her other-regarding payoffs. In this setting, player 2's utility at node/history $h \in H = A_1$ from her action a_2 is

$$U_2(a_2 \mid h = a_1, e, \sigma_1) = \pi_2(a_2, a_1) + \gamma \cdot m_e \cdot g(a_2, a_1, \sigma_1) \tag{1}$$

For example, $m_{dr} > 1 = m_{sm}$ would capture the typical idea that the decision maker is more emotional (hot) in the direct response method while $1 = m_{dr} < m_{sm}$ reflects the case where the decision maker is more aware of the (social) implications of 1's action as she takes a more global perspective of her actions. For our theoretical predictions, we will show that behavior may differ between elicitation *even in the absence of such immediate effects*. In other words, we will assume $m_e = 1$ throughout. However, we will keep these terms for the reader to explore alternative, simpler, ideas in this general framework.[45]

Our main idea is that player 2 may suffer from incomplete contingent reasoning, parameterized by $\lambda_e \in [0, 1]$. In the context of sequential games, incomplete reasoning implies that at some node $h = a_1$, player 2 may not fully realize that a_1 is played and thus her social preferences may be influenced by the alternative choices. Crucially, this mistake will depend on the elicitation method employed. As there is no need for any contingent reasoning whenever player 1's action is fully observed, let $\lambda_{dr} = 0$.[46] This mistake will affect the decision maker whose choices are elicited using the strategy method, however, as she will have to consider various hypothetical choices at the same time. We capture this idea of incomplete contingent reasoning, by the following utility function for player 2:

$$\begin{aligned} U_2(a_2 \mid h = a_1, e, \sigma_1) =& \pi_2(a_2, a_1) \\ &+ \gamma \cdot m_e \cdot \left((1 - \lambda_e) \cdot g(a_2, a_1, \sigma_1) + \lambda_e \cdot \sum_{a_1'} \mathbb{1}(\sigma_1(a_1') \geq 1/2) \cdot g(a_2, a_1', \sigma_1) \right) \end{aligned} \tag{2}$$

The first term in the large social-preference expression, which we call the *expected social payoffs*, represents player 2's true social-preferences from a_2 at node $h = a_1$. The second term captures her (potential) mistake from incomplete contingent reasoning. In the above specification, this term points to the choice of player 1 that she views as most likely due to the indicator function.[47] If the most likely choice for player 1 equals the history player 2 is considering (is located at), the expected social payoffs collapse to the correct social-preference term. However, if player 2 views the alternative action as more likely, then her expected social payoffs are a convex combination of her respective payoffs at the correct node and the alternative node, which alters her "true" preferences. While equation 2 theoretically allows the decision maker to only consider the choice implied by her ex-ante belief, it is reasonable to think that λ is fairly small, and typically $\lambda < 1/2$.

Assumptions. Before we derive our predictions, we need to make assumptions on (1) the material payoffs

[43] In general, defining the notions of treating and being treated well or badly is not as straightforward as our current notation suggests since these concepts are generally a function of (the payoffs of) the whole game. We recognize this, but as our goal is to outline the implications of incomplete conditional thinking in the context of social preferences and not to formalize a complete theory of social preferences with incomplete conditional thinking, we opt for the simplest notation to improve readability and employ various simplifying shortcuts along the way.

[44] For example, an increase in cooperation by player 1 in the Prisoners' Dilemma may capture a stronger norm for social behavior, which in turn would increase the social payoffs from cooperation (relative to defection) in response to cooperation. Alternatively, player 2 may want to reward helpful actions more when she didn't expect such an action to be taken, and so the same increase in cooperation could also result in a decrease in the social payoffs from cooperation.

[45] We will also revisit the predictions of a simple "constant-difference" model that only accounts for differences in m_e to explain behavior between the two elicitation methods at the end of this section.

[46] This is without loss. Our predictions remain qualitatively unchanged for $\lambda_{sm} > \lambda_{dr} \geq 0$.

[47] We use this coarse-thinking over simply averaging the two potential nodes with respect to her belief about player 1's ex-ante choice as it (1) conveniently simplifies the model, and (2) because we deem it a sensible way to capture how attention on some (possibly alternative) node shapes beliefs. The model could also be extend to allow for averaging over nodes at the cost of a more complicated analysis and additional assumptions on preferences.

to represent the games that are being played and (2) on the other-regarding preferences to capture which actions player 2 regards as nice or nasty, and how this shapes her preferences. For the sequential Prisoners' Dilemma (*sPD*), let $A_1 = \{C, D\}$, $A_2 = \{c, d\}$.[48] Player 2 has reciprocity-like preferences, i.e., she wants to treat player 2 well (badly) if he treats her well (badly). Given that cooperation generally improves the payoffs of the other player, this is captured by $g(c, C, \sigma_1) > g(d, C, \sigma_1)$ and $g(c, D, \sigma_1) \leq g(d, D, \sigma_1)$.[49] For the (mini) Ultimatum Game (*mUG*), let player 1's action denotes his offer, which in our case is $A_1 = \{0.5, 0.15\}$. Player 2 either accepts or rejects: $A_2 = \{a, r\}$. Reciprocity-like preferences in this setting are described by $g(r, 0.15, \sigma_1) > g(a, 0.15, \sigma_1)$ and $g(r, 0.5, \sigma_1) \leq g(a, 0.5, \sigma_1)$.

In line with our experimental treatment, we will focus on two particular ex-ante beliefs, namely one where the decision maker believes that defection (unequal offer) is more likely, which represents our *selfish belief treatment*, and one where she believes that cooperation (equal offer) is more likely, which represents our *non-selfish belief treatment*. We use σ_1^S and σ_1^{NS} to denote these selfish and non-selfish beliefs, i.e., where $\sigma_1(C) < 1/2$ and respectively $\sigma_1(C) > 1/2$ for the *sPD*, and $\sigma_1(0.15) > 1/2$ and respectively $\sigma_1(0.15) < 1/2$ for the *mUG*.[50]

Analysis. With no mistake from contingent reasoning, $\lambda_e = 0$, player 2's relative social payoffs point in opposite direction relative to her material payoffs (c vs. d) at $h = C$ and $h = 0.15$, but (weakly) align with her material payoffs at $h = D$ and $h = 0.5$. Consequently, when γ is sufficiently large, she will prefer to take the social action C after cooperation and r after the unequal offer. In turn, this implies that there is a cutoff value for γ, denoted by $\bar{\gamma}(h, e, \sigma_1)$ for these two nodes, at which player 2 is indifferent between both actions. This cutoff captures the minimal sensitivity towards social preferences that is required to induce a social response. Assuming for the population of player 2, γ is distributed with full support on R^+, it follows that a smaller $\bar{\gamma}$ results in a higher frequency of the socially-motivated choices $\sigma_2(c|C)$ and $\sigma_2(r|0.15)$.[51]

We will now solve the model for the *sPD*, which allows us to avoid additional notation and to keep this section succinct. The analysis of the *mUG* is identical. We will summarize the respective predicts for the *mUG* at the end. For the *sPD*, the cutoff in the direct-response ($\lambda_{dr} = 0$) is:

$$\bar{\gamma}(C, dr, \sigma_1) = \frac{1}{m_{dr}} \cdot \frac{\pi_2(d, C) - \pi_2(c, C)}{g(c, C, \sigma_1) - g(d, C, \sigma_1)} \tag{3}$$

We observe that an increase in the sensitivity of the social preference with respect to the elicitation method, $m_{dr} \uparrow$, results in a decrease in the cutoff and thus an increase in $\sigma_2(c|C)$. The equation also shows how ex-ante beliefs about player 1's behavior can directly affect player 2's preferences. Suppose, for example, that player 2 expects more cooperation from player 1, e.g., her belief changes from σ_1^S to σ_1^{NS}. If this increases (decreases) the *relative* social payoffs from c vs. d, e.g., due to social norms (surprise), then $\sigma_2(c|C)$ increases (decreases). Finally, note that for $\lambda_e = 0$ a cutoff for $h = D$ doesn't exist as the decision maker prefers to defect - both for material and (possibly) for social reasons - for any degree of γ.

We now turn to the strategy method. Notice that in this case, the expected social payoffs at $h = C$ may not oppose the material payoff if λ is sufficiently large *and* the decision maker believes that the alternative node to be more likely, $\sigma_1(D) \geq 1/2$. We say that we are at *an interior* (point) if λ is such that the expected social payoffs incentives point in the opposite direction as the material incentives. For now, assume this to be true for $h = C$.[52] The cutoffs for the strategy method at $h = C$ for our two beliefs are:

$$\bar{\gamma}(C, sm, \sigma_1^{NS}) = \frac{1}{m_{sm}} \cdot \frac{\pi_2(d, C) - \pi_2(c, C)}{g(c, C, \sigma_1^{NS}) - g(d, C, \sigma_1^{NS})} \tag{4}$$

48 Note that the use of small/large letters differs slightly in this online appendix to our main text, where we do not use small/large fonts to distinguish between the two players, but aim for better readability in the text.

49 Recall that the choice will be determined by relative material and relative social payoffs between the two actions. In other words, whether $g(a_2, a_1, \sigma_1)$ is positive or negative is irrelevant.

50 Such beliefs are the key drivers of differences that we will observe in the strategy method. Changes within the two coarse categories of C (0.15) or D (0.5) will also have an influence on individual behavior, but will have a relatively minimal impact on the overall population behavior.

51 The full-support assumption on R^+ is made for convenience only. The results remain true if this is true for a reasonable range of γ.

52 Doing so does not affect our analysis is a significant way. We revisit this later below. Aside, notice that this assumption is always true for all $\lambda < 1$ if, for example, $g(c, D, \sigma_1) = g(d, D, \sigma_1)$.

$$\bar{\gamma}(C, sm, \sigma_1^S) = \frac{1}{m_{sm}} \cdot \frac{\pi_2(d,C) - \pi_2(c,C)}{(1-\lambda_{sm}) \cdot [g(c,C,\sigma_1^S) - g(d,C,\sigma_1^S] + \lambda_{sm} \cdot [g(c,D,\sigma_1^S) - g(d,D,\sigma_1^S]} \quad (5)$$

For the non-selfish beliefs, the cutoff equals that of the direct-response method, $\bar{\gamma}(C, sm, \sigma_1^{NS}) = \bar{\gamma}(C, dr, \sigma_1^{NS})$.[53] This is because the ex-ante belief points to the same node as the one the decision maker is currently considering in. When her focus from her ex-ante belief is on the present node of consideration, the two elicitation methods result in the same population behavior, $\sigma_2(c|C, sm, \sigma_1^{NS}) = \sigma_2(c|C, dr, \sigma_1^{NS})$.

If the decision maker's initial belief does not align with her hypothetical position in the game, her preferences are incorrectly swayed by the potential alternative action, which does not warrant a nice response in this particular case. For selfish beliefs, there is thus less cooperation in the strategy-method relative to the direct-response method, however, as $\bar{\gamma}(C, sm, \sigma_1^S) > \bar{\gamma}(C, dr, \sigma_1^S)$ due to $g(c, D, \sigma_1^S) - g(d, D, \sigma_1^S) \leq 0$.

At this point, it should be clear that the assumption that an interior solution exists was without loss. After all, $\bar{\gamma}(C, sm, \sigma_1^{NS})$ always exists regardless. If there are some decision makers in the population whose expected social payoffs are negative for selfish beliefs, then the frequency of cooperation in the selfish belief treatment would be even lower.

Table C.9: Predictions for Player 2 Behavior

1. seq. Prisoners' Dilemma: $\sigma_2(c|h = C, e, \sigma_1)$

	Selfish Beliefs	Non-Selfish Beliefs	$\Delta : (2) - (1)$
Direct Response	$\sigma_2(c\|C)$	$\sigma_2(c\|C) + b$	b
Strategy Method	$\sigma_2(c\|C) - \lambda$	$\sigma_2(c\|C) + b$	$b + \lambda$
$\Delta : (2) - (1)$	$-\lambda$	0	λ

2. Ultimatum Game: $\sigma_2(r|h = 0.15, e, \sigma_1)$

	Selfish Beliefs	Non-Selfish Beliefs	Δ
Direct Response	$\sigma_2(r\|0.15)$	$\sigma_2(r\|0.15) + b$	b
Strategy Method	$\sigma_2(r\|0.15)$	$\sigma_2(r\|0.15) + b - \lambda$	$b - \lambda$
$\Delta : (2) - (1)$	0	$-\lambda$	$-\lambda$

Table C.9 tabulates the predicted behavior by our theory. We use the behavior in the direct response method with selfish beliefs as a baseline, abbreviated by $\sigma_2(c|C)$ for convenience. $b \in R$ describes the direct effect of prior beliefs on preferences (which may vary between games) while $\lambda > 0$ captures the effect from incomplete contingent thinking.

Table C.9 summarizes these predictions, both for the *sPD* and the *mUG* (which can be obtained by simply repeating our previous analysis). For convenience, we use behavior in direct response method for selfish beliefs as a baseline, denoting the respective population behavior by $\sigma_2(c|C)$ and $\sigma_2(r|0.15)$. The direct effects from beliefs and incomplete contingent thinking on behavior are captured by the constant b, which may be positive or negative, and $\lambda > 0$. While we didn't go over the *mUG* in detail, it is worthwhile to point out that for this game, the predicted difference in difference points in the opposite direction as *sPD*. This was part of the reason that motivated us to use these two games, which represent the broader class of "reward" and "punishment" games.

Lastly, we do quickly want to point out what our theory predicts in the *sPD* after D (in similarly for the *mUG* after 0.5). First, observe that since the material and social payoffs align at $h = D$, $\sigma_2(c|D, dr, \sigma_1^S) = \sigma_2(c|D, dr, \sigma_1^{NS}) = 0$. The same is true for the strategy method and with selfish beliefs, $\sigma_2(c|D, sm, \sigma_1^S) = 0$. However, if λ_{sm} is sufficiently large, some decision makers may actually prefer to cooperate in response to defection. To see this, note that the cutoff, assuming λ_{sm} is large enough, is

$$\bar{\gamma}(D, sm, \sigma_1^{NS}) = \frac{1}{m_{sm}} \cdot \frac{\pi_2(d,D) - \pi_2(c,D)}{(1-\lambda_{sm}) \cdot [g(c,D,\sigma_1^{NS}) - g(d,D,\sigma_1^{NS}] + \lambda_{sm} \cdot [g(c,C,\sigma_1^{NS}) - g(d,C,\sigma_1^{NS})}$$

It follows that behavior coincides across all elicitation-methods and belief-treatments at node D except for the strategy-method with non-selfish beliefs, which may have some positive levels of cooperation, that arises

[53] Simply set $\sigma_1 = \sigma_1^N S$ in equation 3 and recall that we assume $m_{dr} = m_{sm}$ throughout.

from confusion due to incomplete conditional thinking. Interestingly, we tend to observe the opposite effect in our experiment, namely less confusion in the strategy method after the node, that we did not expect to be too important when designing the experiment.

Predictions of simple models (hot/cold; increased fairness-considerations, etc). Finally, we want to quickly cover the typical predictions of any simple model. By simple model, we refer to explanations that are based on fundamental difference in social-preferences due to the elicitation method, which, in our mathematical formulation, are captured by difference in m_e. For this analysis, it is irrelevant whether $m_{sm} > m_{dr}$ or $m_{sm} < m_{dr}$.[54] The general cutoff is such a model, focussing again on $h = C$, for elicitation method e is

$$\bar{\gamma}(C, e, \sigma_1) = \frac{1}{m_e} \cdot \frac{\pi_2(d, C) - \pi_2(c, C)}{g(c, C, \sigma_1) - g(d, C, \sigma_1)} \tag{6}$$

It follows that if, for example, $m_{sm} > m_{dr}$, $\sigma_2(c|C, sm, \sigma_1) > \sigma_2(c|C, dr, \sigma_1)$ for any prior belief σ_1, and vice versa for $m_{sm} < m_{dr}$. Since our model captures this direct effect using a multiplicative term m_e, it is generally not true that the difference in behavior between the strategy-method and direct-response method is perfectly constant.[55] That being said, we can think of the predictions of this model as *approximately constant* in the sense that, given a reasonably large sample, the difference between the strategy method and direct-response method is (i) statistically significant regardless of the prior belief σ_1, and (ii) the sign of that difference is always the same. Moreover, (iii) extending this analysis to multiple games, such difference can be found at any node for which the social-preference term, $g(a_2, h, \sigma_1) - g(a'_2, h, \sigma_1)$, is relatively amplified - for any game. In order words, if $m_{sm} > m_{dr}$, then the model also predicts that $\sigma_2(r|0.15, sm, \sigma_1) > \sigma_2(r|0.15, dr, \sigma_1)$ for both belief-treatments in the mUG.

[54] Of course, $\lambda_e = 0$ for any elicitation method.

[55] Indeed, even with an additive term, the difference would not need to be constant as a constant difference in the cutoff does not imply a constant difference in (the frequency of) behavior unless the distribution of γ is uniform in the parameter range of interest.

D. Robustness Checks

D.1. Probit and Logit

In this section, we replicate the main tables from our paper, estimating Probit and Logit regressions. We note that the numbers of observations for these results vary (and hence vary from our OLS-regressions) when controls for individual characteristics are included. This is due to the fact that some of the categories for some control variables feature very few observations. As a result, behavior may be perfectly predicted by their respective category dummy. In table D.10, for example, a single participant said to be between 12-17 years old, and in table D.11, 4 participants, who prefer to not reveal their age, all cooperated after cooperation. For player 2's behavior in the *mUG*, the effect is slightly more pronounced given that most players behave in the same way. Here, the omitted observation again come from categories with very few observations, such those who used "other" for their gender or preferred not to disclose it, those who are "75 years or older" or prefer not to reveal their age, etc. Unlike in OLS-regressions, Probit and Logit regressions (necessary) omit such observations, while their variation is fully explained by the respective dummies in the OLS regression, which essentially "drops" these observations as well. Given that our results are consistent between estimates that control or not control for characteristics, it is clear that the results are unaffected by this (which is unsurprising as these should be orthogonal to the randomization by design).[56]

More generally, this subsection shows that our main results, table 4 - 6 (D.10 - D.12) are fully robust to a change in estimator. The same is true for table 7 (D.13). Signs and significance levels also match fully for table 8 (D.14) in the sPD. For the *mUG*, they also match for the Strategy Method and the Inattentive dummy, and are qualitatively similar for their respective interaction.

Table D.10: Robustness. Player 1's behavior - Probit / Logit

	sPD				mUG			
Dep. Var: P1 cooperates; offers 50-50	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Non-Selfish Belief	0.425***	0.432***	0.683***	0.697***	0.320***	0.322***	0.523***	0.529***
	(0.0460)	(0.0463)	(0.0743)	(0.0751)	(0.0469)	(0.0475)	(0.0769)	(0.0783)
Strategy Method	-0.231***	-0.233***	-0.370***	-0.374***	-0.0971	-0.0826	-0.156	-0.131
	(0.0671)	(0.0681)	(0.108)	(0.110)	(0.0669)	(0.0680)	(0.107)	(0.110)
Non-Selfish × Strategy Method	0.325***	0.326***	0.523***	0.527***	0.177*	0.192*	0.290*	0.307*
	(0.0957)	(0.0966)	(0.156)	(0.158)	(0.0970)	(0.0981)	(0.160)	(0.162)
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Observations	4009	4008	4009	4008	4009	4008	4009	4008
Estimator	Probit	Probit	Logit	Logit	Probit	Probit	Logit	Logit

Notes: this table reports estimates from Probit and Logit regressions, with control variables identical to those in Table 4. Estimates for control variables are not reported. Robust standard errors are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

[56] Needless to say, we prefer to keep the characteristic dummies as they are (provided by the participants) for transparency reason over arbitrarily reclassifying them into some bigger category.

Table D.11: Robustness. Player 2's behavior in sPD - Probit / Logit

	after P1 cooperates				after P1 defects			
Dep. Var: Player 2 cooperates	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Non-Selfish Belief	0.117*	0.131**	0.190*	0.215**	0.115	0.101	0.198	0.170
	(0.0618)	(0.0626)	(0.100)	(0.102)	(0.0784)	(0.0797)	(0.135)	(0.138)
Strategy Method	0.0922	0.108	0.150	0.178	-0.383***	-0.471***	-0.698***	-0.876***
	(0.0755)	(0.0770)	(0.123)	(0.126)	(0.0917)	(0.0971)	(0.170)	(0.180)
Non-Selfish × Strategy Method	-0.130	-0.162	-0.212	-0.263	-0.0888	-0.0460	-0.148	-0.0543
	(0.104)	(0.106)	(0.170)	(0.175)	(0.132)	(0.136)	(0.243)	(0.250)
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Observations	2722	2718	2722	2718	2247	2242	2247	2242
Estimator	Probit	Probit	Logit	Logit	Probit	Probit	Logit	Logit

Notes: this table reports estimates from Probit and Logit regressions, with control variables identical to those in Table 4. Estimates for control variables are not reported. Robust standard errors are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Table D.12: Robustness. Player 2's behavior in mUG - Probit / Logit

	after P1 offers 85-15				after P1 offers 50-50			
Dep. Var: Player 2 rejects	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Non-Selfish Belief	0.156	0.155	0.297	0.292	-0.272**	-0.342***	-0.638**	-0.707***
	(0.102)	(0.105)	(0.195)	(0.198)	(0.113)	(0.120)	(0.266)	(0.270)
Strategy Method	0.208**	0.215**	0.394**	0.421**	-0.641***	-0.746***	-1.593***	-1.816***
	(0.1000)	(0.102)	(0.190)	(0.193)	(0.198)	(0.234)	(0.530)	(0.556)
Non-Selfish × Strategy Method	-0.191	-0.219	-0.363	-0.420	0.666***	0.761***	1.654**	1.813***
	(0.144)	(0.147)	(0.271)	(0.275)	(0.249)	(0.286)	(0.645)	(0.680)
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Observations	1999	1957	1999	1957	2970	2625	2970	2625
Estimator	Probit	Probit	Logit	Logit	Probit	Probit	Logit	Logit

Notes: this table reports estimates from Probit and Logit regressions, with control variables identical to those in Table 4. Estimates for control variables are not reported. Robust standard errors are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Table D.13: Robustness. Player 2's behavior: Mistakes and Beliefs - Probit / Logit

	sPD: after P1 defects				mUG: after P1 offers 50-50			
Dep. Var: Player 2 makes mistake	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Non-Selfish Belief	0.0107	-0.00829	0.0213	-0.0230	-0.0412	-0.0943	-0.130	-0.172
	(0.0955)	(0.0988)	(0.163)	(0.173)	(0.124)	(0.140)	(0.289)	(0.311)
Belief Player 1 cooperates	0.275*	0.268*	0.468*	0.495*				
	(0.144)	(0.150)	(0.245)	(0.262)				
Belief Player 1 offers 85-15					0.575***	0.620***	1.282***	1.248***
					(0.183)	(0.203)	(0.408)	(0.459)
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Observations	1298	1293	1298	1293	2021	1766	2021	1766
Estimator	Probit	Probit	Logit	Logit	Probit	Probit	Logit	Logit

Notes: this table reports estimates from Probit and Logit regressions, with control variables identical to those in Table 4, for player 2s in the direct-response method. Estimates for control variables are not reported. Robust standard errors are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Table D.14: Robustness. Player 2's behavior: Mistakes and Inattention - Probit / Logit

	sPD: after P1 defects				mUG: after P1 offers 50-50			
Dep. Var: Player 2 makes mistake	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Strategy Method	-0.486***	-0.358***	-0.891***	-0.660***	-0.331**	-0.165	-0.789***	-0.488
	(0.0688)	(0.0851)	(0.127)	(0.160)	(0.131)	(0.185)	(0.300)	(0.467)
Inattentive	0.318***	0.443***	0.583***	0.766***	0.548***	0.622***	1.280***	1.379***
	(0.0673)	(0.0846)	(0.120)	(0.145)	(0.110)	(0.124)	(0.251)	(0.281)
Inattentive × Strategy Method		-0.337**		-0.559**		-0.306		-0.482
		(0.139)		(0.254)		(0.248)		(0.606)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2242	2242	2242	2242	2625	2625	2625	2625
Estimator	Probit	Probit	Logit	Logit	Probit	Probit	Logit	Logit

Notes: this table reports estimates from Probit and Logit regressions, with control variables are identical to those in Table 4. Estimates for control variables are not reported. Robust standard errors are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Finally, we also estimate Table 8 (as well as Table D.14) without controls. These estimates were skipped despite all other tables reporting results with and without controls in order to limit the table size in the main part of the paper, and, as we had seen and will again see, because such controls have little effect on the estimates themselves.

Table D.15: Robustness (2). Player 2's behavior: Mistakes and Inattention - no controls

	sPD: after P1 defects						mUG: after P1 offers 50-50					
Dep. Var: P2 makes mistake	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Strategy Method	-0.110***	-0.0767***	-0.446***	-0.342***	-0.815***	-0.632***	-0.0167***	-0.00564	-0.312**	-0.173	-0.763***	-0.454
	(0.0156)	(0.0179)	(0.0670)	(0.0836)	(0.124)	(0.157)	(0.00566)	(0.00524)	(0.122)	(0.174)	(0.295)	(0.463)
Inattentive	0.0841***	0.130***	0.315***	0.417***	0.569***	0.711***	0.0371***	0.0480***	0.572***	0.629***	1.358***	1.454***
	(0.0180)	(0.0273)	(0.0662)	(0.0832)	(0.117)	(0.141)	(0.00758)	(0.0104)	(0.102)	(0.116)	(0.243)	(0.271)
Inattentive × SM		-0.102***		-0.275**		-0.441*		-0.0321**		-0.247		-0.486
		(0.0352)		(0.137)		(0.249)		(0.0140)		(0.238)		(0.597)
Controls	No	No	No	No	No	No	No	No	No	No	No	No
Observations	2247	2247	2247	2247	2247	2247	2970	2970	2970	2970	2970	2970
Estimator	OLS	OLS	Probit	Probit	Logit	Logit	OLS	OLS	Probit	Probit	Logit	Logit

Notes: this table reports estimates from OLS, Probit and Logit regressions for Table 8 and D.14 without controls. Robust standard errors are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

D.2. Order Effects

In this section, we discuss the presence of order effects in our data (which had, for example, been highlighted previously in footnote 21). The general takeaway is that order effects exist and that our analysis restricted to task 1 yields qualitatively similar results despite the reduction in sample size by 50%.

Table D.16 provides the average beliefs and choice frequencies for all four treatment groups by task-order, with significant differences between task 1 and task 2 for a given belief/elicitation group highlighted by the respective star-indicators in the task 1 columns. The null-hypothesis that behavior (and beliefs) is constant across tasks is rejected. Moreover, behavior is neither constant for the direct-response method nor the strategy method for either player.[57] Consequently, we repeat our analysis using behavior of task 1 only in Tables D.17 to D.21 below.

For player 1 (Table D.17), the treatment effect estimates are very similar.[58] If anything, the impact of the strategy method in the *mUG* follows more closely that of the *sPD* when the data is restricted to task 1 - which we previously interpreted as a higher degree of strategic sophistication.

Regarding player 2's behavior in the *sPD* (Table D.18), we see that treatment effects remain similar, with all having the same sign. Crucially, the effect of the strategy method in reducing mistakes remains strongly significant. Somewhat surprisingly, the impact of the non-selfish belief treatment is no longer significant after cooperation (we interpreted this small yet significant effect previously as a 'norms') but is significant after defection. This effect may be interpreted in line with our previous argument regarding mistakes (see our analysis and respective discussion of Table 7).

For player 2's behavior in the *mUG* (Table D.21), the strategy method dummy remains significant after unequal offers (and, as before, with other estimates being insignificant). After equal splits, all estimates are qualitatively similar to before. In light of their small coefficients, and the 50% reduction in sample size, all but the strategy method dummy with controls are no longer significant, however.

We also repeat the analysis of mistakes (table D.20 and D.21), which result in a similar takeaway as before.

Table D.16: Beliefs and Behavior by Task Order

	Task 1				Task 2			
	Selfish		Non-Selfish		Selfish		Non-Selfish	
	DR	SM	DR	SM	DR	SM	DR	SM
Beliefs about Player 1								
Belief Player 1 cooperates	0.34	0.33	0.75**	0.74	0.34	0.32	0.73	0.74
Belief Player 1 offers 50-50	0.32***	0.30	0.75***	0.76***	0.35	0.33	0.73	0.71
Behavior: Player 1								
Player 1 cooperates	0.53**	0.45**	0.67	0.69	0.46	0.35	0.65	0.70
Player 1 offers 50-50	0.62**	0.56	0.77***	0.79**	0.57	0.56	0.66	0.69
Behavior: Player 2								
Player 2 cooperates after C	0.68***	0.71**	0.70**	0.69*	0.57	0.62	0.64	0.62
Player 2 cooperates after D	0.22	0.12	0.30***	0.13	0.19	0.11	0.19	0.11
Player 2 rejects 50-50	0.03*	0.01	0.02	0.01	0.05	0.00	0.03	0.03
Player 2 rejects 85-15	0.10	0.17	0.09**	0.16	0.10	0.12	0.16	0.11

Notes: *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively, testing for differences between the respective columns of among the pre-Covid and Covid samples.

[57] We also note that the direct-respond appears to feature more occurrences of significant difference between tasks for both player 1 (despite the identical set of information for player 1 in both elicitation methods) and player 2. The latter leads us to explore the effect of prior experience in the direct response method at the end of this section (Table D.22). We find no indication for this.

[58] We will usual refer to (two) estimates as being similar when they are (i) both statistically significant (with similar signs), or (ii) the directional effect is at least similar (subject to neither being too close to zero).

Table D.17: Robustness. Player 1's behavior - Task 1

	sPD		UG	
Dep. Var: P1 cooperates; offers 50-50	(1)	(2)	(3)	(4)
Non-Selfish Belief	0.140***	0.144***	0.141***	0.138***
	(0.0248)	(0.0248)	(0.0233)	(0.0234)
Strategy Method	-0.0744**	-0.0798**	-0.0616*	-0.0584
	(0.0374)	(0.0382)	(0.0371)	(0.0371)
Non-Selfish × Strategy Method	0.0942*	0.0937*	0.0898*	0.0906*
	(0.0506)	(0.0510)	(0.0480)	(0.0479)
Controls	No	Yes	No	Yes
Observations	2011	2011	1998	1998

Notes: this table reports estimates from OLS regressions based only on task 1 data. Control variables for individual characteristics include gender, age, income, highest-education, dummies for prior participation in experiments, and Covid sample. Estimates for control variables are not reported. Robust standard errors are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Table D.18: Robustness. Player 2's behavior in sPD - Task 1

	after P1 cooperates		after P1 defects	
Dep. Var: Player 2 cooperates	(1)	(2)	(3)	(4)
Non-Selfish Belief	0.0211	0.0260	0.0728**	0.0663*
	(0.0308)	(0.0307)	(0.0363)	(0.0359)
Strategy Method	0.0296	0.0397	-0.103***	-0.129***
	(0.0380)	(0.0384)	(0.0312)	(0.0319)
Non-Selfish × Strategy Method	-0.0356	-0.0559	-0.0664	-0.0489
	(0.0523)	(0.0529)	(0.0474)	(0.0472)
Controls	No	Yes	No	Yes
Observations	1391	1391	1074	1074

Notes: this table reports estimates from OLS regressions based only on task 1 data, with control identical to those in Table D.17. Estimates for control variables are not reported. Robust standard errors are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Table D.19: Robustness. Player 2's behavior in UG - Task 1

	after P1 offers 85-15		after P1 offers 50-50	
Dep. Var: Player 2 rejects	(1)	(2)	(3)	(4)
Non-Selfish Belief	-0.0121	-0.0129	-0.0140	-0.0139
	(0.0277)	(0.0285)	(0.00918)	(0.00927)
Strategy Method	0.0618**	0.0724**	-0.0168	-0.0229**
	(0.0299)	(0.0299)	(0.0105)	(0.0113)
Non-Selfish × Strategy Method	0.00421	-0.0104	0.0144	0.0208
	(0.0437)	(0.0438)	(0.0137)	(0.0148)
Controls	No	Yes	No	Yes
Observations	953	953	1551	1551

Notes: this table reports estimates from OLS regressions based only on task 1 data, with control identical to those in Table D.17. Estimates for control variables are not reported. Robust standard errors are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Table D.20: Robustness. Player 2's behavior: Mistakes and Beliefs - Task 1

	sPD: after P1 defects		mUG: after P1 offers 50-50	
Dep. Var: Player 2 makes mistake	(1)	(2)	(3)	(4)
Non-Selfish Belief	0.0450	0.0376	0.00123	0.000473
	(0.0452)	(0.0443)	(0.0120)	(0.0124)
Belief Player 1 cooperates	0.0697	0.0750		
	(0.0662)	(0.0652)		
Belief Player 1 offers 85-15			0.0703***	0.0687***
			(0.0241)	(0.0238)
Controls	No	Yes	No	Yes
Observations	602	602	947	947

Notes: this table reports estimates from OLS regressions based only on task 1 data, with control variables identical to those in Table D.17, for player 2s in the direct-response method. Estimates for control variables are not reported. Robust standard errors are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Table D.21: Robustness. Player 2's behavior: Mistakes and Inattention - Task 1

	sPD: after P1 defects		mUG: after P1 offers 50-50	
Dep. Var: Player 2 makes mistake	(1)	(2)	(3)	(4)
Strategy Method	-0.149***	-0.0974***	-0.0118	-0.00162
	(0.0237)	(0.0272)	(0.00727)	(0.00729)
Inattentive	0.0884***	0.163***	0.0255***	0.0356***
	(0.0273)	(0.0425)	(0.00891)	(0.0124)
Inattentive × Strategy Method		-0.157***		-0.0308*
		(0.0532)		(0.0175)
Controls	Yes	Yes	Yes	Yes
Observations	1074	1074	1551	1551

Notes: this table reports estimates from OLS regressions, with control variables identical to those in Table D.17. Estimates for control variables are not reported. Robust standard errors are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Finally, in table D.22, we explore the role of past experience on behavior for task 2. One of the key difference between task 1 and task 2 in the direct-response treatment is that player 2 gets to observe player 1's action in the prior round. If this experience influences her choices in the second round (despite the fact that the person behind player 1 is different for task 2), then order effects will shape player 2's behavior in the direct response method in a different way than in the strategy method for task 2. We find no indication for this when we relate player 2's behavior (cooperate/reject) in task 2 at any node of the game to a dummy that indicates whether player 1 takes the Non-Selfish action (cooperates in the *sPD* or offers 50-50 in the mUG) for task 1.[59]

Table D.22: Regression Table: Player 2 Behavior based on their Experience in first Task

	Player 1's choice in task 2:							
	cooperates		defects		offers 85-15		offers 50-50	
Dep. Var: Player 2 cooperates / rejects	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
P1 takes Non-Selfish Action in t=1	0.0364	0.0351	-0.0210	-0.0108	0.0306	0.0258	0.00769	0.00220
	(0.0386)	(0.0381)	(0.0318)	(0.0317)	(0.0278)	(0.0279)	(0.0135)	(0.0133)
Non-Selfish Belief	0.0660*	0.0754**	-0.000544	-0.0153	0.0577**	0.0542*	-0.0244*	-0.0237*
	(0.0345)	(0.0346)	(0.0305)	(0.0308)	(0.0285)	(0.0289)	(0.0136)	(0.0135)
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Observations	854	854	696	696	574	574	947	947

Notes: this table reports estimates from OLS regressions for player 2 behavior in the second task in the direct response treatment(s) for all possible choices of player 1. Control variables are identical to those in Table 4. Estimates for control variables are not reported. Robust standard errors are reported in parentheses. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

[59] We also checked if player 1's behavior "matched" the belief-treatment, but found no evidence that this to affect player 2's behavior for task 2 (Results available upon request).

E. Sample Comparison: Pre-Covid and Covid Sample

In this section, we compare the pre-Covid (October to November 2019) and Covid sample (October 2021) with regards to the participants' characteristic and behavior.

Differences among Participants. Table E.23 summarizes the participants' (mean) characteristics and highlights their potential differences between our samples. Overall, the Covid sample tends to have slightly more educated participants, with higher incomes. They are also slightly older, which is consistent with more people working from home, using mTurk either as a second source of income or out of interest. More participants have participated in similar experiments before. Despite that, they tend to make more mistakes on the initial understanding test even though they essentially take the same time to complete the experiment.

Differences in Behavior/Beliefs. In terms of behavior, we first want to point the reader's attention back to Table 2 in the paper itself (section 3), which documented a higher rate of rejection in the *mUG* as well as more mistakes in the Covid sample. A more detailed picture is given by Table E.24, which provides beliefs and behavior by treatment for both samples.

The table shows that beliefs are quite similar across samples. While beliefs are significantly different for the non-selfish/direct response treatment at 5% and 10%, the total difference is small at 2 percentage points.

For player 1, there is a significant increase in cooperation in the selfish belief/direct-response treatment over the samples, yet a very significant drop in cooperation for the selfish belief/strategy method group. We also see fewer equal offers for the selfish belief/strategy method treatment in the Covid sample. It is interesting that all these differences arise for the selfish belief treatments, suggesting that behavior may be more susceptible to change in a selfish-framing.

In the *sPD*, player 2's tendency to reward cooperation by cooperating themselves is very stable, with no significant differences between the two samples. Cooperation after defection tends to occur more frequently, yet is only significantly higher for the non-selfish/strategy method treatment. The *mUG* is where we see the largest differences between the two samples. Rejecting offers is significantly higher in the Covid sample (for all treatments but the strategy methods groups in response to 50-50 offers).

Table E.23: (Mean) Characteristics of Participants by Sample

	Pre-Covid	Covid
Participated in experiments before	0.733**	0.756
Gender:		
Female	0.515	0.520
Male	0.478	0.470
Other / Prefer not to say	0.006*	0.010
Age:		
$<$ 12 years	0.000	0.000
12-17 years old	0.000	0.000
18-24 years old	0.081***	0.063
25-34 years old	0.369	0.370
35-44 years old	0.275	0.280
45-54 years old	0.147*	0.162
55-64 years old	0.093	0.088
65-74 years old	0.032	0.029
$\geq$ 75 years	0.002**	0.004
Prefer not to say	0.002	0.003
Income		
Less than 20 000	0.000	0.000
20 000 to 34 999	0.162***	0.137
35 000 to 49 999	0.176	0.182
50 000 to 74 999	0.239	0.244
75 000 to 99 999	0.147**	0.166
100 000 to 140 999	0.113	0.102
over 150 000	0.044*	0.053
Prefer not to say	0.022	0.022
Education:		
No Degree	0.010	0.008
High School Degree	0.298***	0.211
Bachelor Degree	0.487***	0.525
Master Degree	0.133***	0.193
Other Post-Grad Degree	0.033	0.028
Doctorate Degree	0.026	0.024
Prefer not to say	0.012	0.011
Other Game Outcomes		
Total earnings from games (in USD)	1.766**	1.733
Total time (in sec.)	368.737	364.701
Number of mistakes in understanding test	0.300**	0.326

Notes: this table reports the mean of various variables for the Pre-Covid and Covid Sample. Statistically significant differences between pre- and Covid data (based on t-tests) at significance levels of * $p<0.10$, ** $p<0.05$, *** $p<0.01$, are indicated by the respective stars in the Pre-Covid column.

Table E.24: Beliefs and Behavior in the pre-Covid and Covid Sample by Treatments

	Pre-Covid				Covid			
	Selfish		Non-Selfish		Selfish		Non-Selfish	
	DR	SM	DR	SM	DR	SM	DR	SM
Beliefs about Player 1								
Belief Player 1 cooperates	0.34	0.33	0.73**	0.73	0.34	0.31	0.75	0.75
Belief Player 1 offers 50-50	0.34	0.32	0.73*	0.73	0.33	0.31	0.75	0.74
Behavior: Player 1								
Player 1 cooperates	0.47**	0.46***	0.65	0.70	0.53	0.33	0.68	0.69
Player 1 offers 50-50	0.61	0.60**	0.72	0.76	0.58	0.51	0.71	0.71
Behavior: Player 2								
Player 2 cooperates after C	0.63	0.63	0.66	0.67	0.63	0.70	0.68	0.64
Player 2 cooperates after D	0.19	0.10	0.23	0.09**	0.23	0.14	0.26	0.16
Player 2 rejects 50-50	0.02***	0.01	0.01***	0.02	0.07	0.01	0.04	0.03
Player 2 rejects 85-15	0.08**	0.11***	0.10*	0.10***	0.13	0.19	0.17	0.19

Notes: *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively, testing for differences between the respective columns of among the Pre- and Post-Covid Samples.

E.1. Pictures of Experiment

In this section, we present all screenshots from the experiment, starting with the screen that mturk workers see on the platform, followed by our landing page (introduction), and so forth. Some screens are shown from multiple perspectives, i.e., player 1, 2 using different elicitation methods and belief treatments.

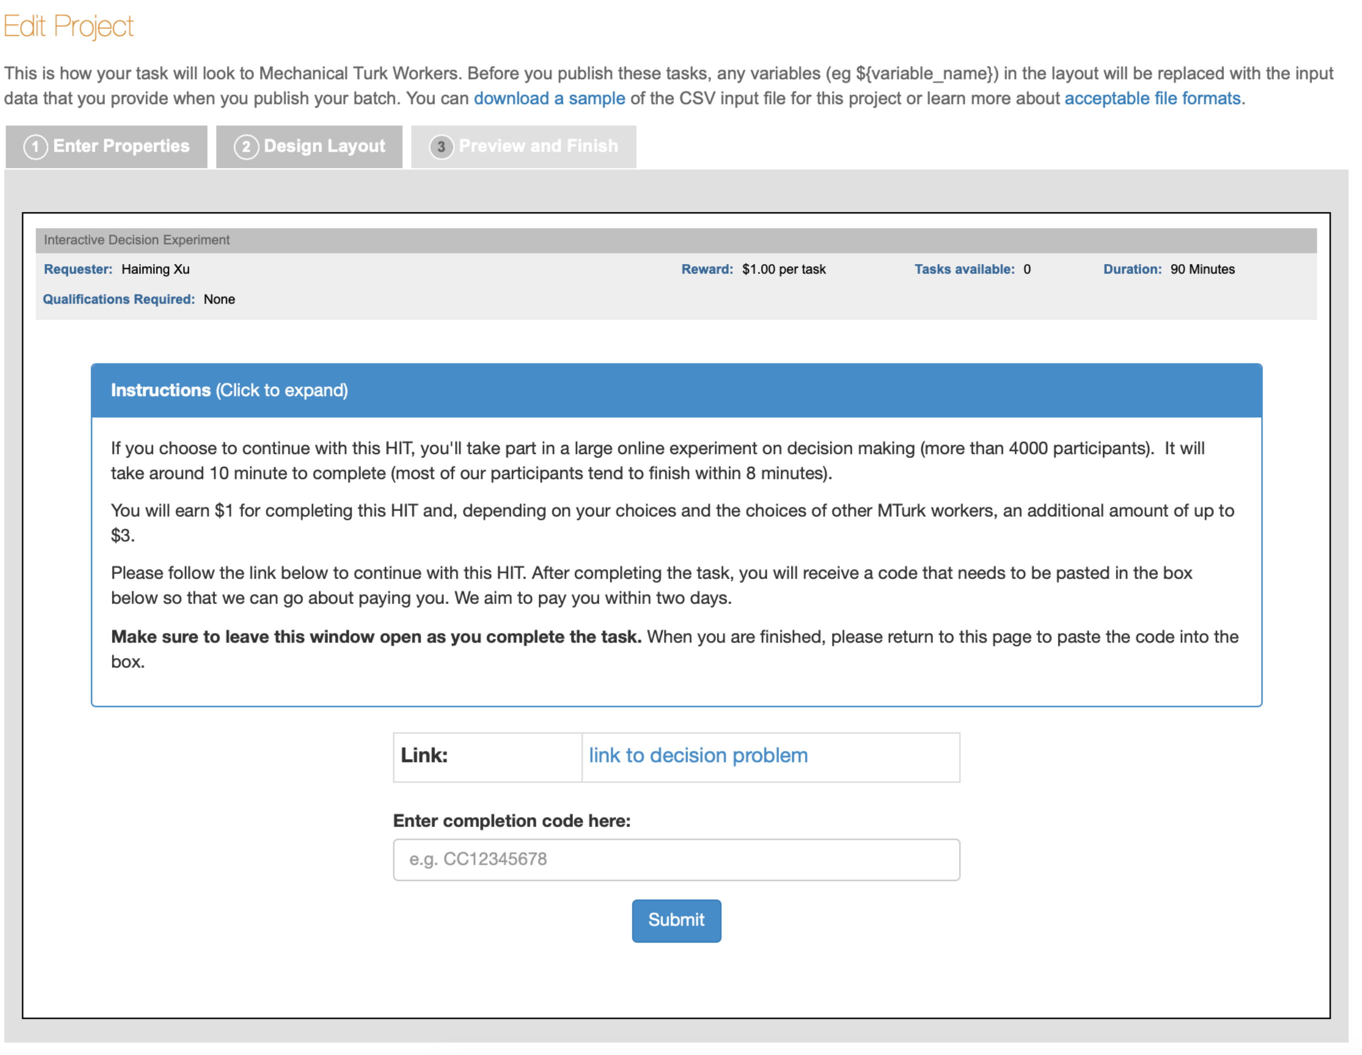

Figure E.4: Job-Ad on MTurk

Introduction

Thank you for accepting this HIT. If you choose to continue with this job, you will participate in an experiment on decision making.

This experiment has three parts: two decision tasks and a short survey with 7 questions. The entire experiment will take 10 minutes to complete. You will earn $1 for completing the HIT and, depending on your choices and the choices of other MTurk workers, an additional amount of up to $3.

In each decision task, you will be randomly matched with another participant. The interaction is completely anonymous. Neither you nor the other worker will know the other person's worker ID.

This experiment follows a no-deception policy. All tasks are implemented exactly as outlined in the instructions. The instructions are the same for all participants that you may interact with. All participants are real MTurk workers. Finally, your earnings and decision in each part of the experiment do not depend on earnings and decisions in other parts.

If you wish to continue with this HIT, please ensure you have sufficient time to complete the whole study.

Please do not close this page during the experiment. If you leave the website during the experiment, you will **not** receive any earnings. Moreover, you will only be able to participate in this experiment once.

By clicking the next button, you consent to taking part in this experiment and promise to do your best to complete the whole experiment.

Before continuing, please enter your mTurk worker ID:

Next

Figure E.5: Landing page / introduction

Instructions

We now explain how the decision tasks work. **Please read these instructions carefully** as we will ask you some simple questions about it on the next page.

Who you interact with

In each of the two decision tasks, you will be randomly matched with another Mechanical Turk worker. The interaction is completely anonymous. Neither you nor the other worker will know the other person's worker ID. Moreover, you will **not** face the same worker twice, i.e. you will interact with one participant in task 1 and another in task 2.

The amount of money that you earn in these tasks will depend on your choice and the other participant's choice. For each task, you will be given a table, similar to the one below, that summarizes your potential earnings. The numbers in the table represent real dollars.

An example of your task (slightly different from the actual task)

We will now walk you through an example to illustrate the finer details. Note that you will not be paid for this particular example and that the earnings associated with the actual tasks will be quite different.

		Other Participant	
		C	**D**
You	**A**	You earn: $2.00 Other earns: $3.00	You earn: $1.00 Other earns: $2.00
	B	You earn: $0.50 Other earns: $0.50	You earn: $6.00 Other earns: $5.00

In this example, you can choose between option **A** and **B** (the rows) while the other participant decides between **C** and **D** (the columns). If, for example, you choose **B** and the other participant chooses **D**, you will earn $6 while the other participant will earn $5.

Figure E.6: Instructions, part 1

Who acts when

Either you or the other participant will move first. You will be randomly assigned to be the **first mover** or the **second mover**. Your role will be the same for both tasks, that is you will be either a first mover for both tasks or a second mover for both tasks.

The difference between these two roles is as follows:

The first mover makes his or her decision first.
Afterwards, the second mover will be informed about the first mover's choice and decides how to respond.

Note: All information that you see as the first or second mover will also be available to the other participant.

Your earnings

Your total earnings from participating in this HIT will be sum of your earnings from the two decision tasks, money earned in the survey, and the participation fee.

Next

Figure E.7: Instructions, part 2 - direct response treatment

Who acts when

Either you or the other participant will move first. You will be randomly assigned to be the **first mover** or the **second mover**. Your role will be the same for both tasks, that is you will be either a first mover for both tasks or a second mover for both tasks.

The difference between these two roles is as follows:

The first mover makes his or her decision first.
The second mover needs to make two choices, one in response to each of the first mover's possible decisions.

For example, if you are the **second mover**, you will make the following choices:
If the first mover chooses **C**, I respond with [select **A** or **B**]
If the first mover chooses **D**, I respond with [select **A** or **B**]

The actual outcome will be determined by the first mover's choice and how the second mover responds to that *particular* choice. For instance, suppose the first mover chose **C** and you, as the second mover, chose **A** in response to **C** and **B** in response to **D**. In this case you earn $2 and the other participant earns $3.

Note: All information that you see as the first or second mover will also be available to the other participant.

Your earnings

Your total earnings from participating in this HIT will be sum of your earnings from the two decision tasks, money earned in the survey, and the participation fee.

Next

Figure E.8: Instructions, part 2 - strategy method treatment

Control Questions

Before we start with task 1, we want to ensure that you have understood the instructions.
In order to continue with this study, you will need to get at least 3 out of 4 questions correct. If you aren't quite sure about your answers, have a look at the instructions at the bottom of this page again.

Please answer the following questions:

Question 1: Do you know the identity, i.e. their MTurk ID or any other personal information, of the participant you are matched with?

◯ Yes ◯ No

Question 2: Imagine you assume the role of the second mover in task 1. Will your role change in task 2?

◯ Yes ◯ No

Question 3. In the two decision tasks, will you interact with the same Mechanical Turk worker?

◯ Yes ◯ No

Question 4: Suppose you are the first mover and earnings are determined by the following table:

		Other Participant	
		C	**D**
You	**A**	You earn: $2.00 Other earns: $3.00	You earn: $1.00 Other earns: $2.00
	B	You earn: $0.50 Other earns: $0.50	You earn: $6.00 Other earns: $5.00

Figure E.9: Control question

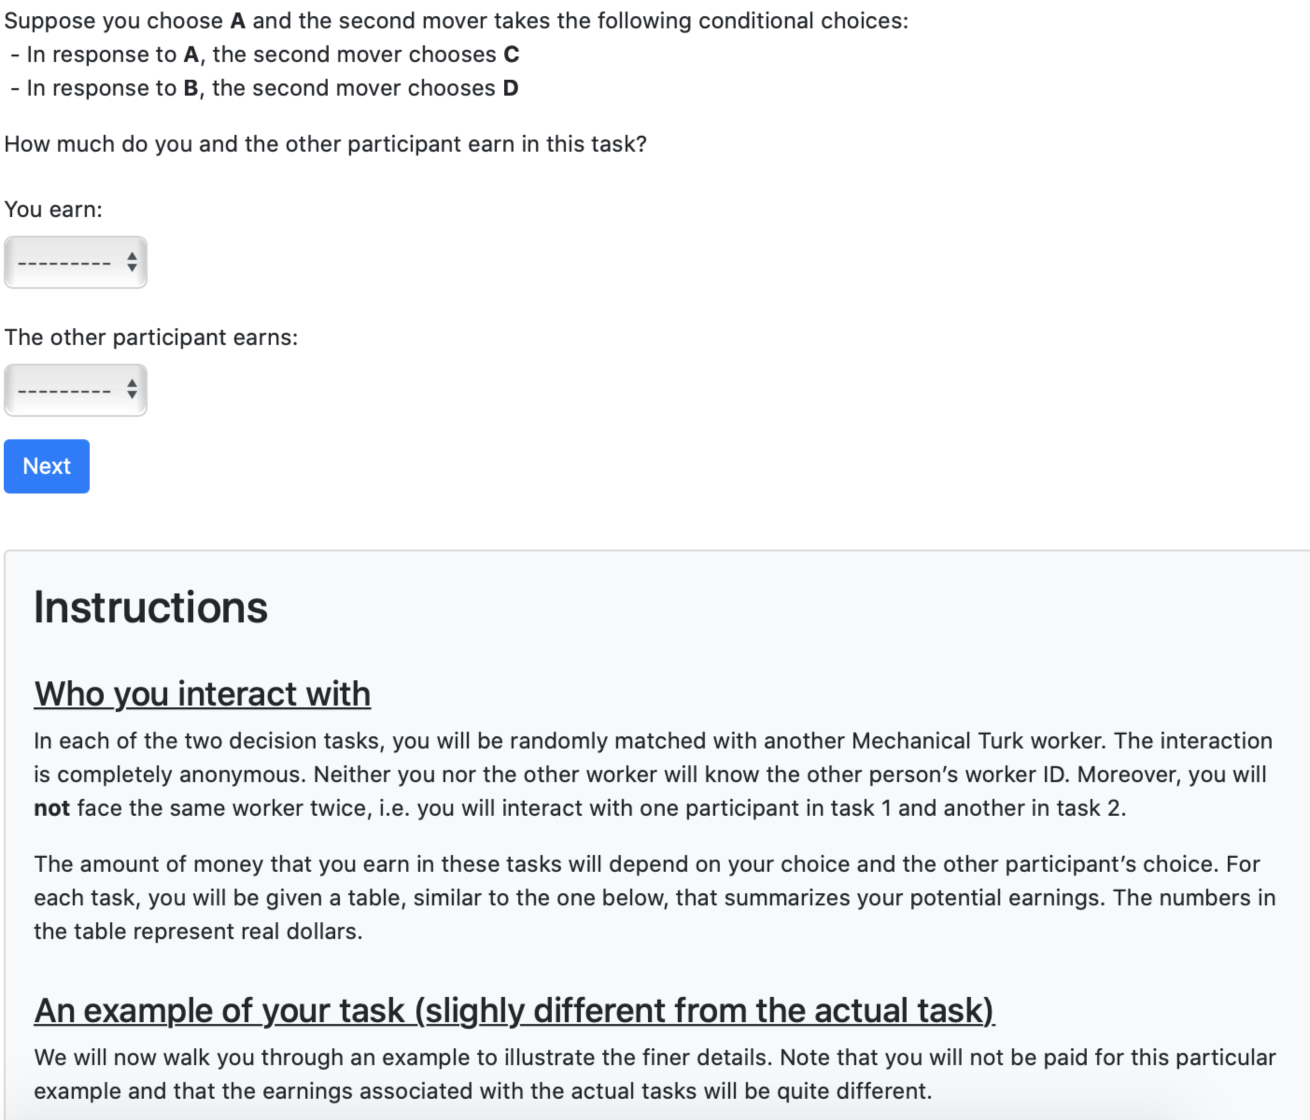

Figure E.10: Control question, continued

Note: in the grey box at the bottom of the page, the full set of instructions (from the prior page) are repeated for the participant, but are cropped for reasons of space in this screenshot.

Decision Task 1

		Other Participant	
		C	**D**
You	**A**	You earn: $1.00 Other earns: $1.00	You earn: $0.00 Other earns: $1.50
	B	You earn: $1.50 Other earns: $0.00	You earn: $0.50 Other earns: $0.50

Your role: you are the **first mover**.

Background Information: In a well-known study of this task by Watabe, Terai, Hayashi, and Yamagishi, published in the year 1996, 82.6% of the first movers chose **A**.

As the first mover, I choose:

- ○ A
- ○ B

Next

Figure E.11: Task 1, player 1, non-selfish belief treatment (sPD)

Decision Task 1

		Other Participant	
		A	**B**
You	**C**	You earn: $1.00 Other earns: $1.00	You earn: $0.00 Other earns: $1.50
	D	You earn: $1.50 Other earns: $0.00	You earn: $0.50 Other earns: $0.50

Your role: you are the **second mover**.

Background Information: In a well-known study of this task by Bolle and Ockenfels, published in the year 1990, 82.7% of the first movers chose **B**.

The other participant chose: **A**

As the second mover, I respond with:

- ○ C
- ○ D

Next

Figure E.12: Task 1, player 2, direct-response, selfish belief treatment (sPD)

Decision Task 1

		Other Participant	
		A	**B**
You	**C**	You earn: $1.00 Other earns: $1.00	You earn: $0.00 Other earns: $1.50
	D	You earn: $1.50 Other earns: $0.00	You earn: $0.50 Other earns: $0.50

Your role: you are the **second mover**.

Background Information: In a well-known study of this task by Watabe, Terai, Hayashi, and Yamagishi, published in the year 1996, 82.6% of the first movers chose **A**.

As the **second mover**

if the first mover chooses **A**, I respond with:

- ○ C
- ○ D

if the first mover chooses **B**, I respond with:

- ○ C
- ○ D

Next

Figure E.13: Task 1, player 2, strategy method, non-selfish belief treatment (sPD)

Decision Task 2

		Other Participant	
		C	D
You	A	You earn: $1.00 Other earns: $1.00	You earn: $0.00 Other earns: $0.00
	B	You earn: $1.70 Other earns: $0.30	You earn: $0.00 Other earns: $0.00

Your role: you are the **first mover**.

Background Information: In a well-known study of this task by Güth, Huck, and Müller, published in the year 2001, 70.6% of the first movers chose **A**.

As the first mover, I choose:

- A
- B

Next

Figure E.14: Task 1, player 2, strategy method, non-selfish belief treatment (mUG)

Survey - page 1/3

The first decision task you completed today was the following interaction:

		Other Participant	
		C	D
You	A	You earn: $1.00 Other earns: $1.00	You earn: $0.00 Other earns: $1.50
	B	You earn: $1.50 Other earns: $0.00	You earn: $0.50 Other earns: $0.50

Your role: you are the **first mover**.

Background Information: In a well-known study of this task by Watabe, Terai, Hayashi, and Yamagishi, published in the year 1996, 82.6% of the first movers chose **A**.

Among the MTurk workers who participated in this experiment with you today, what percentage of first movers do you think will choose **A**?

50

Note: If you are within 5% of the correct answer you will receive an additional $0.25.

Next

Figure E.15: Survey page 1, Belief Elicitation for task 1 of a player 1 in the non-selfish belief treatment (sPD)

Survey - page 2/3

The second decision task you completed today was the following interaction:

		Other Participant	
		C	D
You	A	You earn: \$1.00 Other earns: \$1.00	You earn: \$0.00 Other earns: \$0.00
	B	You earn: \$1.70 Other earns: \$0.30	You earn: \$0.00 Other earns: \$0.00

Your role: you are the **first mover**.

Background Information: In a well-known study of this task by Güth, Huck, and Müller, published in the year 2001, 70.6% of the first movers chose **A**.

Among the MTurk workers who participated in this experiment with you today, what percentage of first movers do you think will choose **A**?

50

Note: If you are within 5% of the correct answer you will receive an additional \$0.25.

Next

Figure E.16: Survey page 2, Belief Elicitation for task 2 of a player 1 in the non-selfish belief treatment (mUG)

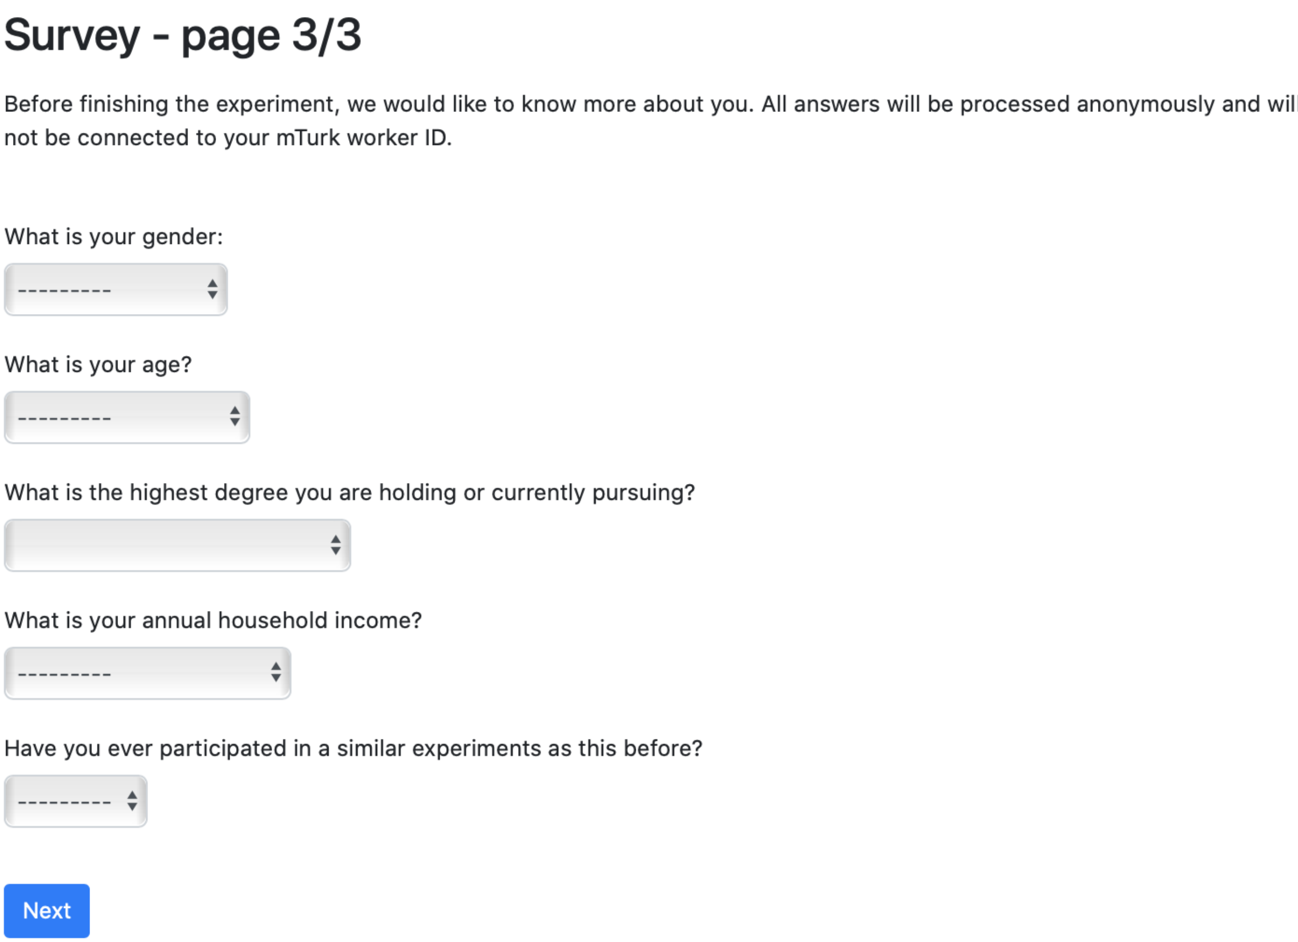

Survey - page 3/3

Before finishing the experiment, we would like to know more about you. All answers will be processed anonymously and will not be connected to your mTurk worker ID.

What is your gender:

What is your age?

What is the highest degree you are holding or currently pursuing?

What is your annual household income?

Have you ever participated in a similar experiments as this before?

Next

Figure E.17: Survey page 3

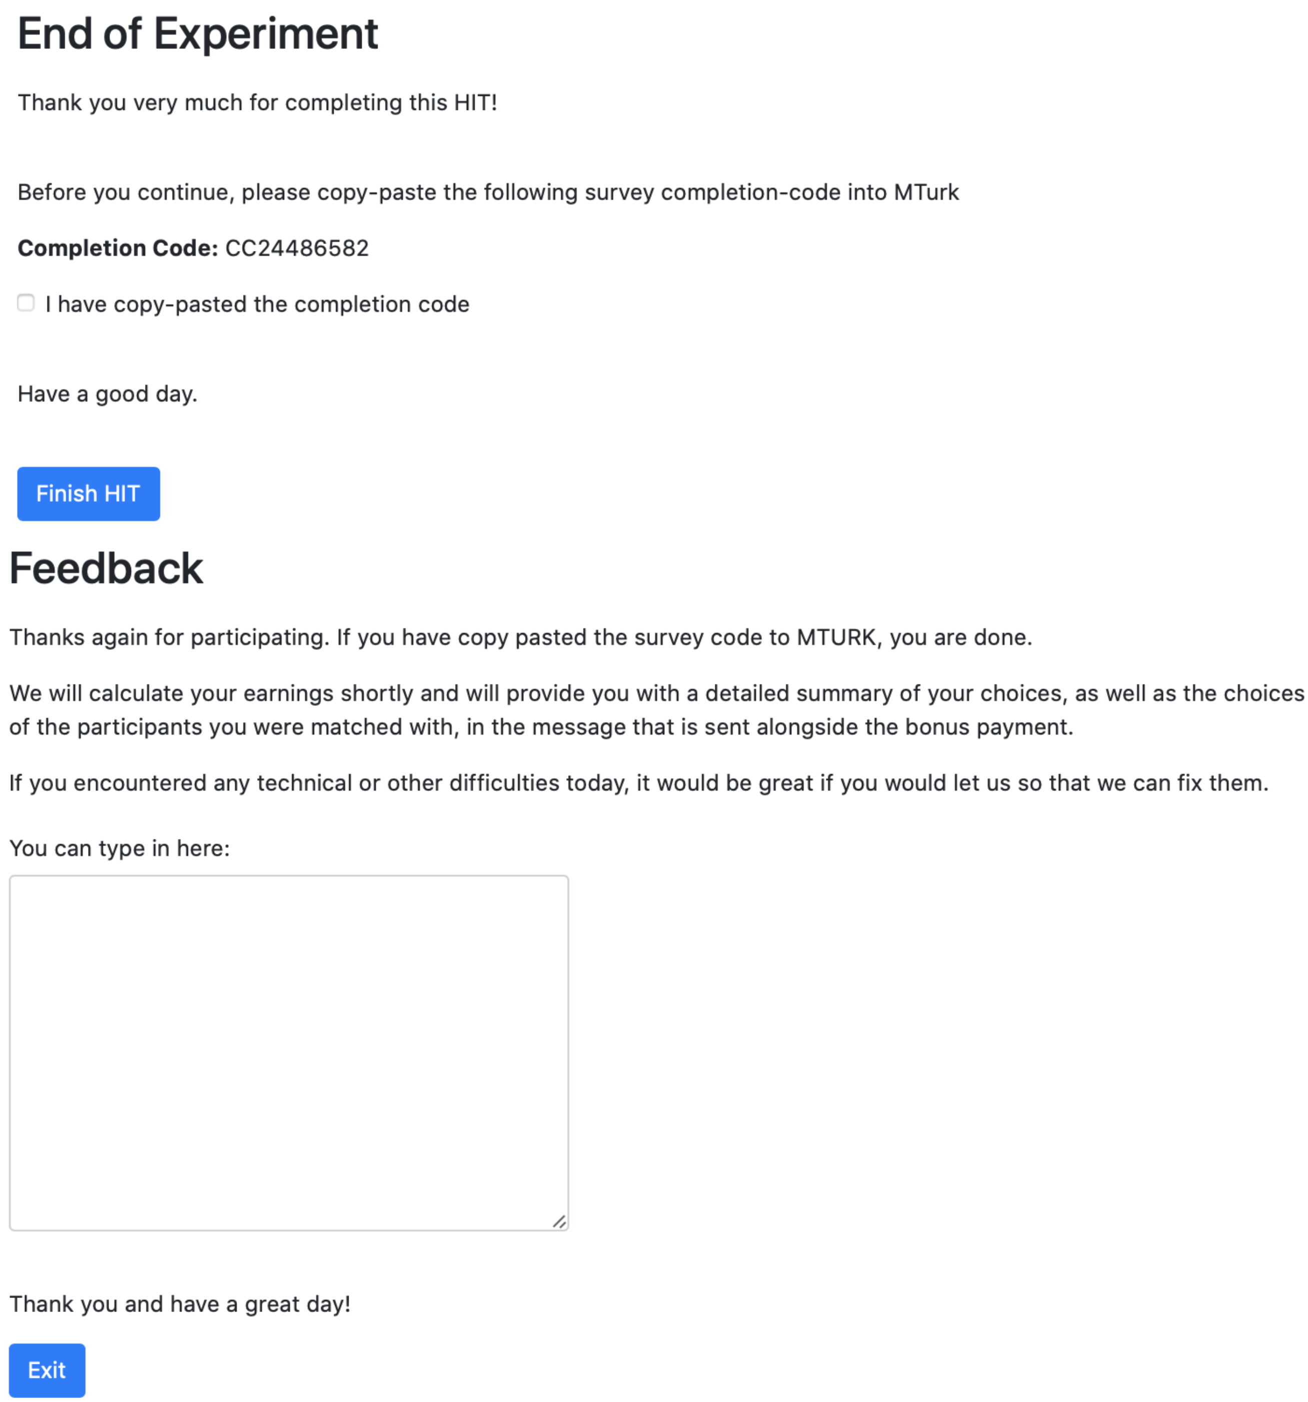

End of Experiment

Thank you very much for completing this HIT!

Before you continue, please copy-paste the following survey completion-code into MTurk

Completion Code: CC24486582

☐ I have copy-pasted the completion code

Have a good day.

Finish HIT

Feedback

Thanks again for participating. If you have copy pasted the survey code to MTURK, you are done.

We will calculate your earnings shortly and will provide you with a detailed summary of your choices, as well as the choices of the participants you were matched with, in the message that is sent alongside the bonus payment.

If you encountered any technical or other difficulties today, it would be great if you would let us so that we can fix them.

You can type in here:

Thank you and have a great day!

Exit

Figure E.18: Completion page and optional feedback Page

End of Experiment

Thank you for attempting this HIT.

We regret to inform you that you are **not allowed to continue** with this experiment as you answered two or more questions (out of 4) incorrectly.

We hope you understand that we cannot provide you with a completion code as a result.

After closing this page, please be so kind and return the HIT in MTurk.

Have a good day.

Finish HIT

Figure E.19: Early termination screen for those who did pass the control questions