Conditionally accepted for publication in ***Experimental Economics***

How do Monetary Incentives Affect the Measurement of Social Preferences?*

Ernst Fehr ⓡ Julien Senn ⓡ Thomas Epper ⓡ Aljosha Henkel

October 1, 2025

Abstract

In this registered report, we investigate (i) whether incentives affect subjects' willingness to pay to *increase*, and to *decrease* the payoff of others, (ii) whether they affect the *distribution* of social preference types, and (iii) whether they affect the *strength* and the *precision* of individuals' structurally estimated social preference parameters. Using an online experiment with a general population sample, we show that the use of monetary incentives, as well as the size of the stakes, have little impact on subjects' modal choices (descriptive analysis), as well as for the distribution of *qualitatively* distinct preference types in the population (clustering analysis). However, monetary incentives affect *quantitative* measures of the strength and the precision of social preferences. Indeed, a structural analysis reveals that the preference elicitation with merely hypothetical stakes leads to an overestimation and a less precise measurement of social preferences. Together, these results highlight that incentivizing the elicitation of social preferences is most useful when interested in *quantitative* estimates. For researchers interested in identifying merely *qualitative* preferences types, however, hypothetical stakes might suffice.

Key Words: Social Preferences, Altruism, Inequality Aversion, Incentives

JEL Codes: C80, C90, D30, D63

*The "ⓡ" symbol indicates that the authors' names are in certified random order. Ernst Fehr acknowledges the support from the Research Priority Program of the University of Zurich on Equality of Opportunity. Thomas Epper acknowledges the support from the Métropole Européenne de Lille (MEL). Epper: CNRS, IESEG School of Management, Univ. Lille, UMR 9221 - LEM - Lille Economie Management F-59000 Lille, France (thomas.epper@cnrs.fr). Fehr: Department of Economics, Zurich University. Blümlisalpstrasse 10, 8006 Zurich, Switzerland (ernst.fehr@econ.uzh.ch). Henkel: ETH Zurich, KOF Swiss Economic Institute. Leonhardstrasse 21, CH–8092 Zurich, Switzerland (henkel@kof.ethz.ch). Senn: University of Paris 1 Panthéon-Sorbonne, Paris, France (julien.senn@univ-paris1.fr).

1 Introduction

Social preferences have been shown to play an important role in a wide variety of domains such as the labor market (e.g., Fehr et al., 1993; Charness, 2000; Bellemare and Shearer, 2007; Dur, 2009; Kube et al., 2012), bargaining decisions (e.g., Camerer and Thaler, 1995; Camerer and Loewenstein, 1993; Camerer, 2011), political economy (e.g., Tyran and Sausgruber, 2006; Kerschbamer and Müller, 2020; Fisman et al., 2017; Fehr ⓡ al., forthcoming; Henkel ⓡ al., 2025), and contract design (e.g., Bierbrauer and Netzer, 2016; Bierbrauer et al., 2017; Schmidt and Ockenfels, 2021; Fehr et al., 2021), among others. Given their theoretical and empirical relevance, it is critical to measure them accurately. In this context, a particularly important question is whether monetary incentives matter for the elicitation of social preferences.

In contrast to other disciplines, the use of monetary incentives has been a pillar of experimental economics (Plott, 1986; Smith, 1982, 1991; Hertwig and Ortmann, 2001). Among others, behaviors measured using incentivized decisions are believed to more accurately capture subjects' "true preferences" than self-reported answers in hypothetical scenarios. In the context of social preferences, where concerns regarding social desirability are particularly high, there are good reasons to believe that monetary incentives might matter for the weight that people put on others' payoff.[1] However, incentivizing decisions also comes at a cost for researchers, who often need to spend thousands of dollars to incentivize the decisions of their participants. This raises the following question: Do monetary incentives improve the measurement of social preferences, or can a non-incentivized elicitation method mimic the properties of an incentivized task?

We study this question using a pre-registered online experiment conducted on Prolific with over 3,000 subjects broadly representative of the US population.[2] Broad population samples have, for example, been used to study questions such as the role of social preferences for political support for redistribution (Fisman et al., 2017; Kerschbamer and Müller, 2020; Fehr ⓡ al., forthcoming; Henkel ⓡ al., 2025), to get a deeper understanding about how people reason about the economy (Andre et al., 2022; Stantcheva, 2021), or to study the role of beliefs for willingness to act against climate change (Dechezleprêtre et al., 2022; Falk et al.,

[1]Consider a standard dicator game where subjects are asked to decide how to split USD 10 with an anonymous recipient. When the decision is hypothetical, sharing the money and behaving in an altruistic way is costless, which might lead to an overestimation of the extent to which individuals are other-regarding. Tying subjects' payment to their decisions might mitigate this issue by forcing decision makers to carefully tradeoff their own material benefit with other-regarding concerns.

[2]This paper was submitted as a registered report. The proposal for the registered report is available at https://www.socialscienceregistry.org/trials/15147.

2021), among others. The question of the role of monetary incentives for the identification of social preferences in broad population samples is particularly relevant given that several survey providers do *not* allow to incentivize decisions. In this context, understanding whether incentivization matters is an important empirical question.

We elicit social preferences using a set of choice situations in which the decision maker has to decide on how to allocate "points" between herself and an anonymous other participant (Fehr ⓡ al., forthcoming). While the literature has often relied on simple dictator games (i.e., choice situations in which subjects can sacrifice resources to *increase* the payoff of others) to identify social preferences, our experimental paradigm also includes several decision situations where the decision maker can pay in order to *decrease* the payoff of others. This allows for the identification of a broader range of social preferences. Indeed, while standard dictator games are well suited to identify altruism and the extent to which individuals are willing to trade-off equality and efficiency (see, e.g., Fisman et al., 2007, 2017), they do *not* allow to identify a broad range of other social preferences. For example, inequality aversion (Fehr and Schmidt, 1999; Charness and Rabin, 2002; Bolton and Ockenfels, 2000) implies that individuals might not only be willing to sacrifice some of their own payoff to increase the payoff of those worse off (aversion to advantageous inequality), but also to decrease the payoff of those who are better off (aversion to disadvantageous inequality). Similarly, envious and spiteful individuals are willing to pay to destroy the payoff of others. Our design solves this identification issue.

To identify the effects of monetary incentives, we exogenously vary whether decisions are incentivized, and by how much. In the *Low-* and the *High-Incentives* treatments, decision makers are paid on the basis of their choice in a randomly drawn choice situation. These two treatments, in which stakes are scaled by a factor of 5, allow us to assess whether the *strength* of monetary incentives matters for the identification of social preferences. In the *Hypothetical* treatment, subjects' decisions were *not* incentivized, but decision makers were encouraged to make a decision "as if" decisions were incentivized. Together, these treatments allow us to cleanly identify whether and how monetary incentives affect the identification of social preferences. We explore these questions in several steps.

First, we investigate whether and how monetary incentives affect subjects' modal choices at the descriptive level. Specifically, we examine the effects of monetary incentives for subjects' willingness to pay to *increase* the other participants' payoff (modal choice on negatively sloped budget lines), and for their willingness to pay to *decrease* the other participants' payoff (modal choice on positively sloped budget lines). In all treatments, we find that subjects' choices are strikingly similar, irrespective of whether decisions are incentivized or not, and irrespective of the size of the monetary stakes.

Second, we assess whether monetary incentives play a role for the *distribution* of social preference types in the population. To answer this question, we uncover the distribution of preferences in each treatment by applying the Dirichlet Process means (DP-means) algorithm, a Bayesian nonparametric clustering algorithm. In all the treatments, we find that the *same* three clusters with a clear behavioral interpretation emerge: an inequality averse type, an altruistic type, and a predominantly selfish type. Moreover, the *distribution* of types is remarkably similar across treatments, with each type comprising about one third of subjects in all treatments. These findings are noteworthy, given that the DP-means algorithm does not impose any assumptions on the behavioral interpretation of types nor on their distribution.

Third, we assess whether monetary incentives affect the *strength* of social preferences. To answer this question, we compare across the different treatments the (distributions of) structurally estimated parameters of a model of social preferences that comprises inequality aversion, altruism and envy as special cases (Fehr and Schmidt, 1999; Charness and Rabin, 2002). In all treatments, the average estimated levels of other-regardingness are sizeable. Indeed, our structural estimates of aversion to disadvantageous inequality (the α-parameter in the Fehr Schmidt model) range from 0.219 in the Low-Incentives to 0.368 in the Hypothetical treatments, corresponding to a willingness to pay of 18 to 26.9 cents to *decrease* the payoff of those ahead by one dollar. Turning to aversion to advantageous inequality (the β-parameter in the Fehr Schmidt model), we estimate parameters ranging from 0.613 in the Low-Incentives to 0.744 in the Hypothetical treatment, corresponding to a willingness to pay to increase the payoff of those worse off by one dollar of 1.58 and 2.90, respectively. Importantly, we find that the size of the monetary stakes (Low vs. High-Incentives) does *not* play a large role for the strength of social preferences. However, social preferences parameters are consistently larger—i.e., likely to be overestimated—when stakes are hypothetical. Interestingly, we show that these treatment differences are mainly driven by subjects identified as being inequality averse by the clustering algorithm. For these subjects, a lack of monetary incentives yields larger aversion to both advantageous and disadvantageous inequality.

We also explore how monetary incentives affect the *precision* of structural estimates. To that end, we compare individuals' posterior standard deviation of the estimated social preference parameters.[3] We find that the precision of the estimates improves substantially when moving from the hypothetical to the incentivized treatments, suggesting that preference parameters recovered under monetary incentives are more reliable and better measures of social preferences. The precision of estimated preference parameters in the low and the high incentives treatment are, however, quite similar.

Together, our results indicate that monetary incentives and stake size have only small effects on subjects' modal choices and on the identification of *qualitatively* distinct preference types. However, they do affect the *quantitative* assessment of the strength and precision of social preferences. In particular, our structural analysis suggests that social preferences tend to be overestimated when elicited using hypothetical stakes. These findings imply that whether or not one should use incentives to elicit social preferences depends on the research objective. If one is interested in obtaining a rough, aggregate, measure of social preferences (e.g., in terms of subjects' modal choices) or in identifying qualitatively distinct preference types, then hypothetical stakes may suffice. If, however, one is interested in quantitative estimates, e.g., to quantitatively calibrate a theoretical model that is then used for predictive purposes, then the use of monetary incentives is advisable, as hypothetical stakes appear to inflate inequality aversion and lead to much more imprecisely estimated preference parameters.

While our paper and the results above are based on a large sample drawn from the general population, it is also interesting to understand the role of monetary incentives for the identification of social preferences in students. Indeed, student samples are still widely used in economic experiments, and insights from the general population might not extend to students. To shed light on this issue, we also collected an additional sample of students carefully prescreened from Prolific.[4] In a nutshell, we find that some of our conclusions extend to students: the use of monetary incentives (and their size) appears to have little effect on the identification of social preferences both at the descriptive level of subjects' modal choices and in the clustering analysis. The results of the structural analysis are a little less clear cut, as we find the largest levels of inequality aversion (and the lowest precision in the estimates) in the High-Incentives treatment. However, these results need to be interpreted with caution

[3]Note that this posterior standard deviation is a measure of the uncertainty of the estimated preference parameters at the individual level.

[4]Using students from Prolific allows to keep the entire experimental protocol constant and enhances the comparability between the results from the general population and the results from the student sample. Note, however, that our aim is *not* to extensively analyze differences in social preferences between students and the general population. We address this specific research question in a separate paper (Epper ⓡ al., 2023).

because our student sample is much less well powered than our general population sample.

Our paper is connected to the literature interested in the effects of using monetary incentives in experimental research. In particular, our paper is linked to studies that have investigated the effects of monetary rewards for behavior in dictator games using student samples.[5] An early contribution to this literature is Forsythe et al. (1994), who compare the effects of incentivized versus hypothetical decisions in a dictator game. They find that dictators donate significantly less when decisions are incentivized. They also study the effects of stake sizes ($5 vs $10) and cannot reject the hypothesis that stakes do *not* affect behavior.[6] More recently, Bühren and Kundt (2015) report on an experiment where participants are asked to make decisions in three games (a prosocial, an envy, and a sharing game). Subjects are randomized into a treatment condition where decisions are incentivized or a condition where decisions are hypothetical. They find that participants in the incentives treatment display more spite and less inequality aversion. Clot et al. (2018) compare the effects of different compensation mechanisms in a dictator game and find that hypothetical stakes lead to fewer egoistic and more egalitarian decisions.[7]

Engel (2011) reviews the literature on dictator games in his meta-analysis, covering over 100 papers published prior to 2010. On the role of incentives, he concludes that dictators' behavior when decisions are incentivized is *not* significantly different than when their decisions are hypothetical. He only finds very weak evidence that stakes size affects giving, despite covering studies with dramatic changes in stakes sizes ranging from $0 to $130.[8] These results are largely confirmed by a more recent meta-analysis by Larney et al. (2019), who find a significant but small effect of stake size on dictator games offers (and no effect of stakes on ultimatum game offers). Note, however, that this meta-analysis excluded studies with hypothetical stakes, and therefore cannot provide evidence on the effects of incentivization *per*

[5] Other papers have studied the effects of alternative incentivization techniques (e.g., probabilistic payment) for giving in dictator games. We do not review these papers here for brevity. For interested readers, we recommend the excellent meta-analyses by Engel (2011) and by Larney et al. (2019). These reviews also cover a handful of other studies which we do not report here due to their very low sample sizes or their lack of randomization of treatment conditions.

[6] In addition, they also study whether the conclusions drawn in dictator games carry over to ultimatum games. Note, however, that these experiments have low statistical power.

[7] In addition, they also show that probabilistic payment does not affect decisions, holding stakes constant. When stakes differ but expected stakes are constant, they report that dictators behave in a more egoistic way when their decision is implemented with certainty.

[8] The relation between stake size and giving is insignificant when he uses all the papers covered in his meta-analysis. However, when he focuses on studies that explicitly manipulated stakes sizes, he finds that higher stakes significantly reduce dictators' willingness to give. However, he qualifies this effect as "*very small*", despite the wide variation in stake size.

se.

More recently, Hufe and Weishaar (2025) studied the effects of incentives for fairness ideals. They find that respondents make similar choices when allocating monthly earnings between two persons in a spectator design, independent of whether decisions are hypothetical or not.[9] Kosfeld et al. (2025) conducted a cross-cultural validation of the (Global) Preference Survey Module (Falk et al., 2018, 2023) and find that quantitative survey items (hypothetical experiments) aimed at capturing social preferences tend to be good predictors of behavior in related incentivized choice experiments.[10] Last, Fitzgerald (2024) shows that traditional hypothetical bias measures are often misleading estimates of hypothetical bias for intervention experiments.

Our results contribute to the existing literature in several ways.

First, our study goes beyond comparing choices in dictator games or ultimatum games across incentives conditions. Instead, we study whether incentives matter (i) for subjects' willingness to pay to *increase* the payoff of others, and for their willingness to *decrease* the payoff of others, (ii) for the *distribution* of social preference types, where we uncover social preferences using state-of-the-art Bayesian nonparametric methods, and (iii) for the *strength* of social preferences and the *precision* of its structural estimates. As such, the scope of this study is—to the best of our knowledge—much broader than the existing studies on the role of monetary incentives for social preferences.

Second, while previous research has investigated the effects of monetary incentives for behavior in dictator games, these studies often rely on student samples of relatively modest size. In contrast, our study addresses this question in a much larger sample drawn from the general population. In our view, this is important given that such samples are being increasingly used by researchers as they are economically much more relevant than student samples.

Finally, the fact that this paper reports on an ex-ante carefully pre-registered experiment and analysis lends further credibility to the empirical findings we document. In this sense, this study also contributes to the broader discussion on best practices in scientific research and the role that pre-registered studies and registered reports can play at increasing transparency in science (Nosek et al., 2018; Munafò et al., 2017; Miguel, 2021).

[9] Note that the incentivized treatment is implemented probabilistically, i.e., the income of a single individual will be determined by the decision of a single, randomly drawn, individual.

[10] Kosfeld et al. (2025) empirically document this relationship for the altruism and positive reciprocity items but not for negative reciprocity.

2 Experimental design

2.1 Measuring distributional preferences

We elicited respondents' distributional preferences using a series of *twelve* incentivized money allocation tasks in which participants had to decide how to allocate experimental currency units (ECUs) between themselves and an anonymous other participant of the study. Figure 1 depicts these 12 budget lines, where the decision maker's own payoff is represented on the x-axis and the recipient's payoff is on the y-axis.[11]

Figure 1: Budget lines used to identify other-regarding preferences

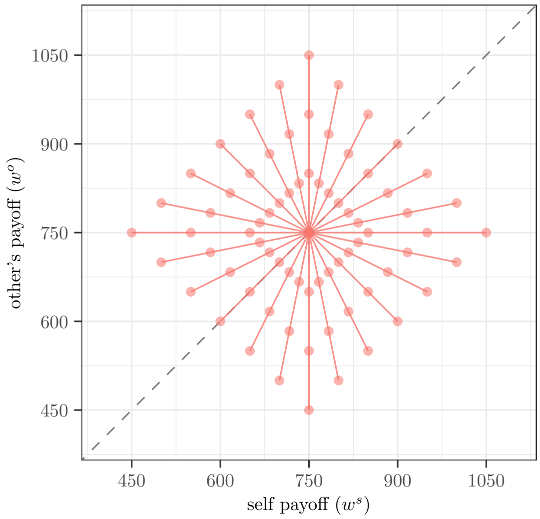

These twelve choice situations systematically vary the cost and the efficiency consequences of redistribution, thereby allowing us to identify a wide range of other-regarding behaviors. Negatively sloped budget lines (where subjects can sacrifice resources to *increase* the payoff of the other participant) allow us to identify behaviors such as altruism and aversion to advantageous inequality. Positively sloped budget lines (where subjects can pay to *decrease* the payoff of the other) allow us to identify behaviors such as envy, spite and aversion to disadvantageous inequality, among others.

Choice situations were presented to subjects in random order directly on subjects' screens.

[11]The design is based on Fehr ⓡ al. (forthcoming, 2023). We provide further details on the various budget lines used for the identification of social preferences in Table 1 in the Appendix B.1.1. In addition, our design also comprises eight more choice situations that can be used to further validate the behavioral interpretation of the types identified. We provide further details on these additional budget lines in Appendix B.1.2.

They were presented in a way that made the distributional consequences of each allocation transparent. Figure 2a illustrates how a typical choice situation was presented to our participants. In each choice situation, subjects were able to choose between seven interpersonal allocations (labeled by 1 to 7)—all of them located on a budget line. Each available allocation consisted of a specific distribution of ECUs between the participant (bars labeled by "You receive") and the other person (bars labeled by "other person receives"). We represented the available choices numerically and graphically in order to make the trade-offs and the associated payoff implications salient. Figure 2b plots the budget line corresponding to the example depicted in Figure 2a in the (ω^s, ω^o)-space ("own payoff", "other's payoff"). In this example, the slope of the budget line is -2, indicating that for every ECU the decision maker gives up, the other player receives 2 ECUs. Here, perfect equality in payoffs can be achieved by choosing allocation 4.

Figure 2: Example choice situation

(a) Decision screen (b) Budget line

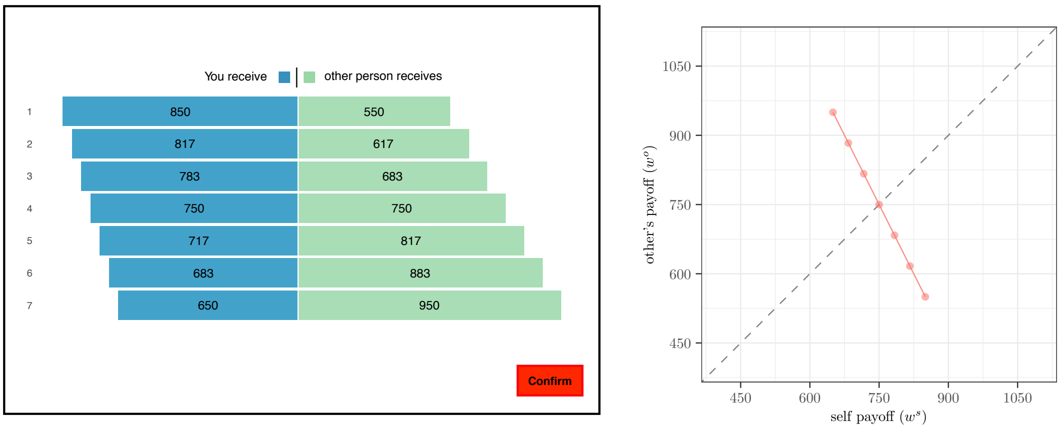

2.2 Treatments

To assess the effects of monetary incentives on measures of social preferences, we randomized individuals into one of the following three treatments in a *between* subjects design.

Low-Incentives. In this treatment, the decisions of participants had real monetary consequences for them and for another, anonymous, participant of the study.[12] The exchange rate

[12]Importantly, our instructions made it clear that the other participant could *not* affect the decision maker's payoff.

between points and USD was set to 500 points = USD 1.

High-Incentives. In this treatment, the stake size was substantially larger than in the Low-Incentives treatment, as we set the exchange rate to 100 points = USD 1, i.e., stakes were *five times larger* than in the Low-Incentives treatment.

Hypothetical. In this treatment, participants choices were hypothetical. Thus, their decisions neither affected their own payoff, nor the payoff of another participant. However, our instructions invited subjects to make decisions *as if* they were incentivized.

Importantly, we made the consequences of subjects' decisions salient by highlighting them in the relevant parts of the instructions (bold fonts). In addition, we included a control question specifically aimed at ensuring that participants understood whether their decisions had real monetary consequences or not (see the control questions at the end of the instructions in Appendix B.2). This was crucial to prevent confusion about the nature of the task, especially in the Hypothetical treatment where incentives were absent. By drawing attention to the presence or absence of incentives and verifying comprehension, we aimed to minimize noise in the measurement of social preferences stemming from misunderstandings. This, in turn, increases our confidence that any observed differences in behavior across treatments are attributable to the incentive structure, rather than differences in understanding.

2.3 Sample

We conducted our study with a general population sample that is broadly representative of the US population with respect to age, gender, and political affiliation. Our final sample comprises of 3,032 subjects, recruited on Prolific.[13] Descriptive statistics on participants' main socio-demographic characteristics can be found in Table C.1 in Appendix C. The average respondent in our sample is 46.2 years old and the share of men is 48.3%. In terms of political leaning, 33.3% of our sample leans towards the Republican party, while 40.8% leans towards the Democratic party. The remaining 25.9% lean towards the Independent or no party. In addition, the last column of Table C.1 also shows that the treatments are generally well balanced among the main observable characteristics.

[13]We discuss the details of the statistical power of our design in Appendix D.

While our main analysis relies on the general population sample, it is also interesting to understand the role of monetary incentives for the identification of social preferences in students. To shed light on this issue, we collected responses from an additional 320 students carefully prescreened from Prolific.[14] The advantage of recruiting students directly from Prolific is that it keeps the entire experimental protocol constant, thereby enhancing the comparability between the results from the general population and the results from the student sample. For the analysis, we pool the 320 subjects recruited in this separate student sample with the 175 students contained in the general population sample, which were identified using the same screening-criteria. Results for this student sample are reported in the Appendix G.

2.4 Data collection and experimental protocol

The experiment was fully computerized using Qualtrics and all the instructions were displayed directly on participants' screens. The study included control questions for the money allocation task. These questions aimed at identifying participants who did not understand the task, or did not pay attention (Berinsky et al., 2014). Participants who failed to pass these control questions were excluded from the final sample. We provide a transcript of the instructions and control questions displayed on participants' screens in Appendix B.2.

All participants were paid a show-up fee of USD 3, provided that they completed the study until the end. In addition, we incentivized respondents' choices in the *Low-Incentives* and the *High-Incentives* treatment by implementing one of their decisions at random.[15]

We pre-registered the study on the AEA RCT registry (AEARCTR-0015147), where we uploaded the (accepted) Stage 1 version of this registered report, along with a detailed pre-analysis plan.[16] Ethics approval was obtained from the Human Subjects Committee of the Department of Economics of the University of Zurich (OEC IRB #2024-043).

[14]Specifically, we recruited students from Prolific who jointly met the following (preregistered) prescreening-criteria: (i) they are located in the USA, (ii) they are currently studying, (iii) they are enrolled in an undergraduate (BA/BSc) or graduate program (MA/MSc/MPhil), and (iv) they are aged between 18 and 30. In our view, subjects who meet these inclusion criteria share the main characteristics of students typically recruited in traditional subject pools for laboratory studies.

[15]Median time to complete the study for subjects in the general population sample was 12.7, and it was 12.0 minutes in the student sample. In the general population sample (student sample), participants in the *Low-Incentives* treatment received an average variable payment of USD 1.69 (USD 1.66), while participants in the *High-Incentives* treatment received an average variable payment of USD 8.35 (USD 8.37).

[16]https://www.socialscienceregistry.org/trials/15147.

3 Hypotheses

Our experimental design allows to shed light on the effects of monetary incentives on measures of social preferences. A key feature of our proposal is that we investigate this question at different levels of analysis. In particular, we examine the *distribution* of qualitatively distinct preference types in a population (clustering analysis), and we assess both the *strength* and the *precision* of these preferences (structural analysis). Thus, our hypotheses consider both of these dimensions.[17]

Our first hypothesis relates to whether relying on incentivized decisions is critical for the measurement of social preferences, or whether hypothetical questions suffice. In the context of social preferences, where concerns regarding social desirability are particularly relevant, individuals might not reveal their true preferences in the absence of real monetary stakes. For example, it is costless for a subject to behave in an altruistic way if decisions are hypothetical. If that is the case, hypothetical stakes might lead to an overestimation of the extent to which individuals are other-regarding. This can be particularly problematic if quantitative estimates are used to make behavioral predictions. Tying subjects' payment to their decisions might mitigate this issue by forcing decision makers to more carefully trade off their own material benefit with other-regarding concerns.

Hypothesis 1. *Monetary incentives do not affect measures of social preferences.*

If our results show that monetary incentives do *not* affect measures of social preferences, researchers could more broadly adopt our elicitation procedure with hypothetical stakes. This would be particularly appealing in contexts where incentivization is logistically difficult to organize (e.g., some online studies) or very costly. If, in contrast, our results show that monetary incentives do affect social preferences, then we will be able to assess how large the mismeasurement related to hypothetical stakes is. This could be particularly useful for researchers who cannot incentivize their subjects but want to know how large the bias due to lack of incentives is likely to be.

Our second hypothesis relates to the effects of stake size. While the social preferences elicited under hypothetical stakes might differ from those elicited with real monetary stakes (Hypothesis 1), it is also possible that the size of the monetary stakes matters for the measurement of social preferences. In particular, larger stakes might affect the distribution of social

[17] These hypotheses correspond to the ones outlined in the Stage 1 version of this registered report. Likewise, our analysis follows the analysis plan that we preregistered.

preferences in a population. This might be the case if, for example, larger stakes induce decision makers to make more selfish decisions. It is also possible that higher stakes increase the precision with which social preferences are estimated, e.g., if higher stakes lead participants to think more carefully about their decisions.

HYPOTHESIS 2. *The strength of monetary incentives has no effect on measures of social preferences.*

If our results suggest that the stake size does *not* matter, then researchers could largely rely on low-powered incentives to reliably elicit social preferences. It is, however, also possible that stake size will affect social preferences. In particular, as discussed above, stakes might affect the precision with which social preferences are estimated. If this is the case, then researchers interested in *precisely* estimating the strength of social preferences may want to rely on using high-powered incentives to elicit preferences, whereas researchers only interested in the qualitative nature of preferences may rely on low-powered incentives.

Finally, it is important to note that answers to the two hypotheses above might vary depending on whether one considers results from the clustering analysis or the structural analysis. For example, it is possible that there are no treatment differences in *all* the dimensions in which we assess social preferences, e.g., that the same behavioral types emerge in the same proportions across treatments *and* that the structurally estimated parameters are identical. However, it is also possible that treatment differences exist in only some dimensions. For example, it is plausible that the distribution of preferences types as identified by the clustering remains stable across treatments, but that there are treatment differences in the structurally estimated parameters. Such a result would imply that the nature of the research questions should determine whether monetary incentives should be used or not: Researchers only interested in assessing the qualitative nature of behavioral types and their prevalence in the population could rely on either hypothetical or real monetary stakes, while researcher interested in the quantitative distribution of types and the strength of these preferences should carefully consider using the appropriate incentivization method.

4 Results

This section closely follows the analysis plan that we pre-registered. In some instances, we also provide additional analyses that were not pre-registered but that are useful to further clarify our findings. For transparency, we list these additional analyses in the Appendix E.1.

4.1 Descriptive analysis

We start our analysis by exploring subjects' choices at the descriptive level. For each budget line, we label the own-payoff maximizing allocation by $z = 1$, the own-payoff-minimizing allocation by $z = 0$, and the payoff-equalizing allocation by $z = 0.5$. The other four available allocations on each budget line are equidistantly placed between 0–0.5 and 0.5–1, respectively. We are interested in the following questions:

- *How do incentives affect subjects choices in the money allocation task at the descriptive level?*
- *Do incentives affect subjects' willingness to pay to decrease the other participant's payoff?*
- *Do incentives affect subjects' willingness to pay to increase the other participant's payoff?*

To answer these questions, we examine whether the treatments affect the distribution of subjects' *modal* choice separately for negatively sloped and for positively sloped budget lines (Figure 1). The distinction between negatively and positively sloped budget lines is important because behavior across budget lines *within* the class of negatively sloped budget lines inform us about how much money individuals are willing to sacrifice to increase another individual's payoff, whereas behavior across budget lines within the class of positively sloped budget lines informs us about how much money individuals are willing to sacrifice to decrease another individuals' payoff. We focus on the mode because it is less susceptible to random responses and to outliers than the mean or the median.[18]

We depict the results of this analysis in Figure 3. The Figure reveals that the distributions of modal choices are strikingly similar across treatments. Among the negatively sloped budget lines, the modal choice of the vast majority of individuals is located at either $z = 0.5$ or $z = 1$. For example, in the Low-Incentives treatment, the modal choice of 42.12% of the participants is $z = 0.5$, whereas it is $z = 1$ for 48.70% of the participants. This means that the majority of participants predominantly choose either the allocation that perfectly equalizes their and the other person's payoff ($z = 0.5$), or the one that maximizes their own payoff ($z = 1$). Likewise, in the High-Incentives treatment, the modal choice of 46.72% of the participants is to equalize payoffs ($z = 0.5$), while it is to maximize own payoffs ($z = 1$) for 43.50% of the participants in this treatment. These shares are remarkably similar to ones observed in the Hypothetical treatment, where the modal choice is $z = 0.5$ for 49.19% of the participants, and $z = 1$ for 41.46%.

[18]For simplicity, we do not consider subjects with no unique mode for this analysis.

Turning to behavior on positively sloped budget lines, own payoff maximization ($z = 1$) is the modal choice of the majority of individuals, while the share of people predominantly implementing payoff equality ($z = 0.5$) is much smaller. For example, in the Low-Incentives treatment, 73.14% of the participants predominantly choose the allocation that is own-payoff maximizing ($z = 1$), while only 23.30% predominantly pick the payoff-equalizing allocation. The distribution of modal choices in the High-Incentives treatment is very similar: 70.56% predominantly pick the own-payoff maximizing allocation, while only 25.57% predominantly pick $z = 0.5$. Likewise, in the Hypothetical treatment, we find that 69.01% mainly choose the own-payoff maximizing allocation, and 27.32% the payoff-equalizing allocation.[19]

These first results suggest that monetary incentives do *not* affect subjects' willingness to pay to increase and decrease others' payoffs. This conclusion is largely confirmed by a series of Kolmogorov-Smirnov tests for pairwise comparisons of distributions of modal choices.[20,21]

While we explored the distribution of modal choices separately for positively and negatively sloped budget lines in Figure 3, it is also instructive to examine their joint distribution. In Figure 4, we depict subjects' modal choice on *both* positively sloped (x-axis) *and* negatively sloped (y-axis) budget in each treatment—where each dot represents a subject. The Figure reveals the existence of *the same three distinct behavioral agglomerations across all three treatments*:

(i) The first behavioral agglomeration is located at $z = 0.5$ for both positively and negatively sloped budget lines. In this behavioral agglomeration, individuals tend to predominantly choose the payoff-equalizing allocation, both on positively and on negatively sloped budget lines. This behavioral pattern, which is suggestive of a preference for equality, characterizes roughly 23% of the subjects in the Low-incentives treatment, 25% in the High-Incentives treatment, and 27% of the individuals in the Hypothetical treatment.

(ii) The second behavioral agglomeration is located at $z = 1$ for positively sloped budget lines and $z = 0.5$ for negatively sloped budget lines. These individuals are characterized by a tendency to choose the payoff-equalizing allocation on negatively sloped bud-

[19] In Appendix E.2, Figure E.1 plots the cumulative distributions and confirms the similarity of modal choices across treatments, both for negatively and positively sloped budget lines.

[20] For these tests, we apply Holm (1979) correction to account for multiple hypothesis testing.

[21] Kolmogorov-Smirnov test p-values for *negatively* sloped budget lines: Low-Incentives vs. High-Incentives ($p = 0.289$), Low-Incentives vs. Hypothetical ($p = 0.045$), High-Incentives vs. Hypothetical ($p = 0.990$). Kolmogorov-Smirnov test p-values for *positively* sloped budget lines: Low-Incentives vs. High-Incentives ($p = 1.000$), Low-Incentives vs. Hypothetical ($p = 1.000$), High-Incentives vs. Hypothetical ($p = 0.996$).

Figure 3: Distribution of modal choices

Low-Incentives treatment

(a) Negatively sloped budget lines

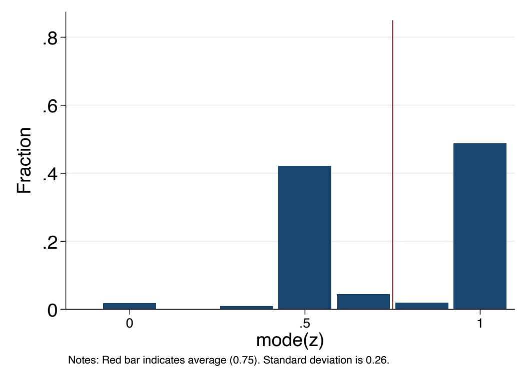

(b) Positively sloped budget lines

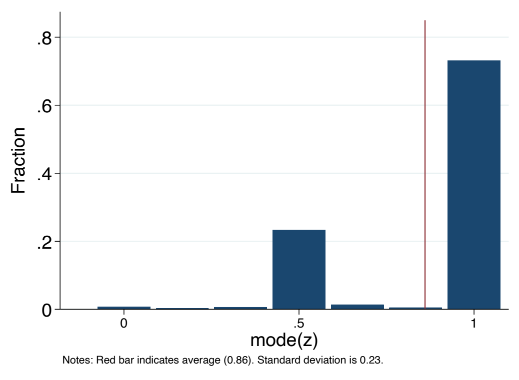

High-Incentives treatment

(c) Negatively sloped budget lines

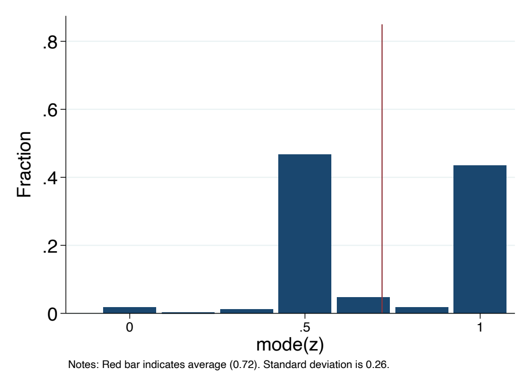

(d) Positively sloped budget lines

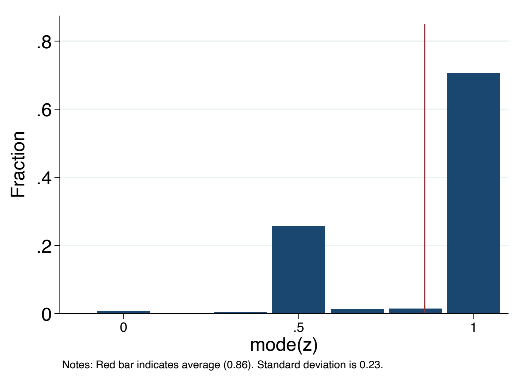

Hypothetical treatment

(e) Negatively sloped budget lines

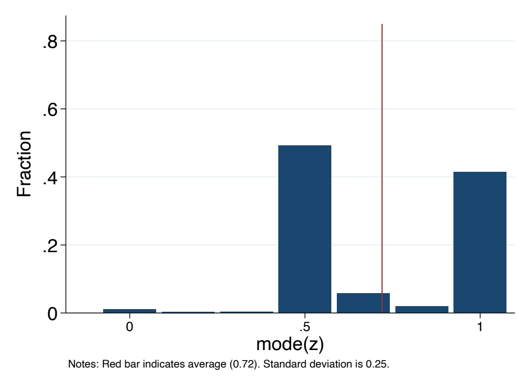

(f) Positively sloped budget lines

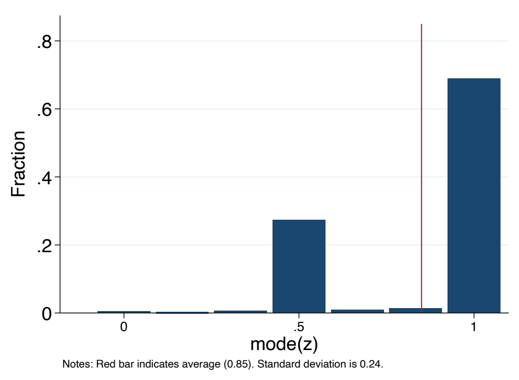

Notes: The figure shows the distribution of individuals' modal choices among negatively sloped and among positively sloped budget lines. For each budget line, $z = 1$ indicates an own-payoff maximizing choice, $z = 0$ indicates an own-payoff minimizing choice, $z = 0.5$ indicates a payoff-equalizing choice. The red vertical line indicates always the average over all modal choices. Panels (a) and (b) are constructed using subjects randomly assigned to the *Low-Incentives* treatment. Panels (c) and (d) are constructed using subjects randomly assigned to the *High-Incentives* treatment. Panels (e) and (f) are constructed using subjects randomly assigned to the *Hypothetical* treatment.

get lines, while at the same time predominantly choosing the own-payoff maximizing allocation on positively sloped budget lines—suggesting that these individuals have altruistic concerns for the worse off but tend to be unwilling to reduce the payoff of others for the sake of equality. This agglomeration comprises roughly 18% of the individuals in the Low-Incentives treatment, 20% in the High-Incentives, and 22% of subjects in the Hypothetical treatment.

(iii) The third behavioral agglomeration is located at $z = 1$ for both positively and negatively sloped budget lines. These participants tend to predominantly maximize their own payoffs, with little considerations for the payoffs of others. These predominantly selfish subjects comprise roughly 49% of the subjects in the Low-incentives treatment, 44% in the High-Incentives treatment, and 42% of the individuals in the Hypothetical treatment.

Altogether, the results suggests that individuals' choices are strikingly similar at the descriptive level, irrespective of whether subjects' choices in the money allocation task are incentivized, and irrespective of the size of the financial stakes.[22] While informative, this descriptive approach also has limitations. For example, it only considers subjects' modal choices and it ignores subjects with non-unique modes. Moreover, this analysis only *suggests* the existence of behavioral types, but it does *not* unambiguously assign all individuals to types. For these reasons, we turn to a more rigorous approach in the next section.

4.2 Cluster analysis: Do monetary incentives affect the distribution of social preferences?

Do incentives affect the distribution of social preference types in the population? To answer this question, we now turn to a more rigorous analysis that assigns subjects to types on the basis of their choices in the money allocation task. To that end, we apply a Bayesian nonparametric approach—the Dirichlet Process (DP) means clustering algorithm (Kulis and Jordan, 2012). This algorithm groups individuals into clusters according to their *behavioral*

[22]In Figure E.2 in Appendix E.2, we also display the average behavior across the different treatments separately for each decision situation (i.e., for each budget line). Like in our analysis of subjects' modal choices, this additional analysis does not reveal any meaningful treatment differences, neither on the main twelve budget lines from the center bundle, nor on the eight additional budget lines from the displaced bundles. These results are confirmed by a series of χ^2 tests, which cannot reject the null hypothesis that the allocation (z) is independent of the treatment for the vast majority of choice situations (see Table E.1).

Figure 4: Descriptive evidence on subjects' modal choices

(a) *Low-Incentives* treatment

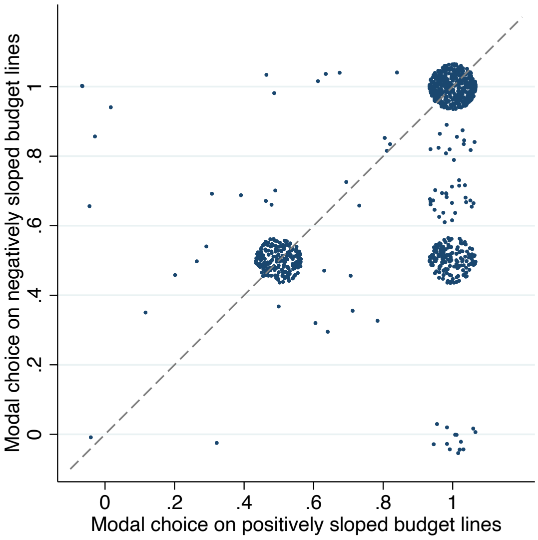

(b) *High-Incentives* treatment

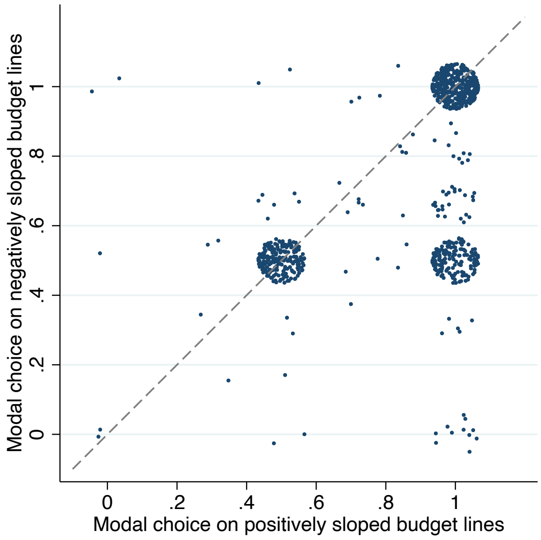

(c) *Hypothetical* treatment

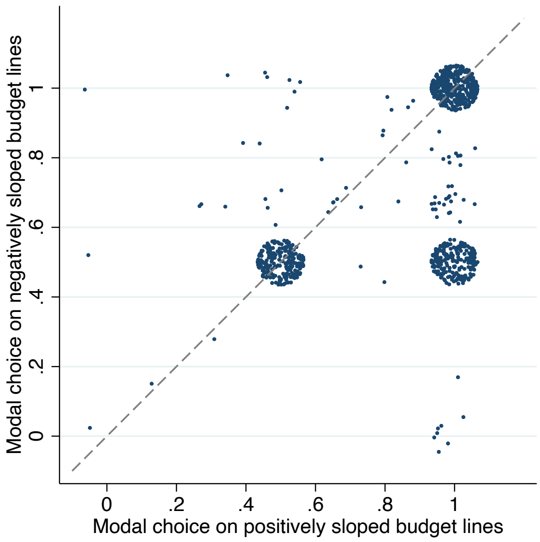

Notes: In all figures, we depict subjects' modal choices among negatively sloped budget lines and among positively sloped budget lines. Each dot represents one individual. Dots are jittered in order to make identical modal choices of individuals visible. For each budget line, $z = 1$ indicates an own-payoff maximizing choice, $z = 0$ indicates an own-payoff minimizing choice, and $z = 0.5$ indicates a payoff-equalizing choice. Panel (a) is constructed using subjects randomly assigned to the *Low-Incentives* treatment. Panel (b) is constructed using subjects randomly assigned to the *High-Incentives* treatment. Panel (c) is constructed using subjects randomly assigned to the *Hypothetical* treatment. Note that if we replace individuals' modal choices by their median choices, very similar behavioral agglomerations emerge.

similarities. In our context, clusters are based on subjects' twelve distributional choices (Figure 1), and similarity is measured by "how close" an individual's allocation profile is to the average allocation of a cluster. Ultimately, individuals' are assigned to the cluster whose centroid—i.e., the mean allocation in the twelve distributional choices—is the closest to their own allocation profile in the twelve-dimensional space of interest. We describe the formalism of the DP-means algorithm in greater details in Appendix F.

An important aspect of the DP-means approach is that it enables the identification of preference types without committing to a pre-specified number of different preference types. Moreover, this approach does neither require an ex-ante specification or parameterization of types, nor does it presume a specific error structure. This means that it remains ex-ante agnostic about key distributional assumptions, and it does not constrain heterogeneity to lie within a predetermined set of models or parameter space.[23] Moreover, the DP-means algorithm allows for all possible type partitions of the data spanning from a representative agent (i.e. a single data-generating process) up to as many types as there are individuals in the population (i.e. n data-generating processes).

We run the DP-means algorithm separately on each treatment. We display the distribution of clusters identified by the DPM in the Table 1 below. Consistent with the results of the descriptive analysis, we find that three clusters emerge in each experimental group. The distribution of these clusters is remarkably similar across treatments, with each cluster comprising roughly one third of the population. This result is confirmed by a χ^2 test, which cannot reject the null hypothesis that the assignment to clusters is independent of the treatments (p = 0.255).

Table 1: Type distributions identified using clustering analysis

	Low-Incentives	**High-Incentives**	**Hypothetical**
Cluster 1: Inequality averse	31.74%	35.09%	32.87%
Cluster 2: Altruistic	30.34%	30.89%	32.67%
Cluster 3: Predominantly selfish	37.92%	34.02%	34.46%

Notes: The table displays the distribution of individuals to the three clusters (in percent) that emerge in our dataset, separately for each treatment. The behavioral interpretation of the clusters (indicated in the left column) is based on the interpretation of each cluster's typical behavior provided in Figure 5.

Importantly, the DP-means algorithm does *not* assign labels to clusters, i.e., it is agnostic

[23] In this regard, our approach differs from previous work (e.g. Bellemare et al., 2008; Fisman et al., 2015, 2017; Bruhin et al., 2018) that characterized preference heterogeneity on the basis of structural assumptions on preferences and error terms.

as to the behavioral interpretation of types. In order to assign a label to the clusters, one needs to inspect the behavioral characteristics of each cluster identified. To that end, we explore the distribution of choices that characterizes each cluster in each treatment in Figure 5.

The Figure reveals that subjects assigned to Cluster 1 mainly select payoff-equalizing allocations both on negatively sloped and positively sloped budget lines. They thus exhibit a willingness to pay to reduce inequality, both when this involves increasing the payoff of those worse off (negatively sloped budget lines) and decreasing the payoff of those better off (positively sloped budget lines). This behavior is consistent with models of inequality aversion (Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000) and we therefore assign the label *"inequality averse"* to this cluster.

Individuals assigned to Cluster 2 display a strikingly different form of other-regarding behavior. They are also willing to pay in order to increase the payoff of those worse off (negatively sloped budget lines), but they are generally unwilling to pay to decrease the payoff of those better off (positively sloped budget lines). They therefore display a tendency to forego own payoff when this benefits the poor, but this willingness to pay vanishes almost entirely when equalizing payoffs would entail reducing income of those who are better off. This behavioral pattern is consistent with altruistic concerns for the worse off (Charness and Rabin, 2002) and with altruistic other-regarding behavior that incorporates an equity-efficiency tradeoff (Fisman et al., 2007, 2015). We therefore label this behavioral cluster as *"altruistic"*.

The remaining subjects, assigned to Cluster 3, are characterized by a making predominantly own-payoff maximizing choices, irrespective of the choice situation. For this reason, we label these individuals as being *"predominantly selfish"*.

Together, this clustering analysis reveals two striking findings. First, the behavioral interpretation of the three clusters is consistent across the three treatments. In all our experimental condition, the same three clusters with a similar behavioral interpretation coexist: inequality averse, altruistic, and predominantly selfish subjects.[24] Second, irrespective of the presence of monetary incentives and irrespective of the stake sizes, the distribution of types is remarkably stable, i.e., the relative share of types remains roughly the same across experimental condition. This finding is noteworthy, given that the DP-means algorithm imposes no prior assumptions regarding the behavioral interpretation of types nor their distribution in the population.

Overall, these results suggest that whether or not monetary incentives are used to elicit

[24]Importantly, note that the fact that the three clusters each have a clear behavioral interpretation rules out the possibility that the DP-means algorithm identified arbitrary or random types.

social preferences does *not* matter for the identification of *qualitatively* distinct preference types.

Figure 5: The distribution of choices for positively and negatively sloped budget lines in each cluster and each treatment

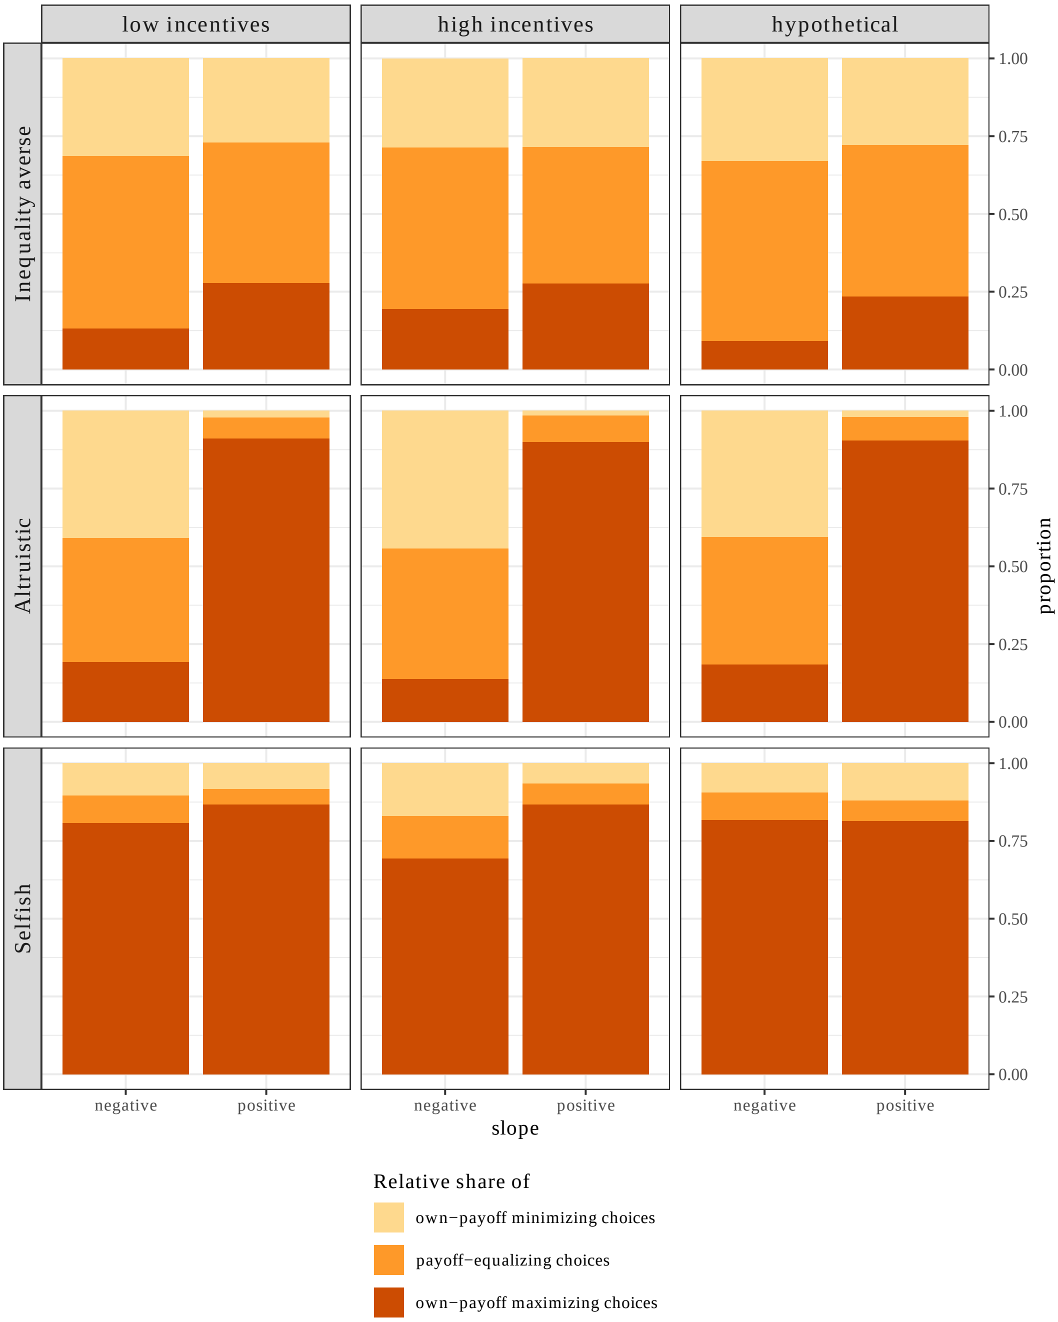

Notes: The figure shows the distribution of choices for positively and negatively sloped budget lines in each cluster and each treatment. Subjects could choose among seven different allocations. A choice is classified as own-payoff minimizing (own-payoff maximizing) if it belongs to the two choices that give the subject the lowest (highest) payoff. It is classified as payoff-equalizing if it implements perfect equality or one of its nearest neighboring allocations.

4.3 Structural Analysis: Do monetary incentives affect the strength of social preferences and the precision of their estimates?

The previous sections revealed that monetary incentives do not affect subjects aggregate willingness to pay to increase or decrease the payoff of others (descriptive analysis), nor the distribution of preference types in the population (clustering analysis). But do they affect the *strength* of individuals' social preferences and the *precision* of their estimates? It is possible, for example, that low monetary incentives (or absence thereof) leads subjects to take decision that make them appear more other-regarding than they are if the monetary stakes were high–which would lead one to overestimate social preferences. Likewise, the absence of incentives might increase the share of participants answering randomly, thereby decreasing the precision of social preference estimates.

To address these questions, we estimate the parameters of a hierarchical Bayesian model of social preferences (Fehr and Schmidt, 1999) *separately* for subjects in different treatments. Specifically, we structurally estimate the following model in each treatment:

$$V_i\left(w_{ij}\right) = w_{ij}^s - \alpha_i \max\left\{w_{ij}^o - w_{ij}^s, 0\right\} - \beta_i \max\left\{w_{ij}^s - w_{ij}^o, 0\right\}$$

where $w_{ij} = \left(w_{ij}^s, w_{ij}^o\right)$ corresponds to individual i's decision on budget line j on how to allocate money between herself (superscript s for self) and the other person (superscript o for other), α_i denotes aversion towards disadvantageous inequality (behindness aversion) and β_i denotes aversion towards advantageous inequality (aheadness aversion).[25]

4.3.1 The effects of monetary incentives on the strength of social preferences

We depict the mean values and standard deviations of the estimated structural parameters, separately by treatment, in Table 2. The average level of aversion to disadvantageous inequality (α) is 0.219 in the Low-Incentives treatment and 0.278 in the High-Incentives treatments, i.e., stake size does not significantly affect average levels of behindness aversion ($p = 0.106$). In contrast, we estimate a considerably higher average α-parameter of 0.368 for individuals in the Hypothetical treatment, i.e., aversion to disadvantageous inequality in the Hypothetical treatment is on average 68.0% higher than under low stakes ($p < 0.01$), and 32.4% higher than

[25]Note that in the absence of restrictions on the alpha and better parameters, the inequality aversion model of Fehr and Schmidt (1999) is equivalent to the two person case in Charness and Rabin (2002). Thus, the model can capture altruistic, envious, or inequality averse preferences.

under high stakes (p = 0.024). These degrees of aversion to disadvantageous inequality are quantitatively sizeable. For example, a subject with an α of 0.278 (High-Incentives) is willing to pay up to 21.75 cents to *decrease* the payoff of those ahead by one dollar. In contrast, the estimates from the Hypothetical treatment imply a maximum willingness to pay of up to 26.9 cents, and of up to 18 cents in the Low-Incentives treatment.[26]

The results are somewhat similar when we turn to aversion to advantageous inequality. There too, the highest average β is in the Hypothetical treatment (0.744), which is 21.4% higher than in the Low-Incentives treatment (0.613, test of the difference: $p < 0.01$) and about 4.5% higher than in the High-Incentives treatment (0.712, test of the difference: $p = 0.487$). In terms of magnitude, a subject with a β of 0.712 (High-Incentives) is willing to pay up to 2.47 dollars to increase the payoff of someone worse off by one dollar, as opposed to a willingness to pay of 1.58 in the Low-Incentives treatment and of 2.90 in the Hypothetical treatment.[27]

Table 2: Summary statistics across treatment

	Low-Incentives		High-Incentives		Hypothetical	
	Mean	S.D.	Mean	S.D.	Mean	S.D.
α	0.219	0.775	0.278	0.843	0.368	0.956
β	0.613	0.965	0.712	0.961	0.744	1.124

Do the distributions of α and β differ accross treatments? We depict their probability density function in Figure 6, and a series of pairwise Kolmogorov-Smirnov (KS) tests of equality in distributions in Table 3.[28]

Consistent with our discussion above, we find that the distributions of α and β-parameters are relatively similar in the Low and High-Incentives treatment, but that the Hypothetical treatment significantly shifts both distributions to the right.[29] Indeed, while stake sizes do not significantly affect the distribution of individuals' estimated α-parameters (Low vs. High-Incentives: p = 0.106), we can reject the null hypotheses that individuals in the Hypothetical treatment have similar distributions than those in the two incentivized treatments (Low-Incentives vs. Hypothetical: p = 0.001; High-Incentives vs. Hypothetical: p = 0.032).

Turning to the distributions of β, we again most strongly reject the null hypothesis of equality in distribution when comparing the Hypothetical with the Low-Incentives treatments

[26] Formally, the willingness to pay to reduce the payoff of someone ahead by one unit is $\alpha/(1+\alpha)$.
[27] The willingness to pay to increase the payoff of someone behind by one unit is $\beta/(1-\beta)$.
[28] For these tests, we apply Holm (1979) correction to account for multiple hypothesis testing.
[29] An analysis of cumulative distribution functions, which we relegate to Figure E.3 in Appendix E, confirms that social preference estimates are larger in the Hypothetical treatment.

(p = 0.001), and the Hypothetical with the High-Incentives treatment (p = 0.001). Interestingly, we can also reject the null hypothesis of equality in distributions between the Low and the High-Incentives treatments, although the significance level of this test is lower (p = 0.034).

These results point towards a common conclusion: while the size of the monetary stakes does not play a large role for the strength of social preferences, whether or not monetary incentives are used matters. Specifically, we find that the social preference parameters are consistently larger—i.e., possibly overestimated—in absence of monetary incentives.

Figure 6: Distribution of structurally estimated parameters

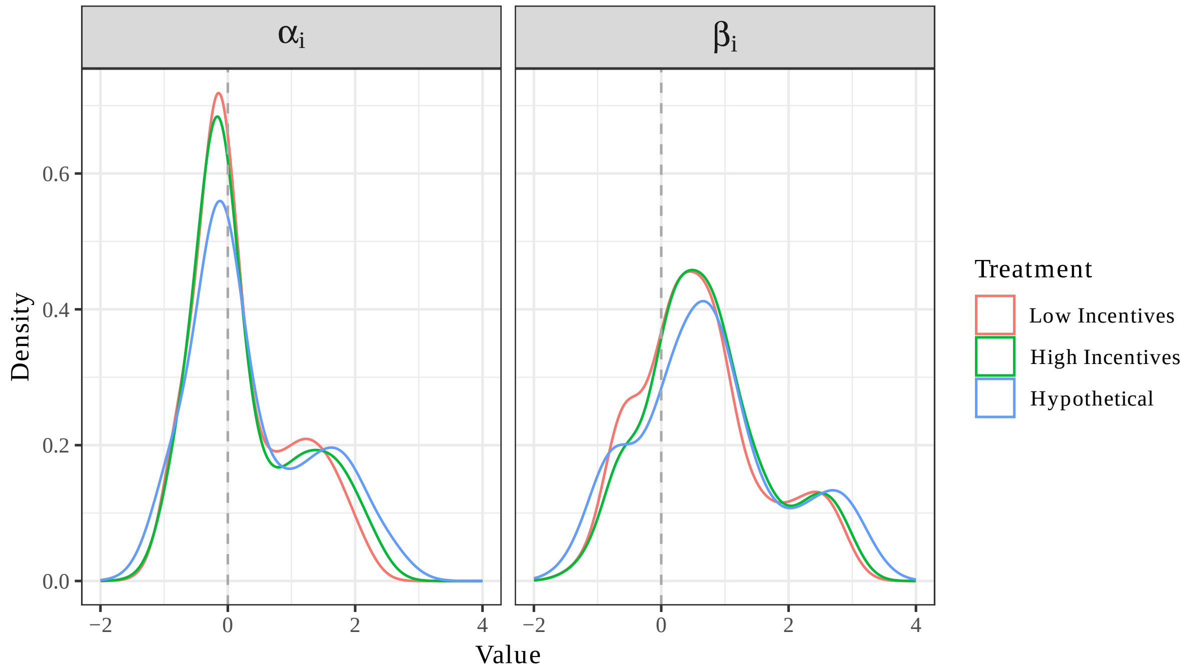

Notes: The left panel depicts subjects' structurally estimated α-parameters by treatment. The right panel depicts subjects' structurally estimated β-parameters by treatment.

Table 3: Kolmogorov-Smirnov Test p-values with Holm (1979) correction

Comparison	α-parameter p-value	β-parameter p-value
Low-Incentives vs. High-Incentives	0.106	0.034
Low-Incentives vs. Hypothetical	0.001	0.001
High-Incentives vs. Hypothetical	0.032	0.001

How do the different preference types identified in Section 4.2 relate to structural parameters of α and β? In Figure E.4 of the Appendix E.3, we shed light on this question by examining the distributions of these parameters separately by preference type.

This analysis reveals that the treatment differences highlighted above are mainly driven by the inequality averse subjects (Cluster 1). Inequality averse subjects display a strong and significant rightward shift in the distribution of social preferences parameters in the Hypothetical treatment. In other words, inequality averse subjects respond to the lack of monetary incentives with *larger* values of α and β. For the other two preference types, we do not find such strong treatment differences.

4.3.2 The effects of monetary incentives on the precision of structural estimates

To assess whether monetary incentives improve the reliability of our preference estimates, we use the parameter estimates from the structural hierarchical Bayesian models of each treatment and extract, for every individual, the posterior standard deviation of the two behavioral parameters α_i and β_i. These within-person dispersions quantify the precision with which the model recovers each individual's preferences: the narrower the posterior, the more informative the data. Comparing the median posterior standard deviation across incentive conditions therefore provides a direct test of whether incentivized tasks yield tighter, higher-quality parameter estimates than non-incentivized ones, independent of the heterogeneity in preferences between individuals. Table 4 presents the results. The precision of the estimates improves substantially when moving from hypothetical to incentivized measures of social preferences. Hence, using financial incentives increases the precision with which social preference parameters are estimated (compared to hypothetical choices).

Table 4: Median posterior standard deviation of individuals' preference parameters by treatment condition

Treatment	α	β
Low-Incentives	0.238	0.221
High-Incentives	0.224	0.196
Hypothetical	0.310	0.290

Overall, this structural analysis reveals an important lesson that was not identified in the descriptive nor in the clustering analyses. With hypothetical stakes, the estimated *strength* of social preferences is larger than with monetary incentives–suggesting that they are overestimated. This result is primarily driven by individuals identified as being inequality averse in the clustering analysis. Moreover, the *precision* with which individuals' preference parameters are estimated is lower under hypothetical stakes. In contrast, the magnitude and the precision of the parameters estimated in the Low and High-Incentives treatments is similar.

4.4 Do monetary incentives affect the measurement of students' social preferences?

While our main interest is to understand whether and how monetary incentives affect measures of social preferences in the general population, it is also instructive to investigate how they affect students. To shed light on this issue, we reproduce the analysis above, focusing on our sample of 495 students (for details on the sample, see Section 2.3). Because this analysis is not our main focus, we only summarize it here and relegate the details to the Appendix G.[30] Moreover, note that this analysis is less well-powered than our analysis for the general population. As such, it should be interpreted with greater caution.

At the descriptive level, we find that the distributions of students' modal choices are essentially unaffected by the treatments: The majority of students predominantly select either the payoff-equalizing ($z = 0.5$) or the own-payoff maximizing allocations ($z = 1$), irrespective of whether their decisions are incentivized (and by how much). Moreover, all the six Kolmogorov-Smirnov tests for pairwise comparisons reject the null hypothesis of equality in distributions between condition.

Turning to the clustering analysis, we again find that heterogeneity in preferences is best captured by three types—an inequality averse, an altruist, and a predominantly selfish type–in the Low-Incentives and in the Hypothetical treatments. However, the DP-means algorithm does *not* identify a stable three-type clustering in the High-Incentives treatment.[31] Interestingly, for those treatments for which we identify three clusters, the assignment of students to types is quite responsive to incentives: in the hypothetical treatment, 83.55% of the subjects are assigned to a one of the other-regarding clusters (i.e., either altruistic or inequality averse), whereas only 55.07% are assigned to such a cluster when money is at stake (Low-Incentives treatment). In other words, students are more likely to be assigned to the selfish type when monetary incentives are used.

The impact of incentives on the structurally estimated preference parameters is less clear cut for students compared to the general population. While monetary incentives reduce the strength of social preferences in the general population, the results for the student sample

[30] Note also that it is not the aim of this paper to provide an extensive analysis of the differences in the social preferences of students and the general population. We address this specific research question in a separate paper (Epper ⓡ al., 2023).

[31] The algorithm's failure to detect distinct types might mean they are absent in this treatment/sample, or that heterogeneity takes a different form here; this is speculative, and larger or more diverse student samples could clarify the matter.

are more mixed. Specifically, while we also find that the distribution of estimated aversion to advantageous and disadvantageous inequality is higher under Hypothetical incentives than under Low-Incentives, we find that it is the highest under High-Incentives. This suggests that stronger monetary incentives yield stronger social preferences in the student sample. These results should be taken with a grain of salt, however, as the precision with which these parameters are estimated is substantially lower than in the general population.[32]

Summarizing, the results from the students sample are broadly consistent with those of the general population. While we identified a few inconsistencies between these two different samples, we cannot rule out that they are due to the smaller sample size (and associated greater estimation uncertainty) in the student sample.

5 Concluding remarks

Using a pre-registered online experiment with a representative sample of the US population, we examined whether the use of monetary incentives and variations in stake size matter for the identification of social preferences. We explored this question in three steps: at the descriptive level of subjects' model choices, at the level of qualitatively distinct preferences types, and at the level of individuals' structurally estimated social preference parameters.

Our results show that the use of monetary incentives, as well as the size of the stakes, have little impact on choices at the descriptive levels, as well as for the identification of *qualitatively* distinct preferences types. They appear to matter, however, for the *quantitative* identification of the strength and the precision of social preferences. In particular, our structural analysis reveals that the social preferences of the general population are likely overestimated when elicited with hypothetical stakes.

These results suggest that whether or not incentives should be used to elicit social preferences depends on the specifics of the research question at hand. If one is solely interested in having a rough, descriptive measure of social preferences at the aggregate level, or if one wants to identify qualitatively distinct preferences types, then relying on hypothetical stakes might suffice. However, if one is interested in making a quantitative assessment of subjects' other-regardingness, e.g., in order to make quantitative predictions, then our result suggest that using monetary incentives is advisable, as hypothetical stakes appear to yield an overes-

[32]For example, the median posterior standard deviation of α and β in the High-Incentives treatment is 0.224 and 0.196 in the general population, while it is 0.441 and 0.319 in the student sample.

timation of inequality aversion. However, we find no evidence that using a larger stake size improves the identification of social preferences, consistent with two meta-analyses showing that larger stakes have little to no effect on decisions in dictator games (Engel, 2011; Larney et al., 2019).

References

Andre, Peter, Carlo Pizzinelli, Christopher Roth, and Johannes Wohlfart (2022) "Subjective models of the macroeconomy: Evidence from experts and representative samples," *The Review of Economic Studies*, 89 (6), 2958–2991.

Bellemare, Charles, Sabine Kröger, and Arthur Van Soest (2008) "Measuring inequity aversion in a heterogeneous population using experimental decisions and subjective probabilities," *Econometrica*, 76 (4), 815–839.

Bellemare, Charles and Bruce Shearer (2007) "Gift exchange within a firm: Evidence from a field experiment."

Berinsky, Adam J, Michele F Margolis, and Michael W Sances (2014) "Separating the shirkers from the workers? Making sure respondents pay attention on self-administered surveys," *American Journal of Political Science*, 58 (3), 739–753.

Bierbrauer, Felix and Nick Netzer (2016) "Mechanism design and intentions," *Journal of Economic Theory*, 163, 557–603.

Bierbrauer, Felix, Axel Ockenfels, Andreas Pollak, and Désirée Rückert (2017) "Robust mechanism design and social preferences," *Journal of Public Economics*, 149, 59–80.

Bolton, Gary E and Axel Ockenfels (2000) "ERC: A theory of equity, reciprocity, and competition," *American Economic Review*, 90 (1), 166–193.

Bruhin, Adrian, Ernst Fehr, and Daniel Schunk (2018) "The many faces of human sociality: Uncovering the distribution and stability of social preferences," *Journal of the European Economic Association*.

Bühren, Christoph and Thorben C Kundt (2015) "Imagine being a nice guy: A note on hypothetical vs. incentivized social preferences," *Judgment and Decision Making*, 10 (2), 185–190.

Camerer, Colin F (2011) *Behavioral game theory: Experiments in strategic interaction*: Princeton University Press.

Camerer, Colin F and George Loewenstein (1993) "Information, fairness, and efficiency in bargaining."

Camerer, Colin and Richard H Thaler (1995) "Anomalies: Ultimatums, dictators and manners," *Journal of Economic Perspectives*, 9 (2), 209–219.

Charness, Gary (2000) "Responsibility and effort in an experimental labor market," *Journal of Economic Behavior & Organization*, 42 (3), 375–384.

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

Clot, Sophie, Gilles Grolleau, and Lisette Ibanez (2018) "Shall we pay all? An experimental test of random incentivized systems," *Journal of behavioral and experimental economics*, 73, 93–98.

Comiter, Marcus Zachary, Miriam Cha, HT Kung, and Surat Teerapittayanon (2016) "Lambda means clustering: automatic parameter search and distributed computing implementation."

Dechezleprêtre, Antoine, Adrien Fabre, Tobias Kruse, Bluebery Planterose, Ana Sanchez Chico, and Stefanie Stantcheva (2022) "Fighting climate change: International attitudes toward climate policies,"Technical report, National Bureau of Economic Research.

Dur, Robert (2009) "Gift exchange in the workplace: Money or attention?" *Journal of the European Economic Association*, 7 (2-3), 550–560.

Engel, Christoph (2011) "Dictator games: A meta study," *Experimental economics*, 14, 583–610.

Epper, Thomas ⓡ Ernst Fehr ⓡ Julien Senn (2023) "Social preferences across subject pools: students vs. general population," *URPP Equality of Opportunity Discussion Paper Series* (46).

Falk, Armin, Peter Andre, Teodora Boneva, and Felix Chopra (2021) "Fighting climate change: The role of norms, preferences, and moral values."

Falk, Armin, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman, and Uwe Sunde (2018) "Global evidence on economic preferences," *The quarterly journal of economics*, 133 (4), 1645–1692.

Falk, Armin, Anke Becker, Thomas Dohmen, David Huffman, and Uwe Sunde (2023) "The preference survey module: A validated instrument for measuring risk, time, and social preferences," *Management Science*, 69 (4), 1935–1950.

Fehr, Ernst ⓡ Thomas Epper ⓡ Julien Senn (2023) "The Fundamental Properties, the Stability, and the Predictive Power of Distributional Preferences."

——— (forthcoming) "Social Preferences and Redistributive Politics," *Review of Economics and Statistics*.

Fehr, Ernst, Georg Kirchsteiger, and Arno Riedl (1993) "Does fairness prevent market clearing? An experimental investigation," *The quarterly journal of economics*, 108 (2), 437–459.

Fehr, Ernst, Michael Powell, and Tom Wilkening (2021) "Behavioral constraints on the design of subgame-perfect implementation mechanisms," *American Economic Review*, 111 (4), 1055–1091.

Fehr, Ernst and K Schmidt (1999) "A Theory Of Fairness, Competition, and Cooperation," *Quarterly Journal of Economics*, 114 (3), 817–868, http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?db=pubmed&cmd=Retrieve&dopt=AbstractPlus&list_uids=9017717075410430622related:nirQ9OZdJX0J.

Fisman, Raymond, Pamela Jakiela, and Shachar Kariv (2017) "Distributional preferences and political behavior," *Journal of Public Economics*, 155, 1–10.

Fisman, Raymond, Pamela Jakiela, Shachar Kariv, and Daniel Markovits (2015) "The distributional preferences of an elite," *Science*, 349 (6254), aab0096.

Fisman, Raymond, Shachar Kariv, and Daniel Markovits (2007) "Individual preferences for giving," *American Economic Review*, 97 (5), 1858–1876.

Fitzgerald, Jack (2024) "Identifying the impact of hypothetical stakes on experimental outcomes and treatment effects,"Technical report, Tinbergen Institute Discussion Paper.

Forsythe, Robert, Joel L Horowitz, Nathan E Savin, and Martin Sefton (1994) "Fairness in simple bargaining experiments," *Games and Economic behavior*, 6 (3), 347–369.

Gallup (2024) "Party Affiliation — Gallup Historical Trends," https://news.gallup.com/poll/15370/party-affiliation.aspx.

Henkel, Aljosha ⓡ Ernst Fehr ⓡ Julien Senn ⓡ Thomas Epper (2025) "Beliefs about inequality and the nature of support for redistribution," *Journal of Public Economics*, 246, 105350.

Hertwig, Ralph and Andreas Ortmann (2001) "Experimental practices in economics: A methodological challenge for psychologists?" *Behavioral and brain sciences*, 24 (3), 383–403.

Holm, Sture (1979) "A simple sequentially rejective multiple test procedure," *Scandinavian journal of statistics*, 65–70.

Hufe, Paul and Daniel Weishaar (2025) "Just cheap talk? Investigating fairness preferences in hypothetical scenarios."

Kerschbamer, Rudolf and D Müller (2020) "Social preferences and political attitudes: An online experiment on a large heterogeneous sample," *Journal of Public Economics*, 182.

Kosfeld, Michael, Zahra Sharafi, Maíra Sontag González, and Na Zou (2025) "Measuring Economic Preferences with Surveys and Behavioral Experiments."

Kube, Sebastian, Michel André Maréchal, and Clemens Puppe (2012) "The currency of reciprocity: Gift exchange in the workplace," *American Economic Review*, 102 (4), 1644–1662.

Kulis, Brian and Michael I. Jordan (2012) "Revisiting k-means: New Algorithms via Bayesian Nonparametrics," *Proceedings of the 29th International Conference of Machine Learning*.

Larney, Andrea, Amanda Rotella, and Pat Barclay (2019) "Stake size effects in ultimatum game and dictator game offers: A meta-analysis," *Organizational Behavior and Human Decision Processes*, 151, 61–72.

Miguel, Edward (2021) "Evidence on research transparency in economics," *Journal of Economic Perspectives*, 35 (3), 193–214.

Munafò, Marcus R, Brian A Nosek, Dorothy VM Bishop et al. (2017) "A manifesto for reproducible science," *Nature human behaviour*, 1 (1), 1–9.

Nosek, Brian A, Charles R Ebersole, Alexander C DeHaven, and David T Mellor (2018) "The preregistration revolution," *Proceedings of the National Academy of Sciences*, 115 (11), 2600–2606.

Plott, Charles R (1986) "Rational choice in experimental markets," *Journal of Business*, S301–S327.

Ruggles, Steven, Sarah Flood, Matthew Sobek et al. (2024) "IPUMS USA: Version 15.0 [dataset]," `https://www.ipums.org/projects/ipums-usa/d010.V15.0`.

Schmidt, Klaus M and Axel Ockenfels (2021) "Focusing climate negotiations on a uniform common commitment can promote cooperation," *Proceedings of the National Academy of Sciences*, 118 (11), e2013070118.

Smith, Vernon L (1982) "Microeconomic systems as an experimental science," *The American economic review*, 72 (5), 923–955.

——— (1991) "Rational choice: The contrast between economics and psychology," *Journal of Political Economy*, 99 (4), 877–897.

Stantcheva, Stefanie (2021) "Understanding tax policy: How do people reason?" *The Quarterly Journal of Economics*, 136 (4), 2309–2369.

Tyran, Jean-Robert and Rupert Sausgruber (2006) "A little fairness may induce a lot of redistribution in democracy," *European Economic Review*, 50 (2), 469–485.

U.S. Bureau of Labor Statistics (2025) "70.1 Percent of Workers Were Full Time, Year Round in 2023," The Economics Daily (online), `https://www.bls.gov/opub/ted/2025/70-1-pergcenInstituteDiscussionPapert-of-workers-were-full-time-year-round-in-2023.htm`.

A Statements and Declarations

DISCLOSURE STATEMENT

With reference to the submission:

"*Do monetary incentives matter for identifying social preferences?*"

Registered Report (Proposal) submitted to the Special Issue of Experimental Economics on Incentivization

co-authored with Aljosha Henkel, Thomas Epper, and Julien Senn.

I declare that:

(1) None of the authors have any relevant, material or financial interests that relate to the research described in this paper.

(2) None of the authors have held any position in organizations that relate to this research.

(3) No party outside of the authors has had the right to review the manuscript prior to submission.

I accept that disclosure statements will be made available upon publication.

Zurich, 10.04.2024

Ernst Fehr

B Background information on the experimental task

B.1 Choice situations in the money allocation task

B.1.1 Center bundle

Table B.1 provides further details on these choice situations. The meaning of the list of variables displayed in the Table is as follows:

- 'choiceId': the unique identifier for each choice situation.
- $(own1, other1)$: represents the payoff combination at the lower end of the budget line (in points).
- $(own2, other2)$: represents the payoff combination at the upper end of the budget line (in points).
- 'bundle': indicates to which bundle the respective choice situation belongs to.
- 'slope': the slope of the budget line in the "own payoff – other payoff" space.

Table B.1: Choice situations in the money allocation task

choiceId	own1	own2	other1	other2	slope
1	450	1050	750	750	0.0
2	500	1000	800	700	-0.2
3	550	950	850	650	-0.5
4	600	900	900	600	-1.0
5	650	850	950	550	-2.0
6	700	800	1000	500	-5.0
7	750	750	1050	450	-Inf
8	700	800	500	1000	5.0
9	650	850	550	950	2.0
10	600	900	600	900	1.0
11	550	950	650	850	0.5
12	500	1000	700	800	0.2

B.1.2 Displaced bundles

Following Fehr ⓡ al. (2023, forthcoming), our design also includes additional choice situations that we do *not* use for type identification. Four of these choice situations lie *above* the 45 degree line (depicted in red) and thus inform us further on the decision maker's willingess to pay to *decrease* the other's payoff. The remaining four choice situations lie *below* the 45 degree line (depicted in blue) and therefore provide us further information on the decision maker's willingness to pay to *increase* the other's payoff. We include these additional budget lines as they can help us validate the behavioral interpretation of the types identified using the center bundle. They can also be used to fine-tune the structural estimation of a model of inequality aversion.

Figure B.1: Additional budget lines

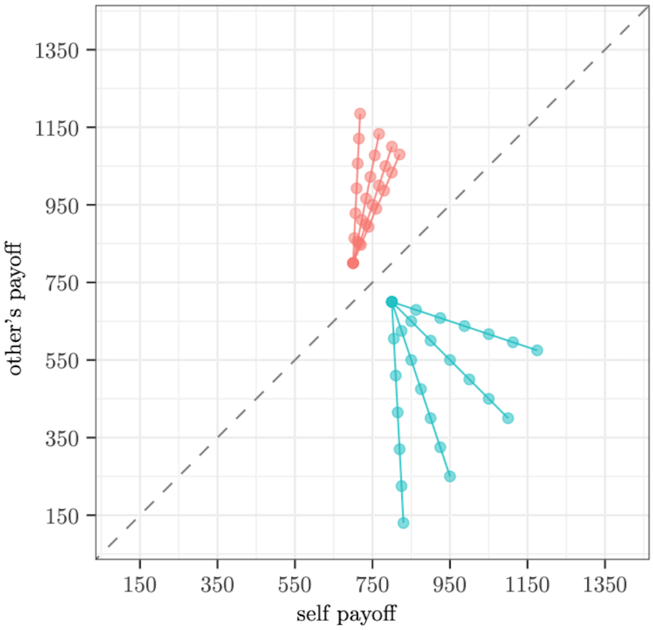

B.2 Instructions

In the following, we reproduce the instructions of the money allocation task for the *Low-Incentives* and the *Hypothetical* treatments. Note that all the instructions were displayed directly on participants' computer screens.

[Social preference task – Low Incentives]

[Instructions]

We now proceed with a task in which you have to take decisions on how to allocate points between yourself and another participant of the study.

In what follows, we describe the instructions for this task. Please read them carefully.

What will you have to do in the following task?

You will be asked to take decisions in different choice situations. In each of these choice situations, you will have to decide how to allocate points between yourself and another participant.

Who is the other participant?

The other participant will take part in another part of the study. Anonymity between yourself and the other participant is guaranteed, i.e. that neither you nor the other person will ever learn about each other's identity.

Moreover, the other participant will not take decisions that affect you, i.e. you will not be affected in any way by the decisions of the other participant.

What will be the consequences of your decisions?

The points gathered during this study will be **converted into US dollars** at the following exchange rate

500 points = $ 1

At the end of this study, the computer will randomly select one of the choice situations **and pay you according to your decision in that choice situation.** This decision-dependent payment will be added to your fixed payment of $3. The other participant will also be paid according to **your** decision in that choice situation.

Since every choice situation has an equal chance of being drawn for payment, it is important that you think carefully about each decision.

What kind of decisions will you have to take?

In each choice situation you will be asked to allocate points between yourself and another participant of the study. You will always have the choice between seven different alternatives, numbered from 1 to 7. Each alternative consists of a distribution of points between you and the other participant.

Example

The figure below illustrates a typical choice situation as it will appear on your screen.

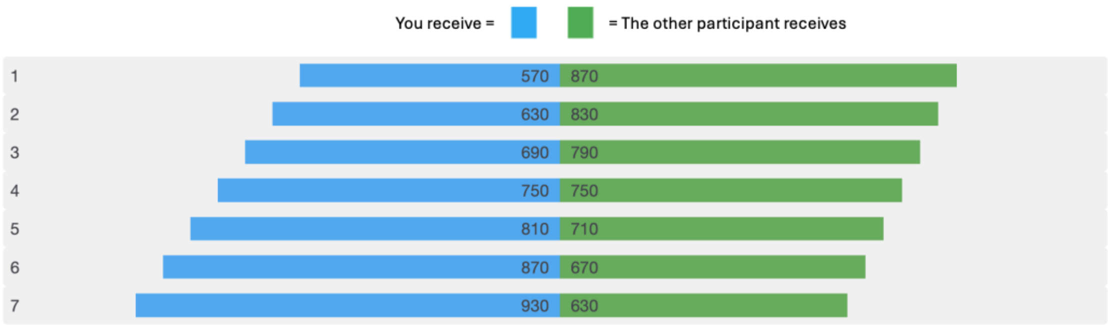

In this example,

- choosing alternative 1 yields you 570 points and the other participant 870 points.
- choosing alternative 7 yields you 930 points and the other participant 630 points.
- the total amount of points to be distributed varies from one alternative to another.

How do you make a choice?

You make your choice by clicking on your preferred alternative. You can change your choice as many times as you want. Once you press the 'Next' button at the bottom right of the screen, your choice is validated and can no longer be reverted. Directly after you press "Next", the next choice situation will appear on the screen. This will be repeated until all the choices have been made.

[Control questions]

Before you start with the task, we would like to make sure that you understand what is asked from you in this task, and what the consequences from your choices are.

To show us that you understand the task, please answer the comprehension questions below. Participants who do not correctly respond to these questions will not be allowed to proceed with the study.

Consider the following example.

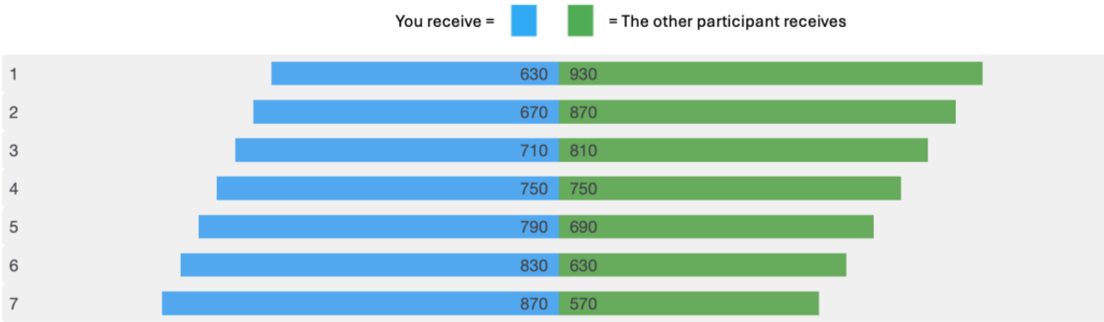

1. How many points do you get if you chose alternative 5? [790]

2. How many points does the other participant obtain if you chose alternative 6? [630]

3. What is the total number of points that you and the other participant receive together if you chose alternative 3? Is it 710, 810, or 1520 points? [1520]

4. Do your choices have real monetary consequences for you and the other participant? [yes, no]

[Success control questions]

You have successfully answered all the control questions. You will now start with the decision task. As of now, your decisions matter for your payment, and for the payment of another participant.

Please think carefully before taking a decision in each choice situation.

[Social preference task: decision screens]

Please choose your preferred alternative.

[DISPLAY CHOICE SITUATIONS]

[Social preference task – Hypothetical]

[Instructions]

We now proceed with a task in which you have to take decisions on how to allocate points between yourself and another participant.

In what follows, we describe the instructions for this task. Please read them carefully.

What will you have to do in the following task?

You will be asked to take decisions in different **hypothetical** choice situations. In each of these choice situations, you will have to decide how to allocate points between yourself and another hypothetical participant.

Who is the other participant?

Imagine that you are paired with a hypothetical participant that participates in another part of the study, and that anonymity between yourself and the other participant is guaranteed, i.e. that neither you nor the other person will ever learn about each other's identity.

Moreover, imagine that the other (hypothetical) participant will not take decisions that affect you, i.e. that you will not be affected in any way by the decisions of the other participant.

What will be the consequences of your decisions?

Your choices will have **no real monetary consequences** for you nor the other participant, but please imagine that you are allocating points that have monetary value between yourself and the other participant.

Imagine, in particular, that the points gathered during this study are converted into US dollars at the following exchange rate:

500 points = $ 1.

Thus, although your choices will have no real monetary consequences for you nor the other participant, please make your choices as if you and the other participant were paid accordingly.

What kind of decisions will you have to take?

In each choice situation you will be asked to allocate points between yourself and another participant of the study. You will always have the choice between seven different alternatives, numbered from 1 to 7. Each alternative consists of a distribution of points between you and the other participant.

Example

The figure below illustrates a typical choice situation as it will appear on your screen.

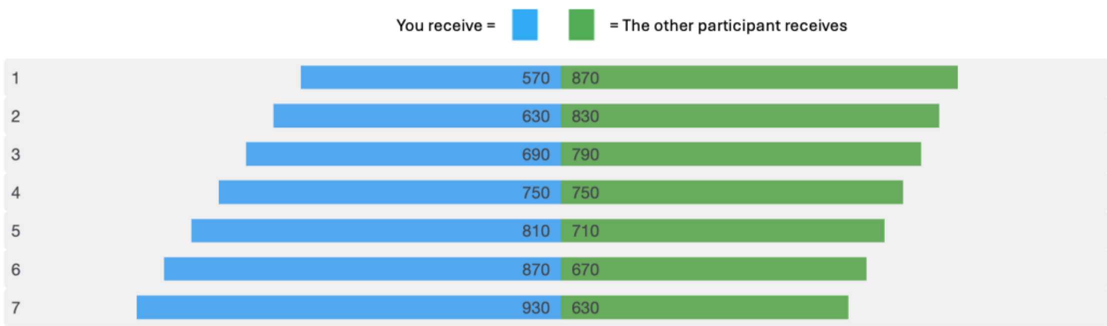

In this example,

- choosing alternative 1 yields you 570 points and the other participant 870 points.
- choosing alternative 7 yields you 930 points and the other participant 630 points.
- the total amount of points to be distributed varies from one alternative to another.

How do you make a choice?

You make your choice by clicking on your preferred alternative. You can change your choice as many times as you want. Once you press the 'Next' button at the bottom right of the screen, your choice is validated and can no longer be reverted. Directly after you press "Next", the next choice situation will appear on the screen. This will be repeated until all the choices have been made.

[Control questions]

Before you start with the task, we would like to make sure that you understand what is asked from you in this task, and what the consequences from your choices are.

To show us that you understand the task, please answer the comprehension questions below. Participants who do not correctly respond to these questions will not be allowed to proceed with the study.

Consider the following example.

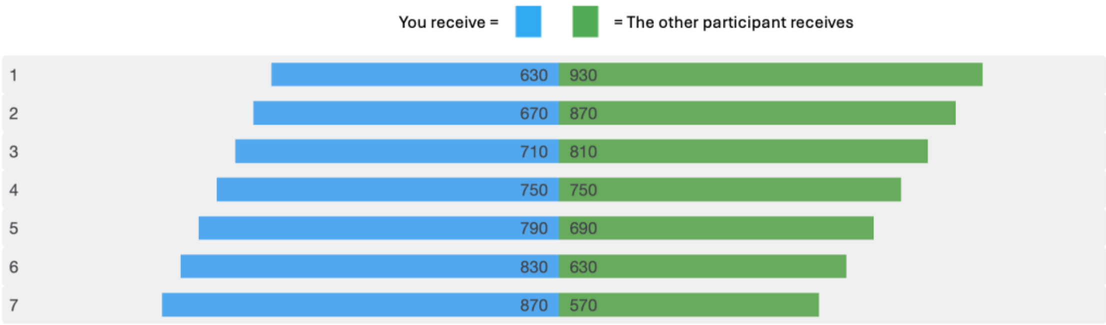

1. How many points do you get if you chose alternative 5? [790]

2. How many points does the other participant obtain if you chose alternative 6? [630]

3. What is the total number of points that you and the other participant receive together if you chose alternative 3? Is it 710, 810, or 1520 points? [1520]

4. Do your choices have real monetary consequences for you and the other participant? [yes, no]

[Success control questions]

You have successfully answered all the control questions. You will now start with the decision task.

Please think carefully before taking a decision in each choice situation.

[Social preference task: decision screens]

Please choose your preferred alternative.

[DISPLAY CHOICE SITUATIONS]

C Demographic characteristics of sample population

We depict the main descriptive statistics in Table C.1, separately for the Low-Incentives, High-Incentives, and Hypothetical treatment. The last column indicates that our treatment is generally well balanced across the main observable characteristics. The table also indicates that our sample is broadly representative of the US population with respect to age, gender, and political affiliation.

Table C.1: Descriptive statistics and balance checks

	Population	Low-Incentives	High-Incentives	Hypothetical	p-value (F-test)
Age (mean)	48.2	46.7	46.7	45.4	.918
18-25 y.o.	0.13	0.12	0.14	0.13	.098
26-35 y.o.	0.17	0.18	0.20	0.15	.001
36-45 y.o.	0.17	0.16	0.16	0.18	.185
46-55 y.o.	0.15	0.17	0.17	0.18	.733
56-65 y.o.	0.16	0.24	0.22	0.24	.385
> 65 y.o.	0.21	0.12	0.11	0.12	.268
Male	0.49	0.47	0.49	0.48	.999
Political spectrum (1 left - 10 right)	-	5.10	5.25	4.99	.186
Leaning towards Republican party	0.28	0.34	0.34	0.31	.663
Leaning towards Democratic party	0.28	0.41	0.38	0.43	.820
Income bracket : ≤ $ 20k	0.08	0.09	0.10	0.10	.263
Income bracket : $ 20-40k	0.11	0.15	0.16	0.15	.298
Income bracket : $ 40-60k	0.12	0.16	0.17	0.17	.614
Income bracket : $ 60-80k	0.11	0.16	0.13	0.16	.016
Income bracket : $ 80-100k	0.10	0.12	0.11	0.10	.124
Income bracket : $ 100-150k	0.19	0.20	0.19	0.18	.245
Income bracket : ≥ $ 150k	0.25	0.10	0.13	0.12	.005
Income bracket : NA	0.03	0.02	0.01	0.01	.000
Have been unemployed in the past	-	0.74	0.76	0.79	.105
Occupation: Full-time worker	0.44	0.43	0.41	0.43	.984
Occupation: Part-time worker	0.19	0.17	0.17	0.15	.079
Occupation: Student	0.08	0.14	0.15	0.15	.373
Occupation: Pensioner	0.14	0.08	0.07	0.08	.052
Occupation: Unemployed	0.03	0.11	0.11	0.14	.001
Occupation: Other	0.12	0.07	0.08	0.05	.000
Education: High school	0.23	0.19	0.23	0.19	.022
Education: Technical college	0.24	0.23	0.21	0.22	.656
Education: Undergraduate degree	0.29	0.35	0.35	0.36	.958
Education: Graduate or doctorate degree	0.13	0.19	0.18	0.19	.619
Education: Other	0.10	0.04	0.03	0.04	.000
Religiosity (0 - 10)	-	4.63	4.60	4.35	.826
Observations		1002	1023	1007	

Note: The table displays descriptive statistics of the US population and of our sample, separately for the Low-Incentives, High-Incentives, and Hypothetical treatment. The descriptive statistics include age (mean), the shares of people falling into each age bracket, and the share of male people. Moreover, they include subjects' political leaning (mean), the shares of people leaning towards the Republican and Democratic parties, and the shares of people falling into each yearly household income bracket. In addition, they include the share of people that have been unemployed in the past, the shares of people falling into each occupation category, as well as the shares of shares of people falling into each highest educational degree category. Lastly, the descriptive statistics include subjects' religiosity (mean). Statistics of the US population were obtained from IPUMS data (Ruggles et al., 2024) from the American Community Survey 2023 and are restricted to the adult US population (i.e., individuals who are at least 18 years old). The proportions of workers who work full-time (part-time) among the working population stem from U.S. Bureau of Labor Statistics (2025). Political affiliation comes from Gallup (2024).

D Power analysis

Prior to the start of the study, we assessed sample size requirements for our main analysis with the general population sample by considering the statistical power of our main hypotheses tests. Specifically, we looked at a two-sided Welch t-test on the difference between inequality aversion parameters α or β between two treatments (e.g., Hypothetical and Low-Incentives). We set the probability of a Type I error to 1%, and the standard deviations of the parameters in the two samples to the values we obtained from a previous (incentivized) study conducted in Switzerland using a similar experimental design. The sample size depicted on the x-axis refers to the per-group sample size (e.g. the Hypothetical treatment). We computed the power of the test under the assumption that the number of participants is identical in both treatments. Thus, a specific number X on the x-axis means that both the hypothetical and the incentivized (e.g. Low-Incentives) treatments contain X participants. The blue curve shows, for different per-group sample sizes, the power we have to detect a difference of 0.25 (using the above assumptions) between the parameters of the two treatments. With a sample size of 1,000 participants per group, we obtain a power which is slightly higher than 80 percent to detect a difference in the parameters α or β of 0.25, i.e., an effect size of about 15 percent of a standard deviation of the structural parameters in the Swiss broad population sample. The green curve depicts the power curve for a larger effect size of 0.35, and the red curve depicts the power curve for an effect size of 0.45.

Figure D.1: Power vs. per-group sample size for various effect sizes

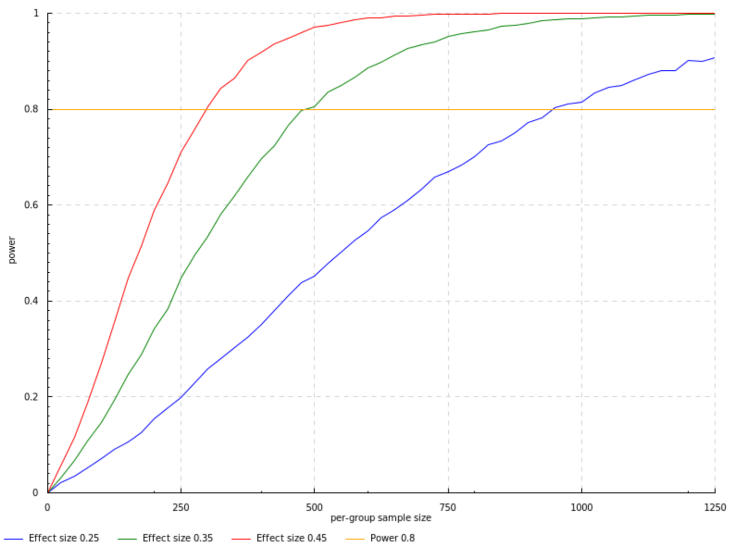

Note: The figure depicts power curves for various effect sizes. The horizontal line indicates a power of 80%.

E Additional tables and figures

E.1 Additional analyses

The analyses and results presented in Section 4 closely follow the analysis plan that we pre-registered. In some instances, we also provide additional analyses that were not pre-registered but that are useful to further clarify our findings. For transparency, we list all these additional analyses below:

- Figure E.1: "Cumulative distribution of modal choices"
- Figure E.2: "Average choices"
- Table E.1: "χ^2-squared test of independence by choice situation (Holm (1979) corrected)"
- Figure 5: "The distribution of choices for positively and negatively sloped budget lines in each cluster and each treatment"
- Figure E.3: "CDFs of structurally estimated parameters"
- Figure E.4: "CDFs of structurally estimated parameters by social preference type"

E.2 Descriptive analysis

Figure E.1 plots the cumulative distributions of modal choices, separately for negatively sloped budget lines (Figure E.1a) and positively sloped budget lines (Figure E.1b). The Figure confirms the similarity of modal choices across treatments for both types of budget lines.

Figure E.1: Cumulative distribution of modal choices

(a) Negatively sloped budget lines (b) Positively sloped budget lines

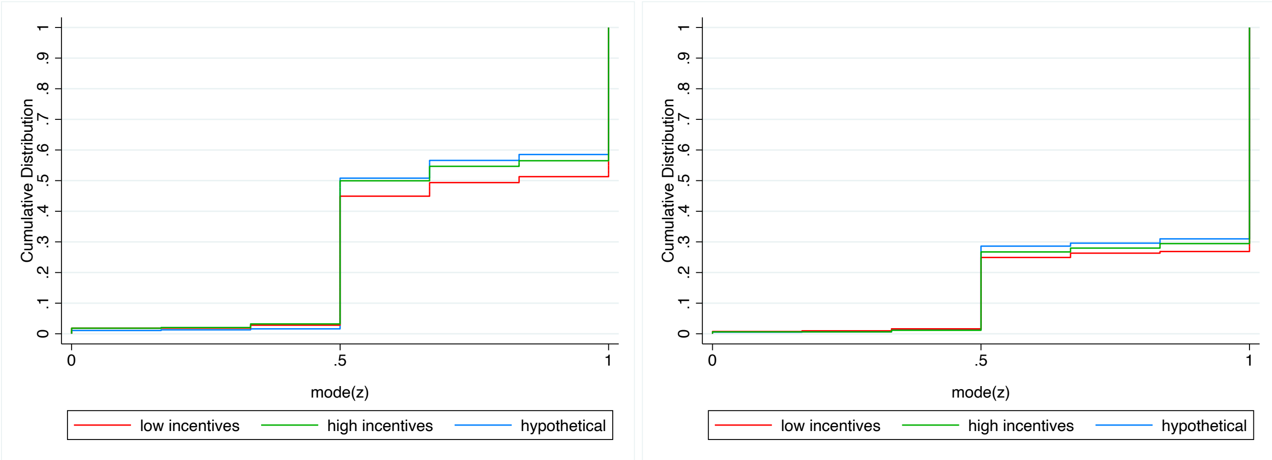

Notes: The figure shows the cumulative distribution of individuals' modal choices among negatively sloped and among positively sloped budget lines. For each budget line, $z = 1$ indicates an own-payoff maximizing choice, $z = 0$ indicates an own-payoff minimizing choice, $z = 0.5$ indicates a payoff-equalizing choice.

Figure E.2 displays the average behavior across the different treatments, separately for each budget line (x-axis) and bundle of budget lines (panels). In line with the results for the modal choices, we do not find any meaningful differences in average behavior across treatments, neither for the twelve budget lines that are centered around the 45-degree line nor for the budget lines that are fully in the disadvantageous domain (”north bundle”) or the budget lines that are fully in the advantageous domain (”south bundle”).[33] Table E.1 shows the χ^2-squared test of independence by choice situation, corrected for multiple testing (Holm, 1979).[34] The table further supports the conclusion that average behavior is similar across treatments in each choice situation.

[33]For details on the budget lines, see Appendix B.1.

[34]In other words, the table provides the results of a series of tests (Holm (1979) corrected) of equality in distributions of implemented choices z (where z ranges from 0 to 1) across the three treatments, by decision situation (choice ID).

Figure E.2: Average choices

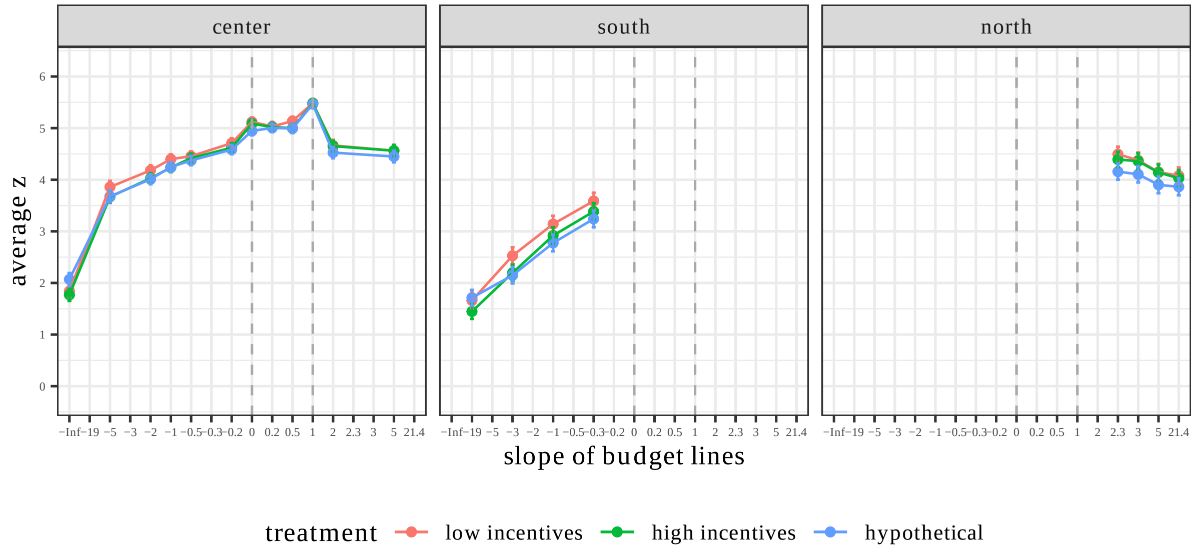

Notes: The figure shows the average implemented choice (z) on the y-axis, by budget line (ordered by their slopes) and by bundle of choice situations (3 panel). For each budget line, $z = 6$ indicates an own-payoff maximizing choice, $z = 0$ indicates an own-payoff minimizing choice, $z = 3$ indicates a payoff-equalizing choice. Budget lines are sorted by slope on the x-axis.

Table E.1: χ^2-squared test of independence by choice situation (Holm (1979) corrected)

	choiceId	statistic	p-value
1	1	33.701	0.020
2	2	17.469	1.000
3	3	12.390	1.000
4	4	18.452	1.000
5	5	14.208	1.000
6	6	18.342	1.000
7	7	25.198	0.330
8	8	9.077	1.000
9	9	22.015	0.520
10	10	6.113	1.000
11	11	21.908	0.850
12	12	27.474	0.100
13	13	17.523	1.000
14	14	26.380	0.190
15	15	10.807	1.000
16	16	21.617	0.660
17	17	31.368	0.040
18	18	13.772	1.000
19	19	15.658	1.000
20	20	30.969	0.050

Notes: This table shows the χ^2-squared test of independence by choice situation, corrected for multiple testing (Holm, 1979). It provides the results of a series of tests of equality in distributions of implemented choices z (where z ranges from 0 to 1) across the three treatments, by decision situation (choice ID).

E.3 Structural analysis

Figure E.3 depicts the cumulative distribution functions (CDFs) of α and β by treatment. Consistent with the PDFs depicted in the main text, this analysis reveals that the distributions under Low and High-Incentives are relatively similar, but that distributions in the Hypothetical treatment are shifted to the right.

Figure E.3: CDFs of structurally estimated parameters

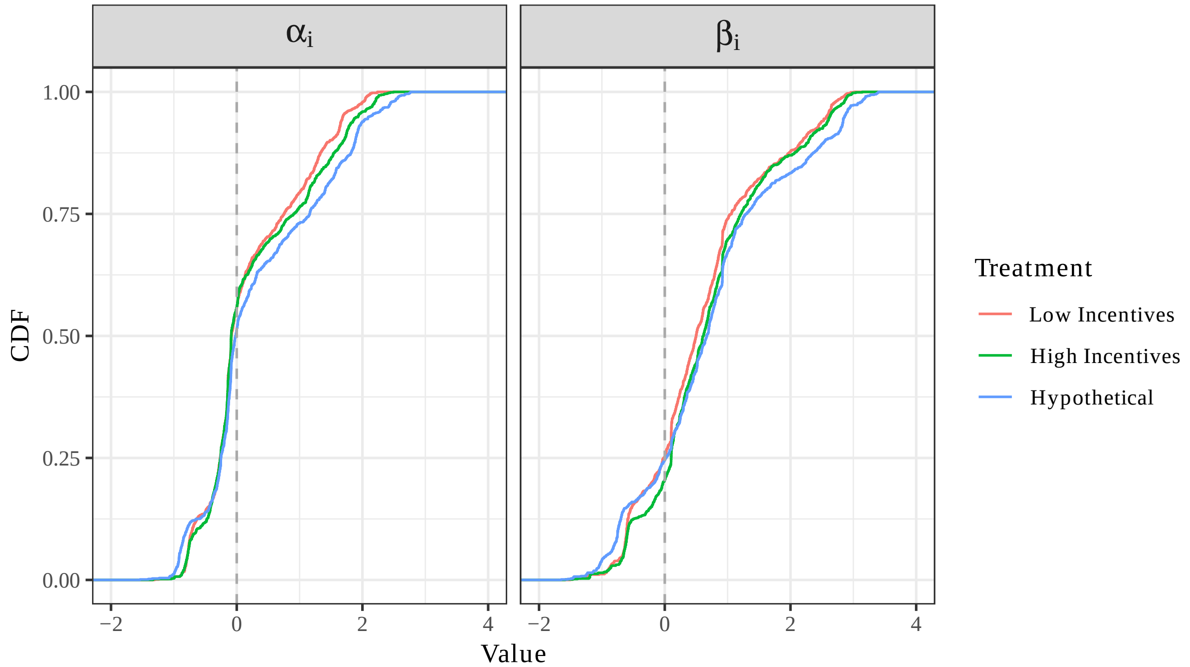

Notes: The CDFs on the left side depict subjects' structurally estimated α-parameters by treatment. The CDFs on the right side depict subjects' structurally estimated β-parameters by treatment.

We break down these distributions by social preference type (as identified using the clustering algorithm) in Figure E.4. This analysis shows that treatment differences are mainly driven by the inequality-averse subjects (Type 1; upper panel). It is for these subjects the distributions of social preferences parameters shifts the most to the right in the Hypothetical treatment. For the other two types, we do not find such strong treatment differences.

Figure E.4: CDFs of structurally estimated parameters by social preference type

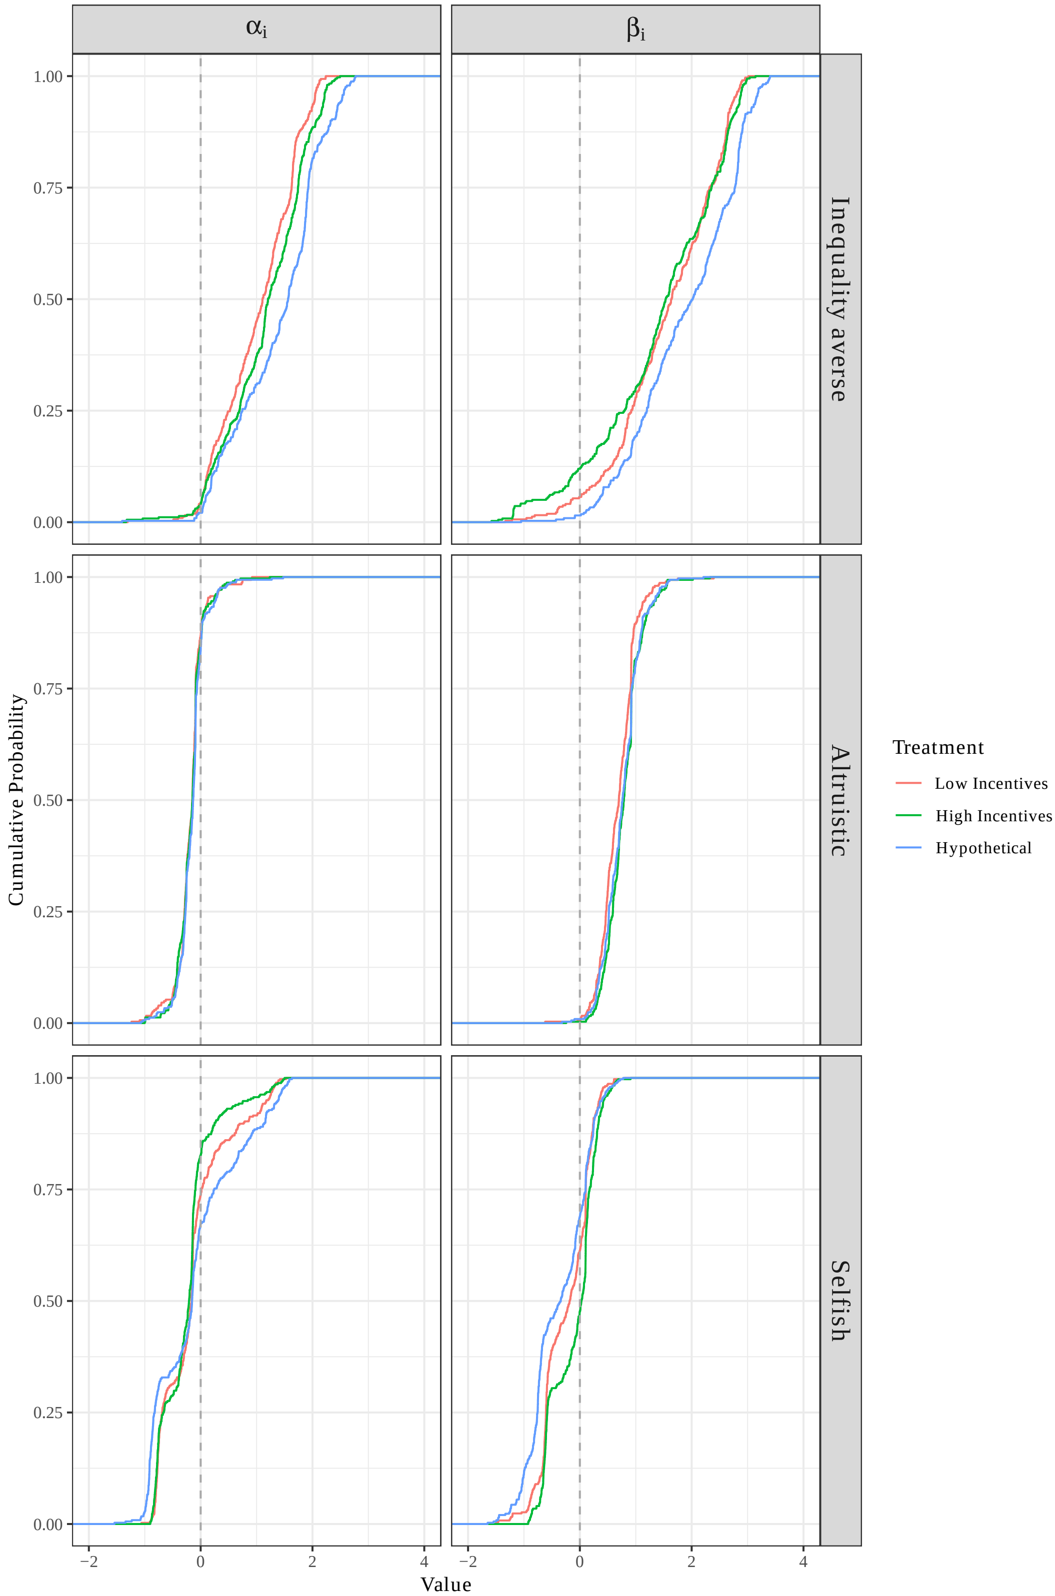

Notes: The CDFs on the left side depict subjects' structurally estimated α-parameters by treatment, separately for each social preference type. The CDFs on the right side depict subjects' structurally estimated β-parameters by treatment, separately for each social preference type.

F Identifying preference types using Dirichlet Process Means

F.1 The method

This appendix provides an overview of the clustering algorithm used to identify the preference types and their distribution in the population. For a more detailed description of the DP-means algorithm and for a discussion of its key differences with other clustering methods such a k-means, see Fehr ⓡ al. (forthcoming, 2023).

Our implementation of the algorithm is based on an iterative refinement. We first span an m-dimensional space, with m denoting the number of budget lines used for the clustering algorithm (in our case, $m = 12$, the twelve budget lines presented in Table 1 in the main paper). Consequently, each individual's choices are represented by a single point in the 12-dimensional space. We then ask how subjects populate this space. Specifically, we are interested in the number of clusters (i.e. types) that emerge and individuals' assignment to clusters. A cluster is characterized by the set of the individuals assigned to the cluster and the associated mean vector of observations (the "centroid"), which – in our case – represents the mean (cluster- representative) behavior of all individuals in m-dimensional space that belong to the cluster.

We initialize the algorithm with a single centroid specified as the global mean vector. At this stage, all data points are assigned to this single centroid. We then refine by iterating over the following two steps: First, we sequentially go through the list of data points in m-dimensional space (i.e. subjects), and check for each subject whether any of the squared Euclidean distances to the centroid exceeds the cluster penalty parameter λ. If this is the case, we open up a new cluster with the actual data point's location vector as the centroid. Otherwise, we assign the data point to its nearest cluster. Second, we collect the subjects assigned to the same clusters and update the centroids by computing the mean vector for each cluster. These two steps are repeated until convergence is reached, i.e. until there is no change in subjects' assignments.

As Kulis and Jordan (2012) demonstrate, this iterative procedure is equivalent to minimizing the objective

$$\min_{\{g_c\}_{c=1}^{k}} \sum_{c=1}^{k} \sum_{\boldsymbol{x}\in g_c} \|\boldsymbol{x}-\boldsymbol{\mu}_c\|^2+\lambda k\,,$$

where $\boldsymbol{x}$ denotes the vector of observations, $\boldsymbol{\mu}$ the vector of centroids, and g the cluster partitioning of $\boldsymbol{x}$. It is straightforward to see that this objective is equivalent to the k-means objective except for the additional penalty term λk.

An important aspect of the DP-means approach is that it enables the identification of preference types without committing to a pre-specified number of different preference types. Moreover, this approach does neither require an ex-ante specification or parameterization of types, nor does it presume a specific error structure. This means that it remains ex-ante agnostic about key distributional assumptions, and it does not constrain heterogeneity to lie within a predetermined set of models or parameter space.[35] The DP-means algorithm allows for all possible type partitions of the data spanning from a representative agent (i.e. a single data-generating process) up to as many types as there are individuals in the population (i.e. n data-generating processes), i.e., it determines the number of preferences types endogenously. Thus, (i) the actual number of types, (ii) the assignment of each individual to one of the types and (iii) the behavioral (preference) properties of the types emerge endogenously.[36]

[35]In this regard, our approach differs from previous work (e.g. Bellemare et al., 2008; Fisman et al., 2015, 2017; Bruhin et al., 2018) that characterized preference heterogeneity on the basis of structural assumptions on preferences and error terms.

[36]The fact that the number of types adapts to the data has important benefits (see Kulis and Jordan, 2012). Most notably, as previous work has shown (see Comiter et al., 2016), this feature of the algorithm yields higher quality type-separation than methods that specify the number of types prior to clustering (such as k-means).

G Student sample

In this Appendix, we present the results for the student sample.

G.1 Demographic characteristics of the student sample

We depict the main descriptive statistics of the student sample in Table G.1, separately for the Low-Incentives, High-Incentives, and Hypothetical groups. The last column indicates that our treatments are generally well balanced across the main observable characteristics.

Table G.1: Descriptive statistics and balance checks

	Low-Incentives	High-Incentives	Hypothetical	p-value (F-test)
Age (mean)	23.1	22.8	22.8	.649
18-25 y.o.	0.75	0.82	0.80	.246
26-35 y.o.	0.25	0.28	0.30	.246
Male	0.52	0.50	0.50	1.000
Political spectrum (1 left - 10 right)	5.46	5.21	5.18	.504
Leaning towards Republican party	0.35	0.34	0.36	.987
Leaning towards Democratic party	0.37	0.41	0.38	.977
Income bracket : ≤ $ 20k	0.16	0.20	0.10	.002
Income bracket : $ 20-40k	0.13	0.10	0.14	.125
Income bracket : $ 40-60k	0.18	0.18	0.24	.283
Income bracket : $ 60-80k	0.15	0.09	0.10	.004
Income bracket : $ 80-100k	0.11	0.10	0.11	.875
Income bracket : $ 100-150k	0.13	0.21	0.13	.011
Income bracket : ≥ $ 150k	0.13	0.10	0.14	.248
Income bracket : NA	0.02	0.02	0.04	.000
Have been unemployed in the past	0.77	0.77	0.74	.801
Education: Technical college	0.74	0.74	0.79	.610
Education: Undergraduate degree	0.26	0.26	0.21	.610
Religiosity (0 - 10)	5.47	5.10	5.27	.918
Observations	158	173	164	

Note: The table displays descriptive statistics of the student sample, separately for the Low-Incentives, High-Incentives, and Hypothetical group. The descriptive statistics include age (mean), the shares of people falling into each age bracket, and the share of male people. Moreover, they include subjects' political leaning (mean), the shares of people leaning towards the Republican and Democratic parties, and the shares of people falling into each yearly household income bracket. In addition, they include the share of people that have been unemployed in the past, as well as the shares of shares of people falling into each highest educational degree category. Lastly, the descriptive statistics include subjects' religiosity (mean).

G.2 Descriptive analysis of students

For the descriptive analysis, we examine whether the treatments affect the distribution of subjects' *modal* choice separately for negatively sloped and for positively sloped budget lines (Figure 1). We depict the results of this analysis in Figure G.1. The Figure reveals that, among the negatively sloped budget lines, the distributions of modal choices are comparable across treatments. While we do see slightly more own payoff maximizing choices in the Low-Incentives treatment, the overall pattern is similar across experimental conditions as the modal choice of the vast majority of individuals is located at either $z = 0.5$ or $z = 1$. Turning to behavior on positively sloped budget lines, own payoff maximization ($z = 1$) is the modal choice for the majority of the individuals, while the share of people predominantly implementing payoff equality ($z = 0.5$) is much smaller. This pattern is similar across treatments, with slightly more payoff-equalizing choices in the High-Incentives treatment. Figure G.2 plots the cumulative distributions and confirms the similarity of modal choices across treatments, both for negatively and positively sloped budget lines. Kolmogorov-Smirnov tests for pairwise comparisons of distributions of modal choices confirm the conclusion that monetary incentives do *not* affect subjects' willingness to pay to increase and decrease others' payoffs.[37,38]

[37] For these tests, we apply Holm (1979) correction to account for multiple hypothesis testing.

[38] Kolmogorov-Smirnov test p-values for *negatively* sloped budget lines: Low-Incentives vs. High-Incentives ($p = 0.253$), Low-Incentives vs. Hypothetical ($p = 0.488$), High-Incentives vs. Hypothetical ($p = 0.951$). Kolmogorov-Smirnov test p-values for *positively* sloped budget lines: Low-Incentives vs. High-Incentives ($p = 1.000$), Low-Incentives vs. Hypothetical ($p = 1.000$), High-Incentives vs. Hypothetical ($p = 0.974$).

Figure G.1: Distribution of modal choices of students

Low-Incentives treatment

(a) Negatively sloped budget lines

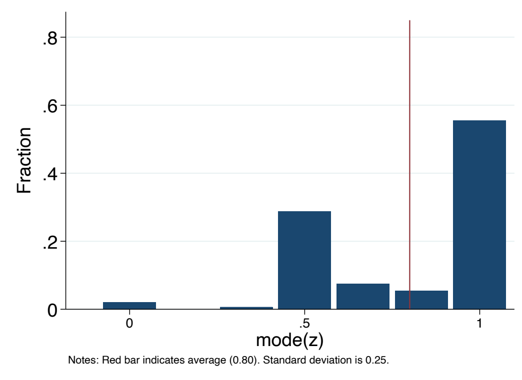

(b) Positively sloped budget lines

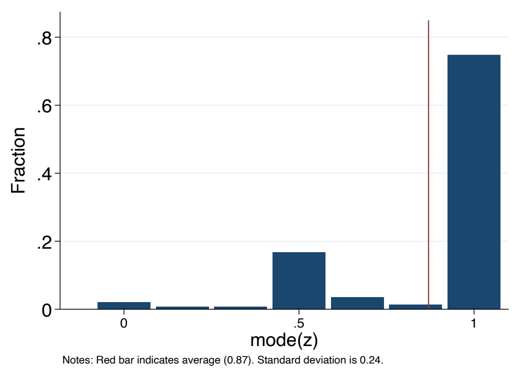

High-Incentives treatment

(c) Negatively sloped budget lines

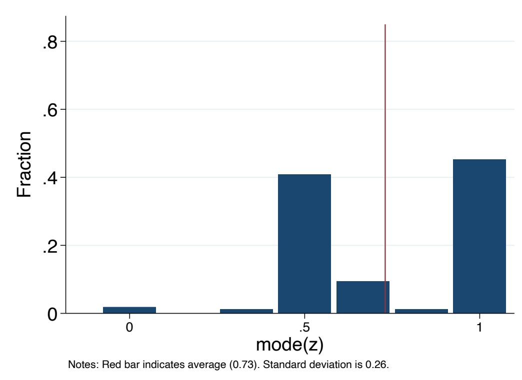

(d) Positively sloped budget lines

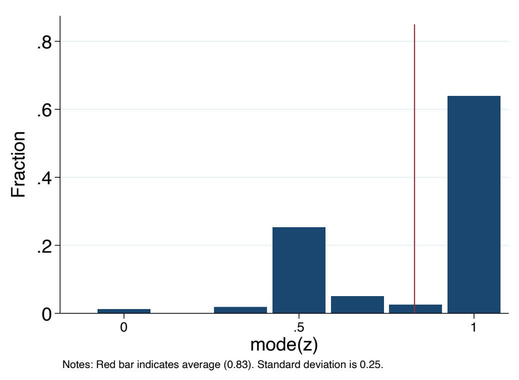

Hypothetical treatment

(e) Negatively sloped budget lines

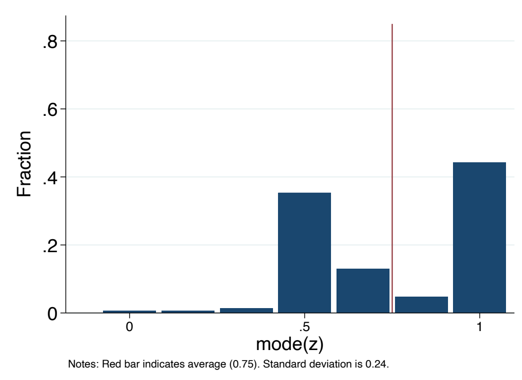

(f) Positively sloped budget lines

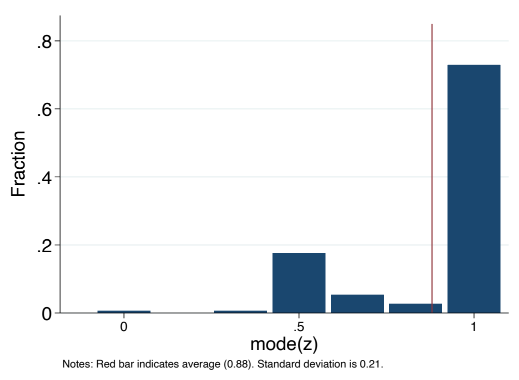

Notes: The figure shows the distribution of individuals' modal choices among negatively sloped and among positively sloped budget lines. For each budget line, $z = 6$ indicates an own-payoff maximizing choice, $z = 0$ indicates an own-payoff minimizing choice, $z = 3$ indicates a payoff-equalizing choice. The red vertical line indicates always the average over all modal choices. Panels (a) and (b) are constructed using subjects randomized into the *Low-Incentives* treatment. Panels (c) and (d) are constructed using subjects randomized into the *High-Incentives* treatment. Panels (e) and (f) are constructed using subjects randomized into the *Hypothetical* treatment.

Figure G.2: Cumulative distribution of modal choices of students

(a) Negatively sloped budget lines (b) Positively sloped budget lines

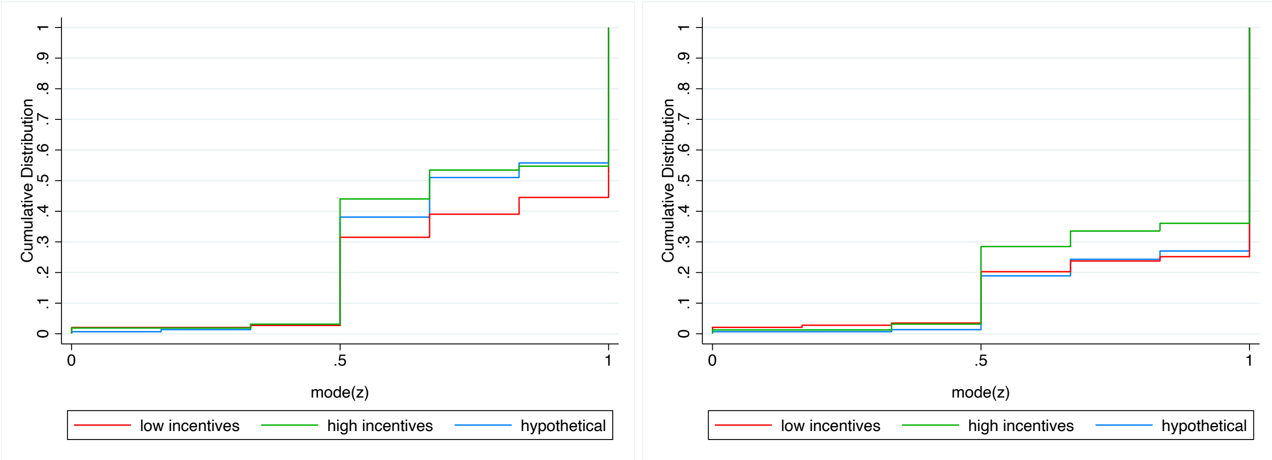

Notes: The figure shows the cumulative distribution of individuals' modal choices among negatively sloped and among positively sloped budget lines. For each budget line, $z = 1$ indicates an own-payoff maximizing choice, $z = 0$ indicates an own-payoff minimizing choice, $z = 0.5$ indicates a payoff-equalizing choice.

We depict subjects' modal choice on *both* positively sloped (x-axis) *and* negatively sloped (y-axis) budget in each treatment in Figure G.3. The Figure reveals the existence of the same three behavioral agglomerations across all three treatments: (i) A first agglomeration located at $z = 0.5$ for both positively and negatively sloped budget lines, i.e., a behavioral pattern suggestive of a preference for equality. (ii) A second agglomeration located at $z = 1$ for positively sloped budget lines and $z = 0.5$ for negatively sloped budget lines, i.e., a behavioral pattern suggestive of altruistic concerns for the worse off but no willingness to reduce the payoff of others for the sake of equality. (iii) A third agglomeration located at $z = 1$ for both positively and negatively sloped budget lines, i.e., a behavioral pattern suggestive of own payoff maximization.

Figure G.3: Descriptive evidence on students' modal choices

(a) *Low-Incentives* treatment

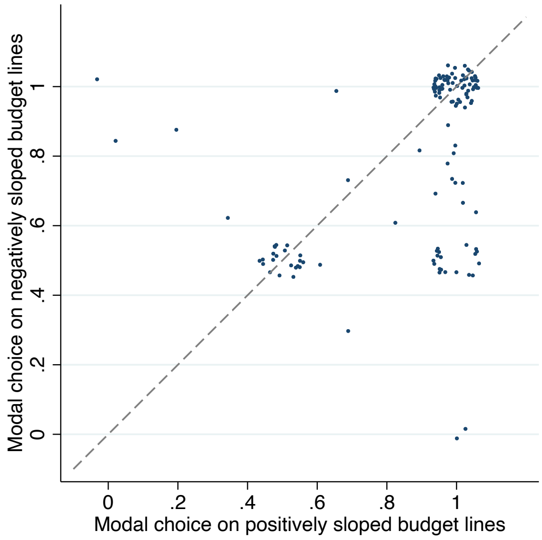

(b) *High-Incentives* treatment

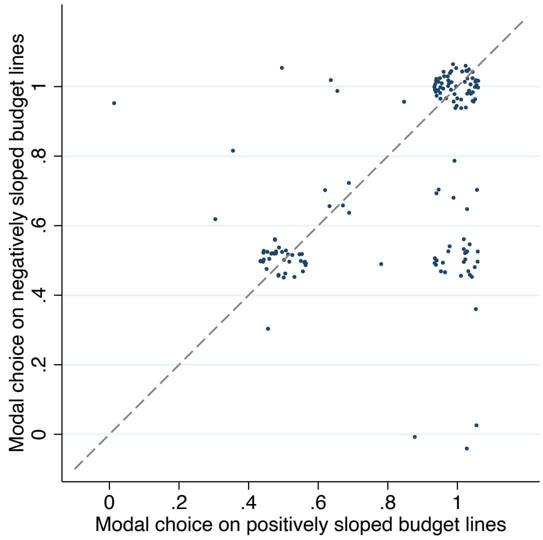

(c) *Hypothetical* treatment

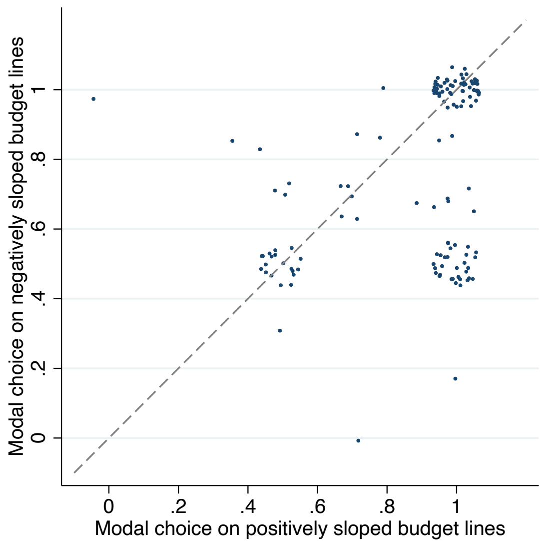

Notes: In all figures, we depict subjects' modal choices among negatively sloped budget lines and among positively sloped budget lines. Each dot represents one individual. Dots are jittered in order to make identical modal choices of individuals visible. For each budget line, $z = 1$ indicates an own-payoff maximizing choice, $z = 0$ indicates an own-payoff minimizing choice, and $z = 0.5$ indicates a payoff-equalizing choice. Panel (a) is constructed using subjects randomized into the *Low-Incentives* treatment. Panel (b) is constructed using subjects randomized into the *High-Incentives* treatment. Panel (c) is constructed using subjects randomized into the *Hypothetical* treatment. Note that if we replace individuals' modal choices by their median choices, very similar behavioral agglomerations emerge.

Figure G.4 displays the average behavior across the different treatments, separately for each budget line. While those in the Low-Incentives appear to chose slightly more often the own payoff maximizing allocation, overall we do not find any meaningful differences in the average behavior across treatments. This holds for the twelve budget lines that are centered around the 45-degree line (”center bundle”, see Figure 1) as well as for the budget lines in the advantageous domain (”south bundle”) and the disadvantageous domain (”north bundle”). For details on the budget lines, see Appendix B.1. Table G.2 shows the χ^2-squared test of independence by choice situation, corrected for multiple testing using Holm (1979). It further supports the notion of similar average behavior across treatments in the different budget lines.

Figure G.4: Average choices of students

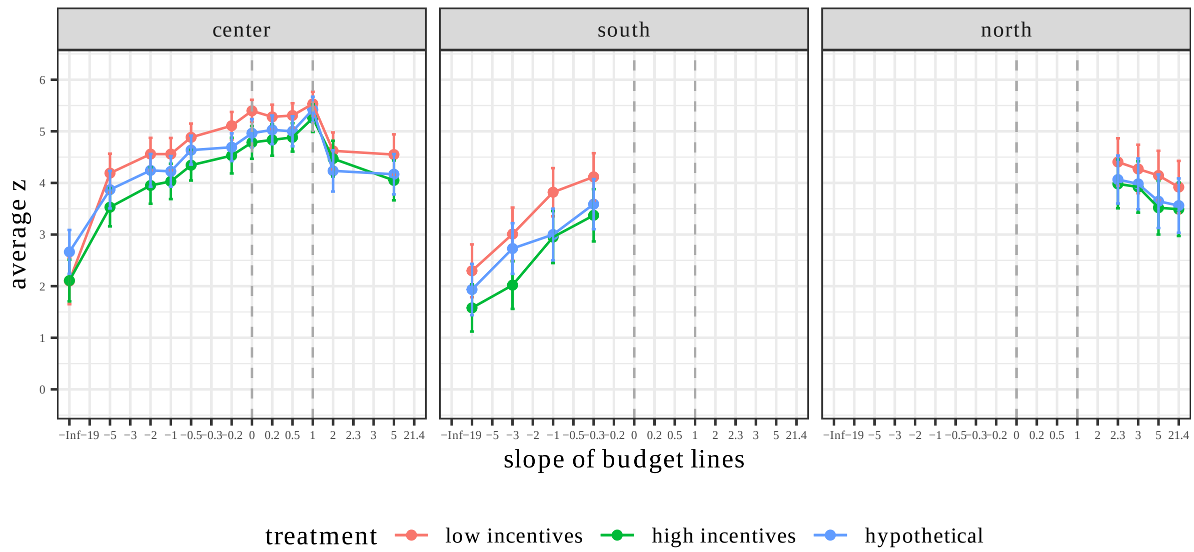

Notes: The figure shows the average implemented choice (z) on the y-axis, by budget line (ordered by their slopes) and by bundle of choice situations (3 panel). For each budget line, $z = 6$ indicates an own-payoff maximizing choice, $z = 0$ indicates an own-payoff minimizing choice, $z = 3$ indicates a payoff-equalizing choice. Budget lines are sorted by slope on the x-axis.

In sum, the descriptive results suggests that students’ choices are strikingly similar at the descriptive level, irrespective of whether their choices in the money allocation task are incentivized, and irrespective of the size of the financial stakes.

Table G.2: χ^2-squared test of independence by choice situation (Holm (1979) corrected) of students

	choiceId	statistic	p-value
1	1	18.380	1.000
2	2	22.898	0.520
3	3	11.305	1.000
4	4	17.223	1.000
5	5	19.227	1.000
6	6	13.692	1.000
7	7	24.084	0.440
8	8	17.138	1.000
9	9	20.836	1.000
10	10	11.672	1.000
11	11	15.497	1.000
12	12	14.022	1.000
13	13	12.388	1.000
14	14	25.278	0.250
15	15	9.640	1.000
16	16	18.964	1.000
17	17	19.565	1.000
18	18	15.888	1.000
19	19	22.321	0.580
20	20	21.831	0.620

Notes: This table shows the χ^2-squared test of independence by choice situation, corrected for multiple testing (Holm, 1979). It provides the results of a series of tests of equality in distributions of implemented choices z (where z ranges from 0 to 1) across the three treatments, by decision situation (choice ID).

G.3 Cluster analysis of students

To study whether incentives affect the distribution of social preferences in the student sample, we use the same approach as for the general population. To that end, we again apply a Bayesian nonparametric approach—the Dirichlet Process (DP) means clustering algorithm (Kulis and Jordan, 2012).

We run the DP-means algorithm separately on each treatment. We display the distribution of clusters identified by the DPM in the Table G.3. In order to assign labels to the identified clusters, we examine the behavioral characteristics of each cluster in Figure G.5, Figure G.6, and Figure G.7.

Like for the general population sample, we find that three clusters with a clear behavioral interpretation emerge in the Low-Incentives and in the Hypothetical treatments: a cluster comprising individuals predominantly equalizing payoffs[39], a cluster of subjects making altruistic choices towards those worse off, and a cluster comprising subjects making predominantly selfish choices. Importantly, however, note that the clustering algorithm does *not* identify a stable 3-types clustering in the High-Incentives treatment.[40] Interestingly, the assignment of students to types seems more responsive to incentives: in the hypothetical treatment, 83.55% of the subjects are assigned to a one of the other-regarding clusters (i.e., either altruistic or inequality averse), whereas only 55.07% are assigned to such a cluster when money is at stake (Low-Incentives treatment). In other words, students are more likely to be assigned to the selfish type when monetary incentives are used.

[39]For example, in the Low Incentives Treatment (Figure G.5) these subjects equalize payoffs by predominantly picking the center allocation ($z = 3$) for the budget lines centered around the 45-degree line. In the North Bundle, which comprises four budget lines with a positive slope above the 45 degree line (where the decision maker is always better off than the other—see Figure B.1), they tend to give up a large portion of their payoff ($z < 3$) in order to decrease the payoff of those better off–thereby achieving greater equality. In the South bundle, which comprises four budget lines with a negative slope under the 45 degree line–see Figure B.1, they implement allocations in which they have to give up a substantial portion of their endowment in order to achieve greater inequality ($z < 3$).

[40]More precisely, it pools almost all subjects into a single cluster that therefore contains a mixture of types, and assigns only one subject into each of the remaining two clusters (for details, see Figure G.6).

Table G.3: Type distributions of students identified using clustering analysis

	Low-Incentives	High-Incentives	Hypothetical
Cluster 1 (Inequality averse)	27.22%	N.A.	34.76%
Cluster 2 (Altruistic)	27.85%	N.A.	48.78%
Cluster 3 (Selfish)	44.94%	N.A.	16.46%

Notes: The table displays the distribution of individuals to the three clusters (in percent) that emerge in our dataset, separately for each treatment. The behavioral interpretation of the clusters (indicated in the left column) is based on the interpretation of each cluster's typical behavior.

Figure G.5: Average choices of students: Low-Incentives

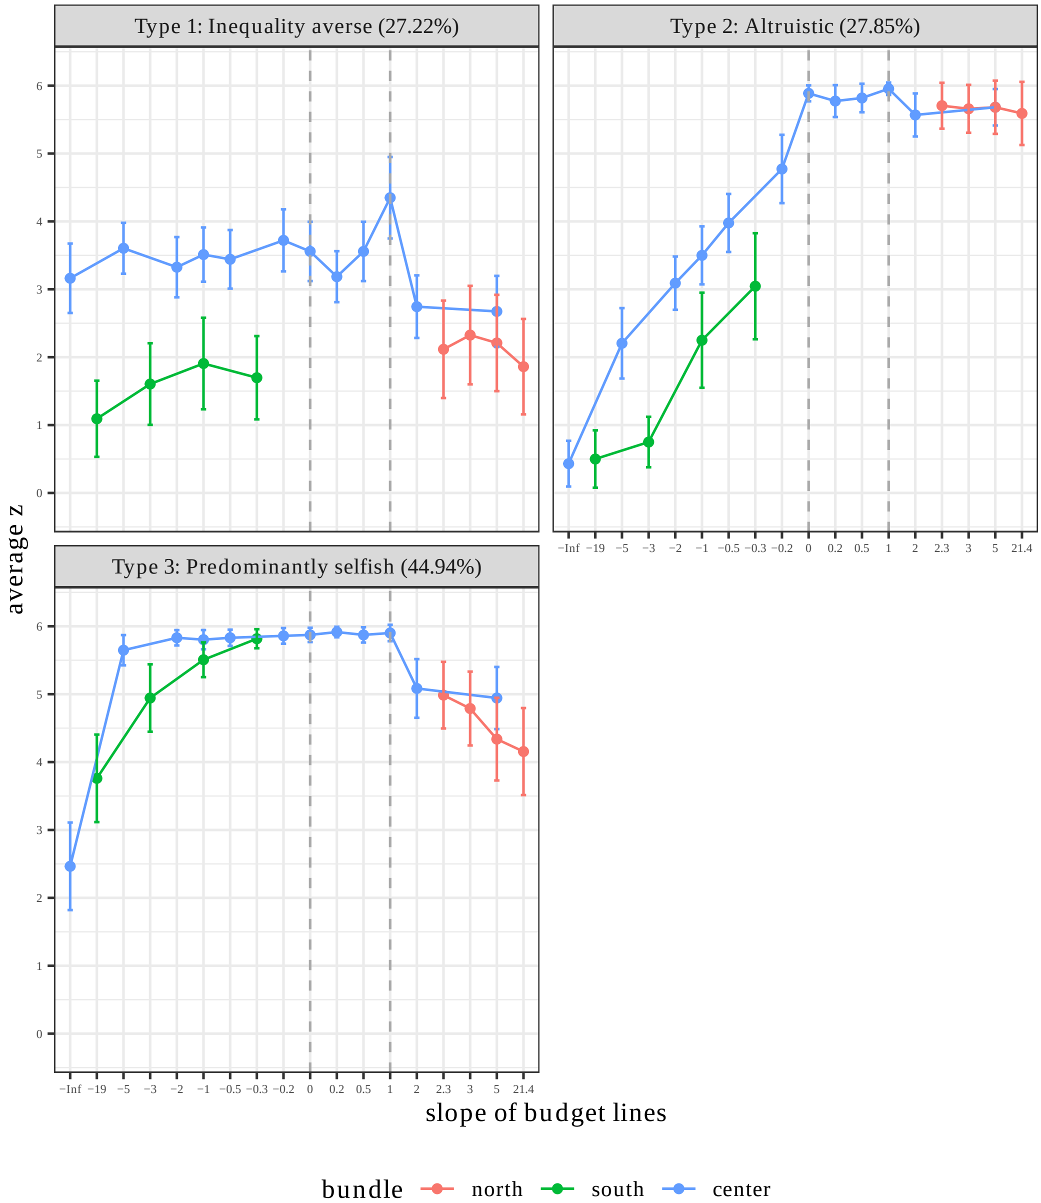

Notes: The figure shows the mean z of students' choices in the Low-Incentives treatment on the y-axis. For each budget line, $z = 6$ indicates an own-payoff maximizing choice, $z = 0$ indicates an own-payoff minimizing choice, $z = 3$ indicates a payoff-equalizing choice. Budget lines are sorted by slope on the x-axis.

Figure G.6: Average choices of students: High-Incentives

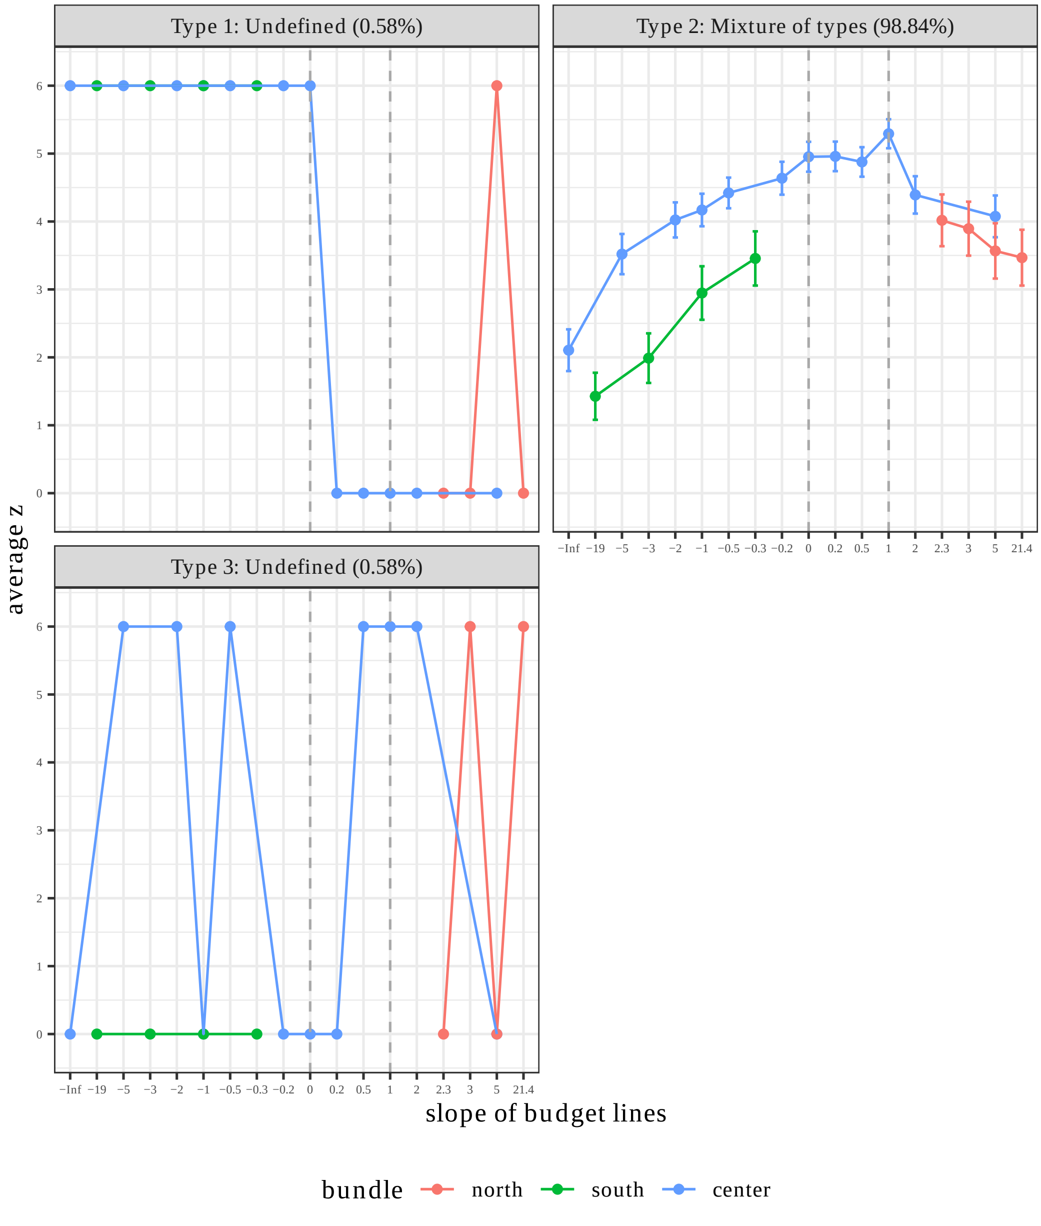

Notes: The figure shows the mean z of students' choices in the High-Incentives treatment on the y-axis. For each budget line, $z = 6$ indicates an own-payoff maximizing choice, $z = 0$ indicates an own-payoff minimizing choice, $z = 3$ indicates a payoff-equalizing choice. Budget lines are sorted by slope on the x-axis. The clustering algorithm does not identify a stable 3-types clustering in the High-Incentives treatment. It pools almost all subjects into a single cluster (Type 2) that therefore contains a mixture of types. It assigns only one subject into the Type 1 cluster and only one subject into the Type 3 cluster. Thus, the subfigures for Type 1 and Type 3 provide no statistically reliable type information.

Figure G.7: Average choices of students: Hypothetical

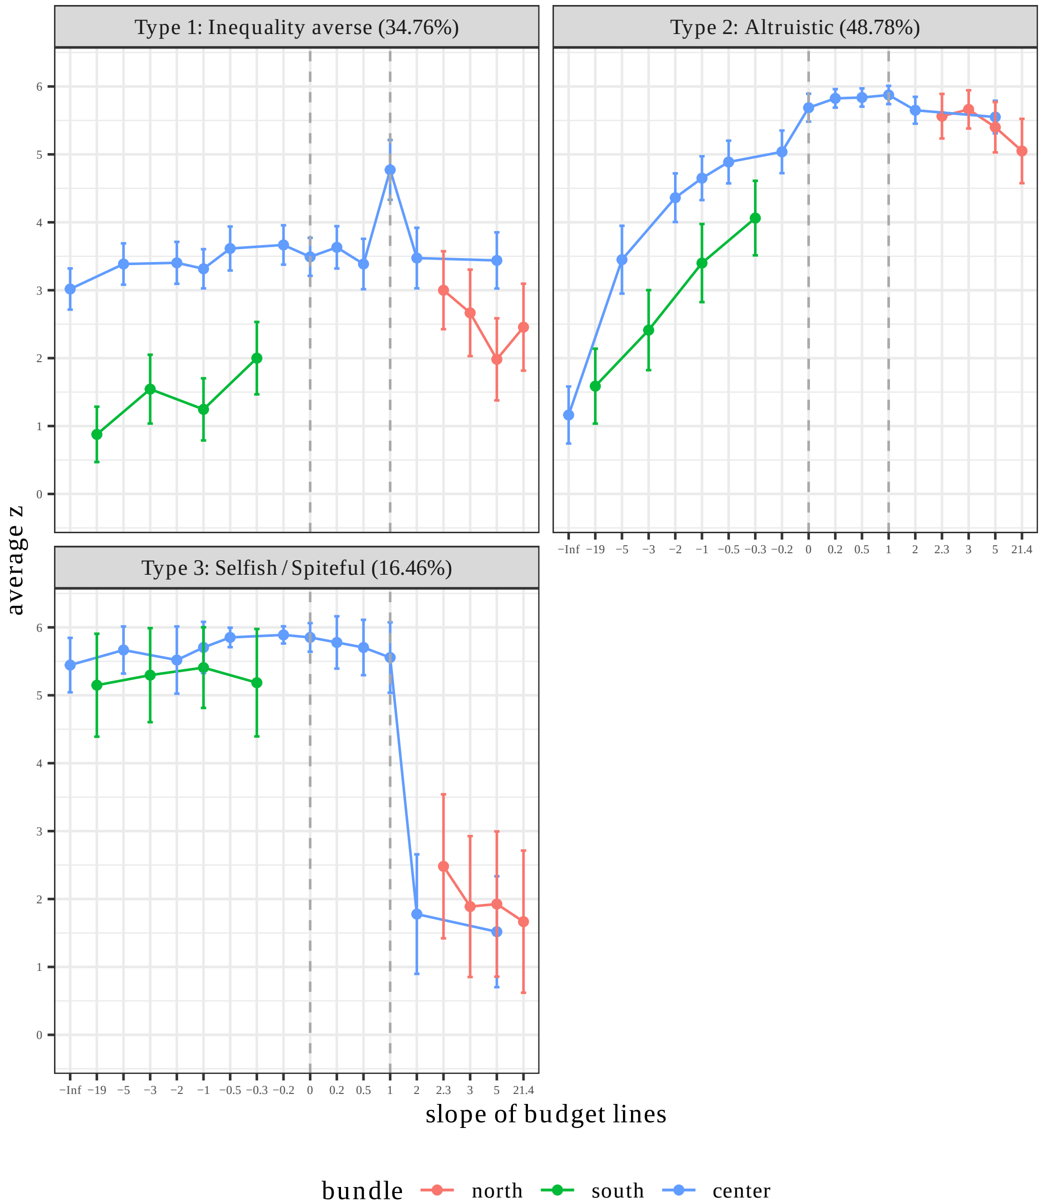

Notes: The figure shows the mean z of students' choices in the Hypothetical treatment on the y-axis. For each budget line, $z = 6$ indicates an own-payoff maximizing choice, $z = 0$ indicates an own-payoff minimizing choice, $z = 3$ indicates a payoff-equalizing choice. Budget lines are sorted by slope on the x-axis.

G.4 Structural analysis of students

We depict the mean values and standard deviations of the estimated structural parameters, separately by treatment, in Table G.4. Surprisingly, and in contrast with our results from the general population, we find the highest estimated parameters of inequality aversion in the High-Incentives treatment.

Table G.4: Summary statistics across treatment of students

	Low-Incentives		High-Incentives		Hypothetical	
	Mean	S.D.	Mean	S.D.	Mean	S.D.
α	0.216	0.752	0.544	1.010	0.426	0.764
β	0.393	0.871	0.675	0.959	0.522	0.945

To study whether the distributions of α and β differ accross treatments, we depict their probability density function in Figure G.8, their cumulative distribution functions (CDFs) in Figure G.9, and a series of pairwise Kolmogorov-Smirnov (KS) tests of equality in distributions in Table G.5.[41] Figure G.8 and Figure G.9 indicate slight differences across treatments that are consistent with the results of the average values discussed above. Again, we find that – compared to the Low-Incentives treatment – the Hypothetical and High-Incentives treatments shift the distributions of α and β-parameters to the right, with the shift being more pronounced for the latter. Despite these slightly visible differences, Table 3 reveals that stake sizes only significantly affect the distribution of individuals' estimated α-parameters (Low vs. High-Incentives: $p = 0.024$). On the contrary, we find no significant differences between the α-parameter distributions of the Hypothetical treatment and the two incentivized treatments. Moreover, we do not find any significant differences between the distributions of the β-parameters.

Turning to the precision of the estimates, we find that precision is the lowest in the High-Incentives treatment (see Table G.6), which is also in contrast with our results from the general population sample.

Overall, these results are slightly less consistent than those established in the general population sample. However, it is important to note that these results might have to be taken with a grain of salt since our student sample is much less well powered than our general population sample.

[41] For these tests, we apply Holm (1979) correction to account for multiple hypothesis testing.

Figure G.8: Distribution of students' structurally estimated parameters

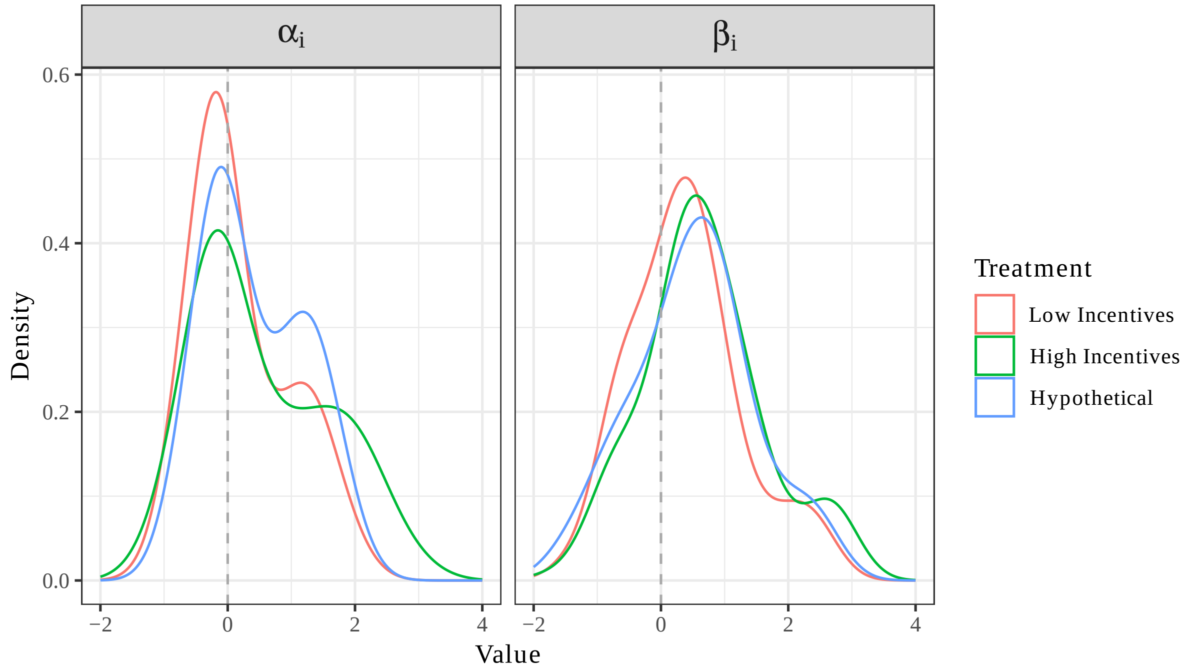

Notes: The distribution on the left side depicts subjects' structurally estimated alpha parameters by treatment. The distribution on the right side depicts subjects' structurally estimated beta parameters by treatment.

Figure G.9: CDFs of students' structurally estimated parameters

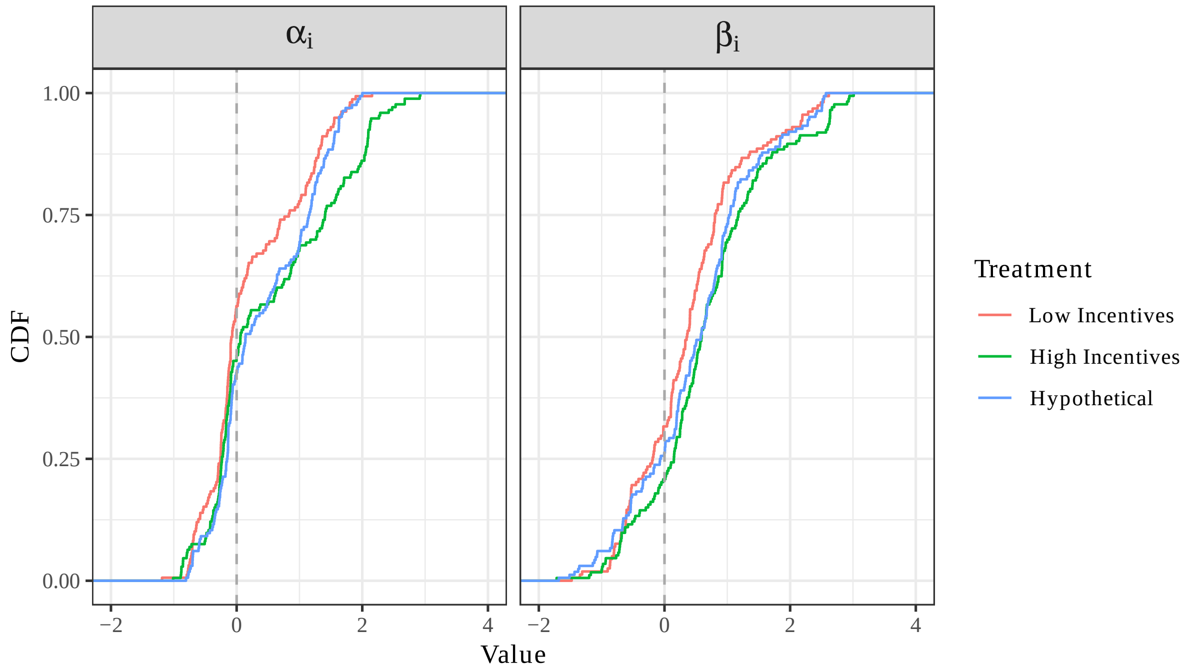

Notes: The CDFs on the left side depicts subjects' structurally estimated α-parameters by treatment. The CDFs on the right side depicts subjects' structurally estimated β-parameters by treatment.

Table G.5: Kolmogorov-Smirnov Test p-values with Holm (1979) correction of students

Comparison	α-parameter p-value	β-parameter p-value
Low-Incentives vs. High-Incentives	0.024	0.055
Low-Incentives vs. Hypothetical	0.073	0.093
High-Incentives vs. Hypothetical	0.073	0.512

Table G.6: Median posterior standard deviation by treatment condition of students

Treatment	α	β
Low-Incentives	0.238	0.241
High-Incentives	0.441	0.319
Hypothetical	0.341	0.308