When Linda Meets Preeti: The Validation of Behavioral Biases in India

WORKING PAPER 2304

AUGUST 2023

Nicholas Owsley

Narges Hajimoladarvish

Aditya Laumas

When Linda meets Preeti: The Validation of Behavioral Biases in India

Nicholas Owsley[*]   Narges Hajimoladarvish[†]   Aditya Laumas[‡]

August 24, 2023

Abstract

This paper presents results from an experiment testing 10 of the canonical biases from the behavioral economics literature amongst two distinct 'non-WEIRD' (Western Educated Industrialized Rich and Democratic) populations: low-income Indians, and university students from an elite Indian university. The study tests for the existence of the 'behavioral bias' for each measure with the pooled 'non-WEIRD' sample, and tests for heterogeneity across each distinct sub-sample. We find that both sub-samples display significant 'bias' in the majority of tests, suggesting that these well-documented behavioral biases are not peculiar to Western samples. We further find that the patterns of bias are the same for each sub-sample for most measures, but that there are notable exceptions for a small subset of measures. In most of these cases, the student sample, closer to typical samples for this type of research, shows stronger bias than the low-income sample[1].

[*]University of Chicago, Booth School of Business, Illinois, USA; Busara Center for Behavioral Economics, Nairobi, Kenya, nwsley@chicagobooth.edu; orcid: 0000-0001-7898-6947

[†]Centre for Social and Behaviour Change, Ashoka University, India

[‡]Centre for Social and Behaviour Change, Ashoka University, India

[1]This research was funded by the Centre for Social and Behaviour Change at Ashoka University. The authors would like to acknowledge Sami Horn's and Pooja Haldea's meaningful contributions to the research design.

1 Introduction

Most research subjects in behavioral economics are culturally and socioeconomically homogeneous (Henrich et al., 2010b). Western populations account for 96% of research subjects in top psychology journal publications (Arnett, 2016) and the bulk of subjects in foundational behavioral economics research (Thaler, 2018; Kahneman, 2011). Even within this peculiar group, the significant majority of research participants are university students, and even more specifically, psychology and economics undergraduates. While this sampling bias does not pose a challenge to the internal validity of individual effects, the accumulation of these findings has produced a skewed body of evidence that is far too often taken as universal and which provides the basis for decision-making models and the direction of research in applied settings (Fehr and Fischbacher, 2002; Cohen et al., 2020; Thaler, 2018; Camerer et al., 2016; Rad et al., 2018). This is particularly problematic given that several studies find considerable heterogeneity in decision patterns across different populations (Henrich et al., 2010b,c; Na et al., 2010; Barrett et al., 2016). Social preferences, risk preferences, and several other areas of decision-making appear to differ sharply across different geographies and socio-economic groups (Falk et al., 2018; Henrich et al., 2010a), and there is some evidence that more fundamental differences in cognition might exist (Nisbett et al., 2001; Tomasello et al., 2005). This highlights the potential risk that behavioral biases described in the behavioral economics literature are specific to Western samples and that their increasing application to policy in non-WEIRD contexts might not be appropriate (Manning et al., 2020).

Recent attempts to replicate core behavioral findings and expand subject pools to less typical samples represent an important shift in the field, but do not fully address the above problems. Most still sample from Western populations, and efforts to sample from non-Western populations are typically concentrated amongst elite university students (McShane et al., 2019; Klein et al., 2018; Camerer et al., 2016). There also remains a specific dearth of evidence amongst low-income samples in developing countries (Rad et al., 2018). Further, many studies rely heavily on measures constructed and validated with Western samples, that in some cases fail tests of reliability and validity when used with non-WEIRD groups (Falk et al., 2018; Laajaj et al., 2019) - using uncontextualized versions of these measures could lead to considerable measurement error and the failure to identify true effects amongst new populations. Finally, most of these studies are not designed to identify heterogeneity in effects across specific relevant sub-samples (McShane et al., 2019). However, identifying heterogeneous effects is both crucial for enriching models of decision-making and determining whether applications of behavioral economics will succeed across different contexts. There is thus considerable room to expand the evidence for behavioral biases amongst non-Western and non-typical samples.

This study presents the results of a randomized controlled trial in an experimental laboratory to test core behavioral economics measures with samples drawn from two non-Western populations of interest - 457 low-income Indians in Haryana state (typical targets of behavioral development interventions), and 457 students from an elite Indian University based in the same state (more similar to typical Western samples and the target of newer theoretical or replication studies expanding to non-WEIRD regions). We elicit responses for each of 10 behavioral biases through 11 individual survey-based measures modeled on original measures. Given several original measures were deemed inappropriate for these groups (due to cultural references and low comprehensibility in this context), we adapted them through a structured contextualization process - this included working with Global South behavioral scientists, prototyping, and piloting with local populations. Pilot and qualitative data show that without contextualization these instruments would likely generate considerable measurement error. We randomly assign treatment for each of our final measures and test for both the existence of the behavioral bias for the 'pooled sample', and for the presence of effect heterogeneity between the distinct sub-samples.

The majority of the measures generated significant effects, indicating the presence of 'behavioral biases' in our 'non-WEIRD' sample: 9 of the 11 tests show a highly significant effect for the pooled sample. For most measures, the effects were not statistically nor qualitatively different between our sub-samples, suggesting low-income and elite student populations show roughly similar levels of bias for these measures. However, the student population alone showed the presence of "mental accounting" (Thaler, 1985), and showed marginally higher rates of bias for the "common ratio effect" (Kahneman and Tversky, 1979). The low-income group showed a marginally higher rate of bias for the cognitive reflection test (Frederick, 2005), a largely unsurprising result given it's correlation with individual education.

The results from this study show that many of the core findings from behavioral economics are robustly observed across two highly distinct non-WEIRD populations. This gives cause for optimism regarding the replicability of these and similar findings across geographic and socio-economic contexts. Nonetheless, given the heterogeneity in certain outcomes, and that some of these measures are foundational for models of decision making and many development interventions, it remains important to continue to gather evidence across diverse populations. The results also show that, when doing such work, researchers and practitioners should adapt their measures and materials if they intend to gather valid evidence from diverse target populations.

This study adds to the behavioral economics literature by testing core measures with non-typical samples, providing replication evidence for 9 behavioral biases and building the body of behavioral economics evidence, specifically in controlled laboratory environments, from the Global South. Furthermore, this study provides details on a process for contextualizing measures which shows promise for generating more valid data in non-WEIRD contexts. Finally, this study tests specifically for heterogeneity in effects between national sub-populations of scientific and policy interest and finds evidence supporting the application of behavioral models across both groups.

2 Study design

2.1 Sample

Three primary criteria guided the selection of target populations. i) We sought populations from a country in the Global South in order to determine if core biases established through studies on Western samples are observed in a Global South context. ii) We sought two highly distinct groups within a Global South country to test for replication and heterogeneity across socio-economic dimensions. iii) We sought two relevant groups regarding the status of the behavioral economics literature: a group of elite university students, a population that replications and theoretical studies in the behavioral economics and judgement and decision-making are now focusing on to expand their reach to 'non-WEIRD' populations; and a group

Summary Statistics by Sample Group

Variable	Pooled Mean	Student Mean	Low-income Mean	T-test
Age	23.52	19.69	27.35	0.00***
		(0.10)	(0.36)	
Female	0.62	0.56	0.68	0.00***
		(0.02)	(0.02)	
Education	12.43	13.23	11.65	0.00***
		(0.07)	(0.08)	
Took-Psych-Course	0.24	0.33	0.14	0.00***
		(0.02)	(0.02)	
Married	0.31	0.00	0.61	0.00***
		(0.00)	(0.02)	
Home-tongue-English	0.41	0.8	0.02	0.00***
		(0.02)	(0.01)	
Home-tongue-Haryani	0.15	0.00	0.29	0.00***
		(0.00)	(0.02)	
Home-tongue-Hindi	0.43	0.19	0.68	0.00***
		(0.02)	(0.02)	
Home-tongue-Other	0.00	0.02	0.00	0.09*
		(0.01)	(0.00)	

of low-income Indians, similar to those typically targeted by behaviorally informed development policy and field research, but which continue to be excluded from more theoretical decision laboratory studies.

The final sample for this study thus consists of two groups of participants, for a full sample of 914 people:

- A sample of 457 low-income people from peri-urban villages in the vicinity of Sonepat, Haryana State, India.
- A sample of 457 students from an elite liberal arts university in Haryana State, India.

To recruit the low-income sample, we identified all distinct villages within a 35km radius of Sonepat, Haryana State, India. We then randomly selected 14 villages from this list for inclusion. We first contacted local village leaders and with their permission and help, actively recruited participants in focal public areas within each village, such as local schools. In several villages, we recruited large proportions of the adult population. We first acquired consent and then recorded baseline data for each potential participant. Participants were only included in the study if they had completed eight years of education, and were above the age of 18 and below the age of 60. A total of 457 participants from this pool completed the full experiment.

To recruit the university sample, participants were recruited from the student administration of an elite liberal arts university in Haryana State; the primary eligibility requirement was being a registered student at the university. We excluded students who were in a psychology program to avoid possible reporting bias from exposure to similar materials, and included a question and control for whether or not students had taken any behavioral economics or psychology courses. A total of 457 participants from this pool completed the full experiment.

The study is powered to detect an Minimum Detectable Effect Size (MDES) of approximately 0.2 standard deviations for within-sub-sample treatment effects, and 0.41 for between-sub-sample differences in treatment effects (described below in the Econometric strategy).

2.2 Procedures and Tasks

2.2.1 Task Selection

The tasks were selected to cover a range of the most well-known and influential biases in the behavioral economics literature (Thaler, 2015). We focus largely on measures derived from the early 'Heuristics and Biases' and 'Prospect Theory' work by Kahneman and Tversky; these measures and the papers in which they were originally published account for some of the most highly cited social science papers of all time (Tversky and Kahneman, 1974; Kahneman and Tversky, 1979), and Prospect Theory has specifically been described as the most influential social science theory since the beginning of the 20th century (Ruggeri et al., 2020; Barberis, 2013). In each case, we derived our final measures from original or commonly used versions of the measure to better compare to the findings with Western samples. In some cases, we underwent a rigorous iterative contextualization process (described below) to make measures culturally appropriate and preserve their ability to identify underlying biases. We finally selected 11 measures to identify 10 different biases. The full list of measures is described below, with details and wording of each measure in the Appendix.

List of Biases and Measures:

1. **Representativeness Heuristic**: Specifically the 'conjunction fallacy', a bias in probabilistic reasoning owing to the 'representativeness' of certain judgement targets leading to overestimating probability judgements (Tversky and Kahneman, 1974).

2. **Cognitive Reflection Test**: A test of the capacity for reflective thinking over wrong intuitive responses. (Frederick, 2005)

3. **Mood Heuristic**: The effect of 'mood' on general well-being assessments. (Strack et al., 1988)

4. **Anchoring**: The effect of arbitrary numerical primes, that can serve as 'anchors', on evaluations and willingness to pay (Tversky and Kahneman, 1974).

5. **Loss Framing**: A test for differences in choices in identical gambles depending on loss and gain framing (Tversky and Kahneman, 1974).

6. **Default - Assessment**: Whether a default selection for an assessment on a numerical scale affects the final assessment.

7. **Defaults - Charity Contribution**: Whether a pre-selected default option affects decisions on which charity to contribute to (Johnson and Goldstein, 2003).

8. **Present-bias**: Whether participants are more 'impatient' over choices that could yield benefits now compared with in the future.

9. **Common Ratio Effect**: Whether participants show inconsistent risk preferences across different elicitation mechanisms (specifically different probabilities) (Kahneman and Tversky, 1979).

10. **Decoy (Atrraction) Effect**: Whether the inclusion an irrelevant 'decoy' option in a menu affects choices (Huber et al., 1982).

11. **Mental Accounting**: Whether different framing of identical monetary endowments changes purchasing decisions (Thaler, 1985).

2.2.2 Task Contextualization

Most tasks (9 out of 11) were closely drawn from the original or most common measures in the literature [2]. Several of the these nine measures contained culturally-specific reference points that were potentially not known to the Indian subject pools. In a related study, Coman et al. (2017) find a high degree of measurement error in responses to several identical original measures in a laboratory study with student and low-income Kenyan samples. As such, in these cases, there was a high risk of measurement error due to potential low comprehension of questions and choice sets, which could attenuate the estimates of biases (and as such, would not allow us to replicate these effects even if they exist). The measures therefore underwent a review and multi-step contextualization process in order to preserve the core components of original measures [3].

First, a team of India- and Kenya-based behavioral science researchers reviewed each measure for contextual appropriateness for the target subject pools; they specifically identified all measures and features that included references or concepts that were distinctly 'Western', or that might appear 'alien' to the local target populations. After determining the inappropriate features, the authors and these researchers identified the fundamental structure and key elements of each measure. This imposed a clear structure for identifying and substituting local cultural references for foreign ones; there was then a thorough review to ensure that the format and features of each original measure was preserved. All contextualized measures were then piloted with Indian field staff, which included a combination of cognitive debriefing, prototyping, and response collection. This was in order to check for comprehension, cultural appropriateness, and for markers of measurement error. Measures were then iterated following several rounds of qualitative feedback and pilot data. This process was then repeated with participants from the target subject pools, and measures were again adapted and finalized until low comprehension and clear patterns of measurement error (such as random responding, or counter-theoretical responses) were minimized. Ultimately, we finalized nine India-specific measures that preserved the structure of the originals.

2.3 Data Collection

Data collection took place between July and September 2019 and was conducted in a decision laboratory for all participants, divided into individual sessions. Each session included 20 participants, who conducted tasks and surveys on individual Windows tablet computers using the oTree software. The full set of instructions were translated and back-translated into Hindi; each screen presented instructions in both

[2] The two exceptions were the two measures of 'default' effects which were straightforward applications of the principle of default selections in economic decisions.

[3] See the Appendix for both originals and contextualized versions, and for an example of the process described below

English and Hindi, and research assistants read instructions out loud for all tasks for the entire set of participants in both languages. Each participant received 400 Indian Rupees for their participation, provided in cash directly after the session. We determined that for all participants, but particularly for the low-income sample, the controlled laboratory setting with clear, translated, and dictated instructions was important for identifying biases and comparing them across sub-samples.

Each of the 11 tasks were administered sequentially after a basic demographic survey and the order of each task was randomized.

2.4 Treatment Assignment

Out of the 11 measures, 7 include between-subject treatment and control conditions - in these 7 cases, participants were randomly assigned with equal probability to each condition, and assignment occurred on the session-task level, such that individuals within the same session were randomly assigned to different conditions 'across tasks' but the same condition 'across individuals'. For each respective task, individuals within the same session are in the same treatment group (See the Appendix for details of treatment conditions for each task). For 4 of the 11 measures, we can identify the bias within-subject (The Appendix describes how each bias is identified).

3 Empirical Strategy

There are two types of tasks we use to identify biases: between subjects' measures, where participants are randomly assigned to a treatment or control condition, and within-subjects' measures. For the within subjects' measures, we identify individual bias for each participant as per each measure's description in the Appendix (in most cases, this is a simple transformation of their responses). The identification of the mean incidence of the bias is equivalent to an intercept-only model for each bias. To identify differences in the bias between each sub-sample for the within-subject measures, we estimate the following model:

$$Y_i = \beta_0 + \beta_1 Student_i + \epsilon_i \tag{1}$$

For biases where the identification relies on a between-subjects' design, and thus assignment to a treatment condition, we use the following set of models:

$$Y_i = \beta_0 + \beta_1 Treatment_i + \epsilon_i \tag{2}$$

$$Y_i = \beta_0 + \beta_1 Student_i + \beta_2 Treatment_i + \beta_3 Student_i * Treatment_i + \epsilon_i \tag{3}$$

Equation 2 simply allows us to identify the overall bias across both sub-samples (what we call the 'pooled sample') through β_1. For all models, Y_i is the outcome of interest for participant i (the response for each measure), $Student_i$ is a dummy variable equal to 1 if the participant is in the student sub-sample, $Treatment_i$ is a dummy variable equal to 1 if the participant was randomly assigned to a treatment group, and $Student_i * Treatment_i$ is the interaction term between student sub-sample and treatment group assignment and is used to identify the differential effect of the treatment for students. β_1 in Equation 2 provides the raw treatment effect (or bias estimate) for the full sample. β_2 in Equation 3 provides the treatment effect for the low-income group alone, while β_3 in Equation 3 provides the differential treatment effect for students compared to the low-income sub-sample. For robustness checks, each of the above models is also estimated with a vector of control variables X_i included in an additional specification.

These models apply to all measures with the exception of the Mood Heuristic, which follows a slightly different analytical procedure owing to the treatment effect being the product of an additional interaction [4].

4 Results

4.1 All Biases - Overall Results

Figure 1: Treatment Effect Sizes: Comparison with Meta-analyses Effect Sizes

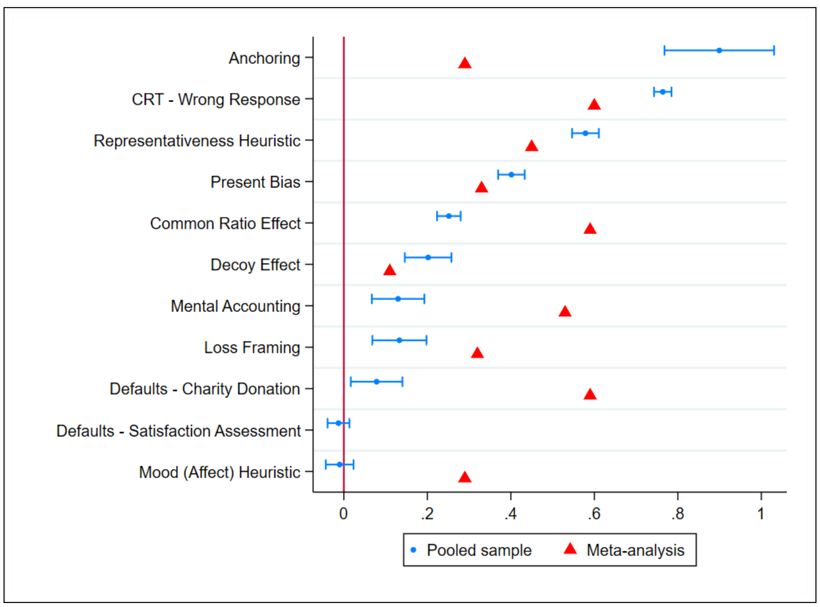

Note: Blue points represent effect sizes from Equation 2 for the main effect for each measure, and blue lines represent the 95% confidence intervals for these estimates. Notably, all effects except for Anchoring are measured in proportional terms (ie. the proportion of the sample showing the bias, or the average proportional bias for the 'Defaults - Assessment'). The Anchoring effect represents the coefficient of the anchor itself on the willingness to pay (also reasonably scalable between 0 and 1; the interpretation allows a value of 1, meaning the anchor perfectly explains the willingness to pay, and 0, that the anchor does not explain any of the willingness to pay). In all cases, the red triangle represents the effect sizes in the same units from meta-analyses conducted with these measures.

Overall, there is clear evidence for the existence of most of the behavioral biases in the pooled Indian sample - 9 out of the 11 measures show significant effects in the expected direction (Table 1, Column 1, Panels A and B; Figure 2), many of which have a p-value less than 0.001. The exceptions are the 'Mood

[4] This is described under the 'Mood Heuristic' in the Appendix.

Table 1: Full Regression Results - All Measures

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A - Within Subjects' Measures						
	Overall Mean		Low-income Mean	Student Mean Difference		
Representativeness	0.579***		0.584***	-0.011		
	(0.022)		(0.037)	(0.044)		
CRT	0.236***		0.071***	0.329***		
	(0.026)		(0.008)	(0.019)		
Common Ratio	0.252***		0.140***	0.223***		
	(0.024)		(0.019)	(0.034)		
Present Bias	0.402***		0.403***	-0.002		
	(0.019)		(0.026)	(0.038)		
Panel B - Between Subjects' Measures						
	(1)	(2)	(3)	(4)	(5)	(6)
	Overall Treatment Effect	Control Group Mean	Low-income Treatment Effect	Student Effect Difference	Low-income Control Group Mean	Student Control Group Mean Difference
Anchoring	719.541***	824.258	720.078***	-1.068	823.948	146.438**
	(46.462)	(40.871)	(71.135)	(93.058)	(47.402)	(60.072)
Loss Framing	0.133***	0.444	0.092	0.081	0.461	-0.083
	(0.043)	(0.032)	(0.068)	(0.086)	(0.035)	(0.051)
Mental Accounting	-0.130***	0.415	-0.027	-0.205**	0.362	0.281***
	(0.041)	(0.041)	(0.057)	(0.077)	(0.049)	(0.064)
Default-Charity	0.079**	0.648	0.080	-0.003	0.647	-0.014
	(0.038)	(0.042)	(0.068)	(0.078)	(0.059)	(0.065)
Default-Assessment	-0.140	10.341	-0.309	0.323	10.399	-3.911***
	(0.177)	(0.098)	(0.211)	(0.344)	(0.092)	(0.254)
Decoy Effect	0.202***	0.707	0.205***	-0.006	0.706	-0.147***
	(0.027)	(0.025)	(0.038)	(0.054)	(0.029)	(0.037)
Mood Heuristic	0.002	8.721	-0.011	0.019	8.729	-1.903***
	(0.046)	(0.200)	(0.073)	(0.093)	(0.269)	(0.347)

Panel A shows the sample means, and mean differences between the Low-income and Student sub-sample for the within-subject bias measures, described in each row. Column 1 represents the mean for the full sample, with standard errors of the means in parentheses, while column 3 represents the same for the Low-income sample specifically, and column 4 for the mean difference between the Low-Income and Student samples (alternatively, columns 3 and 4 can be measured using β_0 and β_1 from Equation 1). Panel B represents output from regressions of the outcomes for all between-subjects biases from Equations 2 and 3. All regressions include a dummy for 'Student', but no other control variables. This table is reproduced in the Appendix but including a set of control variables. Column 1 represents the treatment effect from Equation 2; The right-hand panel represents results from Equation 3. Specifically, column 3 represents β_2, column 4 represents β_3, while columns 5 and 6 represent β_0 and β_1 respectively. Standard errors of coefficients are in parentheses. ***, **, and * indicate significance at the 1, 5, and 10 percent critical level. Standard errors are clustered at the session level (the level of randomization).

Heuristic' and the Default selection for assessment. Figure 1 illustrates the treatment effect sizes, and provides a comparison with effect sizes from meta-analyses for each bias amongst predominantly Western samples[5]

[5]This is intended to benchmark the effects for each sub-sample. Meta-analyses include mostly Western samples except for the CRT, Present Bias, Common Ratio Effect and Mental Accounting. Brañas-Garza et al. (2019) conducted a meta-study of the CRT, which includes studies from Argentina, Brazil, and Colombia. Ashraf et al. (2006) use a field experiment in the Philippines to elicit time preferences and the review of Common Ratio Effects by Blavatskyy et al. (2023) includes a study in South Africa. Priolo et al. (2023) conducted a series of Mental Accounting tests across 21 countries, including Brazil, Egypt, India and Indonesia. See Appendix 1 for more details regarding meta-analyses.

Figure 2: Treatment Effect Sizes by Sub-sample: Comparison with Meta-analyses

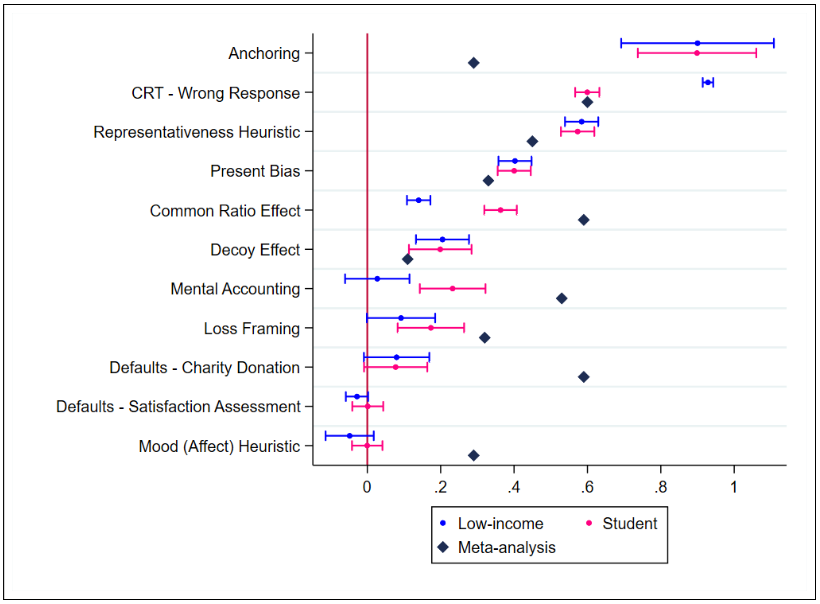

Note: Blue and pink points represent effect sizes from Equation 2, run separately for each sub-sample (Low-income and student samples respectively), for the main effect for each measure, and blue and pink lines represent the 95%a confidence intervals for these estimates. Notably, all effects except for Anchoring are measured in proportional terms (ie. the proportion of the sample showing the bias, or the average proportional bias for the 'Defaults - Assessment'). The Anchoring effect represents the coefficient of the anchor itself on the willingness to pay (also reasonably scalable between 0 and 1; the interpretation allows a value of 1, meaning the anchor perfectly explains the willingness to pay, and 0, that the anchor does not explain any of the willingness to pay). In all cases, the black diamond represents the effect sizes in the same units from meta-analyses conducted with these measures. These are the same meta-analysis effect sizes as in Figure 1.

Second, the patterns of bias are predominantly the same across both the student and low-income sub-samples. Effects are significant (at the 10% level) in the expected direction for both sub-samples in 6 of the 11 measures. There is a significant difference in the effect estimate in only 3 of the 11 tests between students and the low-income sample (Table 1, Column 5, Panels A and B), and in 2 of these cases the bias is still significant in both sub-samples in the same direction, but simply more pronounced in one than the other (for the Cognitive Reflection Test and for the Common Ratio Effect). There are none where the effects are differently signed.

Third, there are notable differences in biases between each sub-sample for 3 of the measures. The low-income sample scored significantly lower for the cognitive reflection test (CRT); this is not unexpected given the correlation between years of education and CRT scores (Frederick et al., 2002). More interestingly, the student sample displayed significantly larger biases in both the common ratio effect and in the measure of mental accounting. For mental accounting, specifically, the low-income sample showed a relatively 'tight null' effect while the student-sample showed a highly significant effect, far more similar to that found in previous studies. Figure 2 depicts treatment effect sizes by sub-sample and compares them with effect sizes from meta-analyses.

4.2 Effect Heterogeneity

We also explored heterogeneity in biases across another set of covariates, to test whether the biases are similarly consistent across other segments of the sample, and whether this exploration could help explain differences between low-income and student samples. We specifically explore differences in effects by gender (male and female participants), CRT score (low and high performers, median split), and by years of education (those with 12 or more years of education or those with less). Table 2 shows the results of this analysis, showing both means of within-subject measures and coefficients from Equation 2 for between-subjects' effects.

Across each dimension, we observe the presence of biases in the within-subjects' measures. However, similarly to the comparison between low-income and student samples, we observe significant differences in CRT scores across some dimensions, with males and participants with higher education achieving significantly higher scores on the test. This is consistent with existing literature that education is correlated with CRT scores and that men perform better on logic, mathematical reasoning, and similar tasks. Participants with high CRT scores also exhibited higher susceptibility to the 'common ratio effect', suggesting that higher ability in 'reflective thinking' is correlated with inconsistencies in risk preferences implied by the common ratio measure.

Regarding biases measured with between-subjects' measures, anchoring and the decoy effect was significant for each sub-sample and did not show any notable differences, while the Defaults-Assessment was not significant in any of the sub-samples; this is consistent with the low-income and student sample comparisons. However, participants with high levels of education and high CRT scores showed significantly greater loss framing and mental accounting effects. This suggests that the differences observed in these biases between low-income and student samples might be partly explained by differences in education and the related trained cognition, though we cannot rule out other correlates of socioeconomic status. Regarding the default effect with charity donations, while all sub-samples show directionally consistent effects, only those in the high education sub-sample were significantly biased. Overall, it appears that education and cognitive reflection measured by the CRT are associated with greater bias across several measures.

4.3 Relationships Between Effects

Given these biases are often explained as the result of the application of decision-making heuristics, it follows that some people might be more prone to applying heuristics across a range of decisions than others (some people may be systematically more 'biased' than others). Moreover, certain biases might be the result of similar heuristic processes, while others might be the result of more distinct ones. We therefore test for relationships between biases. For biases from within-subjects' measures, we simply analyze correlations between each measure. For biases drawn from between-subjects' measures, we use an interaction between individual within-subject responses (biases measured within subject) and treatment status in between subject measures to test whether within-subjects' bias predicts between subjects' bias. We do not measure the relationship between biases that use between-subjects' measures. The results are presented in Table 3. Overall, there is surprisingly little relationship between different biases in our samples. Across 45 tests of pairwise relationships between biases, only 5 are significant at the 5% level. This suggests that these biases are not necessarily capturing subjects' general cognitive tendencies or general inattentiveness, nor can they be explained by some underlying heuristic process. To further explore more general relationships between biases, we create an overall bias measure by summing the Representative

Table 2: Effect Heterogeneity

	CRT		Gender		Education	
	(1)	(2)	(3)	(4)	(5)	(6)
Biases	Low	High	Male	Female	Low	High
Panel A - Within Subjects' Measures						
Cognitive Reflection Test			0.325***	0.181***	0.169***	0.318***
			(0.035)	(0.023)	(0.025)	(0.025)
Representativeness Heuristic	0.579***	0.579***	0.521***	0.614***	0.583***	0.574***
	(0.027)	(0.028)	(0.034)	(0.025)	(0.033)	(0.024)
Common Ratio Effect	0.178***	0.349***	0.281***	0.234***	0.229***	0.280***
	(0.020)	(0.030)	(0.036)	(0.022)	(0.028)	(0.028)
Present Bias	0.398***	0.406***	0.427***	0.386***	0.392***	0.414***
	(0.024)	(0.026)	(0.032)	(0.022)	(0.023)	(0.025)
Panel B - Between Subjects' Measures						
Anchoring	745.777***	691.551***	687.025***	739.839***	701.810***	743.393***
	(68.805)	(59.276)	(67.600)	(60.643)	(73.035)	(52.327)
Loss framing	0.102*	0.175***	0.106**	0.154***	0.086	0.187***
	(0.059)	(0.058)	(0.052)	(0.051)	(0.056)	(0.050)
Mental Accounting	-0.088	-0.183***	-0.221***	-0.071	-0.071	-0.199***
	(0.053)	(0.055)	(0.053)	(0.049)	(0.052)	(0.047)
Defaults - Satisfaction Assessment	-0.278	0.040	-0.322	0.006	-0.023	-0.244
	(0.219)	(0.302)	(0.263)	(0.200)	(0.218)	(0.324)
Defaults - Charity Donation	0.094	0.054	0.072	0.088	0.054	0.107**
	(0.060)	(0.042)	(0.052)	(0.051)	(0.055)	(0.045)
Decoy Effect	0.214***	0.187***	0.176***	0.219***	0.204***	0.197***
	(0.034)	(0.045)	(0.041)	(0.036)	(0.041)	(0.042)
Mood Heuristic	-0.023	0.074*	-0.016	-0.009	-0.026	0.118**
	(0.018)	(0.028)	(0.033)	(0.039)	(0.025)	(0.048)

Panel A shows the sample means within each specified sub-sample for each within-subjects' bias. Panel B reports the coefficient on the treatment variable from a regression for Equation 2 run on each sub-sample with a dummy for 'Student'. Columns 1 and 2 show these outcomes for the sub-sample of low and high CRT scores, respectively; columns 3 and 4 for male and female participants; and columns 5 and 6 for those with low and high education. Standard errors are in parentheses. ***, **, and * indicate significance at the 1, 5, and 10 percent critical level. Standard errors are clustered at the session level (the level of randomization).

Heuristic, Common Ratio Effect, Present Bias, and CRT Intuitive scores for each participant, and then interacting it with treatment for each between-subjects bias measure in the same way. This measure, too, does not meaningfully predict any other individual bias (Table 3, Column 6).

Perhaps somewhat surprising is that the CRT does not predict many biases, though it is often considered as a proxy for more general 'automatic' thinking. However, some biases are correlated. There is a significant negative association between the Common Ratio Effect and both incorrect CRT responding and Intuitive CRT responding (choosing the intuitive but wrong answer in questions in the CRT). Hence, the identified bias in the Common Ratio Effect, which is to display inconsistent risk preferences across question framings, is driven by individuals with *higher* CRT scores (consistent with Section 4.2).

Another exception is a strong significant relationship between the Representative Heuristic and Mental Accounting. This indicates that individuals who make wrong probability judgements due to how 'representative' some description feels, are more likely to use 'mental accounts' when making purchasing decisions (and therefore to treat money as non-fungible).

Finally, we found that those who respond more intuitively to the CRT are more prone to defaulting to the charity donation option. This relationship alone corresponds with the literature arguing that the CRT measures a proclivity towards 'automatic' behaviors.

Table 3: Relationships between Biases

	CRT Wrong	CRT Intuitive	Representativeness Heuristic	Common Ratio Effect	Present Bias	Total Bias
CRT Wrong	1					
CRTIntuitive	0.660***	1				
Representativeness Heuristic	-0.013	-0.031	1			
Common Ratio Effect	-0.202***	-0.143***	0.045	1		
Present Bias	0.009	0.0561	0.043	0.044	1	
Anchoring	-1.757	184.539	36.301	83.765	12.134	47.168
	(120.577)	(140.898)	(104.635)	(129.408)	(90.163)	(37.461)
Loss framing	0.007	0.192*	0.094	0.074	-0.012	0.062
	(0.125)	(0.106)	(0.077)	(0.080)	(0.061)	(0.039)
Mental Accounting	0.080	0.107	-0.141***	-0.087	0.053	-0.037
	(0.116)	(0.108)	(0.051)	(0.069)	(0.059)	(0.023)
Defaults - Satisfaction Assessment	-0.530	-0.330	0.206	-0.257	0.043	-0.070
	(0.552)	(0.400)	(0.323)	(0.334)	(0.255)	(0.152)
Defaults - Charity Donation	0.088	0.329***	0.016	-0.050	0.077	0.052
	(0.106)	(0.096)	(0.069)	(0.059)	(0.082)	(0.041)
Decoy Effect	0.050	0.101	0.087	0.023	-0.089	0.006
	(0.095)	(0.088)	(0.058)	(0.083)	(0.055)	(0.032)
Mood Heuristic	-0.100**	-0.080**	-0.036	0.002	-0.034	-0.021
	(0.042)	(0.030)	(0.029)	(0.087)	(0.029)	(0.014)

The top panel reports Pearson's correlations between within subjects' measures of each bias. For each of these measures, we are able to identify individual bias measures. These estimates thus provide measures of the relationships between within-subjects' biases. The second panel reports coefficients from regressions of each primary outcome for each between-subject measure, using Equation 2, but with the addition of an interaction term between the individual within-subject bias measure (ie. the bias measure for CRT Wrong, CRT Intuitive, Representativeness Heuristic, Common Ratio Effect, Present Bias) and the treatment variable for the relevant between-subject measure. Each regression is run separately, dealing with the interaction of each individual bias measure (shown in the columns) and each between-subject treatment variable (represented by each row) separately. Specifically, we report the coefficient on the interaction term. This coefficient allows us to identify whether those displaying the bias in the within-subjects' measures are more or less likely to display the between subjects' biases. Standard errors are in parentheses. ***, **, and * indicate significance at the 1, 5, and 10 percent critical level. Standard errors are clustered at the session level (the level of randomization).

5 Discussion and Conclusion

Overall, the results of this study suggest that our populations of interest - low-income Indians liable to be targeted by behavioral development policy, and elite Indian university students - predominantly display the behavioral biases from the behavioral economics literature. This suggests that these biases are not exclusively patterns of decision-making of 'WEIRD' samples. The effects are significant for the majority of measures in the expected direction. For most, the direction is the same for each sub-sample, the magnitude of effects is pronounced, and in many cases the point estimates are similar to those for Western samples (Ariely and Jones, 2008; Kahneman and Tversky, 1979; Cohen et al., 2020; Frederick, 2005; Kahneman, 2011). For instance, the proportion of the sample showing present-biased time preferences (40%) over money tightly accords with the seminal literature on this phenomenon, (that finds that between 35% and 45% of typical samples are present-biased; (Cohen et al., 2020)), while the results for the student sample for the cognitive reflection test are nearly identical to that for student samples in the original study (40% correct answers compared to 43% respectively; Frederick (2005)) and in a meta-analyses (40%; (Brañas-Garza et al., 2019)). Moreover, the effect sizes for the representativeness heuristic, the decoy effect and anchoring are larger than in meta-analyses of these measures on Western samples (Yechiam and Zeif, 2023; Heath and Chatterjee, 1995; Li et al., 2021). While the estimates for the loss framing, the common ratio effect, and mental accounting (for students only) are smaller than those in meta-analyses, they remain significantly different from zero, suggesting that these biases still exist within this population, but that perhaps they are somewhat more muted or apply to a smaller proportion of the population (Kühberger, 1998; Blavatskyy et al., 2023; Priolo et al., 2023). Therefore, in spite of the existing evidence of observed differences in reasoning patterns, social preferences, and other differences in cognition and behavior (Henrich et al., 2010c; Falk et al., 2018), the results from this study show encouraging replication potential for some of the more fundamental behavioral biases.

These results are also consistent with some of the large-scale replication work in behavioral economics and similar disciplines. Camerer et al. (2016) replicate 18 behavioral economics experimental effects published in prominent economics journals between 2011 and 2014 and observe a 66% replication rate - slightly lower than but similar to that observed in this study. Klein et al. (2018) run a large replication project across multiple countries, including a number in the Global South, to test several of the seminal social psychology findings. Though they find a lower replication rate than this study (approximately 50%), they find little heterogeneity between WEIRD and 'less-WEIRD' (including middle-income countries) sub-samples, suggesting that the effects that 'survive' replication do not differ systematically by nationality.

In spite of the surprising lack of systematic differences, we still observe clear heterogeneity in a subset of measures. Heterogeneity in effects for the common ratio effect and mental accounting suggest that some biases can vary across different sub-national populations in important ways. Notably, the effects for the elite student sample are practically similar to estimates on Western samples for both measures, while the low-income group shows significantly less or no bias. The result for mental accounting for the student sample is also in sharp contrast to a recent study that uses the exact same measure amongst an elite Indian student sample (Banerjee et al., 2019) and finds no sign of mental accounting. However, the sample for this study, of 67, is far less than ours, of 457, and their use of inference to confirm a null hypothesis is not appropriate. The findings for both the common ratio effect and mental accounting also support some existing evidence that biases regarding small differences in monetary amounts are less pronounced among lower-income groups compared to higher income groups (Shah et al., 2015; Kroft et al., 2020). Mental accounting, specifically, has been more strongly observed for those with higher incomes and less for those facing conditions of financial scarcity (Olsen et al., 2019; Muehlbacher and Kirchler, 2019). Evidence also shows that those facing financial scarcity tend to pay more attention to small differences in monetary outcomes and are therefore more likely to act rationally over them (Mullainathan and Shafir, 2013; Kroft et al., 2020). Taken together, this suggests that research on similar economic preferences should explore socio-economic heterogeneity as a primary component of research designs.

Overall, this study provides evidence that these biases exist in non-WEIRD populations. However, it cannot make any claims about their universality. Future work should explore the replicability of and heterogeneity in these and similar measures. Robust large-sample and cross-country studies are needed to map these biases across a broader range of regions and sub-national groups, preferably with representative samples. 'Multi-lab' studies provide a possible facility for this. However, these studies need to do more to access non-student populations and more funding needs to be made available for theoretical work in the Global South for this to happen. At the least, even if universality is far away, we can support the claim that these biases are not just WEIRD.

References

Maurice Allais. Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école américaine. *Econometrica: Journal of the Econometric Society*, pages 503–546, 1953.

James Andreoni and Charles Sprenger. Estimating time preferences from convex budgets. *American Economic Review*, 102(7):3333–3356, 2012.

Dan Ariely and Simon Jones. *Predictably irrational.* Harper Audio New York, NY, 2008.

Jeffrey J Arnett. The neglected 95%: why american psychology needs to become less american. 2016.

Nava Ashraf, Dean Karlan, and Wesley Yin. Tying odysseus to the mast: Evidence from a commitment savings product in the philippines. *The Quarterly Journal of Economics*, 121(2):635–672, 2006.

Pronobesh Banerjee, Promothesh Chatterjee, Sanjay Mishra, and Anubhav A Mishra. Loss is a loss, why categorize it? mental accounting across cultures. *Journal of Consumer Behaviour*, 18(2):77–88, 2019.

Nicholas C Barberis. Thirty years of prospect theory in economics: A review and assessment. *Journal of Economic Perspectives*, 27(1):173–196, 2013.

H Clark Barrett, Alexander Bolyanatz, Alyssa N Crittenden, Daniel MT Fessler, Simon Fitzpatrick, Michael Gurven, Joseph Henrich, Martin Kanovsky, Geoff Kushnick, Anne Pisor, et al. Small-scale societies exhibit fundamental variation in the role of intentions in moral judgment. *Proceedings of the National Academy of Sciences*, 113(17):4688–4693, 2016.

Pavlo Blavatskyy, Valentyn Panchenko, and Andreas Ortmann. How common is the common-ratio effect? *Experimental Economics*, 26(2):253–272, 2023.

L. V. Hedges J. P. Higgins Borenstein, M. and H. R. Rothstein. *Introduction to meta analysis.* Wiley Online Library, England, 2009.

Pablo Brañas-Garza, Praveen Kujal, and Balint Lenkei. Cognitive reflection test: Whom, how, when. *Journal of Behavioral and Experimental Economics*, 82:101455, 2019.

Colin F Camerer, Anna Dreber, Eskil Forsell, Teck-Hua Ho, Jürgen Huber, Magnus Johannesson, Michael Kirchler, Johan Almenberg, Adam Altmejd, Taizan Chan, et al. Evaluating replicability of laboratory experiments in economics. *Science*, 351(6280):1433–1436, 2016.

Jonathan Cohen, Keith Marzilli Ericson, David Laibson, and John Myles White. Measuring time preferences. *Journal of Economic Literature*, 58(2):299–347, 2020.

Alan Coman, Johannes Haushofer, Chaning Jang, Elizabeth Levy Paluck, and Oliver Schweickart. Behavioral biases across cultures: Evidence from kenya and the usa. *Unpublished manuscript*, 2017.

Isaac Dinner, Eric J Johnson, Daniel G Goldstein, and Kaiya Liu. Partitioning default effects: why people choose not to choose. *Journal of Experimental Psychology: Applied*, 17(4):332, 2011.

Thomas Dohmen, Falk Armin, Huffman David, and Sunde Uwe. Dynamic inconsistency predicts self-control problems in humans. *Working Paper*, 2006.

Armin Falk, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman, and Uwe Sunde. Global evidence on economic preferences. *The Quarterly Journal of Economics*, 133(4):1645–1692, 2018.

Ernst Fehr and Urs Fischbacher. Why social preferences matter–the impact of non-selfish motives on competition, cooperation and incentives. *The economic journal*, 112(478):C1–C33, 2002.

Shane Frederick. Cognitive reflection and decision making. *Journal of Economic perspectives*, 19(4): 25–42, 2005.

Timothy B Heath and Subimal Chatterjee. Asymmetric decoy effects on lower-quality versus higher-quality brands: Meta-analytic and experimental evidence. *Journal of Consumer Research*, 22(3): 268–284, 1995.

Pamela W Henderson and Robert A Peterson. Mental accounting and categorization. *Organizational Behavior and Human Decision Processes*, 51(1):92–117, 1992.

Joseph Henrich, Jean Ensminger, Richard McElreath, Abigail Barr, Clark Barrett, Alexander Bolyanatz, Juan Camilo Cardenas, Michael Gurven, Edwins Gwako, Natalie Henrich, et al. Markets, religion, community size, and the evolution of fairness and punishment. *science*, 327(5972):1480–1484, 2010a.

Joseph Henrich, Steven J Heine, and Ara Norenzayan. Most people are not weird. *Nature*, 466(7302): 29–29, 2010b.

Joseph Henrich, Steven J Heine, and Ara Norenzayan. The weirdest people in the world? *Behavioral and brain sciences*, 33(2-3):61–83, 2010c.

Joel Huber, John W Payne, and Christopher Puto. Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of consumer research*, 9(1):90–98, 1982.

Taisuke Imai, Tom A Rutter, and Colin F Camerer. Meta-analysis of present-bias estimation using convex time budgets. *The Economic Journal*, 131(636):1788–1814, 2021.

Eric J Johnson and Daniel Goldstein. Do defaults save lives?, 2003.

Daniel Kahneman. *Thinking, fast and slow*. Macmillan, 2011.

Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292, 1979.

Daniel Kahneman and Amos Tversky. Choices, values, and frames. *American Psychologist*, 39(4):341, 1984.

Richard A Klein, Michelangelo Vianello, Fred Hasselman, Byron G Adams, Reginald B Adams Jr, Sinan Alper, Mark Aveyard, Jordan R Axt, Mayowa T Babalola, Štěpán Bahník, et al. Many labs 2: Investigating variation in replicability across samples and settings. *Advances in Methods and Practices in Psychological Science*, 1(4):443–490, 2018.

Kory Kroft, Jean-William P Laliberté, René Leal-Vizcaíno, and Matthew J Notowidigdo. Salience and taxation with imperfect competition. Technical report, National Bureau of Economic Research, 2020.

Anton Kühberger. The influence of framing on risky decisions: A meta-analysis. *Organizational behavior and human decision processes*, 75(1):23–55, 1998.

Rachid Laajaj, Karen Macours, Daniel Alejandro Pinzon Hernandez, Omar Arias, Samuel D Gosling, Jeff Potter, Marta Rubio-Codina, and Renos Vakis. Challenges to capture the big five personality traits in non-weird populations. *Science advances*, 5(7):eaaw5226, 2019.

David Laibson. Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 112(2): 443–478, 1997.

Lunzheng Li, Zacharias Maniadis, and Constantine Sedikides. Anchoring in economics: a meta-analysis of studies on willingness-to-pay and willingness-to-accept. *Journal of Behavioral and Experimental Economics*, 90:101629, 2021.

Lauren Manning, Abigail Goodnow Dalton, Zeina Afif, Renos Vakos, and Faisal Naru. Behavioral science around the world volume ii: Profiles of 17 international organizations. *eMBeD report. Washington, D.C.: World Bank Group*, 2020.

Blakeley B McShane, Jennifer L Tackett, Ulf Böckenholt, and Andrew Gelman. Large-scale replication projects in contemporary psychological research. *The American Statistician*, 73(sup1):99–105, 2019.

Stephan Meier and Charles Sprenger. Present-biased preferences and credit card borrowing. *American Economic Journal: Applied Economics*, 2(1):193–210, 2010.

Stephan Muehlbacher and Erich Kirchler. Individual differences in mental accounting. *Frontiers in Psychology*, 10, 2019.

Sendhil Mullainathan and Eldar Shafir. *Scarcity: Why having too little means so much*. Macmillan, 2013.

Jinkyung Na, Igor Grossmann, Michael EW Varnum, Shinobu Kitayama, Richard Gonzalez, and Richard E Nisbett. Cultural differences are not always reducible to individual differences. *Proceedings of the National Academy of Sciences*, 107(14):6192–6197, 2010.

Richard E Nisbett, Kaiping Peng, Incheol Choi, and Ara Norenzayan. Culture and systems of thought: holistic versus analytic cognition. *Psychological review*, 108(2):291, 2001.

Jerome Olsen, Matthias Kasper, Christoph Kogler, Stephan Muehlbacher, and Erich Kirchler. Mental accounting of income tax and value added tax among self-employed business owners. *Journal of Economic Psychology*, 70:125–139, 2019.

Giulia Priolo, Stablum Federica, Vacondio Martina, D'Ambrogio Simone, Caserotti Marta, Conte Beatrice, and De Roni Prisca. The robustness of mental accounting: a global perspective. *Science advances*, 2023.

Mostafa Salari Rad, Alison Jane Martingano, and Jeremy Ginges. Toward a psychology of homo sapiens: Making psychological science more representative of the human population. *Proceedings of the National Academy of Sciences*, 115(45):11401–11405, 2018.

Kai Ruggeri, Sonia Alí, Mari Louise Berge, Giulia Bertoldo, Ludvig D Bjørndal, Anna Cortijos-Bernabeu, Clair Davison, Emir Demić, Celia Esteban-Serna, Maja Friedemann, et al. Replicating patterns of prospect theory for decision under risk. *Nature human behaviour*, 4(6):622–633, 2020.

Ulrich Schimmack and Shigehiro Oishi. The influence of chronically and temporarily accessible information on life satisfaction judgments. *Journal of Personality and Social Psychology*, 89(3):395–406, 2005.

Anuj K Shah, Eldar Shafir, and Sendhil Mullainathan. Scarcity frames value. *Psychological science*, 26 (4):402–412, 2015.

N Craig Smith, Daniel G Goldstein, and Eric J Johnson. Choice without awareness: Ethical and policy implications of defaults. *Journal of Public Policy & Marketing*, 32(2):159–172, 2013.

Fritz Strack, Leonard L Martin, and Norbert Schwarz. Priming and communication: Social determinants of information use in judgments of life satisfaction. *European journal of social psychology*, 18(5):429–442, 1988.

Richard Thaler. Mental accounting and consumer choice. *Marketing science*, 4(3):199–214, 1985.

Richard H Thaler. *Misbehaving: The making of behavioral economics*. WW Norton Company, 2015.

Richard H Thaler. From cashews to nudges: The evolution of behavioral economics. *American Economic Review*, 108(6):1265–87, 2018.

Michael Tomasello, Malinda Carpenter, Josep Call, Tanya Behne, and Henrike Moll. Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and brain sciences*, 28(5):675–691, 2005.

Amos Tversky and Daniel Kahneman. Judgment under uncertainty: Heuristics and biases. *Science*, 185 (4157):1124–1131, 1974.

Amos Tversky and Daniel Kahneman. Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological review*, 90(4):293, 1983.

Eldad Yechiam and Dana Zeif. The effect of incentivization on the conjunction fallacy in judgments: A meta-analysis. *Psychological Research*, pages 1–9, 2023.

Ning Zhao, Liu Xin, Li Shu, and Rui Zheng. Nudging effect of default options: A meta-analysis. *Advances in Psychological Science*, 30(6):1230, 2022.

Appendix

Appendix 1 - Measures

Cognitive Reflection Test

The Cognitive Reflection Test (CRT) is a 3-item test designed to create a conflict between an automatic intuitive response and a reflective reasoned response on all three items. Together, they present a measure of a participant's engagement with their reflective capacities. All three questions present a relatively simple mathematical question which has an intuitive response that is wrong, and a correct response which requires some reasoning. In the original study, participants got 1.23 out of 3 correct on average (Frederick, 2005). In this study, all participants will complete the same three questions.

- **Instruction**: "On the next pages, you will be asked several questions that vary in their difficulty. Please answer as quickly and as accurately as you can."
- **Questions / Tasks**:
  - A bat and a ball cost INR 110 in total. The bat costs INR 100 more than the ball. How much does the ball cost? (answer in rupees)
  - If it takes 5 machines 5 minutes to make 5 buttons, how long would it take 100 machines to make 100 buttons? (answer in minutes)
  - In a lake, there is a patch of lotus flowers. Every day, the patch doubles in size. If it takes 48 days for the patch to cover the entire lake, how long would it take for the patch to cover half of the lake? (answer in days)
- **Proposed Analytical Procedure**: The main variable of interest is the number of correct responses. A secondary variable of interest is an indicator for responding Rs10 for the first question, 100 minutes for the second question, and 24 days for the third question.

Brañas-Garza et al. (2019) report the results of a meta-study of 118 Cognitive Reflection Test studies comprising of 44,558 participants across 21 countries including Argentina, Brazil and Colombia from the Global South. While 32% of participants answered the first task correctly, the fraction rises to 48% for the third question. The second task was answered correctly by 40% of the sample. Hence, on average, 40% of responses were correct. Looking at the total number of correct answers, they find that a third of the population lack reflective, or cognitive, abilities. Meanwhile, the remaining 62% have at least some, including 18% that provide all correct answers.

Representativeness Heuristic

Tversky and Kahneman (1983) gave a group of undergraduates the following set of descriptions about an individual named Linda: "Linda is 31 years old, single, outspoken and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations." participants were then asked to rank the probability that Linda was involved with certain organizations or occupations including the likelihood that Linda was a “bank teller”, “active in the feminist movement”, and that Linda was both “a bank teller and active in the feminist movement.” Tversky and Kahneman found that undergraduates rated the probability that Linda was both a bank teller and active in the feminist movement as higher than that of Linda being a bank teller, despite the fact that the former category of individuals is a subset of the latter. We use a modified version of this that applies the same principles in generating both the profile and the possible outcomes to be ranked. In our version, we created the profile of a ‘typical’ liberal and educated Indian woman, and a list of current occupations that capture a range of probabilities. Importantly, we combine an option of an unintuitive outcome (“Preeti is a garment retailer”), with an option that has intuitive association with the profile of Preeti - ie. the combined options appear to better ‘represent’ her, though is by definition less likely that the unintuitive outcome alone. We also reduce the total options from 8 to 5, allowing for variation in ranks but reducing the computational and logistical difficulties of sorting 8 options. This measure was developed through qualitative exercises with participants and piloted.

- **Instruction**: “On the following page, you will be asked to evaluate a profile of a hypothetical individual. Please click through to continue.”
- **Questions / Tasks**: Preeti is 30 years old; she did very well at school, is independent, and always speaks her mind. At university, she studied a degree in Politics, and was the president of the student union. She cares deeply about her community and wants them to grow and succeed. Please rank the following possibilities by how likely they are to describe Preeti, where 1 is the most likely to describe her, and 5 is the least likely to describe her.
  1. Preeti is a teacher in primary school
  2. Preeti volunteers for a local NGO
  3. Preet is a garment retailer
  4. Preeti works in a call center
  5. Preeti is a garment retailer and volunteers for a local NGO
- **Proposed Analytical Procedure**: The variable of interest is a dummy variable taking the value of 1 if the participant has selected option 5 to be more likely than option 3.

Yechiam and Zeif (2023) run a meta-analysis to investigate the effect of incentivization on the conjunction fallacy (representativeness heuristic) and found that, on average, the rate of correct choices in the non-incentivized condition (n=1895) was 45%. All studies included in Yechiam and Zeif (2023) meta-analysis are from Western samples (Europe, US and UK).

Mood Heuristic

The mood heuristic is a cognitive shortcut by which an individual’s subjective impressions about the goodness or badness of an event, object, or idea informs their judgments and decision-making. Specifically, individuals allow how they ‘feel’ at a given moment to affect more universal assessments. We model our measure on that used by Strack et al. (1988) and adapt it to the local context as it includes culturally specific reference points. The original measure uses a personal numerical prime - asking participants ‘how many dates they went on in the past month' - that is correlated with positive affect. They are then asked about their general life satisfaction. The mood prime induces a positive affective state which influences life satisfaction assessments. The control condition receives the same questions with the life satisfaction question asked first. In the original study, the 'number of dates' correlates strongly with life satisfaction when asked first and not when asked second. We substitute a similar, and more contextually appropriate mood prime by asking "How many times did someone you know do a favour for you in the last week?".

- **Instruction**: "You will now be asked to think about a situation, followed by several questions. Please answer as carefully and accurately as possible"
- **Questions / Tasks**:
  - No prime condition:
    - How happy are you with life in general? Please rate your general happiness on a scale from 1 to 11, where 1 is 'not so happy' and 11 is 'extremely happy' - you can choose any number between 1 and 11.
    - How happy do you feel right now? Please rate your happiness right now on a scale from 1 to 11, where 1 is 'not so happy' and 11 is 'extremely happy' - you can choose any number between 1 and 11.
    - How many times did someone you know do a favour for you in the last week?
  - Mood prime condition:
    - How many times did someone you know do a favour for you in the last week?
    - How happy are you with life in general? Please rate your general happiness on a scale from 1 to 11, where 1 is 'not so happy' and 11 is 'extremely happy' - you can choose any number between 1 and 11.
    - How happy do you feel right now? Please rate your happiness right now on a scale from 1 to 11, where 1 is 'not so happy' and 11 is 'extremely happy' - you can choose any number between 1 and 11.
- **Proposed Analytical Procedure**: Given that this analysis, unlike others, relies on the relationship of the numeric prime response to the life satisfaction outcome conditional on treatment, rather than simply treatment condition, it warrants it's own specification. We use the following OLS models:

No student condition:

$$Y_i = \beta_0 + \beta_1 Favors_i + \beta_2 Favors_i Treat_i + \beta_3 Treat_i + \epsilon_i \tag{4}$$

Controlling for student condition:

$$\begin{aligned} Y_i = \beta_0 + \beta_1 Favors_i + \beta_2 Treat_i + \beta_3 Student_i + \beta_4 Favors_i Treat_i + \beta_5 Student_i Treat_i \\ + \beta_6 Student_i Favors_i + \beta_7 Student_i Favors_i Treat_i + \epsilon_i \end{aligned} \tag{5}$$

Y_i is the outcome variable for participant i. In this measure, the outcome variable is the answer to the general life happiness question. The treatment variable, $Treat_i$, is a dummy variable equal to 1 for the Mood prime condition and 0 otherwise. The $Favors_i$ variable is equal to the answer to the priming question: the number of times that people did favors for participant i. To analyze this data, we will examine whether or not the coefficients for $Favors_iTreat_i$(β_2 in equation 4) and for $Student_iFavors_iTreat_i$ (β_7 in equation 5) are significant. If β_2 is significant, the treated participants are significantly different. If β_7 is significant, there is a significant difference between students and non-students within the treated participant pool.

Schimmack and Oishi (2005) review studies that experimentally manipulated the order of domain satisfaction and life satisfaction items. They computed effect sizes (d) for 16 item-order comparisons, and the 16 effect sizes were subjected to a meta-analysis. This analysis confirmed a statistically significant effect of item-order and an effect size in the small to moderate range ($d = 0.29, r = 0.15$). All studies reviewed by Schimmack and Oishi (2005) are from Western samples (Europe and the US) with the exception of one study from China.

Anchoring

Anchoring refers to the tendency of individuals to overweight the first piece of information available to them in making decisions (Tversky and Kahneman, 1974). In the classic experiment, Tversky and Kahneman asked participants to compute the product of the numbers 1 through 8 in five seconds. Because participants were unable to solve the task in five seconds, they had to estimate the rest of the product using the information available to them – namely, the value of the product up to the point where they had stopped. Subject who were asked to multiply 1 through 8 produced much smaller estimates of the final product than subjects who were asked to multiply 8 through 1. Because subjects in the first condition began by multiplying much smaller numbers together, they relied on a smaller anchor to produce their final estimate. We use a modified version of this task in which we ask participants to give a price estimate for a particular product after being exposed to either a high or low anchor. Procedure: A random half of the sample is given a low-anchor to estimate the price of a Snuggie Microplush product. participants are subsequently asked to give an exact price for the product. The other half of the sample is given a high-anchor, and subsequently asked to give an exact price for the product.

- **Instruction**: "In the next phase, you will be asked to evaluate a product. Please click through to continue."
- **Questions / Tasks**:
  - Low Anchor Condition:"What is your best estimate of the price of the product below?" (Picture shown of Snuggie Microplush)
    - More than INR 700 (1)
    - Less than INR 700 (2)
  - Low Anchor Condition:"What is your exact estimate of the price?"

- High Anchor Condition:"What is your best estimate of the price of the product below?" (Picture shown of Snuggie Microplush)
  - More than INR 1500 (1)
  - Less than INR 1500 (2)
- High Anchor Condition:"What is your exact estimate of the price?"

- **Proposed Analytical Procedure**: The variable of interest is the exact estimate of price. The "high anchor condition" will be considered the treatment condition.

Li et al. (2021) fixed-effect estimate of the overall correlation coefficient between anchor number and elicited valuation is 0.286, with 95% confidence interval [0.263, 0.309]. We can use the standard formula (Borenstein and Rothstein., 2009) to convert this effect size to Cohen's d by $d = \frac{2r}{\sqrt{(1-r^2)}} = 0.596$. Most studies included in Borenstein and Rothstein. (2009) meta-analysis are from Western samples (Europe and US) with a few studies from Turkey and Taiwan.

Loss and Gain Framing

Kahneman and Tversky (1984) showed that framing the same problem as either a gain or a loss can result in subjects changing their preferences. A well-known example is the Asian Disease Problem, which we use here. In the original study, participants were asked to "imagine that the U.S. is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed..." One group of participants was given the choice to either save 200 people in a group of 600 people, or choose a probabilistic option by which there is a 1/3 probability that 600 people will be saved, and a 2/3 probability that no people will be saved. In the original study, 72 percent of participants preferred the safe option. Another group of participants were confronted with an identical problem (in terms of expected values), but were told to choose between an option in which 400 people would die, or a probabilistic scenario in which there is a 1/3 probability that nobody will die, and a 2/3 probability that 600 people will die. When framed as a loss, 78 percent preferred the risky option. Kahneman & Tversky later formalized these findings as risk aversion in the gain domain and risk-seeking in the loss domain. We modify this version slightly to make the probability as $\frac{1}{2}$ instead of $\frac{3}{4}$ and adjust the safe amount to 300 in order to make this computationally as simple as possible but maintain the equivalence of expected value - we also slightly modify the framing of probability to make it more simple, but maintain its accuracy.

- **Instruction**: "On the following page, you will be asked to read a scenario and determine a response. Please read the scenario and options carefully. Click through to continue"

- **Questions / Tasks**:
  - Gain Frame Condition:"Imagine that India is preparing for the outbreak of an unusual African disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimate of the consequences are as follows: If Program A is adopted, 300 people will be saved. If Program B is adopted, there is

a 1 in 2 chance that 600 people will be saved, and a 1 in 2 chance that no people will be saved - this is the same as if we flipped a coin - if the coin landed heads 600 people will be saved, if the coin landed tails no people will be saved.” Which of the two programs would you favor?”

  * Program A (1)
  * Program B (2)

- Loss Frame Condition: “Imagine that India is preparing for the outbreak of an unusual African disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimate of the consequences are as follows: If Program A is adopted, 300 people will die. If Program B is adopted, there is a 1 in 2 chance that nobody will die, and a 1 in 2 chance that 600 people will die - this is the same as if we flipped a coin - if the coin landed heads nobody will die, if the coin landed tails 600 people will die. Which of the two programs would you favor?

  * Program A (1)
  * Program B (2)

• **Proposed Analytical Procedure**: This should change to: The variable of interest is a dummy indicating whether a participant chose Program A in the gain framing and Program B in the loss framing.

Kühberger (1998) calculate the framing effect in risky choices by conducting a meta-analysis consisting of 136 reports with 230 single effect sizes. The overall framing effect shows that the difference between framing conditions in terms of Cohen’s d is 0.329. Most studies included in Kühberger (1998) meta-analysis are from Western samples (Europe, the US and the UK) with a few studies from China, Singapore, Japan and Israel.

Common Ratio Effect

The Common Ratio Effect describes a famous violation of the expected utility model of decision-making under risk. It was first discussed by Allais (1953) and popularized by Kahneman and Tversky (1979) in the following example. Imagine a participant is confronted with the following choice problems:

Problem 1

1. Receive $4,000 with probability 0.8 (and $0 with probability 0.2)
2. Receive $3,000 with certainty

Problem 2

1. Receive $4,000 with probability 0.20 (and $0 with probability 0.8)
2. Receive $3,000 with probability 0.25 (and $0 with probability 0.75)

Because the amounts are identical and the probabilities in Problem 2 are exactly one-fourth of those in Problem 1, expected utility predicts consistent choices in the two problems, i.e. (1.1 and 2.1) or (1.2 and

2.2). Instead, experimental evidence shows humans robustly prefer choice set (1.2 and 2.1). We test a variant of this task. The variable of interest is a dummy indicating whether a participant chose (1b and 2a).

- **Instruction**: "On the following pages, you will be asked to make hypothetical financial decisions between lotteries that have different chances of being implemented. Please pick the option that you would prefer the most."
- **Questions / Tasks**:
  1. Which option do you prefer?
     (a) 25% chance of receiving INR 250,000.
     (b) 50% chance of receiving INR 100,000.
  2. Which option do you prefer?
     (a) 2.5% chance of receiving INR 250,000.
     (b) 5% chance of receiving INR 100,000.
- **Proposed Analytical Procedure**: The variable of interest is a dummy indicating whether a participant chose (1b and 2a). A secondary variable of interest is a dummy for choosing (1a and 2b)

Blavatskyy et al. (2023) re-examine data from 39 articles with 143 experimental parameterizations of the common-ratio effect (14,909 observations). Out of these 143 designs, the standard common-ratio effect was found in 85 designs (59.4%), no common-ratio effect was found in 43 designs (30.1%), and the reverse common-ratio effect was found in the remaining 15 designs (10.5%). Their data shows that in experiments using hypothetical choices, on average, 40% of choices show the standard common-ratio effect (n=4680). Studies included in Blavatskyy et al. (2023) meta-analysis are from Europe, the US, the UK, Australia, South Africa and Thailand.

Intertemporal Choice and Present Bias

Studies on intertemporal choice have found that humans show a systematic preference for rewards to materialize sooner rather than later. In the canonical discounted expected utility model, humans discount future rewards exponentially, i.e. at a constant rate. Numerous experiments have shown that intertemporal preferences are more accurately described by a hyperbolic model (Laibson, 1997), in which the value of a reward falls rapidly in the near future, but more slowly in the distant future. For example, humans tend to prefer $10 today over $11 tomorrow, but prefer $11 in 101 days over $10 in 100 days. In other words, the initial preference for immediacy is attenuated or reversed when rewards are delayed further out into the future. We ask a similar pair of questions to all participants.

- **Instruction**: "On the following pages, you will be asked to make hypothetical financial decisions in which you would receive payments at different points in time. Please pick the option that you would prefer the most. Imagine that this money will be transferred to, or taken out of, your bank account automatically, so you do not have to keep track of it or perform any actions to get it. You simply have to decide. Please respond as if you were deciding about real money."

- **Questions / Tasks**:
  1. Which option do you prefer?
     (a) Receive INR 2000 today.
     (b) Receive INR 2500 in one month.
  2. Which option do you prefer?
     (a) Receive INR 2000 in ten months.
     (b) INR 2500 in eleven months.
- **Proposed Analytical Procedure**: The variable of interest is a dummy that is equal to 1 if a participant shows "decreasing impatience", i.e. chooses "Receive INR 2000 today" and "Receive INR 2500 in eleven months". For any other combination, the value of the dummy variable is set to zero. Secondary variables will be a dummy for "increasing impatience", and "Consistent patience".

The literature shows that, on average, people are present-biased. The mean value of the present bias parameter from meta-analyses is between 0.95 and 0.97 (Imai et al., 2021). Andreoni and Sprenger (2012) state that Ashraf et al. (2006), Dohmen et al. (2006) and Meier and Sprenger (2010) found approximately 30-35% of subjects to display present bias preferences. Hence, for the percentage of present bias subjects from meta-analyses, we use a measure of 33%. Using a similar MPL method, Andreoni and Sprenger (2012) classified 14 of 84 subjects (16.7%) as present biased. We use this figure for comparisons with the original effect sizes. Meier and Sprenger (2010) uses a field experiment in the Philippines. Thus, meta-analyses include a country from the Global South.

Defaults - Charity Selection

A default describes an option that is pre-chosen, and ultimately selected if the decision-maker does not specify an active choice. Research shows that defaults can generate meaningful differences in choices, in spite of its lack of relevance to the decision-makers' preferences or the information on the choice parameters. Changing the default option in a choice context has been known to dramatically increasing willingness to donate organs after death (Johnson and Goldstein, 2003; Dinner et al., 2011). We design our own measure modelled on those described in Smith et al. (2013), which describes a set of highly common default options for online purchases, and present a simple incentivised binary choice between two options where participants are likely to have different yet meaningful preferences to test the effect of defaults in the Indian context.

- **Instruction**: "On the following page, you will be asked to choose between two options. Please mark your choice clearly on the page. Busara and the Center for Social and Behavioral Change will be making a donation to a charity based on which charity you prefer. The charity with the most votes from people involved in this study will receive a donation at the end of this study. Which of the following two charities would you prefer to receive this donation?"

- **Questions / Tasks**:
  - Oxfam Default
    - Oxfam International - a charity focused on reducing poverty around the world [pre-selected]
    - Pratham - an Indian charity focused on improving education in India
  - Pratham Default
    - Oxfam International - a charity focused on reducing poverty around the world
    - Pratham - an Indian charity focused on improving education in India [pre-selected]
- **Proposed Analytical Procedure**: The dependent variable is selection of Pratham, and the treatment condition is the 'Pratham default' condition.

Zhao et al. (2022) conducted a meta-analysis covering 56 articles, 92 default studies, with analysis revealing that opt-out defaults lead to more pre-selected decisions than opt-in defaults do with Cohen's d of ($d = 0.59$), and 95% confidence interval of $[0.47, 0.71]$, indicating that default nudging has a considerable effect. Most studies included in Zhao et al. (2022) meta-analysis are from Western samples (Europe and the US) with a few studies from Australia and one study from the UK, Israel, Singapore and Hong Kong.

Defaults - Assessment

The inclusion of this measure follows a similar motivation to that of the previous measure. In this case, the default is not one of a binary choice, but a default selection on one end of a numerical scale. We include this because a default on a scale end could induce an effect through alternative mechanisms to the default in a binary choice setting, such as anchoring to one end of the scale.

- **Instruction**: "[Completed at the end of the session] On a scale from 1-11, where 1 means you have not enjoyed this at all, and 11 means you greatly enjoyed the session, how much have you enjoyed this session? You can choose any number between 1 and 11."
- **Questions / Tasks**:
  - 1 pre-selected/default
    - 1- Did not enjoy at all [**Pre-selected**],
    - 2,
    - 3,
    - 4,
    - 5,
    - 6 - Moderately enjoyed,
    - 7,

    * 8,
    * 9,
    * 10,
    * 11 - Greatly enjoyed
  - <u>11 pre-selected/default</u>
    * 1- Did not enjoy at all,
    * 2,
    * 3,
    * 4,
    * 5,
    * 6 - Moderately enjoyed,
    * 7,
    * 8,
    * 9,
    * 10,
    * 11 - Greatly enjoyed [**Pre-selected**]
- **Proposed Analytical Procedure**: The dependent variable is the enjoyment rating selection (between 1 and 11), and the treatment condition is the '11 pre-selected' condition.

Mental Accounting

Mental accounting describes the concept where people assign money to different 'mental accounts'. Money, however, is fungible and standard economic theory suggests it is most efficient to allocate money without restriction across all categories of spending. Kahneman and Tversky (1984) demonstrated this in a study where participants were more willing to spend $10 on a movie ticket if they had just lost $10 than if they had to replace a lost $10 movie ticket. The explanation for this is that people have separate mental accounts for transactions specifically related to the movie, and would behave differently if the value lost can be attributed to this 'account'. Henderson and Peterson (1992) repeated this measure with significant results. Banerjee et al. (2019) also used an adapted version of this measure with a very small sample in India. We use their version, which closely mirrors Kahneman and Tversky's measure, with price adjustments and a larger sample size.

Procedure: A randomly selected half of participants will be assigned a hypothetical scenario where they 'lost Rs. 200 cash' and the other half randomly assigned a hypothetical scenario where they "lost the ticket'. In each case participants are then asked if they would still purchase another ticket.

- **Instruction**: "On the following page, you will be asked to read a scenario and determine a response. Please read the scenario and options carefully. Click through to continue."
- **Questions / Tasks**:
  - Condition 1 Imagine that you have decided to see a movie where admission is Rs. 200 per ticket. As you enter the theater, you realize that you have lost Rs. 200 cash. Will you still pay Rs. 200 for a ticket for the play?
    - Yes
    - No
  - Condition 2 Imagine that you have decided to see a movie where admission is Rs. 200 per ticket. As you enter the theater, you realize that you have lost the ticket. Will you still pay Rs. 200 for a ticket for the play?
    - Yes
    - No
- **Proposed Analytical Procedure**: The dependent variable of interest is a dummy variable taking the value 1 if the participant chooses "Yes" to purchase the new ticket. The treatment variable is being assigned to the "Lost Ticket" condition.

Priolo et al. (2023) conducted a test to assess the replicability of seven studies focusing on mental accounting, which included the original lost ticket experiment. The study encompassed 5,589 participants from 21 countries including Brazil, Egypt, India and Indonesia from the Global South. Through hierarchical Bayesian meta-analysis, the results demonstrated a remarkable 100% replication rate across all seven studies. Notably, in the case of the lost ticket experiment referred to as 'Play' in Priolo et al. (2023), the effect size was found to be smaller than that reported in the original study, with the log of odds ratio being 0.53 and a 95% confidence interval of [0.42, 0.63]. Based on their data, we calculated the effect size in terms of Cohen's d to be 0.28, with a 95% confidence interval of [0.24, 0.31].

Decoy (Attraction) Effect

The Attraction Effect (also known as the Decoy Effect) was first identified by Huber et al. (1982) and describes the phenomenon by which the introduction of an irrelevant third option changes an individual's preference between two other options. We use the most popular example of this effect, due to Ariely and Jones (2008) which uses a decoy option of "print subscription" when offering a choice between a "web subscription" and a "print web subscription" to the Economist magazine. The premise is that the "print subscription" option is weakly dominated by the "print web subscription" option, and so should never be chosen, as such acting like a 'decoy'. Procedure: Half of the participants are given a set of possible preferences with a decoy option, while the other half of participants are given a set of possible preferences with no such decoy.

- **Instruction**: "On the following page, you will be asked to make a simple purchasing decision. Please pick the option that you prefer the most."

- **Questions / Tasks**:
  - Decoy Condition Which option do you prefer?
    1. A money holder from a local retailer for INR 300
    2. A bag from a local retailer for INR 800
    3. A bag and a money holder from a local retailer INR 800
  - No Decoy Condition Which option do you prefer?
    1. A money holder from a local retailer for INR 300
    2. A bag and a money holder from a local retailer INR 800
- **Proposed Analytical Procedure**: The variable of interest is a dummy variable that takes the value of 1 if a participant has selected "A bag and a money holder from a local retailer INR 800" as their preferred option. The "decoy condition" is considered the treatment group.

Heath and Chatterjee (1995) meta-analysis shows that lower-quality dominated decoy in non-durable goods increases the target share by 14.83%. Given the 79.30.% target share in the experimental group (n=399) and 64.48% target share in the control group (n=443), the Cohen's d is equal to 0.41 with 95% confidence interval of $[0.24, 0.58]$. All studies included in Heath and Chatterjee (1995) meta-analysis are from the US.

Appendix 2 - Alternative Effect Size Figures

Figure A1: Treatment Effect Sizes: Comparison with Original Effect Sizes

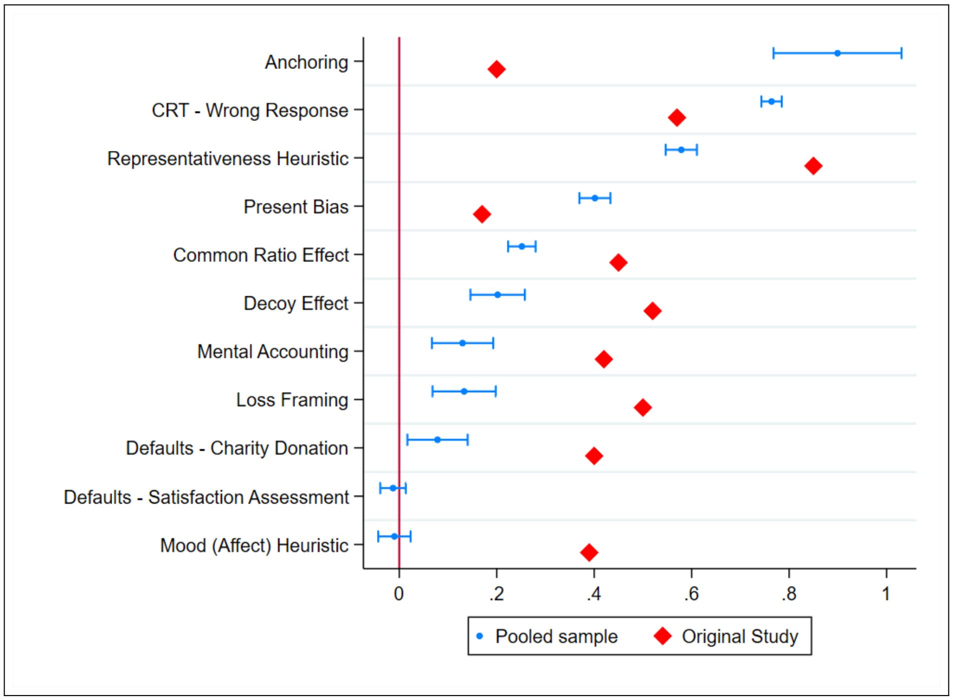

Note: Blue points represent effect sizes from Equation 2 for the main effect for each measure, and blue lines represent the 95% confidence intervals for these estimates. Notably, all effects except for Anchoring are measured in proportional terms (ie. the proportion of the sample showing the bias, or the average proportional bias for the 'Defaults - Assessment'). The Anchoring effect represents the coefficient of the anchor itself on the willingness to pay (also reasonably scalable between 0 and 1; the interpretation allows a value of 1, meaning the anchor perfectly explains the willingness to pay, and 0, that the anchor does not explain any of the willingness to pay). In all cases, the red diamond represents the effect sizes in the same units from the original studies (see Appendix 1 for the relevant references). This is intended to benchmark these effects; however, we expect that small sample, originating studies are likely to have higher effect sizes.

Appendix 3 - Individual Bar Graphs

Figure A2: Representativeness Heuristic

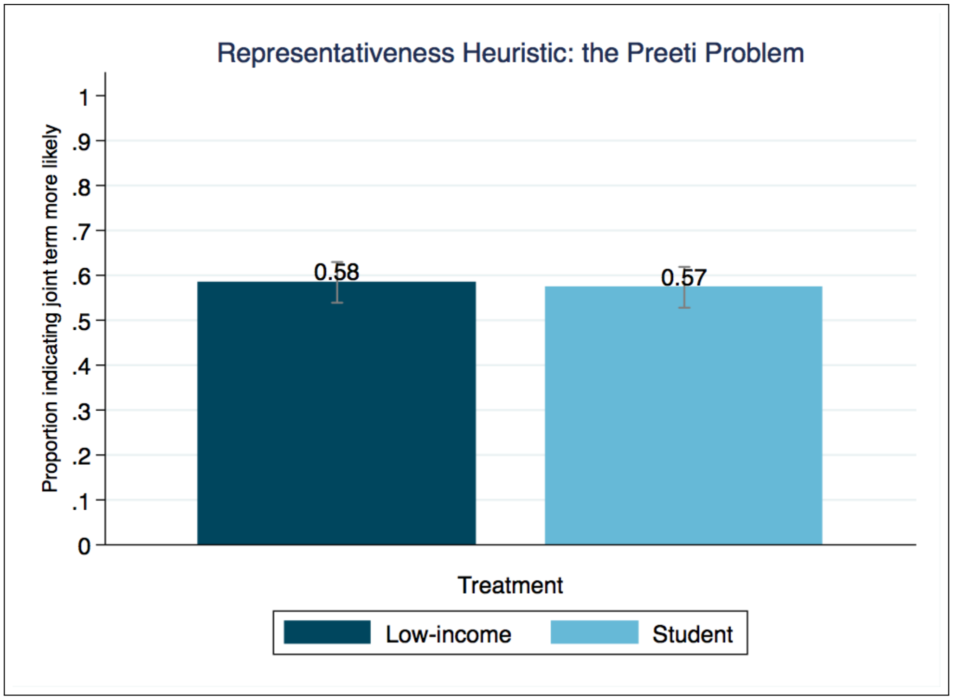

Bars represent the proportion making the error

Figure A3: Cognitive Reflection Test

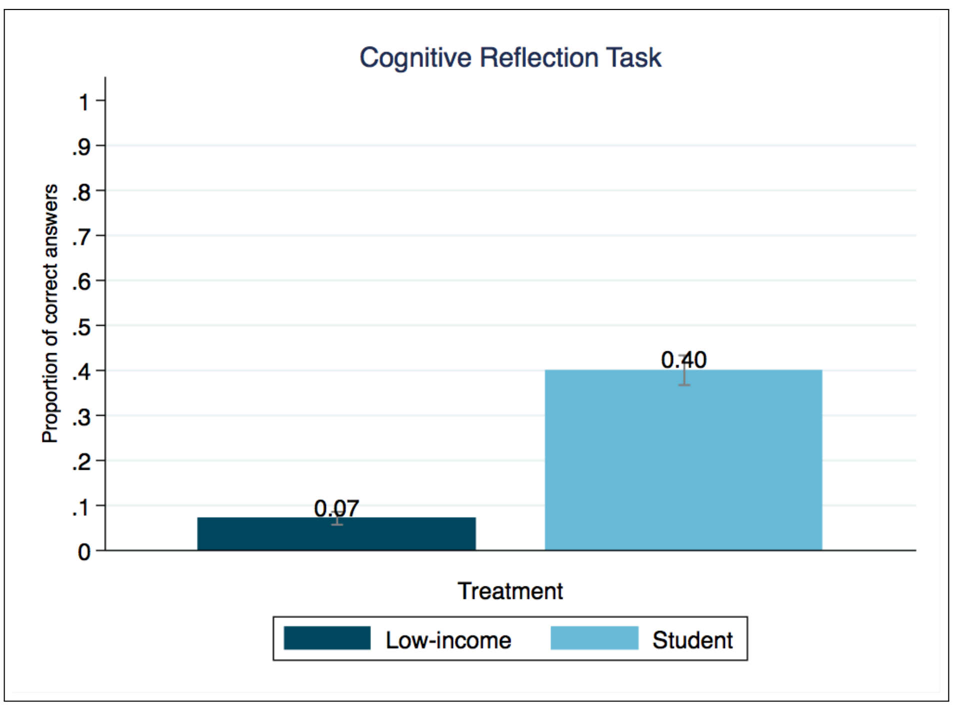

Bars represent the proportion of correct answers

Figure A4: Framing: Loss Framing

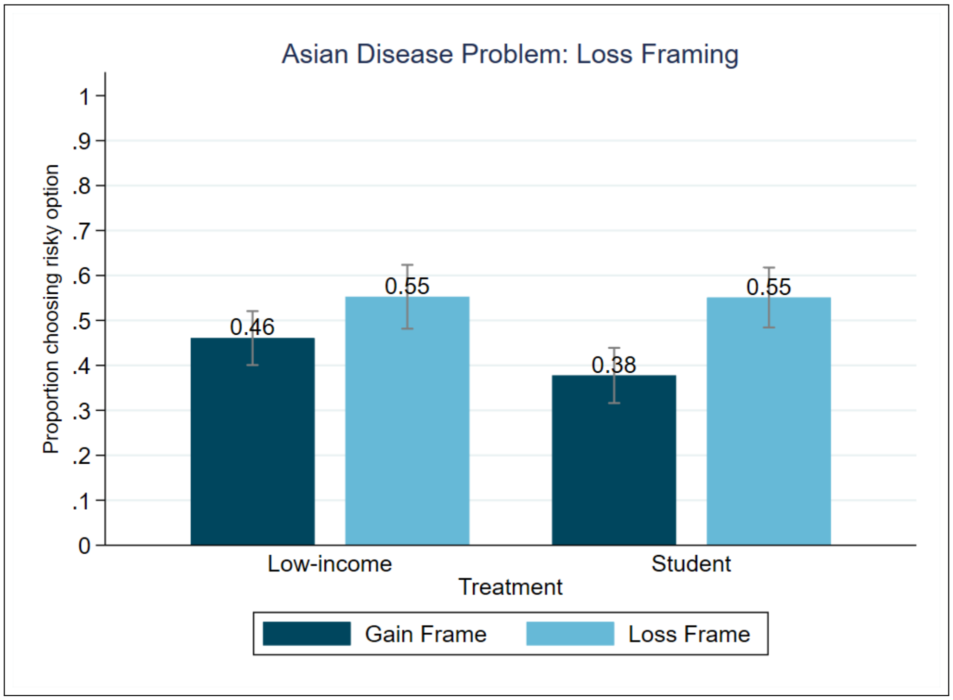

Bars represent the proportion choosing the riskier option

Figure A5: Defaults - Assessment

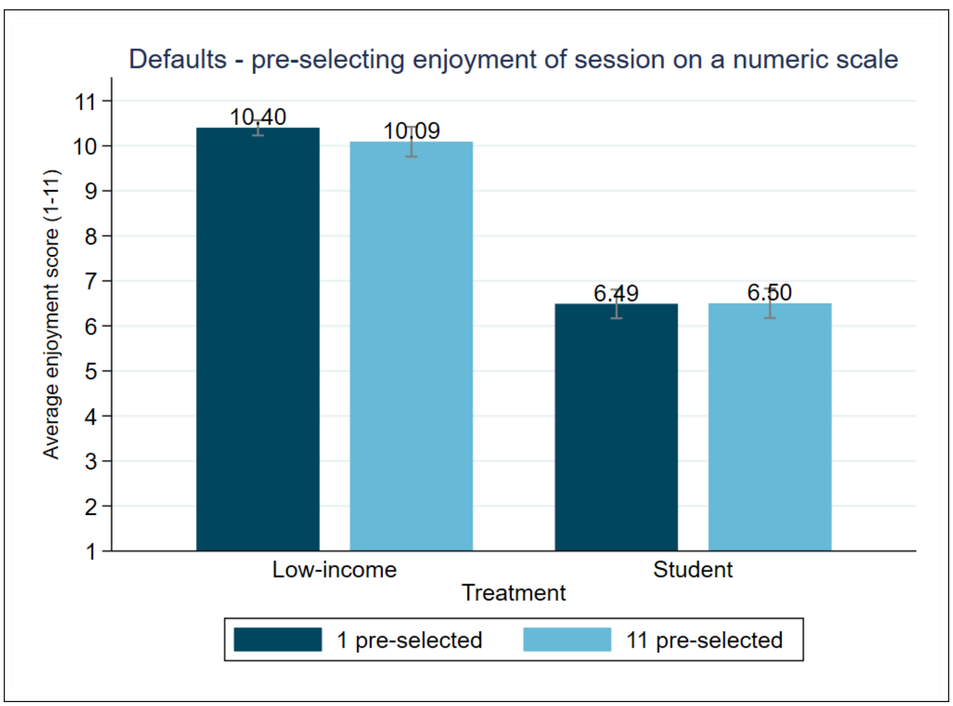

Bars represent the average assessment score for enjoyment of session (out of 11)

Figure A6: Anchoring

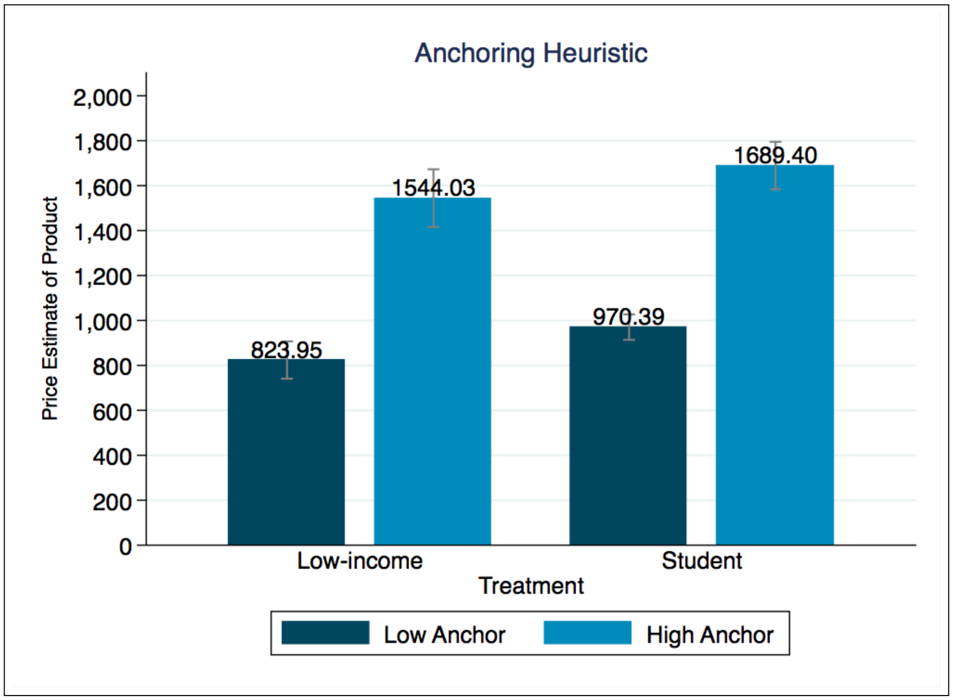

Bars represent the average predicted price of the microplush product

Figure A7: Decoy Effect

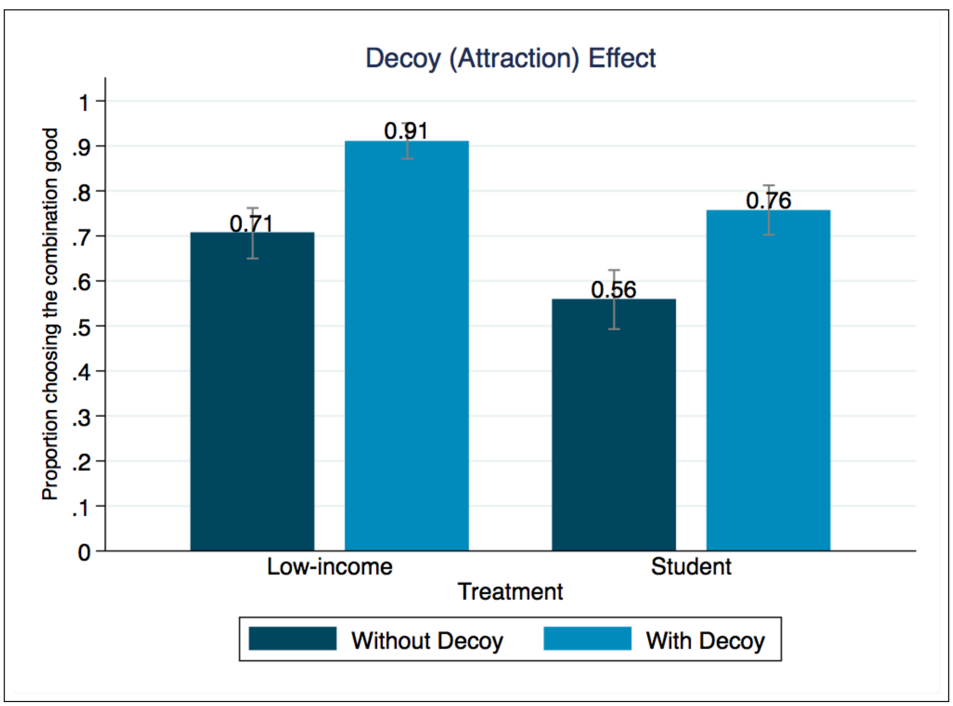

Bars represent the proportion selecting the ‘combination good’

Figure A8: Mental Accounting

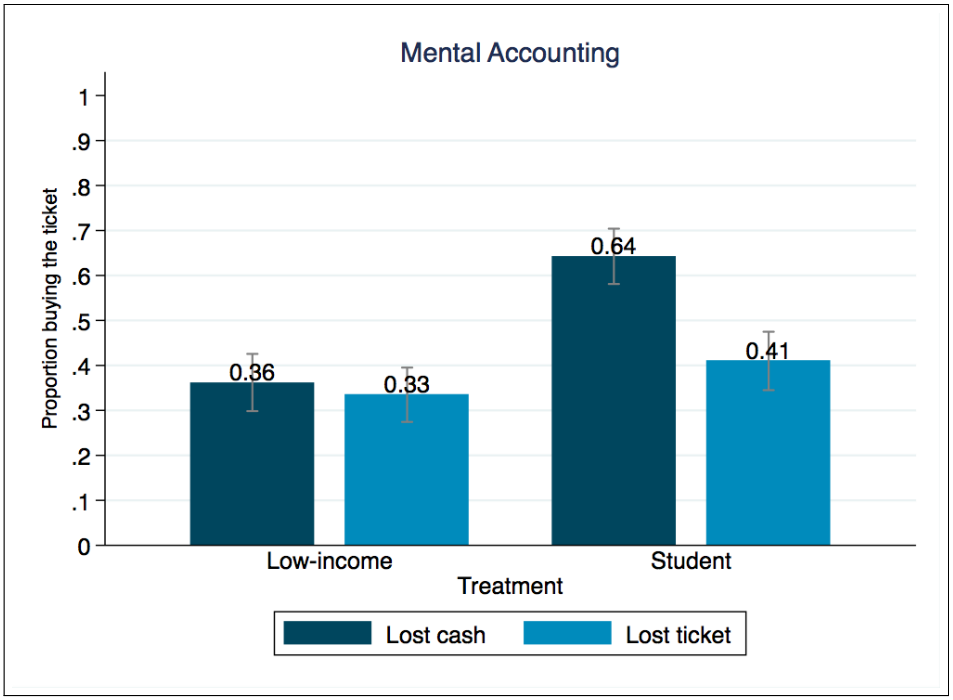

Bars represent the proportion choosing to buy the ticket again

Figure A9: Present Bias

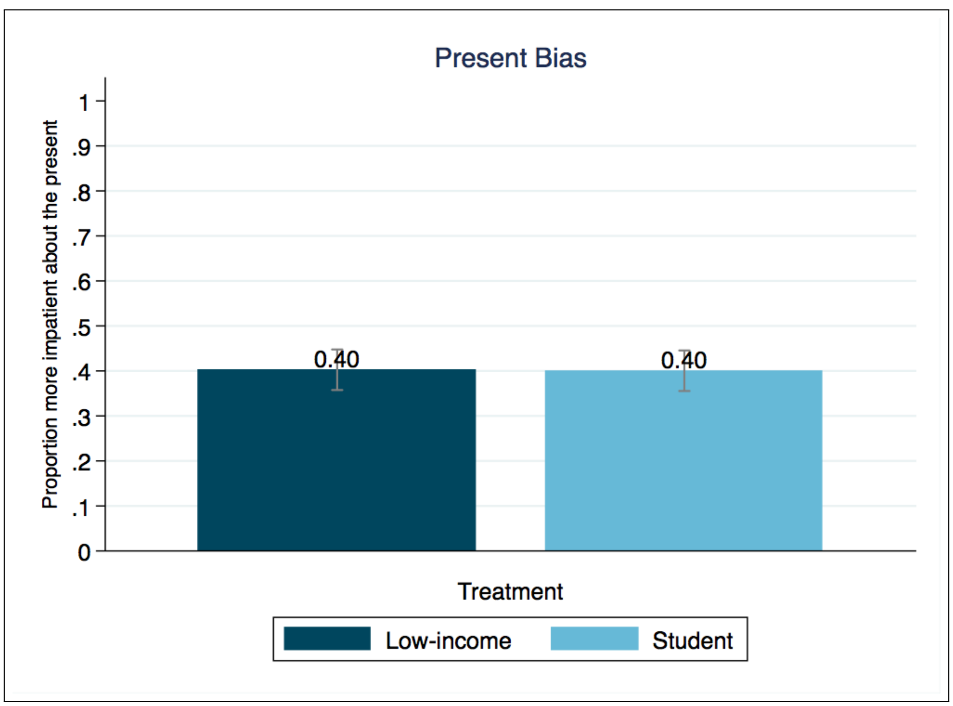

Bars represent the proportion making ‘present-biased’ choices

Figure A10: Common Ratio Effect

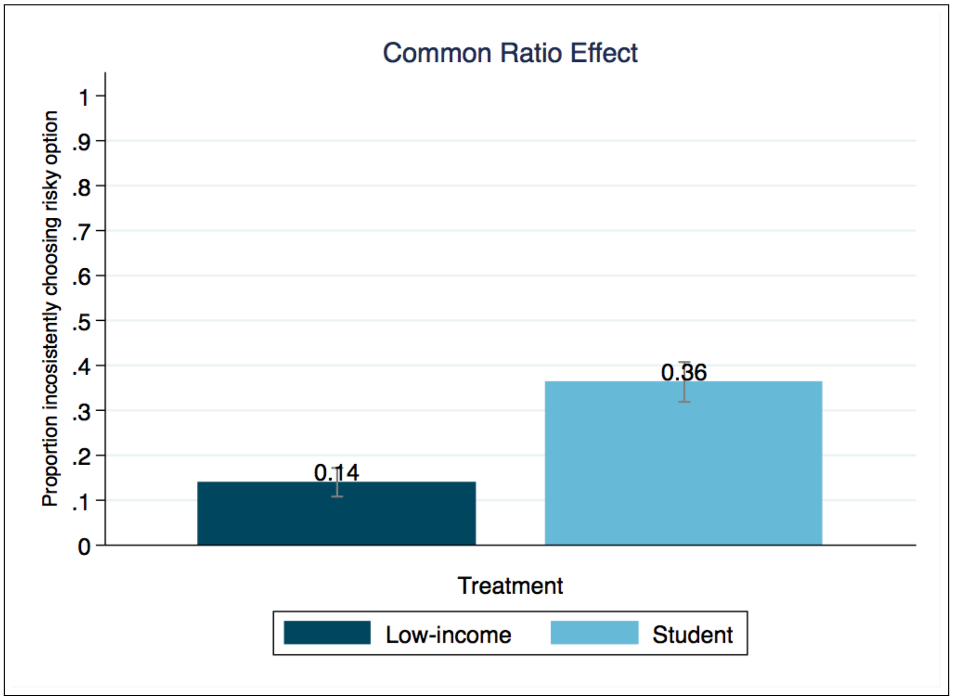

Bars represent the proportion making inconsistent choices in terms of risk preferences and the curvature of their utility function in the selections

Figure A11: Defaults - Charity Selection

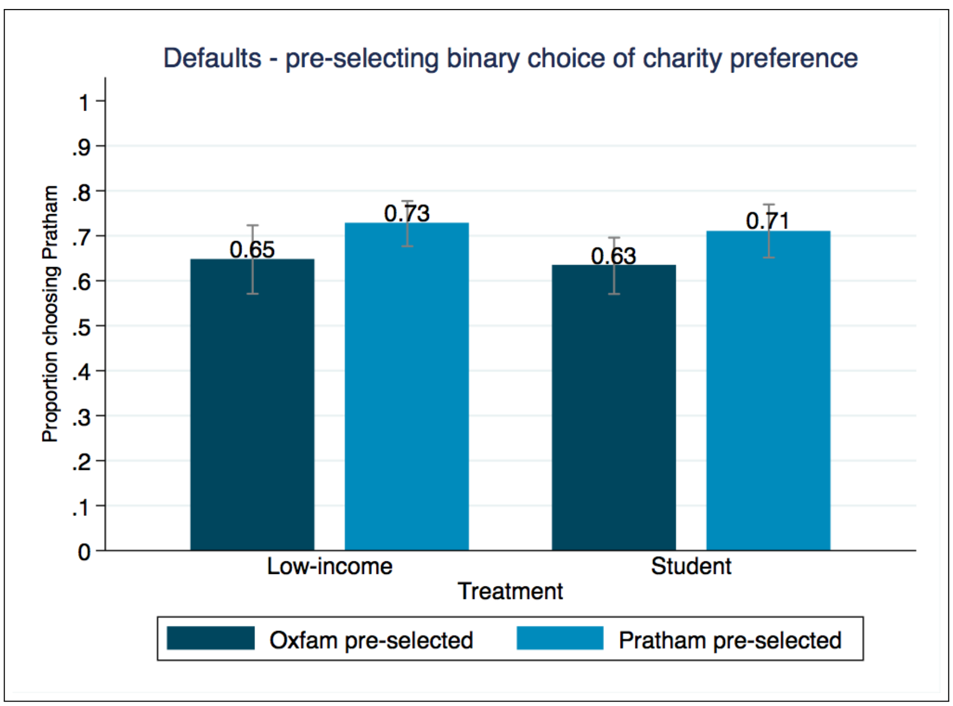

Bars represent the proportion choosing 'Pratham' charity in deciding which charity to donate money to

Figure A12: Mood Heuristic

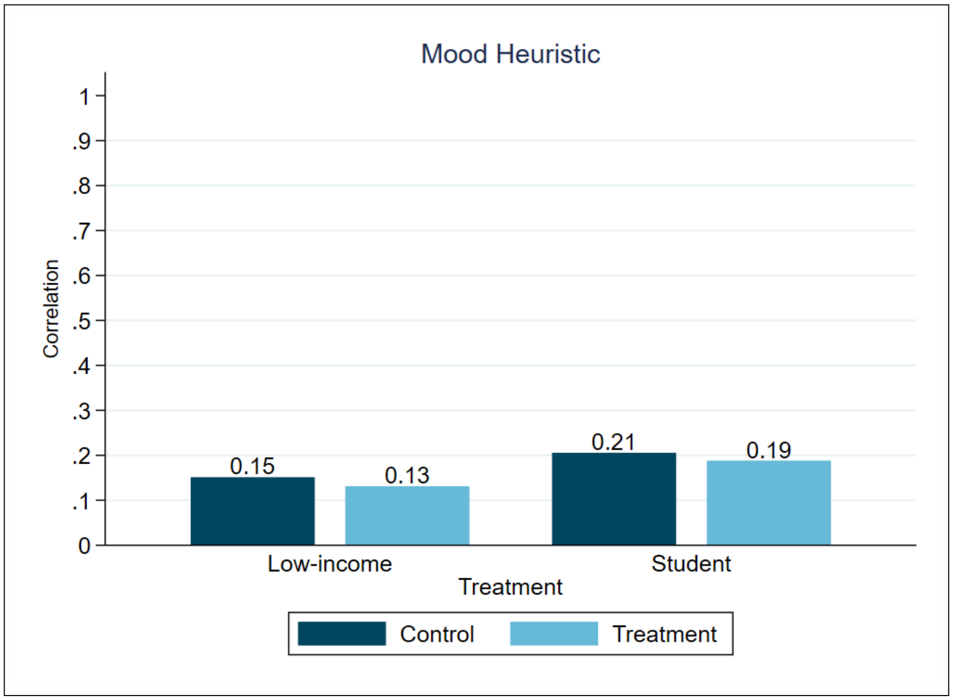

Bars represent the correlation between life satisfaction and mood prime that asks 'How many times did someone you know do a favour for you in the last week?' 13 participants are excluded as they indicated more than 20 favours.

Appendix 4 - Contextualization example

To illustrate the contextualization process, the 'Linda Problem', an original measure of the 'Representativeness Heuristic' (Tversky Kahneman, 1974), was identified and contextualized (see Appendix 1 for both versions). The original measure relies on research subjects understanding culturally-embedded characteristics of a young woman's profile such that certain selection options, even though statistically less likely due to the conjunction fallacy, appear more likely because they are 'representative' of the target. Thus, subjects should be cued that Linda is a certain 'type' of person that they are familiar with and should therefore have somewhat predictable expectations of her likelihood of engaging in different occupations on an accompanying list. Given this, the reference points in the original 'Linda Problem' are distinctly Western. They require understanding of the type of Western woman Linda might be and the likelihood of her engaging in each of a range of typically Western occupations. If the key occupation options do not seem representative due to a lack of knowledge of the context surrounding Linda, then the measure will not identify the conjunction fallacy. In pilot data in Kenya and India, responses to this measure showed signatures of a high degree of noise. Notably, there was a largely uniform distribution of selections, suggesting that subjects chose at random even when ex ante there should be options that clearly dominate others. In debriefing, participants suggested that many of the occupations did not seem 'representative' of Linda even when they were supposed to be. While a conceptual replication would allow the development of a separate measure that included the core elements of representativeness, we attempted to more closely recreate the measure to compare effect sizes to other similar measures. The core elements of the measure that we identified for preservation were: i) a short profile of a recognizable 'type' of person in the Indian context (for comparability, we decided to also use a socially engaged young woman); ii) a list of plausible occupations that could associate with that person; iii) variation in the likelihood of each occupation; iv) two occupations on the list — one that is likely and one that is unlikely - that seem incongruous, but the combination of which increase the 'representativeness'. Within this structure, a profile for 'Preeti' and a list of occupations that an Indian woman 'of this type' could reasonably undertake were developed. See Appendix 1 for the exact measure.

Appendix 5 - Main regression with controls

Table A1: Full Regression Results with controls

Panel A - Within Subjects' Identification

	(1) Overall Mean	(2)	(3)	(4) Low-income Mean	(5) Student Mean Difference	(6)	(7)
Representativeness	0.580**			0.574*	0.008		
	(0.240)			(0.251)	(0.073)		
CRT	0.230*			0.042	0.237***		
	(0.124)			(0.116)	(0.041)		
Common Ratio	0.891***			0.770*	0.177*		
	(0.296)			(0.292)	(0.069)		
Present Bias	0.205			0.208	-0.004		
	(0.250)			(0.246)	(0.078)		

Panel B - Between Subjects' Identification

	(1) Overall Treatment Effect	(2) Low-income Control Group Mean	(3) Student Control Group Mean Difference	(4) Low-income Effect	(5) Student Effect Difference	(6) Low-income Control Group Mean	(7) Student Control Group Mean Difference
Anchoring	731.339***	81.035	114.269	742.762***	-22.748	73.605	127.14
	(48.455)	(401.411)	(89.475)	(76.658)	(95.661)	-400.631	-91.767
Loss Framing	0.138**	0.672*	-0.062	0.098	0.077	0.692*	-0.102
	(0.043)	(0.286)	(0.091)	(0.070)	(0.087)	-0.284	-0.1
Mental Accounting	-0.129**	0.416	0.117	-0.018	-0.221**	0.309	0.220*
	(0.040)	(0.275)	(0.073)	(0.056)	(0.075)	-0.276	-0.088
Default-Charity	0.086*	0.441	-0.053	0.092	-0.010	0.436	-0.047
	(0.040)	(0.236)	(0.066)	(0.070)	(0.082)	-0.244	-0.092
Default-Assessment	-0.135	10.932***	-3.570***	-0.313	0.342	11.070***	-3.732***
	(0.180)	(1.479)	(0.410)	(0.219)	(0.347)	-1.493	-0.431
Decoy Effect	0.205***	0.933**	-0.220***	0.211***	-0.012	0.932**	-0.215***
	(0.027)	(0.324)	(0.049)	(0.037)	(0.054)	-0.324	-0.056
Mood Heuristic	-0.011	10.142***	-1.636***	-0.035	0.035	10.109***	-1.616***
	(0.045)	(1.863)	(0.266)	(0.071)	-0.089	-1.922	-0.414

Table reproduces Table 1 with a set of control variables including Age, Age square, Female, Education, Took-Psych-Course, Married and Home-tongue. Panel A shows the sample means, and mean differences between the Low-income and Student sub-sample for the within-subject bias measures, described in each row. Column 1 represents the mean for the full sample, while column 4 represents the same for the Low-income sample specifically, and column 5 for the mean difference between the Low-Income and Student samples. Panel B represents output from regressions of the outcomes for all between-subjects biases from Equations 2 and 3 with control variables and a dummy for 'Student'. Standard errors of coefficients in parentheses. ***, **, and * indicate significance at the 1, 5, and 10 percent critical level. Standard errors are clustered at the session level (the level of randomization).