Journal of Development Economics 120 (2016) 41–52

Contents lists available at ScienceDirect

Journal of Development Economics

journal homepage: www.elsevier.com/locate/devec

The effect of saving on risk attitudes and intertemporal choices☆

Leandro S. Carvalho [a], Silvia Prina [b,*], Justin Sydnor [c]

[a] University of Southern California, United States
[b] Case Western Reserve University, United States
[c] Wisconsin School of Business, U.W.-Madison, United States

ARTICLE INFO

Article history:
Received 18 February 2015
Received in revised form 16 December 2015
Accepted 2 January 2016
Available online 14 January 2016

Keywords:
Savings
Financial access
Time preferences
Risk preferences
Intertemporal choices

ABSTRACT

We investigate whether access to savings accounts affects choices individuals make about financial risk and intertemporal tradeoffs. We exploit a field experiment that randomized access to savings accounts among a largely unbanked population of Nepalese villagers. One year after the accounts were introduced, we administered lottery-choice and intertemporal-choice tasks to the treatment and control groups. We find that the treatment is more willing to take risks in the lottery-choice task and is more responsive to changes in experimental interest rates in the intertemporal-choice task. The results on time discounting are less conclusive, but suggest that the treatment group is more willing to delay receiving money. These results suggest that access to formal savings devices has a positive feedback loop for poor families by increasing their willingness to take risks and to delay gratification.

1. Introduction

Providing poor households access to savings accounts is becoming a priority in the development agenda (Karlan et al., 2014). While the poor typically engage in some level of informal saving (Banerjee and Duflo, 2007), a growing body of empirical literature shows that access to formal savings accounts often leads to increases in asset accumulation, which in turn can increase investment activity and facilitate consumption smoothing (Ashraf et al., 2006; Brune et al., 2014; Burgess and Pande, 2005; Dupas and Robinson, 2013; Prina, 2015).

However, it is still an open question whether access to savings accounts affects how the poor make decisions about financial risk and intertemporal tradeoffs. We study this question by conducting risk and intertemporal choice tasks with participants from a previous field experiment that randomized access to formal savings accounts for poor villagers in Nepal.

The subjects for our study were participants in a field experiment (Prina, 2015) where 1236 poor households were randomly assigned into a control group or a treatment group that gained access to formal savings accounts. For most of the treatment households, this account represented their first access to a formal saving product. The treatment group used these new accounts at very high rates (making on average 2 deposits every 3 weeks), accumulating modest but meaningful account balances.[1] While these households saved some informally prior to the experiment, the treatment created a fundamental shift in the way that they could engage with saving. The savings program generated weekly visits of deposit collectors, providing not only a secure and convenient savings instrument but also distinct moments each week when participants made active and conscious decisions about how much to save.

We administered to these same control and treatment groups from the savings experiment: a) an incentivized lottery-choice task, typically used to measure risk attitudes; b) survey questions about hypothetical intertemporal choices, frequently used as a convenient measure of time discounting; and c) an incentivized intertemporal-choice task adapted from the Convex Time Budget (CTB) methodology proposed by Andreoni and Sprenger (2012).[2]

A number of mechanisms suggest the treatment group might make more risk-neutral and patient choices in these tasks. First, savings increases assets and reduces liquidity constraints, which may increase the willingness to take risks or to delay gratification. Second, work

☆ This research would not have been possible without the outstanding work of Yashodhara Rana who served as our project coordinator. This paper benefited from comments from David Atkin, Shane Frederick, Xavier Giné, Jessica Goldberg, Glenn Harrison, Mireille Jacobson, Dean Karlan, Dan Keniston, David McKenzie, Andy Newman, Nancy Qian, Laura Schechter, Dan Silverman, Matt Sobel, Charles Sprenger, Chris Udry and Dean Yang. Carvalho thanks the Russell Sage Foundation and the RAND Roybal Center for the Financial Decisionmaking, Prina thanks IPA-Yale University Microsavings and Payments Innovation Initiative and the Weatherhead School of Management, and Sydnor thanks the Wisconsin School of Business for the generous research support.

* Corresponding author at: Case Western Reserve University, Weatherhead School of Management, 11119 Bellflower Road, room 273, Cleveland, OH 44106, United States. Tel.: +1 216 368 0208; fax: +1 216 368 5039.
E-mail address: silvia.prina@case.edu (S. Prina).

[1] One year after the introduction of the savings account, the median account balance was stable at 35–40% of a household's weekly income.

[2] See Giné et al. (2012) for an alternative but similar field adaptation of the CTB in a development setting.

http://dx.doi.org/10.1016/j.jdeveco.2016.01.001
0304-3878/

on endogenous preference formation (Becker and Mulligan, 1997; Bowles, 1998) raises the possibility that access to savings accounts may change "risk preferences" and "time preferences". Work in psychology suggests behaviors like regularly saving – and in particular the sort of routine savings decisions the treatment group was asked to make – may change mental processes and cause people to spend more time envisioning future outcomes and setting consumption priorities (Strathman et al., 1994). Several recent studies argue for a combination of these types of mechanisms, suggesting that increased asset accumulation may lead poor families to behave *as if* they have lower discount rates and higher self-control.[3] Throughout we use the terms "risk attitudes" and "intertemporal choices" to denote the patterns of choices we observe in these tasks, noting that they may reflect not only "risk preferences" and "time preferences" but also other factors such as background economic circumstances.[4]

We find that access to savings accounts led to changes in risk attitudes (broadly defined). The treatment group was more willing to take risks in the lottery-choice task; those offered access to savings accounts were 4 percentage points less likely to choose the risk-free option. We also find that the treatment group was more responsive to changes in the experimental interest rate in the CTB task. These two results are consistent with the notion that those with access to savings accounts experienced less rapidly diminishing utility over the experimental rewards.

Our results on intertemporal choices are less conclusive, but suggest that the treatment group was more willing to delay receiving money.[5] In the hypothetical choice task, the treatment group was more likely to choose a larger, more delayed payment rather than a smaller, more immediate payment. In the CTB task, the point estimates suggest that the treatment group is more patient; however, these differences have large standard errors.[6] Finally, neither the control nor the treatment group is present-biased in their CTB choices, which is consistent with the findings in Andreoni and Sprenger (2012) and Augenblick, Niederle and Sprenger (2015).

To help us quantify the observed treatment–control differences in experimental choices, we use our choice data to estimate a structural utility model, similar to the approach in Andreoni and Sprenger (2012).[7] Following Andreoni et al. (2013), we assume that participants were "narrowly bracketing" when they made their experimental choices. This assumption implies that the differences in choices will be attributed to changes in underlying preference parameters (i.e., discount factors and utility curvature). It is quite possible instead that any differences in observed behavior stem from differences in liquidity constraints or un-modeled decision processes between groups. We are open to those possibilities, but nonetheless see value in using this structural estimation as a way of providing some sense of how the observed behaviors could map into economic parameters of interest. We estimate that under the assumption of narrow bracketing, access to savings accounts is associated with a decline in relative risk aversion of 5% to 7% and an increase in the annual discount rate of 2 percentage points. However, these structural estimates have large standard errors and are not statistically significant. For the control group, estimates of preference parameters show relative risk aversion in line with previous experimental studies and an annual discount rate of 26.1%.

This study contributes to the growing literature on how economic circumstances and life experiences affect attitudes toward risk and intertemporal choices. Studies have documented that growing up during the Great Depression (Malmendier and Nagel, 2011), experiences with civil war and violence (Callen et al., 2014), financial education (Lührmann et al., 2014), or experiencing a large natural disaster (Callen et al., 2014; Cameron and Shah, forthcoming; Eckel et al., 2009) affect risk attitudes and intertemporal choices. There is a more mixed literature on the effects of moderate income and spending shocks (Brunnermeier and Nagel, 2008; Carvalho et al., forthcoming; Chuang and Schechter, 2014; Dean and Sautmann, 2014; Brune et al., 2014; Meier and Sprenger, forthcoming; Tanaka et al., 2010). Our study of the causal effect of gaining access to savings accounts integrates well with this broader literature, as access to savings accounts has both a component of altering the economic circumstances for individuals and also a life-experience component coming through the practice of saving. Our study and Lührmann et al. (2014) are the only ones in this literature to take advantage of a randomized experiment.

We also add to the growing literature in development economics exploring how access to financial products shapes the lives of the poor (e.g., Bruhn and Love, 2009; Burgess and Pande, 2005; Dupas and Robinson, 2013; Kaboski and Townsend, 2005; Karlan and Zinman, 2010a, 2010b; Prina, 2015; Schaner, 2015). An interesting comparison to our work comes from Schaner (2015), who measures discount factors of couples in an experiment offering households the opportunity to open savings accounts. Schaner finds that households where couples have very different intertemporal preferences are more likely to use individual accounts with lower interest rates over joint accounts with higher interest rates. The two studies are similar in exploring the link between savings and discounting for households, but differ in that Schaner (2015) studies how time preferences affect savings decisions, while our study is focused on how savings opportunities affect time preferences. Our work suggests that there are likely at least some modest feedback loops between access to effective financial products and risk attitudes and intertemporal choices.

An open question from this paper is whether the modest differences in choices we observe are related more to how savings changes constraints (e.g., wealth effects) or to more fundamental preference changes. These different mechanisms are difficult to disentangle and we cannot offer any conclusive evidence about their relative importance. However, in our final section we discuss a few pieces of evidence that touch on this distinction and may provide some thoughts for future research into this question.

The remainder of the paper is organized as follows. Section 2 describes the background of the savings accounts experiment conducted by Prina (2015) and outlines the design of our choice tasks. Section 3 presents the reduced-form results. Section 4 discusses the potential mechanisms behind the effects and presents the structural utility estimates. Section 5 concludes.

2. Background and experimental design

2.1. The savings accounts field experiment in Nepal

Formal financial access in Nepal is very limited: only 20% of households have a bank account (Ferrari et al., 2007). That access is concentrated in urban areas and among the wealthy. In the randomized field experiment run by Prina (2015), GONESA bank made savings accounts available to a random sample of poor households in 19 slums surrounding Pokhara, Nepal's second largest city. In May 2010, a baseline survey of 1236 female household heads was conducted.[8] Then, separate public lotteries were held in each slum to randomly assign these female household heads to treatment and control groups: 626 randomly assigned to the treatment group were offered the option to open a savings account

[3] Banerjee and Mullainathan, 2010; Bernheim, Ray, and Yeltekin, 2015; Carvalho et al., forthcoming; Haushofer and Fehr, 2014; Mani et al., 2013; Shah et al., 2012; Mullainathan and Shafir, 2013; Spears, 2011; Ubfal, 2014.

[4] See Frederick et al. (2002) for an in depth discussion of this point related to time discounting and time preferences.

[5] Ogaki and Atkeson (1997) document cross sectional patterns consistent with our findings that asset accumulation may affect the intertemporal elasticity of substitution more than time discounting.

[6] Wilcoxon rank-sum tests also show differences in the distribution of overall CTB allocations between groups with a p-value of 0.04.

[7] Significant papers in the development of structural utility modeling from experimental data include Harrison et al., 2002; Andersen et al., 2008; Tanaka et al., 2010; Andreoni and Sprenger, 2012.

[8] Here female household head is defined as the female member who is taking care of the household. Based on this definition, 99% of the households living in the 19 slums were surveyed by the enumerators.

at the local bank-branch office; the rest, assigned to the control group, were not given this option. After the baseline survey was done, between the last two weeks of May and the first week of June 2010, GONESA bank progressively began operating in the slums.

These accounts have all the characteristics of any formal basic savings account offered by other commercial banks in Nepal at the time of the intervention. The bank does not charge any opening, maintenance, or withdrawal fees and it pays a 6% nominal yearly interest, similar to the average alternative available in the Nepalese market (Nepal Rastra Bank, 2011).[9] Nor do these savings accounts have a minimum balance requirement.[10] Customers can make transactions at their local bank-branch offices in the slums, open twice a week for 3 h, or at the bank's main office, located in downtown Pokhara, during regular business hours.

Table 1 shows the summary statistics of baseline characteristics. The last column in the table shows the p-values on a test of equality of means between the treatment and control groups. It reveals that randomization led to balance along all background characteristics (Prina, 2015). The women in the sample on average have two years of schooling, and they live in households with weekly income averaging 1600 Nepalese rupees (henceforth, Rs.) (~$20) and with Rs. 50,000 (~$625) in assets. On average households have 4.5 members with 2 children. Only 15% of the households had a bank account at baseline. Most households save informally, via microfinance institutions (MFIs) and savings-and-credit cooperatives, storing cash at home, or participating in Rotating Savings and Credit Associations (ROSCAs).[11] Monetary assets account for 40% of their total assets while non-monetary assets, such as durables and livestock, account for the remaining 60%. Finally, 88% of households had at least one outstanding loan (most loans are taken from ROSCAs, MFIs, and family and friends).

As Prina (2015) documents, this experiment generated exogenous variation in access to savings accounts and saving behavior. At baseline, roughly 15% of the control and treatment groups had a bank account. A year later, 82% of the treatment group had a savings account at the GONESA bank, and 77% used it actively making at least two deposits within the first year of being offered the account.[12] The average savings balance in the account after 55 weeks is almost one and a half times the household weekly income at baseline. In Prina (2015), the ITT estimate of the effect on monetary assets (in levels) is positive but not statistically significant.[13] Measures of assets are inherently noisy; consequently, the standard errors are large. Nevertheless, a visual inspection of the cumulative distribution of monetary assets suggests that the treatment group accumulated more assets than the control group — see Fig. 3A in Prina (2015). Indeed, the Kolmogorov–Smirnov and Wilcoxon rank-sum tests for equality of distributions reject the null that the asset distributions of control and treatment groups are drawn from the same population distribution.[14]

The treatment group actively used the savings account. Fig. 1 shows the average number of deposits and withdrawals in the 52 weeks prior to the administration of the experimental tasks. Over this period, accounts holders on average made 34.7 deposits and 3.7 withdrawals. These figures indicate that the accounts were used with high frequency over this period: on average, account holders made 2 deposits every 3 weeks. Fig. 1 also shows that the typical account holder accumulated and maintained a median balance of around 600 rupees over the 52 weeks prior to the administration of the experimental tasks.[15]

Table 1
Descriptive statistics by treatment status at baseline.

	Control		Coefficient on treatment dummy		Equality of means
	Means	SD	Coefficient	SE	p-Value
Characteristics of the female head of household					
Age	36.6	11.50	0.2	0.69	0.73
Years of education	2.8	2.93	0.1	0.18	0.59
Household characteristics					
Household size	4.5	1.65	0.0	0.10	0.68
Number of children	2.2	1.30	0.0	0.08	0.82
Total income last week	1.6	5.33	0.1	0.35	0.73
% of entrepreneurs	15%	0.36	2%	0.02	0.36
% owned house	83%	0.38	1%	0.02	0.55
% owned land on which house was built	77%	0.42	2%	0.03	0.44
Experienced negative income shock	40%	0.49	3%	0.03	0.26
Assets					
Total assets	43.7	50.68	4.2	3.35	0.21
Total monetary assets	13.6	37.12	3.9	2.60	0.14
% with money in a bank	16%	0.37	2%	0.02	0.46
Total money in bank accounts	4.4	23.88	2.6	1.88	0.17
% with money in a ROSCA	18%	0.38	1%	0.02	0.59
Total money in ROSCA	2.2	8.87	1.2	0.85	0.16
% with money in an MFI	55%	0.50	−3%	0.03	0.34
Total money in MFIs	4.0	19.83	−0.1	1.02	0.96
Total amount of cash at home	1.9	4.32	0.2	0.30	0.41
Total non-monetary assets	30.1	29.34	0.3	1.74	0.84
Consumer durables	25.5	25.50	0.0	1.49	1.00
Livestock	4.6	12.47	0.4	0.77	0.65
Liabilities					
Total amount owed	54.5	282.54	−5.9	12.83	0.65
% with outstanding loans	89%	0.32	2%	0.02	0.37
Net Assets	−10.7	279.55	10.1	12.79	0.43

Notes: N = 1105. Columns 1 and 2 report summary statistics for the control group. Columns 3 and 4 display the coefficient on the treatment dummy and its standard error from regressions of the variables listed in the rows on the treatment dummy and a constant. The last column reports the p-value of two-way tests of the equality of the means across the two groups. All monetary values are reported in 1000 Nepalese rupees.

In thinking about exactly what the treatment effect was in this savings experiment, it is important to move beyond just the accumulation of assets. The shift to formal savings for the treatment group involved not just a safe account with the ability to earn interest, but also regular contact with deposit collectors.

These deposit collectors arrived at scheduled days and times during the week. This regular schedule meant that treatment group engaged with savings on a regular and conscious basis. As such, it is likely that the treatment could have fundamentally changed patterns of thinking about savings relative to the control group for whom informal savings may have been less of a conscious decision and more the residual left over after consumption decisions.

2.2. Data

We use data from three household surveys: the baseline survey (N = 1236) and two follow-up surveys conducted in June and September of 2011. The first follow-up survey, conducted one year after the beginning of the intervention, include the hypothetical intertemporal-choice task (N = 1118). It also repeated the modules that were part of the baseline survey and collected additional

[9] The International Monetary Fund Country Report for Nepal (2011) indicates a 10.5% rate of inflation during the study period.

[10] The money deposited in the savings account is fully liquid for withdrawal; the savings account requires no commitment to save a given amount or to save for a specific purpose.

[11] A ROSCA is a savings group formed by individuals who decide to make regular cyclical contributions to a fund in order to build a pool of money, which then rotates among group members, being given as a lump sum to one group member in each cycle.

[12] The percentage of control households with a bank account remained at 15%.

[13] Prina (2015) also shows that the ITT estimate of the effect on monetary assets calculated using survey data is similar in magnitude to the average balance that the treatment group had in the savings account (calculated using bank administrative data).

[14] Prina (2015) also finds reallocation of expenditures across categories (e.g. more spending on education and meat and fish, and less on health and dowries), and higher ability to cope with shocks. Finally, on qualitative outcomes, she finds households report that their overall financial situation has improved.

[15] Even though the average number of deposits is larger than the average number of withdrawals, the balance stabilized around 600 rupees because most account holders deposited small amounts on a regular basis and made occasional withdrawals of larger sums.

information on household expenditures.[16] In the second follow-up survey we administered the lottery-choice and the CTB tasks (N = 1105). In Appendix Table 2 we show that attrition is not different across treatment and control.

2.3. Risk aversion and the lottery-choice task

In the lottery-choice task, subjects were asked to choose among five lotteries, which differed on how much they paid depending on whether a coin landed on heads or on tails. The lottery-choice task is similar to that used by Binswanger (1980); Eckel and Grossman (2002) and Garbarino et al. (2011). Based on a coin flip, each lottery had a 50–50 chance of paying either a lower or higher reward. The five (lower; higher) pairings were (20; 20), (15; 30), (10; 40), (5; 50) and (0; 55). The choices in the lottery task allow one to rank subjects according to their risk aversion: subjects that are more risk averse will choose the lotteries with lower expected value and lower variance.[17] Given the low level of literacy of our sample, we opted for a visual presentation of the options, similar to Binswanger (1980). Each option was represented with pictures of rupees bills corresponding to the amount of money that would be paid if the coin landed on heads or tails (see Appendix Fig. 1 for a reproduction of the images shown to subjects).

2.4. Hypothetical intertemporal choice task

In the first follow-up survey, we measured willingness to delay gratification by asking individuals to make hypothetical choices between a smaller, sooner monetary reward and a larger, later monetary reward (Benzion et al., 1989; Tversky and Kahneman, 1986). Study participants were asked to choose between receiving Rs. 200 today or Rs. 250 in 1 month. Those who chose the Rs. 200 today then were asked to make a second choice between Rs. 200 today or Rs. 330 in 1 month. Those who had chosen Rs. 250 in 1 month were asked to make a second choice between Rs. 200 today or Rs. 220 in 1 month. These hypothetical choices in the intertemporal choice task allow us to rank subjects according to their willingness to delay gratification: the more impatient subjects will be less willing to wait to receive a larger reward. We also asked a second set of questions varying the time frame (that is, in one or two months ahead choices) in order to investigate hyperbolic discounting (see Appendix Figs. 2 & 3).

We note that when possible it is clearly preferable to use incentivized tasks, which helps to motivate our use of the incentivized task described next. However, we also note that hypothetical choices have a long history in investigating attitudes to time discounting (see Frederick et al., 2002) and are a useful and convenient way of collecting additional information for this study.

2.5. Incentivized intertemporal choice task

We adapted an experimental procedure developed by Andreoni and Sprenger (2012) called the "Convex Time Budget" method (henceforth, CTB) for our sample. In the CTB, subjects receive an experimental budget and must decide how much of this money they would like to receive sooner and how much they would like to receive later. The amount they choose to receive later accrues an experimental interest rate. In practice, subjects are solving a two-period intertemporal allocation problem, choosing an allocation along the intertemporal budget constraint determined by the experimental budget and the interest rate.[18]

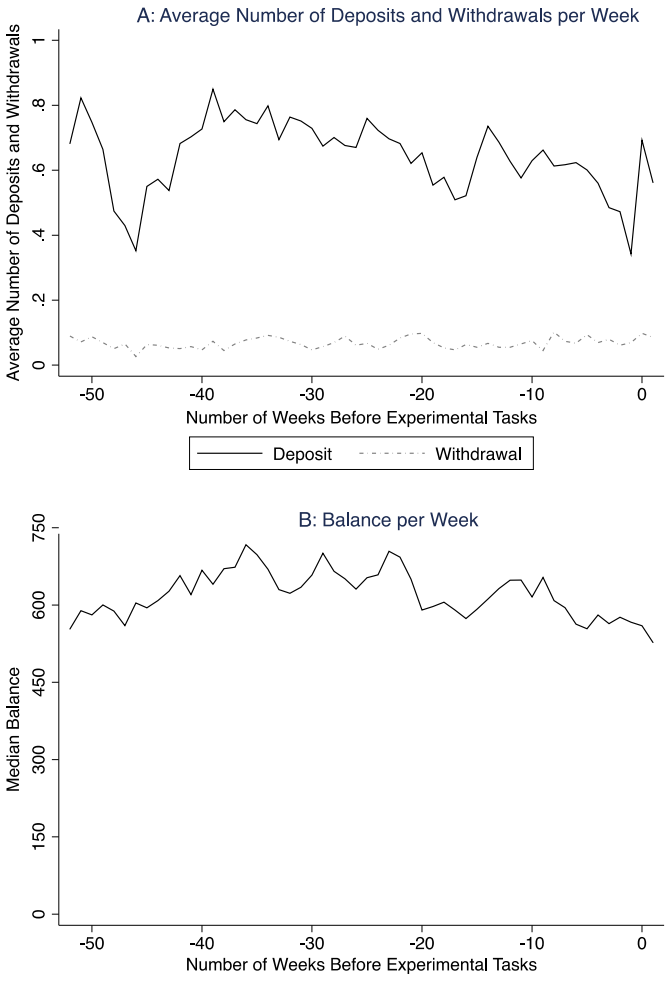

Fig. 1. Saving behavior in savings accounts. *Notes*: Panel A shows the average number of deposits (solid line) and the average number of withdrawals (dashed line) for the 52 weeks preceding the administration of the experimental tasks. Panel B shows the median balance in the savings account for the 52 weeks preceding the experimental tasks. The sample is restricted to treatment women who opened GONESA bank accounts (N = 491).

In our adaptation of the task, participants were asked to choose among three options, which corresponded to three (non-corner) allocations along an intertemporal budget constraint. The experimental endowment was Rs. 200 and the implicit experimental interest was either 10% or 20%. Subjects then were asked to make four of these choices (henceforth, games) in which we varied the time frame and the experimental interest rate. One of the four games was randomly selected for payment.

Table 2 lists the parameters of the four games and the three possible allocations in each game. In game 1, the interest rate was 10%, the earlier date was "today", and the later date was "in 1 month", so the time delay was one month. Game 2 had the same interest rate and time delay as game 1, but the earlier date in game 2 was "in 1 month". Comparing game 1 and 2 outcomes allows us to explore the possibility of present bias. Games 2 and 3 had the same time frame, but the interest rate was 10% in game 2 and 20% in game 3. Finally, in games 3 and 4 the interest rate was 20% but the time delay was 1 month in game 3 and 5 months in game 4 (in both, the earlier date was "in 1 month").

Limiting the decision in each game to a choice among three options greatly simplified the decisions subjects had to make and allowed for a visual presentation with pictures of rupee bills (see Appendix Figs. 4–7 for a reproduction of the images shown to study participants). As with the lottery-choice task, visual presentation of the options was crucial because of the low level of literacy and limited familiarity with interest

[16] Of the 1236 households interviewed at baseline, 91% (1118) were found and surveyed in the first follow-up survey. Attrition for completing the follow-up survey is not correlated with observables or treatment status (see Prina, 2015).

[17] The least risky lottery option involved a sure payout of Rs. 20, while the most risky option (0; 55) was a mean-preserving spread of the second-most risky, and thus should only be chosen by risk-loving individuals.

[18] Andreoni and Sprenger (2012) used a computer display that allowed for a quasi-continuous choice set.

Table 2
Choices for adapted Convex Time Budget (CTB) task.

Game	Interest rate	Dates		Monetary rewards (in Nepalese rupees)					
				Allocation A		Allocation B		Allocation C	
		Sooner	Later	Sooner	Later	Sooner	Later	Sooner	Later
1	10%	Today	1 month	150	55	100	110	50	165
2	10%	1 month	2 months	150	55	100	110	50	165
3	20%	1 month	2 months	150	60	100	120	50	180
4	20%	1 month	6 months	150	60	100	120	50	180

Notes: This table shows the parameters of the intertemporal choice task. Each row corresponds to a different choice ("game") participants would make among three different allocations (A, B, and C). The allocations differed in how much they paid at a sooner and a later date. The sooner and later dates and the interest rate varied across games.

rates among our sample. In addition, the enumerators were instructed to follow a protocol to carefully explain the task to participants and to have subjects practice before making their choices.[19] It is also important to note that our setup mitigates the concern that the treatment and control groups might behave differently because the treatment group has a better understanding of interest or ability to make interest calculations. The visual presentation of choice options did not require individuals to understand interest; instead, it simply offered them choices between different sums of money at different dates. Hence, while the interest rate was manipulated across choice tasks, the individuals did not have to process the interest rate themselves.

One interesting feature of the CTB method is that it allows us to investigate whether treatment and control groups respond differently to changes in the experimental interest rate or the time frame. Moreover, as we explain in greater detail in Section 4, the variations in the time frame and the interest rate permit us to estimate utility-function parameters that better quantify the observed differences in behavior across the two groups.

For both the lottery-choice and the CTB tasks, payments were made using vouchers that the participant could redeem at GONESA's main office. Each voucher contained the earliest date the money could be received. Each participant received two vouchers from the CTB task, one for her "sooner" payment and one for her "later payment"; she received another for the lottery-choice task (which could be redeemed a month later). The earnings from the two tasks were determined – according to a coin toss and a roll of a dice – only at the end of the experiment, after the participants had completed both tasks.

Because the majority of the treatment group owned GONESA bank accounts, one may worry that the transaction costs to redeem the vouchers could have been lower for the treatment group. However, there are factors that mitigate such concerns. First, all bank accounts were opened in the local bank-branches that operated in the villages/slums, not in the bank's main office where the vouchers could be redeemed. 99% of the transactions (i.e., deposits and withdrawals) over the first year took place in the local bank-branches and fewer than 25% of account holders made any transaction at the main office. Finally, concerns about GONESA having a different reputation across treatment and control groups are mitigated by the fact that both control and treatment groups were very familiar with GONESA at baseline because the NGO provides free-of-charge kindergarten in the 19 slums that participated in the study.

2.6. Experimental choices and behavior outside the experimental task

The use of experimental tasks to study attitudes toward risk and intertemporal choices raises some natural concerns about external validity but also provides for controlled measures. An alternative would be to look for real-world decisions where these attitudes are relevant. While there is clearly value in that type of analysis, real-world choices also come with identification problems because not all relevant variables are observed. Frederick et al. (2002), for example, argue that estimation of discount rates from real-world behaviors "*are subject to additional confounds due to the complexity of real-world decisions and the inability to control for some important factors*". By contrast, the controlled environment of an experimental task enables the researcher to control the constraints and the incentives in order to isolate individual differences in preferences (there is of course a concern that treatment-control differences outside the experimental task may lead to treatment-control differences in experimental choices, an issue we discuss in Section 4). Moreover, manipulations in the experimental tasks are designed to disentangle differences in time discounting from differences in the curvature of the utility function. All experimental tasks that we administered are well-established in the experimental literature.

The existing evidence suggests that experimental choices in these types of tasks predict real-world behavior (see Jaminson et al., 2012 for a review). Time preferences measures are associated with a wide-range of outcomes, such as cigarette smoking (Bickel et al., 1999), occupational choice (Burks et al., 2009), credit card borrowing (Meier and Sprenger, 2010), BMI and physical exercise (Chabris et al., 2008), and demand for commitment (Ashraf et al., 2006). Measures of risk aversion are associated with the share of financial wealth in stocks (Kimball et al., 2008), stock participation (Hong et al., 2004), and risky behaviors such as smoking, drinking, and not having insurance (Barsky et al., 1997).

Finally, there is a concern that experimental choices may not reflect subjects' preferences if they do not understand what their experimental choices entails. The protocol of the CTB task was particularly designed to mitigate this concern. As discussed above, the enumerators were instructed to carefully explain the task to subjects, who were given the opportunity to practice before making their actual choices. Second, as we discuss in Section 3.2, the evidence suggests that participants understood the experimental task; on average they were more willing to delay gratification when the interest rate was increased and less willing to delay when the waiting time was increased. More importantly, we expect that any mistakes in identifying and implementing one's preferred experimental choice to be orthogonal to treatment status.

3. Reduced-form results

3.1. Incentivized lottery choices

Fig. 2 presents the distribution over the five possible choices in the lottery-choice task, separately for the control and treatment groups. The bars are indexed by the *lower* × *higher* amounts that subjects would be paid if a coin landed on *heads* × *tails*. For example, the first bar on the left shows the fraction of subjects who chose the risk-free option that paid Rs. 20 irrespective of the coin toss. Similarly, the second bar shows the fraction who chose the lottery that paid Rs. 30 if the coin landed on heads and Rs. 15 if it landed on tails. Thus, the bars further to the right correspond to the lotteries with higher expected value and higher variance.

[19] The protocol of the experiment is described in the Appendix. Giné et al. (2012) also adapted the CTB method into an experiment in the field with farmers in Malawi. Their procedure is closer to the original CTB; they asked subjects to allocate 20 tokens across a "sooner dish" and a "later dish." Our population is less educated than the Malawi sample and thus required an even simpler design.

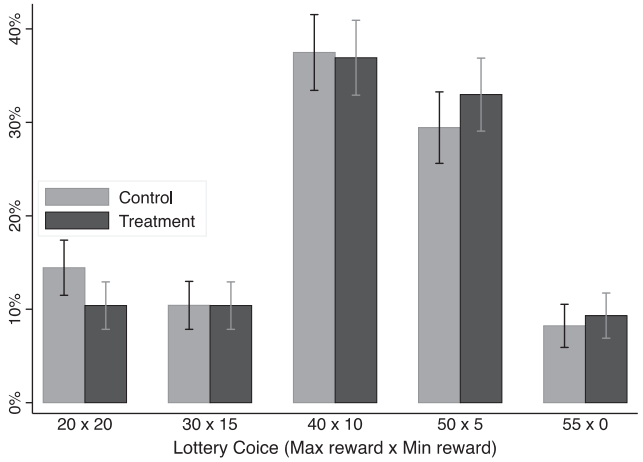

Fig. 2. Distribution of choices in lottery-choice task by treatment status. *Notes*: N = 1105. This figure shows the distribution of choices in the lottery choice task by treatment status. The two values shown below each bar correspond to the amounts subjects would get if the coin landed on heads or tails.

Fig. 2 shows that the treatment group was more willing to choose riskier lotteries. The distribution of the treatment group is shifted to the right relative to the distribution of the control group, that is, the treatment group was more likely than the control group to choose options with higher expected value and higher variance.

Table 3 complements Fig. 2 by showing choice frequencies for the treatment and control groups. The treatment group was 4 percentage points less likely to choose the risk-free option that paid Rs. 20 irrespective of the coin toss.[20] In Section 3.4 we test the statistical significance of these differences.

Later in the paper we turn to a formal structural estimation, but it is also possible to generate a rough calculation of the difference in risk-aversion parameters across the two groups. The risk choice implies bounds on the relative risk aversion from a CRRA model (that considers only experimental earnings); this can be regressed on a treatment dummy (and a constant) using an interval regression. This estimation exercise yields a CRRA parameter of 0.58 for the treatment group and 0.68 for the control group. To put this difference in perspective, we can compare it to the well-documented gender differences in lottery-choice tasks of this type. Studies such as Garbarino et al. (2011) find that women tend to have CRRA parameters around 30% higher (on average) than men; we observe a 17% difference between the treatment and control groups. Thus, the effect of the savings accounts experiment is about half of the size of the observed gender differences often discussed in the experimental literature on risk preferences.

3.2. Hypothetical intertemporal binary choices

Fig. 3 presents the distribution of the answers that subjects gave when asked to make hypothetical choices between Rs. 300 in 1 month and a larger amount in 2 months. It shows the fraction of participants who selected each of the 4 possible answers to the question. The bars are indexed by the delayed amount that subjects would require to be willing to wait. Thus, the bars further to the right correspond to responses of participants who were more willing to delay gratification.[21]

Fig. 3 and Appendix Fig. 8 (which shows the same patterns for the today vs. 1 month condition) show that the treatment group was more willing than the control group to accept delayed payments in the hypothetical intertemporal choice task. In both figures, the mass of distribution of the treatment group is shifted to the right relative to the distribution of the control group.

Table 4 echoes these results. The treatment group is roughly 5 percentage points more likely than the control group to be willing to give up Rs. 300 in 1 month in exchange for Rs. 330 in 2 months. In Section 3.4 we test the statistical significance of these differences.

Table 3
Distribution of choices in the lottery-choice task.

Choices		Distribution	
Payment conditional on coin toss		Control	Treatment
Heads	Tails		
20	20	14.4%	10.4%
30	15	10.4%	10.4%
40	10	37.5%	36.9%
50	5	29.4%	33.0%
55	0	8.2%	9.3%

Notes: N = 1105. This table reports the distribution of choices in a lottery-choice task in which subjects chose one of five lotteries that paid different amounts depending on a coin toss. The first set of columns show the contingent payments of each lottery.

3.3. Incentivized CTB choices

Fig. 4 shows the distribution of choices in the CTB experimental task for each game, separately for the control and treatment groups. It presents four sets of two bars: each set corresponds to one of the four games. The left bar in each set corresponds to the distribution of choices among the control group while the right bar corresponds to the distribution of choices among the treatment group. Each bar has two parts: a black part above the x-axis and a gray part below the x-axis. The black part corresponds to the fraction of participants who were most willing to delay gratification, choosing to delay the maximum amount of Rs. 150 (Rs. 50 sooner). The gray part corresponds to the fraction of participants who were least willing to delay gratification, delaying the minimum amount of Rs. 50 (Rs. 150 sooner).[22] Thus, an increase in the willingness to delay gratification corresponds to an increase in the black bar and/or a reduction in the gray bar.

The comparison of choices across games suggests that participants broadly understood this more complicated task. For example, subjects re-allocated significantly more money to the later date when the experimental interest rate increased from game 2 to game 3. Subjects also reallocated more money to the sooner date when the delay time increased from game 3 to game 4. Interestingly, we find no evidence of present bias. The choices in games 1 and 2 are very similar, even though the sooner date is “today” in game 1 and “in 1 month” in game 2. Andreoni and Sprenger (2012) also find no evidence of present bias when they conduct the CTB task. Augenblick et al. (2013) find that tasks involving choices over monetary rewards may be less suited to capturing present bias than tasks involving choices over real-effort-tasks.

Fig. 4 shows that while the choice patterns were broadly similar, the treatment group showed somewhat more willingness to delay gratification. The treatment group was more likely to delay the maximum amount possible of Rs. 150 and less likely to delay the minimum amount possible of Rs. 50 (with the exception of game 2). In the next section we test the statistical significance of these differences.

3.4. Statistical tests for treatment differences

In Table 5 we present regressions to estimate average treatment differences in our different tasks. For each task we first put choices into standard-deviation units, by subtracting off the mean of the outcome for the control group and dividing by the control group's standard

[20] We note that the stakes in the lottery choices task were small, roughly one-tenth of a day's wage, which mitigates the concern that the treatment group may have chosen riskier lotteries because they had a safe place – the savings accounts – to keep the task's rewards.

[21] Appendix Fig. 8 presents the distribution over the four possible choices when subjects had to choose between Rs. 200 today and a larger amount in 1 month.

[22] The fraction choosing the middle allocation can be inferred from the other two fractions.

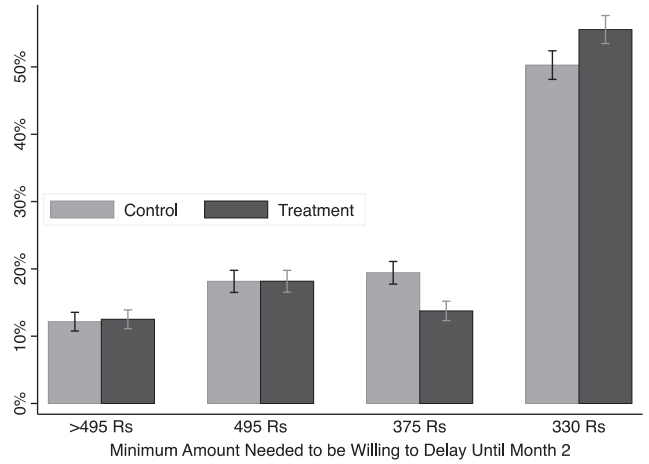

Fig. 3. Distribution of hypothetical choices between 300 Rs. in 1 month and larger amount in 2 months by treatment status. *Notes*: N = 1118. This figure shows the distribution of choices in a task in which subjects hypothetically chose between 300 Rs. in 1 month and a larger amount in 2 months. The horizontal axis shows the amount that was required for subjects to be willing to delay receiving 300 Rs.

deviation for that outcome. Because each of the tasks involves choices over a discrete number of alternatives, for this analysis we have to make a decision on a reasonable unidimensional outcome measure for each task. For the risk task we use as the outcome the expected value of the selected lottery. For the hypothetical intertemporal choice we use an indicator for being willing to accept delayed payment at the lowest interest rate as the outcome.[23] Finally, for the CTB task we use the sooner rewards allocated. The first three columns show treatment differences for the risk task and hypothetical intertemporal choices respectively.

We find that the treatment group selects lotteries that have an expected value 0.12 standard deviations higher on average than those selected by the control (p-value = 0.03). Similarly, the treatment group exhibits a willingness to delay in exchange for the lowest interest rate (10%) of 0.11–0.12 standard deviations higher than the control group (p-values = 0.05 and 0.08).

In the last column of Table 5 we estimate treatment differences for the CTB task. We combine the data from the four games and run a regression of the sooner reward on 1) a dummy for whether the sooner date is in 1 month; 2) a dummy for whether the experimental interest rate is 20%; 3) a dummy for whether the time delay between the sooner and later dates is 5 months; 4) a constant; and the interaction of these four variables with the treatment dummy. Consistent with the raw choice patterns shown in Fig. 4, we find that the control group responds to changes in the interest rate and in the delay time in the expected directions. In general, the treatment–control differences are small and not statistically significant.[24]

Table 4
Distribution of choices in the hypothetical intertemporal choice task.

Choices	Control	Treatment
Panel A: 300 Rs. in 1 month vs. larger amount in 2 months		
Willing to delay for 330 Rs.	50.3%	55.6%
Willing to delay for 375 Rs.	19.4%	13.8%
Willing to delay for 495 Rs.	18.2%	18.2%
Unwilling to delay for 495 Rs.	12.2%	12.5%
Panel B: 200 Rs. today vs. larger amount in 1 month		
Willing to delay for 220 Rs.	50.1%	55.9%
Willing to delay for 250 Rs.	23.2%	19.2%
Willing to delay for 330 Rs.	13.3%	10.9%
Unwilling to delay for 330 Rs.	13.4%	13.9%

Notes: N = 1118. This table reports the distribution of choices in two hypothetical intertemporal choice tasks. Panel A reports the choices when subjects chose between receiving 300 rupees in 1 month and a larger amount in 2 months. Panel B reports the choices when subjects chose between receiving 200 rupees today and a larger amount in 1 month. The choices in this intertemporal task allow us to rank subjects according to their willingness to delay gratification. For example, in Panel A subjects who chose 300 in 1 month versus 495 in 2 months were the least willing to accept a delayed payment. Those who chose 330 in 2 months versus 300 in 1 month were the most willing to accept a delayed payment.

We find, however, that the treatment group is more responsive than the control group to an increase in the experimental interest rate. When the experimental interest rate increases from 10% to 20%, the control group reduces the sooner reward by 0.21 standard deviations while the treatment group reduces the sooner reward by 0.33 standard deviations (p-value = 0.07).[25]

This suggests that the treatment group may be more willing to delay gratification because it has a higher intertemporal elasticity of substitution.[26] That is also consistent with the evidence that the treatment group is more likely to choose riskier options in the lottery choice task. In fact, in models with constant-relative-risk-aversion (CRRA) risk preferences, which are commonly used in the literature, a higher intertemporal elasticity of substitution corresponds to a less concave and more risk-neutral utility function.

The results in Table 5 do not control for baseline covariates. As one would expect, controlling for baseline covariates does not change our point estimates much — see Appendix Table 1. In all three experimental tasks, the differences in the average choices of the treatment and control groups have the expected sign (with some exceptions in the CTB task where there is not a strong ex-ante prediction) but are at times only marginally statistically significant. These effects likely represent a combination of moderate effect sizes and rather large standard errors. The moderate effect sizes in this experiment, which randomized access to savings accounts, are not particularly surprising considering that there may well be a range of influences beyond saving that affect risk and intertemporal-choice attitudes. Also, the need for simplicity led us to keep the choice tasks to a relatively limited set of discrete options that could be displayed visually; that may also affect our ability to detect average choice differences. It is also worth noting that the estimated treatment effects here are intent-to-treatment estimates; the difference in magnitudes would be even larger if one took into account that one-fifth of the treatment group declined the offer to open a savings account.

[23] The hypothetical intertemporal choices provide the greatest challenge to establishing a univariate outcome measure. The challenge is that the choice was "unbounded" in the sense that for a subset of subjects we do not uncover a minimum interest rate they would require to delay. If one wanted to use an alternative univariate measure such as "minimum interest rate required to delay" it would not be clear how to assign an outcome for these subjects. For this reason, we use an indicator for whether or not the subjects were willing to delay for the minimum interest rate, as a feasible univariate measure. We note, however, that this measure is imperfect as it fails to capture any differential variation in choices beyond that initial level of willingness to delay. The Wilcoxon tests presented in Table 6 do not have this problem and allow for a comparison of the full distribution of choices across treatment and control.

[24] There is some weak evidence that the control group may have more of a present bias than the treatment group. In particular, the control group decreased the sooner reward in response to a change from immediate to delayed payments while the treatment group increased, but this difference is not statistically significant. If we use an indicator for whether subjects chose a sooner reward in game 1 higher than in game 2 as our measure of present bias, we find that 24% of subjects displayed behavior consistent with present bias. However, the treatment–control difference in this measure is smaller than a tenth of a percentage point and is not statistically significant.

[25] The response to the change in interest rate from game 2 to game 3 in the CTB also provides a test for the rationality of choices by subjects. There is a clear demand prediction that the amount allocated to sooner rewards should weakly fall as the experimental interest rate rises (Giné et al., 2012; Carvalho et al., forthcoming). We find that 86% of subjects satisfy that (weak) rationality test. Interestingly, however, the treatment group is 5 percentage points more likely to satisfy this rationality test (p-value = 0.01). The observed result that the treatment group responds more to the increase in experimental interest rates than the control group is partly explained by the treatment–control difference in the likelihood of violating the law of demand. Notice that we cannot conduct similar tests of rationality for the lottery-choice or the hypothetical intertemporal choice tasks.

[26] To see this formally, we refer the reader to equation (6) in Andreoni and Sprenger (2012).

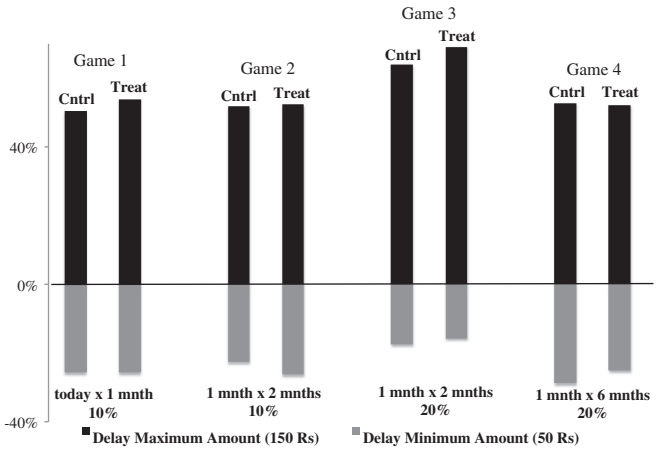

Fig. 4. Choices in the CTB task by treatment status. *Notes*: N = 1105. This figure shows the distribution of choices in the CTB experimental task, separately for the control and treatment groups. Four sets of two bars are presented, corresponding to the different games. The left bar in each set corresponds to the distribution of choices among the control while the right bar corresponds to the distribution of choices among the treatment. The black portion of each bar corresponds to the fraction of participants who were the most willing to delay gratification, choosing to delay the maximum amount of 150 rupees (50 rupees sooner). The gray area corresponds to the fraction of participants who were the least willing to delay gratification, delaying the minimum amount of 50 rupees (150 rupees sooner).

To address the broader question of whether access to savings accounts has *some effect* on attitudes toward risk and intertemporal tradeoffs, one can move from looking at differences in average choice frequencies to considering the distribution of choices more broadly. Imbens and Wooldridge (2009) argue that combining rank-sum tests with randomization-inference for the p-values (á la Rosenbaum, 2002) is one important method for determining whether observed patterns in randomized experiments imply that the treatment had an effect on the outcome of interest. In Table 6, we show the p-values from Wilcoxon rank-sum tests of differences between treatment and control for each task and for combinations of the different experimental tasks. Combining all the tasks, we see a p-value of 0.01 on the test of equality between treatment and control. That provides clear evidence of differential overall choice patterns for those given access to savings accounts.

Table 6
p-Values for randomization-inference Wilcoxon rank-sum tests.

Tests of equality in single tasks		Tests of equality across multiple tasks	
Experimental task	p-Value	Combined tasks	p-Value
Risk game	0.02	Hypothetical intertemporal 2 delays combined	0.03
Hypothetical intertemporal — today vs. 1 month	0.02	CTB (all 4 games combined)	0.04
Hypothetical intertemporal — 1 month vs. 2 months	0.05	Risk + hypothetical intertemporal	0.01
CTB game 1	0.13	Risk + CTB	0.02
CTB game 2	0.20	Hypothetical intertemporal + CTB	0.01
CTB game 3	0.03	All tasks combined	0.01
CTB game 4	0.21		

Notes: This table reports the p-values for one-sided Wilcoxon rank-sum tests (Wilcoxon, 1945) computed using (nonparametric) randomization inference (Rosenbaum, 2002). The left-hand columns show p-values for individual tasks. The right-hand columns show p-values for combined tasks. The sharp null hypothesis is that the outcomes of every study participant would have remained the same if the participant's treatment status was switched. The null hypothesis is rejected with a confidence level of 1-α if the observed Wilcoxon statistic is in the α% upper tail of the distribution (variables in which the observed ranks of treatment were smaller than the observed ranks of control were multiplied by −1). In the tests across multiple tasks, the rank-sum is calculated separately for each task and then aggregated over tasks (Rosenbaum, 1997).

3.5. Magnitudes

In Table 7 we present the results from estimation of a structural model of preferences that help us to quantify how large treatment–control differences in preference parameters would have to be to justify the reduced-form differences we observe. The derivation of the structural model follows the exposition in Andreoni and Sprenger (2012) with an adaptation to the discrete choice setting we use. This derivation is provided in the Appendix.

In particular, we follow Andreoni et al. (2013) for this exercise in assuming that subjects were "narrowly bracketing" when making these

Table 5
Standardized treatment–control differences.

	Risk	Intertemporal hypothetical		CTB
	Expected	{Willing to delay for 10% interest}		Sooner
	return	Today × 1 month	1 month × 2 months	Reward
Treatment	0.12	0.12	0.11	−0.04
	[.058]**	[.060]*	[.060]*	[.061]
{Interest rate = 20%} * treatment				−0.12
				[.067]*
{Delay time = 5 months} * treatment				0.05
				[.072]
{Sooner date = in 1 month} * treatment				0.08
				[.072]
{Interest rate = 20%}				−0.21
				[.048]***
{Delay time = 5 months}				0.27
				[.051]***
{Sooner date = in 1 month}				−0.05
				[.050]
N choices	1105	1118	1118	4408
N subjects	1105	1118	1118	1105

Notes: This table estimates effect sizes of treatment–control differences in terms of standard deviations. In the first column the outcome is the expected value of the lottery selected in the lottery-choice task. In the second and third columns the dependent variable is an indicator variable for whether the subject was willing to delay gratification (in the hypothetical intertemporal choice task) in exchange for a 10% monthly interest rate. All of these outcomes are standardized by subtracting the mean for the control group and then dividing by the standard deviation for the control group. In the last column the dependent variable in the sooner reward (minus the average sooner reward for the control group in game #1 divided by the standard deviation of the sooner reward for the control group in game #1). The omitted categories in the last column are {Sooner date = today}, {Interest rate = 10%}, and {Delay time = 1 month}. Regressions include constants but coefficients are not shown. Robust standard errors in the first three columns. Standard errors clustered at the individual level in the last column. Statistically significant coefficients are indicated as follows: *10%; **5%; ***1%.

Table 7
Maximum Likelihood estimation of preference parameters.

Parameter estimates	Convex Time Budget		Lottery choice
Annual discount factor control (δ)	0.79 [0.022]	0.79 [0.022]	–
Discount factor treatment/discount factor control	1.02 [0.037]	1.02 [0.037]	–
Risk aversion control (ρ)	0.11 [0.007]	0.11 [0.008]	0.40 [0.017]
Risk aversion treatment/risk aversion control	0.93 [0.066]	0.96 [0.075]	0.95 [0.060]
Present bias control (β)	1.00 [0.009]	–	–
Present bias treatment/present bias control	1.01 [0.013]	–	–
Includes intertemporal choice with immediate rewards?	Yes	No	–
N choices	4420	3315	1105
N subjects	1105	1105	1105

Notes: This table shows Maximum Likelihood estimates of preference parameters. The first two columns report results estimated using choices in the Convex Time Budget task while the last column reports results estimated using the choices in the lottery-choice task. The second column excludes intertemporal choices in game #1 that involved immediate rewards. The estimates correspond to the "narrow bracketing" case and assume zero background consumption incorporated in the CTB and risk choices. Standard errors are clustered at the individual level in the first two columns. Robust standard errors in the last column.

experimental choices. This assumption implies that any differences in choice patterns will be reflected in different structural preference parameters. Of course, choices could be driven instead by differences in economic circumstances or other un-modeled decision processes. We discuss these issues more in the next section. At this point we simply note that this structural-estimation exercise provides one common way of estimating the potential economic magnitude of choice patterns.

The first column shows the estimates of the annual discount factor (δ), relative risk aversion (ρ), and present bias (β) based on choices in the four CTB games. The second column shows the estimates of the annual discount factor (δ) and the relative risk aversion (ρ) when we exclude game #1 that involved immediate rewards. The last column shows a separate estimate of relative risk aversion (ρ) from the lottery-choice task. In each case, we show the parameter estimate obtained for the control group and the ratio of the treatment group's estimate to that of the control group.

We estimate the control group to have an annual discount factor of 0.79 (and an annual discount rate of 26.1%).[27] That suggests that this population strongly discounts the future. We note, however, in interpreting the discount rates that annual inflation in Nepal was above 10% during the study period (IMF, 2011), which would generally be expected to increase discount rates relative to lower-inflation environments. Interestingly, our estimates suggest less discounting of the future by the Nepalese villagers than Andreoni and Sprenger (2012) observed when they conducted the CTB with undergraduate students in the United States. We obtain a CRRA parameter for the control group in the narrow-bracketing case of 0.12, which is similar to the estimates Andreoni and Sprenger (2012) provide for their sample. This corresponds almost exactly to the original curvature that Tversky and Kahneman (1992) estimated for the value function in gains for prospect theory.

On these structural parameter estimates the standard errors are sizeable; the treatment–control differences discussed below are not statistically significant. This likely reflects a combination of: the discrete choice set we used in the CTB task, which reduced the variation available for parameter estimation relative to the continuous version; moderate effects; and inherent noise in the experimental data.

Our point estimates indicate that the treatment group is more patient than the control group. The estimated discount factor for the treatment group is 2 percent higher than that of the control group. Alternatively, the treatment group has an annual discount rate that is 2 percentage points lower than the control group's.

We find no present bias for either group, which is consistent with the choice patterns shown in Section 3.3 and with prior studies using the monetary CTB (e.g., Andreoni and Sprenger, 2012; Giné et al., 2012). However, Augenblick et al. (2013) document that while the monetary CTB tends not to reveal present bias, that it can be identified in the same populations using consumption-based tasks. As such, the lack of evidence for present bias in the monetary CTB does not necessarily imply an underlying lack of present bias.

In the second column of Table 7 we re-estimate the parameters from the CTB dropping game 1 involving immediate rewards and hence eliminating the present-bias coefficient from the estimates. The estimates on the discounting and risk-aversion parameters are unchanged when we do this.

Our point estimates also suggest that the treatment group is less risk averse than the control group. In the CTB task, the estimated (coefficient of) relative risk aversion for the treatment group is 7% lower than that of the control group. The estimates from the lottery-choice task imply similar treatment–control differences in percentage terms.[28] In the lottery-choice task, the estimated (coefficient of) relative risk aversion for the treatment group is 5% lower than that of the control group. Again, these results are consistent with the choice patterns that suggested more linear utility for the treatment group.

4. Discussion of mechanisms

Section 3 documented that the treatment and control groups made different experimental choices. In this section, we discuss two broad mechanisms through which access to savings accounts could affect risk-taking and intertemporal choice behavior. One possible mechanism is the relative *change in economic circumstances* for the treatment group (i.e., a "wealth effect"). As discussed in Section 2.1, the savings account may have enabled the treatment group to accumulate more wealth than the control group, which could affect their experimental choices. A second possible mechanism is that gaining access to savings accounts may have *changed preferences* more broadly.

As Dean and Sautmann (2014) discuss, it is challenging to disentangle these mechanisms in choice data.[29] In particular, understanding these forces depends on: 1) how subjects integrate their experimental choices with their background economic situation, and 2) how the background economic situation differed between the treatment and control groups. If participants narrowly bracket and do not consider their background consumption when making experimental choices, then differences in experimental choices can be considered to reflect differences in preferences. However, if participants integrate their choices with background consumption, then it is difficult to establish how those choices reflect preferences versus differential background economic situations.

In the behavioral economics literature, the role of "narrow bracketing" is discussed extensively and many observations of decisions in experimental tasks suggest that subjects are narrowly bracketing

[27] It is important to notice that discount rates estimated using the Convex Time Budget method depend on how subjects respond to changes in the time interval between the two payment dates; that is why the estimated discount rates are effectively lower than the experimental interest rates.

[28] However, the estimates of the level of risk-aversion are different across tasks. From the lottery-choice task, we estimate a coefficient of relative risk aversion of 0.40 for the control group, substantially higher than the estimate from the CTB task. This difference could reflect the challenges of fitting the simple CRRA functional form over varying stakes, because the CTB task had outcomes that were 5 to 10 times the size of the lottery task. Andreoni and Sprenger (2012) find the same pattern, with higher risk aversion measures, in a multiple price list lottery task than in the CTB. They posit that this may suggest that prospects with underlying risk are governed in part by an additional force beyond simple utility-of-outcome curvature.

[29] See Andersen et al. (2008) and Andreoni and Sprenger (2012) for relevant discussions on these issues.

(e.g., Carvalho et al., forthcoming; Rabin and Weizsacker, 2009; Tversky and Kahneman, 1981). One exception is Dean and Sautmann (2014), who provide evidence against narrow bracketing, showing that repeated measures of the marginal rate of intertemporal substitution of subjects in an experiment in Mali vary systematically with income, consumption, savings, and especially expenditure shocks.

In what follows, we present different pieces of evidence on these issues, but we note at the outset that *we cannot conclusively disentangle these two mechanisms*. We first show that to the extent that the control group has a lower level of background consumption, which could in principle explain why the treatment group is more willing to take risks, the treatment–control difference in background consumption is actually small. Fig. 5 shows the cumulative distribution of total household expenditures (in logs) at the time of the first follow-up survey.[30] Although the mode of the distribution of the treatment group (solid line) is shifted slightly to the right relative to the mode of the control group (dashed line), the treatment–control difference in average (log) expenditures is not statistically significant (p-value of 0.38). There is also no evidence that a buffer stock of wealth helps the treatment group to shield consumption from negative income shocks: we cannot reject that the variance of expenditures for the treatment group is equal to the variance of expenditures for the control group (p-value of 0.48). More generally, we cannot reject the null of a Wilcoxon rank-sum test, or of a Kolmogorov–Smirnov test, that the samples are drawn from the same population distribution.[31]

Second, we note that the small-stakes risk aversion observed in the experimental tasks suggests subjects were narrowly bracketing. The lottery-choice task presented subjects with a risky choice over stakes that were small relative to their income (around 3% of weekly income). If subjects were not narrowly bracketing, then they would be expected to be essentially risk neutral over these small stakes (Rabin, 2000; Schechter, 2007). Instead, less than half of the subjects chose the two lotteries with the highest expected-value.

Our subjects also failed to take advantage of a simple arbitrage opportunity, which indicates that they were not perfectly integrating: the experimental interest rate was much higher than the prevailing market interest rate (and the rate of interest the treatment group earned on their savings accounts). If individuals were integrating their background consumption into their decisions, they should have allocated all money in the CTB to the future in order to take advantage of the higher experimental interest rates and adjusted other savings opportunities accordingly. However, a substantial fraction of participants made less-than-perfectly-patient choices in the CTB, even those from the treatment group with substantial savings.

Finally, we present some evidence that choices in the CTB did not respond to variation in background circumstances. Our experimental tasks happened to fall around the *Dashain*, Nepal's most important national holiday. Because households incur major expenses in preparation for these festivities, we would expect the holiday to generate reductions in background consumption in the days leading up to the festivities, and to cause potential liquidity constraints for households without savings.[32] In Fig. 6A we show the relationship between average consumption of chicken and poultry (measured in number of days in the previous week in which household members ate chicken or poultry) and the date at which the experimental tasks were administered.

We observe a strong negative relationship between consumption and proximity to the *Dashain*: over a roughly 30-day period, households reduced their chicken and poultry consumption from approximately 2.5 days per week down to 0.5 days per week.[33,34] This pattern presumably reflects households forgoing consumption in order to finance higher consumption during the *Dashain* festivities. If differing background constraints are an important driver of experimental choices, we might expect to see different patterns of choices depending on when relative to the *Dashain* subjects completed the tasks. However, there is no evidence that proximity to *Dashain* was correlated with willingness to delay gratification in our CTB task. Fig. 6B plots the fraction of participants who in game 1 chose to receive the largest sooner reward of Rs. 150, which they could redeem on the same day, against the interview date.

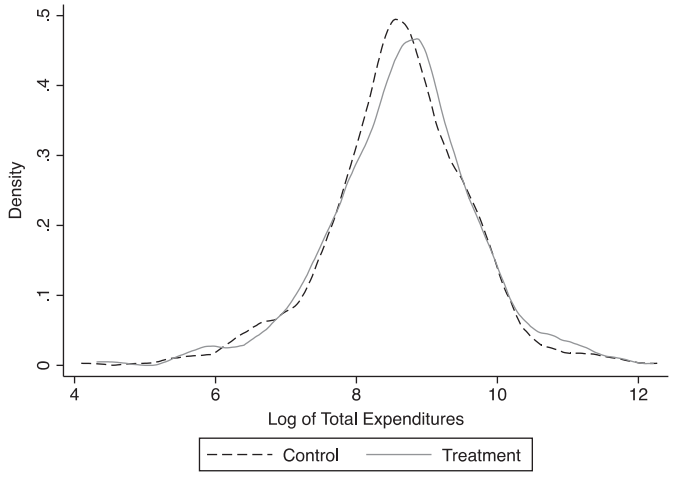

Fig. 5. Distribution of total expenditures. *Notes*: N = 1110. This figure shows the cumulative distribution of log total expenditures measured in the first follow-up survey, separately for the control and treatment groups.

5. Conclusion

We exploited a field experiment that randomized access to savings accounts in order to investigate whether attitudes toward risk and intertemporal choice are affected by the act of saving. Because the majority of the sample had never had a savings account, the experiment generated random variation in saving behavior. A year later, we administered a lottery-choice and intertemporal-choice tasks. Our findings on lottery choices and the responsiveness to interest rates in the CTB task seem to indicate that the group offered savings accounts acts *as if* they have "more linear" utility over money. The results on intertemporal tradeoffs are less conclusive, but the patterns suggest that the treatment group is more patient than the control group.

Understanding the exact mechanisms behind these differences is difficult and, as Section 4 highlighted, we can only provide suggestive evidence about them. We suspect therefore that there may be some value in more closely marrying research in economics with work in psychology that has explored how the ability to "imagine the future" affects preferences (e.g., Strathman et al., 1994; Taylor et al., 1998). For example, it seems plausible to us that the act of saving regularly may change one's frame of reference when making a whole range of choices. It may be that individuals who save regularly appear less risk averse in experimental tasks because they are more able to envision uses for larger

[30] Data on expenditures were collected only in the first follow-up survey. The module with the experimental tasks included only a few questions about how many days of the previous week household members had eaten chicken or poultry, goat or lamb, beef or buffalo, fish, or pork.

[31] p-Value on the Kolmogorov–Smirnov test is 0.26 and on the Wilcoxon rank-sum is 0.32.

[32] A household would spend money for example on new clothes and on animals, like goats and chickens, to be slaughtered as religious sacrifices.

[33] In Appendix Fig. 9 we show that there is a corresponding negative relationship between reported (average) savings at the time of the experimental tasks and proximity to the *Dashain* — even if we control for baseline reported savings (Appendix Fig. 10).

[34] We did not randomize when each participant was administered the experimental tasks, so there is a concern that the relationship in Fig. 6A could reflect baseline differences between subjects who participated in the experimental tasks at different times. Appendix Fig. 11 suggests that this is not the case. If we graph the consumption of chicken and poultry at the time of the first follow-up survey (which was in the field until approximately one month before the experimental tasks were administered) against the date of the experimental tasks, we observe no clear relationship.

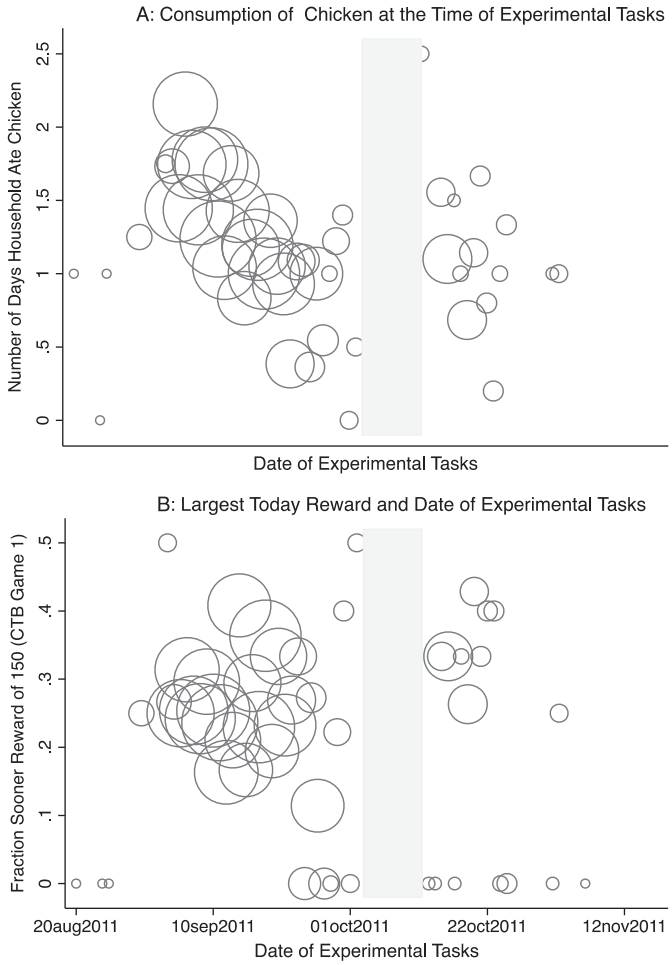

Fig. 6. A. consumption of chicken at the time of experimental tasks. B. Largest today rewards and date of experimental tasks. *Notes*: Fig. 6A plots the average consumption of chicken and poultry at the time of the experimental tasks. Fig. 6B shows the fraction of participants who chose the largest today reward of Rs. 150. The balls' circumferences correspond to the mass of participants surveyed at that given day.

sums of money. Hence, they experience less diminishing marginal utility over experimental earnings. Or, it could be that those with access to savings anticipate smoothing out experimental rewards over time in a way that those without savings do not. It could also be that access to savings affects the perception of scarcity and that those perception in turn affect cognitive function in ways that are related to risk attitudes and intertemporal choices.

Ultimately, we hope that the results of this study will motivate future research focused on better understanding the economic and psychological links between asset accumulation and economic preferences. In particular, there may be important policy implications gained with a better understanding of the potential mechanisms at play in how saving affect risk attitudes and intertemporal choices. For example, if these effects derive principally from wealth effects, then they could be replicated with one-time exogenous shocks to wealth, or with wealth transfers from the rich to the poor. However, if the effects of saving come primarily through such mechanisms as an ability to imagine the future, then the *act of saving* may be important for changing attitudes toward risk and intertemporal tradeoffs.

Appendix A. Supplementary data

Supplementary data to this article can be found online at http://dx.doi.org/10.1016/j.jdeveco.2016.01.001.

References

Andersen, S., Harrison, G.W., Lau, M.I., Rutström, E.E., 2008. Eliciting risk and time preferences. Econometrica 76 (3), 583–618.
Andreoni, J., Sprenger, C., 2012. Estimating time preferences from convex budgets. Am. Econ. Rev. 102 (7), 3333–3356.
Andreoni, J., Kuhn, M.A., Sprenger, C., 2013. On measuring time preferences. NBER Working Papers 19392.
Ashraf, N., Karlan, D., Yin, W., 2006. Tying Odysseus to the mast: evidence from a commitment savings product in the Philippines. Q. J. Econ. 121 (2), 635–672.
Augenblick, N., Niederle, M., Sprenger, C., 2015. Working over time: dynamic inconsistency in real effort tasks. Q. J. Econ. 130 (3), 1067–1115.
Banerjee, A.V., Duflo, E., 2007. The economic lives of the poor. J. Econ. Perspect. 21 (1), 141–167.
Banerjee, A., Mullainathan, S., 2010. The shape of temptation: implications for the economic lives of the poor. NBER Working Paper 15973.
Barsky, R.B., Juster, F., Thomas, K., Miles, S., Shapiro, M.D., 1997. Preference parameters and behavioral heterogeneity: an experimental approach in the health and retirement study. Q. J. Econ. 112 (2), 537–579.
Becker, G., Mulligan, C., 1997. The endogenous determination of time preference. Q. J. Econ. 112 (3), 729–758.
Benzion, U., Rapoport, A., Yagil, J., 1989. Discount rates inferred from decisions: an experimental study. Manag. Sci. 35 (3), 270–284.
Bernheim, D.B., Ray, D., Yeltekin, S., 2015. Poverty and self-control. Econometrica 83 (5), 1877–1911.
Bickel, W.K., Odum, A.L., Madden, G.J., 1999. Impulsivity and cigarette smoking: delay discounting in current, never, and ex-smokers. Psychopharmacology 146, 447–454.
Binswanger, H.P., 1980. Attitudes toward risk: experimental measurement in rural India. Am. J. Agric. Econ. 62 (3), 395–407.
Bowles, S., 1998. Endogenous preferences: the cultural consequences of markets and other economic institutions. J. Econ. Lit. 36, 75–111.
Bruhn, M., Love, I., 2009. The economic impact of banking the unbanked: evidence from Mexico. World Bank Policy Research Working Paper 4981.
Brune, L., Giné, X., Goldberg, J., Yang, D., 2014. "Facilitating Savings for Agriculture: Field Experimental Evidence from Malawi" Unpublished.
Brunnermeier, M., Nagel, S., 2008. Do wealth fluctuations generate time-varying risk aversion? Micro-evidence on individuals' asset allocation. Am. Econ. Rev. 98 (3), 713–736.
Burgess, R., Pande, R., 2005. Do rural banks matter? Evidence from the Indian social banking experiment. Am. Econ. Rev. 95 (3), 780–795.
Burks, S.V., Carpenter, J.P., Goette, L., Rustichini, A., 2009. Cognitive skills affect economic preferences, strategic behavior, and job attachment. Proc. Natl. Acad. Sci. U. S. A. 106 (19), 7745–7750.
Callen, M., Isaqzadeh, M., Long, J.D., Sprenger, C., 2014. Violence and risk preference: experimental evidence from Afghanistan. Am. Econ. Rev. 104 (1), 123–148.
Cameron, L., Shah, M., 2016. Risk-taking behavior in the wake of natural disasters. J. Hum. Resour. (forthcoming).
Carvalho, L., Meier, S., Wang, S.W., 2016. Poverty and economic decision-making: evidence from changes in financial resources at payday. Am. Econ. Rev. (forthcoming).
Chabris, C.F., Laibson, D., Morris, C.L., Schuldt, J.P., Taubinsky, D., 2008. Individual laboratory-measured discount rates predict field behavior. J. Risk Uncertain. 37 (2), 237–269.
Chuang, Y. and L. Schechter. 2014. "Stability of Experimental and Survey Measures of Risk, Time, and Social Preferences Over Multiple Years." Unpublished.
Dean, M. and A. Sautmann. 2014. "Credit Constraints and the Measurement of Time Preferences." Unpublished.
Dupas, P., Robinson, J., 2013. Savings constraints and microenterprise development: evidence from a field experiment in Kenya. Am. Econ. J. Appl. Econ. 5 (1), 163–192.
Eckel, C.C., Grossman, P.J., 2002. Sex differences and statistical stereotyping in attitudes towards financial risks. Evol. Hum. Behav. 23 (4), 281–295.
Eckel, C.C., El-Gamal, M.A., Wilson, R.K., 2009. Risk loving after the storm: a Bayesian-network study of hurricane Katrina evacuees. J. Econ. Behav. Organ. 69 (2), 110–124.
Ferrari, A., Jaffrin, G., Shreshta, S.R., 2007. Access to Financial Services in Nepal. The World Bank, Washington, D.C.
Frederick, S., Loewenstein, G., O'Donoghue, T., 2002. Time discounting and time preference: a critical review. J. Econ. Lit. 40 (2), 351–401.
Garbarino, E., Slonim, R., Sydnor, J., 2011. Digit ratios (2D:4D) as predictors of risky decision making. J. Risk Uncertain. 42 (1), 1–26.
Giné, X., J.A. Goldberg, D. Silverman, and D. Yang. 2012. "Revising Commitments: Field Evidence on Adjustment of Prior Choices." Unpublished.
Harrison, G.W., Lau, M.I., Williams, M.B., 2002. Estimating individual discount rates in Denmark: a field experiment. Am. Econ. Rev. 92 (5), 1606–1617.
Haushofer, J., Fehr, E., 2014. On the psychology of poverty. Science 344 (6186), 862–867.
Hong, H., Kubik, J.D., Stein, J.C., 2004. Social interaction and stock-market participation. J. Financ. 59 (1), 137–163.
Imbens, G.W., Wooldridge, J.M., 2009. Recent developments in the econometrics of program evaluation. J. Econ. Lit. 47 (1), 5–86.
International Monetary Fund, 2011. Nepal Country Report No. 11/319. Asia and Pacific Department.
Jaminson, J., Karlan, D., Zinman, J., 2012. Measuring risk and time preferences and their connections with behavior. Russell Sage Foundation Working Paper.
Kaboski, J., Townsend, R., 2005. Policies and impact: an analysis of village-level microfinance institutions. J. Eur. Econ. Assoc. 3 (1), 1–50.
Karlan, D., Zinman, J., 2010a. Expanding credit access: using randomized supply decisions to estimate the impacts. Rev. Financ. Stud. 23 (1), 433–446.

Karlan, D., and J. Zinman. 2010b. "Expanding Microenterprise Credit Access: Using Randomized Supply Decisions to Estimate the Impacts in Manila." Unpublished.
Karlan, D., Ratan, A.L., Zinman, J., 2014. Savings by and for the poor: a research review and agenda. Rev. Income Wealth 60 (1), 36–78.
Kimball, M.S., Sahm, C.R., Shapiro, M.D., 2008. Imputing risk tolerance from survey responses. J. Am. Stat. Assoc. 103 (483), 1028–1038.
Lührmann, M., M. Serra-Garcia, and J. Winter. 2014. "Education and Intertemporal Choice: Can Interventions Reduce Time Inconsistency?" Unpublished.
Malmendier, U., Nagel, S., 2011. Depression babies: do macroeconomic experiences affect risk taking? Q. J. Econ. 126 (1), 373–416.
Mani, A., Mullainathan, S., Shafir, E., Zhao, J., 2013. Poverty impedes cognitive function. Science 341 (6149), 976–980.
Meier, S., Sprenger, C., 2010. Presente-biased preferences and credit card borrowing. Am. Econ. J. Appl. Econ. 2 (1), 193–210.
Meier, S., Sprenger, C., 2016. Temporal stability of time preferences. Rev. Econ. Stat. (forthcoming).
Mullainathan, S., Shafir, E., 2013. Scarcity: Why Having Too Little Means So Much. Times Books, New York.
Nepal Rastra Bank, 2011. Quarterly Economic Bulletin — Mid October 2011.
Ogaki, M., Atkeson, A., 1997. Rate of time preference, interemporal elasticity of substitution, and the level of wealth. Rev. Econ. Stat. 79 (4), 564–572.
Prina, S., 2015. Banking the poor via savings accounts: evidence from a field experiment. J. Dev. Econ. 115, 16–31 (C).
Rabin, M., 2000. Risk aversion and expected-utility theory: a calibration theorem. Econometrica 68 (5), 1281–1292.
Rabin, M., Weizsacker, G., 2009. Narrow bracketing and dominated choices. Am. Econ. Rev. 99 (4), 1508–1543.
Rosenbaum, P., 1997. Signed rank statistics for coherent predictions. Biometrics 53 (2), 556–566.
Rosenbaum, P.R., 2002. Observational Studies. Springer-Verlag, New York.
Schaner, S., 2015. Do opposites detract? Intrahousehold preference heterogeneity and inefficient strategic savings. Am. Econ. J. Appl. Econ. 7 (2), 135–174.
Schechter, L., 2007. Risk aversion and expected-utility theory: a calibration exercise. J. Risk Uncertain. 35 (1), 67–76.
Shah, A., Mullainathan, S., Shafir, E., 2012. Some consequences of having too little. Science 338 (6107), 682–685.
Spears, D., 2011. Economic decision-making in poverty depletes behavioral control. B.E. J. Econ. Anal. Policy 11 (1), 1–44.
Strathman, A., Gleicher, F., Boninger, D.S., Edwards, S.C., 1994. The consideration of future consequences: weighing immediate and distant outcomes of behavior. J. Pers. Soc. Psychol. 66 (4), 742–752.
Tanaka, T., Camerer, C.F., Nguyen, Q., 2010. Risk and time preferences: linking experimental and household survey data from Vietnam. Am. Econ. Rev. 100 (1), 557–571.
Taylor, S.E., Pham, L.B., Rivkin, I.D., Armor, D.A., 1998. Harnessing the imagination: mental simulation, self-regulation, and coping. Am. Psychol. 53 (4), 429–439.
Tversky, A., Kahneman, D., 1981. The framing of decisions and the psychology of choice. Science 211 (4481), 453–458.
Tversky, A., Kahneman, D., 1986. Rational choice and the framing of decisions. J. Bus. 59 (4), S251–S278.
Tversky, A., Kahneman, D., 1992. Advances in prospect theory: cumulative representation of uncertainty. J. Risk Uncertain. 5, 297–323.
Ubfal, D. 2014. "How General Are Time Preferences? Eliciting Good-Specific Discount Rates." Unpublished.
Wilcoxon, F., 1945. Individual comparisons by ranking methods. Biometrics 1, 80–83.