Credit Constraints and the Measurement of Time Preferences*

Mark Dean[†] Anja Sautmann[‡]

January 17, 2020

Abstract

Incentivized experiments are often used to identify the time preferences of households in developing countries. We argue theoretically and empirically that experimental measures may not identify preference parameters, but are a useful tool for understanding financial shocks and constraints. Using data from an experiment in Mali we find that subject responses vary with savings and financial shocks, meaning they provide information about credit constraints and can be used to test models of risk sharing.

JEL: O16, D90, D14, C90

Experimental methods have become an important part of the applied economist's toolkit. They are regularly used to identify individual-specific preferences, in particular those that

*We would like to thank in particular Andrew Foster and Demian Pouzo for their helpful suggestions and comments, as well as David Atkin, Daniel Björkegren, Yoram Halevy, Kyle Hyndman, Greg Kaplan, Supreet Kaur, Aprajit Mahajan, Dilip Mokherjee, Sriniketh Nagavarapu, Suresh Naidu, Andriy Norets, Nancy Qian, Marta Serra Garcia, Jesse Shapiro, Duncan Thomas, Martin Uribe, David Weil and numerous seminar and conference participants. We also thank Seunghoon Na for excellent research assistance. All errors are exclusively ours. We recognize support from the Aga Khan Foundation, the Economic and Social Research Council (Development Frontiers Grant ES/K01207X/1), the Brown University Seed Fund award, and the Population Studies and Training Center at Brown University in the completion of this project. Dean gratefully acknowledges the support of NSF grant 1156090. Special thanks go to Mali Health for their help in the design and implementation of the project.

†Department of Economics, Columbia University. Email mark.dean@columbia.edu.

‡Corresponding author: Massachusetts Institute of Technology, 400 Main Street, Cambridge, MA 02141, USA. Tel +1-617-258-0667, email: sautmann@mit.edu.

govern intertemporal choices.[1] Typically, such time preference experiments measure the subject's relative valuation for money received in two different periods.[2] In order to identify underlying personal discount factors directly from experimental choices it must be assumed that these choices are divorced from outside conditions. As many have pointed out (for example Frederick et al. (2002)), without this 'narrow bracketing' assumption, experimental trade-offs may be affected by prevailing credit market conditions.

This paper makes two contributions to this literature. First, we develop a model that integrates experimental decisions with the subject's broader intertemporal optimization problem, allowing us to understand what time preference experiments tell us if narrow bracketing fails. We show that, in this case, experimental choices do not directly identify time preference parameters, but instead measure the Marginal Rate of Intertemporal Substitution (MRS) for consumption. This makes them a useful tool for understanding many other questions of interest to economists, such as the relative importance of different types of financial shocks affecting households, and their ability to cope with these shocks through insurance or intertemporal consumption smoothing. Second, we examine the model's implications in a novel panel data set of experimental choices and financial variables from poor households in Mali. We show that subject choices are correlated with financial shocks and savings, in line with our model. The findings contradict the narrow bracketing assumption and support the presence of partial credit constraints, complicating the identification of time preference parameters. The implications of our model are thus of practical, as well as theoretical interest.

The model we propose assumes a decision maker with quasi-hyperbolic preferences who may suffer income shocks as well as preference shocks that affect their marginal utility from consumption expenditure (due, for example, to loss or destruction of household assets). The effective interest rate at which households can borrow and save depends negatively on their current savings stock. This reduced-form model of 'soft' or 'partial' credit constraints is easily tractable, and can accommodate many existing models of credit rationing.

[1]See for example Ashraf et al. (2006); Tanaka et al. (2010); Mahajan and Tarozzi (2011); Schaner (2015).
[2]See Frederick et al. (2002) for a comprehensive overview.

The model predicts that, if subjects take into account their broader economic circumstance when making experimental choices, experimental trade-offs reflect the individual interest rate at the optimal level of savings in each period, which in turn equals the ratio between marginal utility of consumption today and expected discounted marginal utility of wealth tomorrow (i.e. MRS). Importantly, this conclusion does not require that subjects arbitrage the experimental payments, only that they adjust outside consumption optimally to non-experimental shocks, and take the resulting change in their "real-world MRS" into account when making choices in the experiment.

This finding suggests that experimental time-preference measures may be unsuitable for learning about time preference parameters, but instead provides a tool to learn about credit constraints and financial shocks, and to test models of consumption smoothing. For example, the covariance of experimentally measured MRS with other financial variables helps to identify the credit regime under which a household is operating. The partial credit constraints model predicts that positive income shocks decrease measured MRS, that preference shocks increase it, and that a higher stock of savings is directly linked with a lower interest rate and therefore MRS. The same relationships do not hold in the extreme cases of a household without credit constraints (the 'no constraints' model) or one that is completely unable to borrow and save (the 'complete constraints' model). The partial-constraints model implies that the relationship between MRS and measured consumption (i.e. spending) can be positive or negative, depending on the relative importance of preference and income shocks.

Our results have relevance for the large literature which uses experimental choices over dated monetary amounts to measure underlying preference parameters in both experimental and applied settings (e.g. Ashraf et al. (2006); Mahajan and Tarozzi (2011); Augenblick et al. (2015)). Our model shows that MRS measurements cannot be used to directly identify time preference parameters, and individual preference reversals are not generally indicative of time inconsistency, but can be the result of financial shocks. Moreover, in the presence of partial credit constraints, choices which exhibit present bias on average can result from time-inconsistent preferences ($\beta < 1$) only under some conditions, and conversely, their absence does not indicate time consistency. Average choices reveal underlying preference parameters only in the presence of complete constraints.

We apply our model in a unique panel data set from 1013 household in Mali, which contains three consecutive weeks of financial data and week-to-week measures of intertemporal trade-offs. We find that measured MRS responds as predicted to exogenous preference and income shocks. This rules out the narrow bracketing and 'no constraints' model. Moreover, we find a negative correlation of MRS with savings, ruling out the 'complete constraints' model in favor of partial credit constraints. To our knowledge, our paper is the first to document the simultaneous correlation of MRS with income shocks, preference shocks, and savings, allowing us to identify the credit regime that best describes our sample.

Finally we show that households are *more* impatient in periods in which they are spending more, implying that preference shocks play an important role in determining expenditure. This positive correlation is driven by expenses on adverse events and on food and necessities, identifying them as important sources of uninsured risk. Quantitatively important preference shocks not only have potential policy implications, they also make expenditure a poor proxy for household income or (marginal) consumption utility, as is for example needed in parameterized Euler equation estimates (see Deaton and Zaidi (2002)).

Our theoretical work is related to that of Pender (1996), which we extend in a number of ways, for example by allowing for preference shocks, time inconsistency and endogenous labor responses, and by showing that the results do not rely on active arbitrage (see also Cubitt and Read (2007)). Concurrent to our paper, Epper (2017) shows how a liquidity constrained subject with positive income expectations can exhibit many observed behavioral anomalies, even with standard exponential time preferences. In contrast, we focus on what can and cannot be learned from experimental choices in the face of soft credit constraints and income and preference shocks, and allowing for potential present bias.

On the empirical side, several papers investigate the variation of individual's time preference measures over time (see Chuang and Schechter (2015) for an overview). Five studies correlate time preferences with some measure of subjects' outside financial situation in a range of populations and find a relationship.[3] What sets our data apart is its level of detail, which allow us to use the relationships between MRS and financial variables to differentiate between models of credit constraints and learn what types of financial shock hit our sample population. Three other studies do not find a correlation between financial variables and measured discount rates.[4] These differ from our work in sample population, frequency of data collection and type of question asked. We provide a detailed discussion of the related literature in appendix I.

1 Integrated Choices in Time Preference Experiments

Consider the sequences of decisions shown in table 1.

[3] Harrison et al. (2005), Krupka and Stephens (2013), Carvalho et al. (2016), Ambrus et al. (2015) and Cassidy et al. (2018).

[4] Meier and Sprenger (2015), Chuang and Schechter (2015) and Giné et al. (2018).

Table 1: A Multiple Price List Experiment.

Set A		Set B	
Today (a_0)	in 1 week (a_1)	in 1 week (b_1)	in 2 weeks (b_2)
CFA 50	CFA 300	CFA 50	CFA 300
CFA 100	CFA 300	CFA 100	CFA 300
CFA 150	CFA 300	CFA 150	CFA 300
CFA 200	CFA 300	CFA 200	CFA 300
CFA 250	CFA 300	CFA 250	CFA 300
CFA 300	CFA 300	CFA 300	CFA 300
CFA 350	CFA 300	CFA 350	CFA 300
CFA 400	CFA 300	CFA 400	CFA 300

In Set A, the subject is asked to make a series of choices between receiving money today (a_0) and receiving money in one week's time (a_1), here denominated in CFA, or West African Francs (CFA 300 equal approximately USD 0.60 at market exchange rates and USD 1.60 in PPP terms). In Set B, the subject makes choices between money in one week's time (b_1) and money in two weeks' time (b_2). 'Multiple price list' (MPL) experiments of this kind have been used in many experimental investigations into time preferences. From the top to the bottom of each list, the earlier payout becomes more attractive. The parameter of interest is the point at which the subject switches to choosing the early over the late payment.

Typically, behavior in MPL experiments has been understood in the context of 'narrow bracketing' models, in which decisions in the laboratory are treated in isolation from the outside world (see for example Ashraf et al. (2006); Andersen et al. (2008); Benhabib et al. (2010)). It is assumed that subjects ignore both changes in their current outside consumption and their cost of saving and borrowing.[5] However, if narrow bracketing breaks down, outside conditions could intrude on experimental decisions. This may be especially true in a developing-country context, where households are poor and markets are incomplete, meaning that financial shocks are salient because they substantially affect the household's utility from consumption.

We propose an integrated model of experimental choices with quasi-hyperbolic discounting (in the manner of Laibson (1997)), 'soft' credit constraints, and income and preference shocks. While some of the ideas incorporated here have been previously discussed (see appendix I), as far as we are aware, ours is the first model that combines all these elements. The tractability of the model under the assumption of 'smooth' credit constraints allows us to make clear predictions about what can be learned about time preferences from MRS measurements, how measured MRS should covary with other financial variables, and what this tells us about the financial constraints and shocks that affect the household (we note, however, that the key results in section 1.5 do not depend on a differentiable interest rate function). With

[5] See Cubitt and Read (2007) for a discussion. See also Schechter (2007) on the importance of related assumptions for the measurement of risk preferences.

soft credit constraints, MRS experiments can convey more interesting information than in the case of no credit constraints (in which they simply report the market interest rate) or complete credit constraints (in which the household makes no dynamic allocation decisions). Allowing for quasi-hyperbolic discounting enables us to discuss the identification of present bias from the data.

1.1 A Motivating Example

Before developing the full model, we illustrate the main points in a simplified example. Consider a subject who lives for two periods $t \in \{0, 1\}$ and has chosen consumption in period 0 to maximize $u(c_0) + \delta u(c_1)$, subject to $c_1 = y_1 + R(s_0)$, where $c_t \geq 0$ is consumption, $y_t > 0$ is income, $u(c)$ is utility and $R(s_0)$ is the return to savings $s_0 = y_0 - c_0$. For simplicity, assume first that $R(s_0) = (1 + r)s_0$, i.e. there is a fixed interest rate. Standard first-order conditions imply that (at an interior solution) the optimal c_0 satisfies $\frac{u_0'(c_0)}{\delta u_1'(c_1)} = R'(s_0) = (1 + r)$.

Now suppose the subject is offered experimental choices from set A above. It is well known that, if the subject is allowed to arbitrage, they will prefer the earlier payment if and only if $\frac{a_1}{a_0} < (1 + r)$ - i.e. the gross interest offered in the experiment is less than the market rate. Our first result is to show that this remains true for small payments even if subjects do *not* arbitrage, as long as they take into account the true utility value of experimental payments. If experimental payments are consumed in the period they are received, the subject will prefer the earlier payment if the associated utility gain is higher, or approximately if $a_1 \delta u'(c_1) < a_0 u'(c_0)$ (where the linear approximation is close for small enough payments). Thus, the subject chooses the earlier payment if $\frac{a_1}{a_0} < \frac{u_0'(c_0)}{\delta u_1'(c_1)} = (1 + r)$ and the later payment otherwise.

The above means that choices from set A can approximately identify the MRS, regardless of whether the subject arbitrages the experimental payment. However, they do not identify the time preference parameter δ. To illustrate, consider the case in which $r = 0$, $u_i(c) = \ln(c)$ and $y_0 = y_1 = 3$. If $\delta = 1$, the subject chooses $c_0 = c_1 = 3$; if $\delta = \frac{1}{2}$, they choose $c_0 = 4$ and $c_1 = 2$. However, in both cases the MRS is $\frac{u_0'(c_0)}{\delta u_1(c_1)} = \frac{c_1}{c_0} = 1$, so experimental choices will remain the same: δ could only be recovered if the researcher also knew c_0, c_1 and u'.

In the extreme case of constant R, experimental choices are also independent of y_0, y_1, or shocks to u_0. However, we argue that credit market frictions lead to a decreasing marginal return function R'; for example, interest rates may be higher for borrowing than saving. In this case, measured MRS $\frac{u'(c_0)}{\delta u'(c_1)} = R'(y_0 - c_0)$ still does not depend directly on δ. However, an increase in y_0 (under mild assumptions) will lead to an increase in consumption and savings in $t = 0$ and reduce both the MRS and the effective interest rate R'. A similar argument applies to shocks that increase the marginal utility of consumption: they will decrease savings and so increase the interest rate. This means that the experimentally measured MRS will be positively related to preference shocks and negatively related to income and savings. We now formalize these results for the infinite time period model.

1.2 Set-Up

We model a decision maker i whose preferences are described by

$$u_i(c_{i0}, \rho_{i0}) + \beta_i E_0 \sum_{t=1}^{\infty} \delta_i^t u_i(c_{it}, \rho_{it}).$$

Utility is time-separable and given by the instantaneous utility function $u_i(c, \rho)$, where c is period consumption, and ρ is a real-valued stochastic preference parameter, drawn independently from a distribution $F_{i\rho}$ in each period. We use ρ to model the effect of preference shocks and assume that the marginal utility of consumption $\frac{\partial u_i}{\partial c}$ is everywhere increasing in ρ (see below). δ_i is the discount factor, and the parameter $0 < \beta_i \leq 1$ indexes present bias. Preferences are quasi-hyperbolic and time-inconsistent if $\beta_i < 1$.

The resource constraint is given by

$$\begin{aligned} c_{it} &= w_{it} - s_{it} \\ w_{it} &= y_{it} + R_i(s_{i,t-1}) \\ w_{i0} &\ \text{given.} \end{aligned}$$

The stock of savings at the end of period t, s_{it}, can be positive or negative. y_{it} is i's current income, drawn independently from a distribution F_{iy} in each period, and w_{it} is cash-on-hand in t. $R_i(s_{it})$ describes the gross returns to saving and thus the intertemporal budget constraint. From here on out we will suppress the person index i to ease notation.

We assume that R is strictly increasing, concave, and continuously differentiable. This implies that the resources available in period $t+1$ are increasing in s_t, but the marginal returns to saving fall as savings s_t increase, or equivalently that the cost of borrowing rises with the amount of credit. We refer to this as the *partial (credit) constraints* model.

The shape of R is a reduced-form way of modeling the (potentially individual-specific) credit and savings constraints that households in developing countries face. Decreasing returns to savings can arise from diminishing returns to capital in household production, capital market imperfections, or a finite supply of financial assets. The 'classic' liquidity-constraint model with a hard borrowing constraint but unrestricted savings (which includes "storing money under the mattress") is a limit case of this class of return function. Moreover, it can be shown that the key predictions of the model are robust to a piece-wise linear R (i.e. with points of non-differentiability) and infinite slope.

In some settings, the assumption of decreasing returns may not hold: high-return durable assets or starting a business may require a minimum investment, or formal financial instruments may offer better terms for borrowing or saving larger amounts. We do not believe these to be important factors in our data, not least because our population has little access to formal financial instruments, but they may be relevant in other settings. Appendix A discusses our justification of the shape of the return function in more detail.

The curvature of R indexes the degree to which the consumer is credit constrained: the more concave the function, the more the rate of return varies with the amount saved or borrowed. At one extreme, R is globally linear and equal to $1+r$. We call this the *no-constraints case*. At the other extreme, as the second derivative $R''(0) \to -\infty$, the cost of borrowing goes to infinity, while the rate of return on savings goes to zero. In the limit, no borrowing or savings are possible. We call this the *complete-constraints case*. While there is some evidence that savings constraints exist (see appendix A), our primary aim in focusing on the extreme case of complete constraints is expository: as we discuss in section 1.6.1, the complete constraints model is unique in that it allows for the identification of time preferences using average experimental choices.

1.3 The Euler Equation and Marginal Rate of Intertemporal Substitution

We use the results of Harris and Laibson (2001) to identify an Euler equation for the quasi-hyperbolic consumer of our model. These authors characterize the set of perfect equilibria in stationary Markov strategies of the game between the different 'selves' of the consumer in different periods. We assume that consumers are sophisticated about the behavior of their future selves, though our results are essentially unchanged if people are instead naive (see appendix D.4). Because shocks are independent over time, the only state variables at time t are cash on hand w_t and the realization of the preference shock ρ_t.

Harris and Laibson (2001) provide a set of conditions under which the equilibrium of such a game can be described by what they call the Strong Hyperbolic Euler Equation (SHEE). We assume the corresponding conditions hold here.[6]

Definition: *A consumption function $c : \mathbb{R} \times \mathbb{R} \to \mathbb{R}_+$ satisfies the Strong Hyperbolic Euler Equation with credit and savings constraints if the following holds for every $w_t, \rho_t \in \mathbb{R}$:*

$$\begin{aligned} u'(c(w_t,\rho_t),\rho_t) &= R'(s_t)\beta\delta E_t\left[V'(w_{t+1})\right] \\ &= R'(s_t)E_t\left[\left(\beta\delta\frac{\partial c_{t+1}}{\partial w_{t+1}} + \delta\left(1-\frac{\partial c_{t+1}}{\partial w_{t+1}}\right)\right)u'(c(w_{t+1},\rho_{t+1}),\rho_{t+1})\right] \end{aligned} \quad (1)$$

[6] These are: (1) The utility function u is strictly increasing and twice continuously differentiable on $[0,\infty)$; (2) Relative risk aversion is bounded away from zero and below infinity, i.e. $0 < \underline{x} \leq \frac{-cu''(c,\rho)}{u'(c,\rho)} \leq \overline{x} < \infty$, on $[0,\infty)$ for all ρ in the support of $F\varrho$; (3) The distribution function f_y is twice continuously differentiable and has a support that is bounded away from zero and below infinity; (4) The distribution f_ρ is twice continuously differentiable; (5) $\max\{\delta, \delta R(s)^{1-\underline{x}}\} < 1$ for $s > 0$; and (6) The hyperbolic discounting factor satisfies $\beta \in [0,1]$ and the model is parameterized such that the equilibrium consumption function is Lipschitz continuous (β is close to 1).

This version of the SHEE has the familiar Euler equation interpretation. $u'(c(w_t, \rho_t), \rho_t)$ is the marginal utility of consumption at t, while $R'(s_t)$ is the rate at which money today is converted into money tomorrow. The expectation term on the right hand side of equation (1) is the expected discounted marginal value of cash on hand in period $t+1$ from the point of view of the agent at time t (denoted $V(w_{t+1})$). The optimal allocation equalizes the marginal value of consuming funds right away and handing them to one's next-period self.

Equation (1) differs from the standard Euler equation in two ways. First, the marginal utility of income tomorrow from the perspective of time t is discounted by

$$d_{t+1} \equiv \beta\delta \frac{\partial c(w_{t+1}, \rho_{t+1})}{\partial w_{t+1}} + \delta \left(1 - \frac{\partial c(w_{t+1}, \rho_{t+1})}{\partial w_{t+1}}\right)$$

rather than a time-invariant discount factor. This effective discount factor is a weighted average of the short-run discount factor $\beta\delta$ and the long-run discount factor δ, where the (time-variant) weight is given by the future propensity to consume. Second, the standard market interest rate term $1+r$ is replaced by the savings-dependent rate of return $R'(s_t)$.

Rearranging, we can express the MRS of the consumer as:

$$MRS_t \equiv \frac{u'(c(w_t, \rho_t), \rho_t)}{E_t\left[d_{t+1} u'(c(w_{t+1}, \rho_{t+1}), \rho_{t+1})\right]} = R'(s_t). \tag{2}$$

Note that MRS is constant and equal to $1+r$ in the no-constraints case. In the complete-constraints case, the marginal propensity to consume equals one, $d_{t+1} = \beta\delta$, and the identity $w_t = y_t = c_t$ holds in every period, so

$$MRS_t = \frac{u'(y_t, \rho_t)}{\beta\delta E_t u'(y_{t+1}, \rho_{t+1})}.$$

1.4 Choice in MPL Experiments

Assume that the subject has optimized her consumption plan given her current period's income y_t and realized shock ρ_t. Her current level of consumption is c_t^* and her savings are

s_t^*. Now she is offered the experimental choice of a payoff of a_1 one period ahead vs. a_0 immediately. We will now show that her experimental choices reveal her MRS, as long as these payoffs are small. This is regardless of whether experimental payouts in the current period must be consumed immediately, or current period consumption can be adjusted.

Proposition 1. *Consider the decision maker's preferences between receiving $a_0 > 0$ immediately and $a_1 > 0$ in the next period, and the set of all such payments with the ratio $\frac{a_1}{a_0} = \hat{R}$. There exists a non-zero bound on a_0 (and $a_1 = \hat{R}a_0$) below which the decision maker will strictly prefer the earlier payment if*

$$R'(s_t^*) = \frac{u'(c_t^*, \rho_t)}{E_t\left[d_{t+1}u'(c(w_{t+1}, \rho_{t+1}), \rho_{t+1})\right]} > \hat{R}, \tag{3}$$

regardless of whether a_0 must be consumed immediately or if the decision maker can adjust her consumption and saving decision in period t. The later payment will be strictly preferred if the inequality is reversed.

Proof. See appendix B. □

The proposition shows that the pairwise choices in the MPL experiment provide an interval estimate of the MRS. At the point of indifference between earlier and later payments the relative marginal value of money in the two periods is (approximately) equal,[7] both in terms of its consumption value, and its investment value. The subject's experimental choices approximate the slope of the budget constraint $R'(s_t^*)$ when they can arbitrage the payoffs, and the slope of the indifference curve at s_t^* without arbitrage, and at the optimum these are equal.

[7]To give an idea of how this approximation performs, consider the utility function $u(\omega + a) = \frac{(\omega+a)^{1-r}}{1-r}$ with $r = 0.741$, as in Andersen et al. (2008) (with payments converted to Danish Krona) and consider a one-standard deviation shock to consumption. With ω equal to mean consumption in our data (roughly \$14), and ω' mean consumption plus one standard deviation (roughly \$28), $\frac{u'(\omega)}{u'(\omega')} = 1.67$. We can compare this with $\frac{\frac{u(\omega)-u(\omega+a)}{a}}{\frac{u(\omega')-u(\omega'+a)}{a}}$ for different experimental payments a. For $a = 1$ USD this ratio is 1.65, an error of about 1.2%. For $a = 0.15$ USD (our mean immediate payment) this ratio is 1.67 to 2 decimal places, an error of about 0.2%.

A similar argument can be used to determine the subject's choice between future payments at $t = 1$ and $t = 2$, evaluated at period $t = 0$.

Proposition 2. *Consider the decision maker's preferences between receiving $b_1 > 0$ at $t+1$ and $b_2 > 0$ at $t+2$, and the set of all such payments with the ratio $\frac{b_2}{b_1} = \hat{R}$. There exists a non-zero bound on b_1 (and $b_2 = \hat{R}b_1$) below which the decision maker will strictly prefer the earlier payment if*

$$\frac{E_t\left[d_{t+1}u'(c_{t+1}, \rho_{t+1})\right]}{E_t\left[d_{t+1}d_{t+2}u'(c_{t+2}, \rho_{t+2}) + O_b\right]} > \hat{R} \tag{4}$$

where O_b is an 'approximation error' that equals zero if either the decision maker is time-consistent, or the interest rate is fixed (the no-constraints case). The later payment will be strictly preferred if the inequality is reversed.

Proof. See appendix B. □

Note that d_{t+2} is the discount rate that the $t+1$ self applies in trade-offs between periods $t+1$ and $t+2$. We will discuss implications of this result in section 1.6.1.

1.5 Predictions of the Partial Credit Constraints Model

Next, we use the model to make predictions about the relationships between measured MRS and savings, income shocks, and preference shocks, and show how these can differentiate between credit regimes. All proofs from this section appear in appendix C.

First, we consider exogenous variation in income. It is straightforward to show that, all else equal, higher income is associated with higher savings, and therefore lower measured MRS.

Prediction (Income shocks and MRS): *Consider a decision maker who holds savings from the previous period s_{t-1} and has preference parameter ρ_t. For any two possible income realizations y_t, y'_t and associated MRS_t, MRS'_t, $y_t > y'_t$ implies $MRS_t < MRS'_t$.*

Next, we consider the preference parameter ρ. The notion that the derivative of u may vary randomly for a given level of c is motivated by the observation that measured consumption spending and true "value of consumption" do not always perfectly line up. In particular, if we think of the c in the utility function as total consumption expenditure, we have to account for variation in spending that does not translate into immediate utility gains. For example, the expenditure to "undo" an adverse event such as the theft of a productive asset, illness of a family member, or damage to one's house does not actually increase the decision maker's utility in the same way as, say, buying a meal would. A household that is subject to such an event has a higher marginal utility of consumption than a household with the same level of c but without this event. Such a preference shock will lead to an increase in measured consumption and a reduction in savings.

Prediction (Preference shocks and MRS): *Consider a decision maker with cash on hand w_t. For any two realizations of the preference shock ρ_t, ρ'_t and the associated MRS_t, MRS'_t, $\rho_t < \rho'_t$ (and therefore $\frac{\partial u(c,\rho_t)}{\partial c} < \frac{\partial u(c,\rho'_t)}{\partial c}$ for all c) implies $MRS_t < MRS'_t$.*

Note that these predictions refer to exogenous changes in income and preferences (shocks) but not endogenous (chosen) changes in income, e.g. from increased labor supply (see appendix D.1 for a discussion). We exploit the prediction that income sources with greater exogenous variation should be more strongly negatively related to MRS in section 3.1.

Our third prediction uses the fact that an increase in s_t is directly associated with a fall in MRS_t through the shape of the returns function R. Note that, even though the level of savings s_t is endogenously chosen by the household and may depend on the shape of R, $MRS_t = R'(s_t)$ holds with equality in each period.

Prediction (Savings and MRS): *For any two possible savings levels s_t, s'_t and associated MRS_t, MRS'_t, $s_t > s'_t$ implies $MRS_t < MRS'_t$.*

Also of interest is the relationship between household spending and MRS. In the full-constraints model, consumption and income are the same (save for reporting and measurement error), and so spending will be negatively related with MRS. In the no-constraints or narrow-bracketing models, the MRS is unaffected by spending. In the partial-constraints model, the relationship depends on the relative importance of income and preference shocks. If there are few or no preference shocks, then spending is determined mostly by income.

Subjects will consume more when income is high, and so spending and MRS will be negatively correlated. However, if preference shocks dominate, then MRS and spending will be positively correlated: for example if an asset used in household production breaks and has to be replaced, then spending will be high in that week, but utility-relevant consumption (e.g. food) will be low, and so the MRS will be high. In appendix C.1, we formalize this claim, and show how the relationship between MRS and spending could, under simplifying assumptions, be used to bound the relative variance of income and preference shocks.[8]

All predictions are summarized in table 2.

[8] We note that shocks to expected future income could also give rise to a positive relation between spending and MRS.

Table 2: Predicted Relationships between MRS and Financial Variables

	Expected Relationship with MRS			
	Savings	Inc. shocks	Pref. shocks	Spending
	s_t	y_t	ρ_t	c_t
Narrow bracketing (R irrelevant)	0	0	0	0
No credit constraints ($R' = 1 + r$)	0	0	0	0
Full credit constraints ($R' = 0/-\infty$)	0	$-$	$+$	same as income
Partial credit constraints ($R'' < 0$)	$-$	$-$	$+$	indeterminate

The predictions regarding preference and income shocks differentiate our model with credit constraints from narrow bracketing and the no-constraints version of the model. Both predict no relationship between shocks and MRS, in the first case by assumption, in the second because such shocks do not affect the effective interest rate faced by the household. The relationships of MRS with savings and spending serve to differentiate the partial- and complete-constraints version of our model. Under complete constraints, savings are zero, so that any difference between income and spending is due to measurement error. In this case we would not expect a relationship between savings and MRS, while spending and income must have the same relationship with MRS.

The above results pertain to the sign of the relationship between MRS and alternative realizations of other variables (income, spending, and preference shocks) in the same period. We show in appendix C that the covariance between MRS and each of the variables calculated from a T-length sample will, under mild conditions, have the same sign in expectation.

Our model of experimental decisions makes strong rationality assumptions, especially given that the experimental payments are small and therefore errors are not very costly. However, if outside consumption is chosen optimally, our predictions go through as long as subjects recognize changes in the value of money today relative to expected value of money in the future (e.g. through changes in the interest rate), and adjust their experimental choices on average in the right direction.

If outside consumption is not chosen optimally, experimentally measured MRS may still respond to shocks as we predict, as long as the subject feels relatively rich after a positive shock, and relatively poor after a negative one. For example, consider a heuristic model in which the subject applies their utility function to current consumption, but evaluates cash on hand in the future according to some constant function. This subject would behave in line with our predictions of experimental choices under the partial credit constraints model outlined above. Whether the correlations in table 2 would still be able to differentiate between credit regimes in such a case depends on the specifics of the model, for example whether an unconstrained household succeeds in keeping their marginal utility constant.

1.6 Implications

As we show below, the predictions of our proposed model for the relationship of MRS with shocks and savings are supported by the data. We can therefore use the model to ask what can be learned from the experimental measurement of intertemporal trade-offs. We first discuss what conclusions can be drawn about a subject's time preference parameters. Then we show how such experiments can be exploited for empirical research on the effect of financial shocks on household finances and the availability of intertemporal consumption smoothing and (self-) insurance.

1.6.1 Implications for the Measurement of Time Preference Parameters

Section 1.4 shows that, if our model is correct, measured MRS from decision A will typically not directly reflect the discount factor, as is often assumed in time preference experiments. Moreover, decision A and B together do not generally give us information about the (level of) present bias or the value of β.

To see this, substitute the Euler equation into our expression for $\frac{b_2}{b_1}$ to get

$$\begin{aligned} \frac{b_2}{b_1} &\approx \frac{E_t\left[d_{t+1}\cdot E_{t+1}\left[d_{t+2}u'(c_{t+2},\rho_{t+2})\right]\cdot R'(s_{t+1})\right]}{E_t\left[d_{t+1}d_{t+2}u'(c_{t+2},\rho_{t+2})+O_b\right]} \\ &= \left\{E_t\left[R'(s_{t+1})\right]+\frac{Cov\left(d_{t+1}\cdot E_{t+1}\left[d_{t+2}u'(c_{t+2},\rho_{t+2})\right],R'(s_{t+1})\right)}{E_t\left[d_{t+1}d_{t+2}u'(c_{t+2},\rho_{t+2})\right]}\right\} \\ &\times \frac{E_t\left[d_{t+1}d_{t+2}u'(c_{t+2},\rho_{t+2})\right]}{E_t\left[d_{t+1}d_{t+2}u'(c_{t+2},\rho_{t+2})+O_b\right]} \end{aligned} \tag{5}$$

where O_b is the 'approximation error' defined in Proposition 2.

The first term in this expression is the expected future interest rate, which, in a stationary economy, is (approximately) *equal* to the expectation of $\frac{a_1}{a_0}\approx R'(s_t)$. This is the rate at which money can be transferred between $t+1$ and $t+2$ outside the experiment. This equivalence occurs here because the subject at time t can only choose when payments are received, not when they are consumed. This means to a first approximation that the best they can do is maximize expected discounted income.

The covariance term arises from the fact that self t must predict both future consumption utility and the future interest rate. Consider for example the exponential discounting case, where $d_t = \delta$, and assume there are no preference shocks. In this case, the covariance term is positive and the term in brackets is *greater* than $\frac{a_1}{a_0}$ on average, since both u' and R' vary negatively with s_{t+1}. This argument continues to hold with quasi-hyperbolic discounting if the marginal propensity to consume does not respond too strongly to financial shocks. The covariance term disappears only if either there are no credit constraints (R' is constant) or if credit constraints are very high and savings vary little with income (c_{t+2} becomes independent of s_{t+1}).

The last term is a multiplier that equals one if O_b equals zero. From Proposition 2, this is the case if either there are no credit constraints, or the decision maker is not present biased (i.e. $\beta = 1$). We show in appendix B that the term O_b will be *positive* if the decision maker is present biased (as long as β is not too far from one) and the interest rate varies with savings, as in our model.

One approach to identifying time inconsistency in the literature has been to use preference reversals between decisions A and B to conclude that there is present bias, without necessarily identifying β exactly. Note that, without narrow bracketing, any individual preference reversal may be due to financial shocks. Moreover, the covariance term tends to bias any estimate of $\frac{b_2}{b_1}$ upwards, making decision B on average *less* patient than decision A. However, if O_b is positive, which can only be the case if $\beta < 1$, the term in brackets in the expression above is multiplied by a number less than one. Thus, assuming that the economy is stationary, and our model of partial credit constraints is correct, we would observe decision B to be on average more patient than decision A *only if* there is present bias, either on the individual level (when observing many decisions for one person) or on the population level. Due to the covariance term, however, the converse is not true, that is, the presence of present bias does not imply that decision B must be more patient than A on average.

Can there be any further progress on identifying β and δ from decisions A and B? The answer is yes, but only in some special cases. Specifically, equation (2) shows that the MRS identifies the (effective) discount factor if $u'(c(w_t, \rho_t), \rho_t) = E_t\left[u'(c(w_{t+1}, \rho_{t+1}), \rho_{t+1})\right]$. This holds for example if the household has globally linear utility or if consumption is constant in all periods because there are no shocks.[9] The expression also holds in expectation if u' is stable over time and the decision maker is subject to complete credit constraints, so that marginal utility in each period is determined only by realized income and preference shocks. Moreover, we have

$$\begin{aligned} E\left(\frac{a_1}{a_0}\right) &= E\left(\frac{u'(y_t, \rho_t)}{\beta\delta E_t u'(y_{t+1}, \rho_{t+1})}\right) = \frac{1}{\beta\delta}, \quad \text{and} \\ E\left(\frac{b_2}{b_1}\right) &= E\left(\frac{E_t\left[u'(y_{t+1}, \rho_{t+1})\right]}{\delta E_t\left[u'(y_{t+2}, \rho_{t+2})\right]}\right) = \frac{1}{\delta} \end{aligned}$$

where $E(\cdot)$ denotes the unconditional expected value. As a result, the difference between decision A and B can be used to identify time inconsistency on average.[10]

We return to possible pathways for the identification of time preferences from experimental data in the conclusion.

1.6.2 Implications for the Measurement of Consumption Smoothing and Insurance

While our results are somewhat pessimistic about identifying time preference parameters from experimental measures of MRS, they suggest that these can instead help us understand the financial shocks and constraints that affect a household. Repeated MPL experiments can be used to measure the variance of individual MRS over time and between subjects, and the covariance between MRS and other financial variables. Measuring MRS in this way is

[9] In the case of linear utility, the decision maker would adjust savings until reaching the point at which $R'(s_t) = \frac{1}{E_t(d_{t+1})}$.

[10] Note that adding more choices over payments in different time periods does not help. For example, even under the partial constraints model, transitivity implies that the MRS between periods 0 and 2 should equal the MRS between 0 and 1, multiplied by the MRS between period 1 and 2. Thus, such questions would not provide new information that would help identify time preference parameters (although they might help pin down measurement error).

significantly easier than (for example) inferring changes in marginal utility from the variance of consumption, and is unaffected by the problem of preference shocks (which we show to be important in section 3).

This methodology has many potential applications. As we have shown in section 1.5, the relationship between MRS and other financial variables can help to determine the credit regime faced by a household. Furthermore, the better a household's ability to smooth financial shocks, the lower should be the overall variance of its MRS, as well as its MRS response to exogenous shocks (e.g. promised future payments). This could be used to test the impact of programs designed to improve household consumption smoothing, for example of the type evaluated in Karlan et al. (2014).

MRS measurements can also be used to test predictions about the first-order conditions for intertemporal consumption allocation over time. Starting with Hall (1978), a large literature has examined systematic deviations of observed consumption choices from the path prescribed by (a linear approximation of) the Euler equation due to factors such as credit constraints (see e.g. Zeldes (1989); Runkle (1991) for early examples). Other models make predictions for the effects of incentive constraints in problems of risk sharing on (inverse) MRS, and test them using implications for consumption allocations over time (e.g. Rogerson (1985); Green and Oh (1991); Ligon (1998); Golosov et al. (2003); Kocherlakota and Pistaferri (2009); Attanasio and Pavoni (2011); Karaivanov and Townsend (2014); Kinnan (2019)). Yet it has long been recognized that the estimation of (log-linearized) Euler equations is hampered by approximation bias (e.g. Ludvigson and Paxson (2006); Carroll (2001)), whereas nonlinear GMM approaches lead to inconsistent estimates when there is measurement error (see Alan et al. (2009) for a discussion). Correlated measurement error can bias the results when studying consumption and income over time (Runkle (1991)), which is a concern given that these variables are often difficult and costly to measure (Grosh and Glewwe (2000)).

Experimental measures of MRS may be able to address some of these issues, assuming that the measurement error in them is independent from concurrent measures of financial variables. As a simple, illustrative example, consider the classical test of full insurance (Townsend (1994); Deaton (1997); Mace (1991)). Without a savings technology, the Pareto optimal choice by a social planner in period 0 will allocate consumption in any period t and any state of nature s such that weighted marginal utility is equalized across individuals. This means that the full-insurance model in its purest form predicts that MRS is the same for all individuals in any given period, and this prediction can be tested with experimental MPL data. Weaker predictions, such as whether or not MRS is related to individual-specific shocks (in addition to group-level shocks) can also be tested.[11]

Townsend (1994) and others conduct equivalent tests with consumption and income data only, by using a specific utility function (typically CARA or CRRA) to predict the co-movement of individual and group consumption or to test the residual effect of individual income on consumption. Early applications of these tests show problems with this approach; for example, Mace (1991) carries out the test for both a power and an exponential utility function and rejects full insurance in one case but not the other. As Kinnan (2019) and many others have pointed out, measurement error in right-hand side variables or correlated measurement error in individual consumption and income[12] may also lead to a spurious effect of individual income on MRS. By contrast, the use of experimental data does not require estimating MRS from consumption, and measurement error in the experimental data is less likely to be correlated with measurement error in the data on financial shocks. This is particularly important in the presence of preference shocks, which drive a wedge between consumption expenditure and utility. As we discuss below, our data suggests such shocks are quantitatively important.

[11] In an earlier version of this paper we discuss the predictions of the Townsend mutual insurance model for measured MRS more formally (Dean and Sautmann (2014), section 5.2) and argue that the correlations of individual shocks with MRS are a rejection of the full-insurance hypothesis. One could perform further analysis by studying the co-movement of the MRS measures of the risk-sharing group, or relate MRS changes to aggregate shocks.

[12] Introduced for example by errors in pricing household production and consumption, see Deaton (1997).

1.7 Extensions of the Basic Model

Our baseline model makes a number of simplifying assumptions. In appendix D we discuss the implications for the results of section 1.5 of four generalizations: endogenous sources of income, intertemporally correlated shocks, temporary shocks to the individual return function R_i, and naivety on the part of the household. Broadly speaking we find our results to be robust. If some income is under the control of the household, MRS will be negatively related to exogenous income shocks, but *positively* related to endogenous income changes. We utilize this fact in section 3.1. Serial correlation in income does not overturn any of the results of section 1.5, as long as shocks have a larger effect on current income than they do on future income. Exogenous shocks to the return function R_i could potentially lead to a positive relationship between savings and measured MRS. The fact that we find a negative relationship in section 3.2 suggests that in our sample such shocks are less important than changes to the interest rate caused by decreasing returns. Our results also go through under the assumption that our subjects are naive, and believe that in the future they will behave in a manner consistent with their current preferences. However, naivety opens up a new channel that could lead to present bias in measured discount rates.

2 Data

We now apply the insights from the model to data from MPL experiments that were carried out as part of a larger panel survey in Fall 2012 in Bamako, Mali. The survey was the baseline of a randomized control trial for a health care program for children. We collected demographic information at the start of the survey, and household members answered detailed questions on income and spending every week. The head of the household participated in multiple price list time preference experiments in four consecutive visits.

Table E.1 in appendix E shows summary statistics for the population of 1013 subjects. The sample is fairly characteristic for the area, but there is selection to the degree that survey participants were chosen according to the criteria of the NGO providing the health care program. All households have children under five and had to pass a proxy-means test for income. If a household member had a savings account or was holding a salaried job at the time of the proxy test, the household was not eligible for the program and did not participate in the survey.

The time preference experiment consists of a set of multiple price list choices over payoffs at different points in time as shown in table 1. These MPLs measure trade-offs between money in the current week and the next (A), and next week and one week after (B).[13] Households were asked to make choices from set A and B in three consecutive weeks. All households were interviewed in the same three week period.

Each decision in the MPL is a choice between a payment of CFA 300 (about US$ 0.60) at the later point in time, and a payment varying from CFA 50 to CFA 400 (US$ 0.10-0.80) at the earlier point in time. The experimental design follows the standard MPL procedure used in the literature, with the exception that we allow for negative interest rates by offering trade-offs between a higher payoff earlier and a lower payoff later. This is motivated by the idea that a severely savings-constrained household would actually prefer to exchange high amounts today for lower amounts tomorrow, and indeed we see a number of households choose this option (see below).

One decision from all MPL choices during the current visit was selected for payout, using a random draw at the end of the experiment. Subjects then either received their immediate monetary payout, or a written receipt that stated the date and amount of any future payout the subject was owed. In the following weeks, the surveyors used their own notes and the subjects' receipts to make payouts due from past decisions. As the surveyors visited the household every week, transaction costs were the same for current and future payments. In order to establish subjects' trust, the first time-preference experiment consisted only of choices over payouts in the future to make salient that the surveyors actually return and make payments owed in later weeks. These choices are not used here.

[13] In one week, an additional MPL experiment was carried out which is not used here, concerning choices between payouts two and three weeks away.

Table 3 shows a summary of the remaining three weeks of MPL choices. The top row shows the number of observations. As is fairly typical in these types of experiments, 10-14% of subjects were recorded as making inconsistent decisions within a price list, with repeated switches between earlier and later payoffs. Education and literacy are associated with consistency; for instance, illiterate subjects make on average 15.4% inconsistent choices, but literate subjects only 8.7% (different at the 1% significance level). The other demographic variables have no effect on consistency. In the remainder of the table, only consistent choices are reported (we use the inconsistent choices in our conditional logit estimates, see below).

The row labeled "Avg. switch to earlier payment" reports the lowest earlier payoff that was chosen on average. The lowest possible value is therefore CFA 50, the highest value was set to 450 (for individuals who chose the later payment always). Due to the discrete experimental choices, we cannot report exact "indifference points" between earlier and later payments. We will discuss this issue in more detail below.

Table 3: Experimental choices.

	Week 1		Week 2		Week 3	
Observations	973		969		965	
Decision	A	B	A	B	A	B
Consistent	830 (85.3%)	836 (85.9%)	871 (89.9%)	856 (88.34%)	858 (88.9%)	864 (89.5%)
Avg. switch to earlier payment (CFA)	157.2	155.8	153.6	152.5	158.2	154.4
Implied average MRS	4.78	4.7	4.73	4.65	4.53	4.56
Paying negative interest rate (MRS<1)	9.64%	8.25%	7.35%	5.49%	7.34%	6.83%
Equal choice in A and B	69.85%		70.37%		76.31%	
More patient (lower MRS) in A	15.14%		14.02%		10.09%	
More patient (lower MRS) in B	15.01%		15.61%		13.61%	

The next rows in the table show the week-to-week correlations of decisions in A and in B, and the proportions of subjects who made the same, more patient, or less patient decisions in A compared to B. Subjects' decisions in the different MPL experiments are clearly related: a sizable proportion choose the same switch point in both decision A and decision B, and across weeks. However, there is also significant variation in choices both across and between weeks; up to 30% of subjects choose different switch points in A and B in the same week, and the correlation of choices between weeks is high, but far from perfect at 0.67-0.72.[14] In 10-15% of cases subjects make a more "patient" choice in decision A than decision B. The table also shows that at least 7% of subjects are willing to pay a weakly negative interest rate, that is, they choose CFA 300 in one week over CFA 350 right now. None of these patterns can be explained by the quasi-hyperbolic model in the standard "narrow bracketing" framework, but they are possible in the presence of financial shocks.

Table E.2 in appendix E shows the distribution of switch points in decision A for consistent subjects by week. There is bunching at the most patient and most impatient choice, with a large proportion of subjects choosing the earlier payment always. 14% of subjects each week choose always the higher of the two payments in each of the eight choices, implying an interest rate between 1 and 1.167. Aside from these three most frequently observed choices, there is significant and varying dispersion in choices across the three weeks.

The lack of (average) present bias in our data, in the sense of more impatient choices in A over B, may seem surprising, but is consistent with other studies that take care to minimize differences in transaction costs and risk between present and future payments. Andreoni and Sprenger (2012) and Augenblick et al. (2015) estimate present-bias parameters β for money that are never significantly less than 1. Halevy (2015), who repeatedly visited subjects in class in order to make payments, also found little present bias. Repeated visits – in class, or at home as in our study – may not only eliminate transaction costs for the subjects but also reduce self-selection into experiment participation, based, for example, on current financial

[14]In comparison, Halevy (2015) found between 43% and 60% of subjects make identical choices in two "A" type decisions five weeks apart.

need. However, based on the analysis of section 1.6.1, the lack of present biased choices in our data does not mean that we can conclude that households are not time-inconsistent.[15]

For the remainder of the paper we focus on the data from decision A, to which the predictions of section 1.5 refer.[16]

Finally, we collected weekly income and spending data (table 4). Income data was collected

[15] We only point out that ours is not the only study to find no average difference between A and B. The cited studies are carried out on US student populations, and our model makes no clear prediction about how their behavior should compare to our study population. As equation 5 in section 1.6.1 shows, credit-constrained populations with $\beta < 1$ may appear more or less present biased than those that are unconstrained (but note that there is evidence that students are also credit-constrained, Halevy (2015)).

[16] See Dean and Sautmann (2014) for analysis of the relationship between decision B and financial variables. These relationships are not discussed in the current paper because the relevant theoretical predictions are not robust to the presence of serially correlated shocks.

Table 4: Weekly income by source, consumption, adverse events, and the resulting change in savings (income minus spending).

	min	median	mean	max	std. dev.
Income	0	31.00	59.46	1309	92.96
Labor income	0	28.00	51.33	952	80.06
Nonlabor income	0	0	8.02	1001	41.55
- of which "exogenous" sources:	0	0	4.99	1001	34.99
- of which "endogenous" sources:	0	0	3.02	500	21.68
Spending	0	66.37	98.45	1210	106.42
Spending on food and household necessities	0	21.40	27.16	1040	32.30
Adverse event spending	0	0	5.80	600	23.57
Adverse event occurred?	0	0	33.20%	1	47.10%
Savings increase (income - spending)	-1176	-27	-38.59	1147	82.07

All amounts converted to US$. Exogenous sources of nonlabor income: formal transfers, rent payments received, and loan repayment received. Endogenous sources of nonlabor income: informal transfers, sales revenue of an item owned, tontine payouts, and gifts after an adverse event.

by source, and can be broadly categorized into labor and non-labor income. As described in the table, we also break out non-labor income into "endogenous" and "exogenous" categories, according to the degree by which the household can affect the size and timing of payments (see appendix E for more on this breakdown). A typical experimental payment is about 2% of the weekly median household income.[17]

Spending includes any monetary outlays of the household.[18] Of particular importance for our analysis is the expenditure category of "adverse events". Subjects were asked whether they had incurred any unexpected expenditure since the surveyor's last visit due to "damage to an item your household owns; damage to a building; loss, theft, or destruction of a good; loss or theft of animals; or illness to a household member". If they answered yes, they were asked how much money was spent on repairs, replacement or (for illness) treatment. We use such events to proxy for preference shocks of the type discussed in section 1.5.

Some notes on data quality and the match with the model variables are in order. First, savings as reported here are a flow variable. The stock of savings s_t is unobserved, because our survey did not collect information on cash and other liquid assets held from week to week.[19] We discuss this issue in appendix F. Second, spending does not directly correspond to consumption, but rather represents the outflow of cash, whereas 'true' consumption is unobserved. The model addresses this by allowing for preference shocks which do not directly contribute to consumption utility. Third, we may be concerned that households selectively participate in the survey depending on their financial outcomes in a given week, or that individuals who make inconsistent choices differ from those who do not. Comparing households that have some weeks of missing or inconsistent data with households that do not (411

[17]A regression of measured MRS on the receipt of experimental payments yields small and insignificant coefficients, suggesting that these payments are indeed 'small', as required for our theoretical results.

[18]Including purchases of food and household goods, spending on fuel, rent, electricity, and heat, personal expenses of the household head, transfers to other households, business expenses including labor cost, and payments into a savings club or to pay off a debt.

[19]Information on household wealth, while available, is noisy and only includes relatively illiquid assets. Subject are generally reluctant to give information on cash and other savings in the house. Within the time constraints of the health survey in which this data was collected, we expect that, even if liquid asset information had been gathered, it would likely not be precise enough to reflect week-to-week variation in the relevant s_t accurately.

out of 2559 observations), we find that they have on average lower spending, income, and savings. The largest difference is in income (significant at the 14% level); households with missing MRS data report on average \$54 compared to \$61 weekly income. Occurrence of and spending on adverse events is nearly identical for both types of households, so they seem to be subject to similar shocks. Lastly, using our information on consumption and income, we calculate flow savings to be negative on average. While it is possible that our sample of households as a whole is dissaving, this discrepancy is not atypical for household surveys and commonly interpreted as a sign of under-reported income (see e.g. Deaton (1997)).[20] The income distribution is also more skewed than the spending distribution, suggesting that households have rare high income realizations that were not observed in our short panel. In general, it is likely that our financial data exhibits measurement error. We will address this again when discussing individual empirical tests.

3 Analysis

In this section we test the predictions from section 1.5.[21] The aim of our analysis is to use the data to differentiate between different models of financial constraints.[22] The empirical model we use to test the effects of income and preference shocks is

$$MRS_{it} = \alpha_i + \lambda X_{it} + \gamma_t + \epsilon_{it},$$

where X_{it} is the financial variable of interest converted to US\$ 100. MRS_{it} represents the marginal rate of substitution measured by the MPL experiment.

The individual fixed effect α_i implies that we are looking at deviations of MRS_{it} and X_{it} from their individual-specific averages. This accounts for ex-ante differences between households in income and spending levels (which determine for example what constitutes a

[20] Spending on large (durable) purchases, which could be a form of savings, can only account for 10% of the difference.

[21] Prior to data collection, our plan was to estimate the relationship between measured MRS and total income, adverse events, and savings. The additional results we report are exploratory.

[22] While using structural methods to estimate the parameters of the underlying model might be possible in principle, in practice the data requirements are extreme, as we discuss in section 4.

"positive" or "negative" shock), the savings stock, and the returns function R. It is likely that R is stable for a given household over the relatively short span of the experiment, but there may be variation in the interest rate function between households. For example, if household 2 faces a higher interest rate than household 1 at all savings levels, it may induce them to save more, leading to a positive inter-household correlation between savings and MRS. Similarly, the savings stock is endogenous to past shocks and may also differ for individuals with same R but different β or δ. For all these reasons we focus on within-subject variation in absolute terms, although we estimate the average response of the MRS to shocks.

In some specifications we also include period fixed effects γ_t to control for potential period-specific preference changes and time trends, for example due to festivals, holidays, weather changes, changing financial market conditions, or other sample-wide events. The error term ϵ_{it} captures the measurement and approximation error in the experimentally measured MRS, as well as any variance in intertemporal trade-offs not explained by the financial variables.

Since we only have discrete brackets given by the nine possible switch points in the list, we take two different approaches to estimating this model. First, we estimate OLS and IV specifications with errors clustered at the individual level, where we approximate the subject's MRS by calculating the midpoint between the ratios of the later over the earlier payment at which the subject switches from choosing the late to choosing the early payment.[23] The MRS for individuals who always choose the earlier payment within a given decision set may lie anywhere on the interval $(6, \infty)$, and for those who always choose the later payment, it may be anywhere on $(0, 0.75)$. The regression results reported here use 0.708 as the lowest and 8 as the highest MRS; we verified the robustness of our estimates to values between 0.3 and 0.75 at the lower end and between 6 and 10 at the upper end (not shown) as well as to drawing a random value for the MRS from the intervals identified by the experimental choices (discussed in more detail in appendix G).[24]

[23]In the case of preference shocks we instrument spending on the adverse event with a dummy for the event occurrence, see below.

[24]0.708 is the next lower switch point if the MPL had included the choice between CFA 450 earlier vs. CFA 300 later. No equivalent "extension" by CFA 50 is available at the lower end; 8 is the midpoint of the interval

These checks aside, the approach cannot account for (surveyor or subject) errors within a choice list, and we must exclude inconsistent choice lists in which there is more than one switch. Moreover, any OLS specification deals in an ad-hoc manner with both the discreteness and the truncation inherent in the data. Thus, for our second approach we estimate a discrete choice model that is more in line with what we actually observe. We assume that the (latent) MRS is a linear function of the financial variables as above, plus an additive logistic error term. In each of the (up to) 24 binary MPL choices the subject makes, the probability of choosing the later payment is given by the probability that the MRS is lower than the ratio of the later to the earlier payment. This can be used to construct a conditional log likelihood and to estimate the coefficient on the financial variable in the MRS, along with the (inverse) standard deviation of the logistic error term. The conditional likelihood method can accommodate person fixed effects and inconsistent choices within a choice set (see appendix H for details).

3.1 Income and Preference Shocks and MRS

We first examine the relationship of MRS and income. Columns (1), (2), and (5) of table 5 report a significant negative relationship between total income and MRS, with column (5) reporting the conditional logit (CL) estimates (note that in CL each binary choice in the experiment constitutes one observation). This rejects narrow bracketing or fixed interest rates, and supports a model with credit constraints.

Within this model, the coefficients we report may in fact underestimate the effect of exogenous income changes on MRS, if households are able to affect their income to some degree in response to shocks. Our predictions regarding the relation between income shocks and MRS continue to hold under mild conditions when some income is endogenous (see appendix D.1), but an endogenous component to total income leads to a downward bias in the estimates.

---

if an additional choice had been included between CFA 30 earlier vs. CFA 300 later.

Table 5: Effect of income (in US$100) on MRS_t, total income in columns (1), (2), and (5), and by income source in columns (3), (4), and (6).

	OLS	OLS	OLS	OLS	CL	CL
Total income	-0.176 *	-0.187 **			-0.226 **	
	(0.094)	(0.095)			(0.102)	
Labor income			0.0178	-0.005		-0.112
			(0.115)	(0.116)		(0.120)
Nonlabor income			-0.310	-0.299		-0.306
"endogenous"			(0.255)	(0.262)		(0.286)
Nonlabor income			-0.414 ***	-0.415 ***		-0.395 **
"exogenous"			(0.146)	(0.152)		(0.190)
1/(sd)[(a)]	-	-	-	-	0.906 ***	0.908 ***
					(0.043)	(0.043)
Ind FE	yes	yes	yes	yes	yes	yes
Time FE		yes		yes	yes	yes
Observations[(b)]	2484	2484	2484	2484	13208	13208

*Standard errors clustered at the individual level (in parentheses). Significance levels * p<0.10, ** p<0.05, ***p<0.01. (a) Reciprocal of the standard deviation of the error term in the conditional logit model. (b) OLS: max. one observation per week per household. CL: max. eight binary choices per week per household.*

In a partial solution we therefore classify each separately recorded income source according to the level of control that the household likely has over that source (see appendix E). Our model predicts that the negative relationship between MRS and income will be strongest for the most "exogenous" income sources. Columns (3), (4), and (6) of table 5 estimate the effect of income split into its different sources. The results support this assumption: whereas the effect of labor income on MRS is small and insignificant, non-labor earnings have a larger effect, and for those income sources least under the household's control, the effect is strongest and significant at the 1% level. Note that the reported coefficients in the CL estimates must be rescaled by σ (1/sd in the table), the standard deviation of the error, giving an effect size of -0.436 for exogenous non-labor income in column (6) (see appendix H). For a sense of the magnitudes of these effects, note that the standard deviation of the de-meaned measured MRS in this sample equals 1.462, so a \$100 increase in exogenous non-labor income lowers the MRS on average by 0.28 standard deviations.

Since this approach does not use truly exogenous income variation, and the groupings above are to some extent ad hoc, we carry out some robustness checks. First, we examine the correlation of the three income categories with the occurrence of adverse events and find that the correlation is overall low, but highest for labor income (0.052), followed by endogenous non-labor income (0.020) and finally exogenous non-labor income (-0.003). Only the correlation with labor income is significant (at the 1% level), consistent with the idea that labor income responds endogenously to consumption needs. Second, the degree to which the MRS correlates with income may be driven by the (lack of) overall variation of the different income types, rather than different degrees of endogeneity. However, the frequencies of positive income observations in the three categories suggest that, due to many zero-income observations, the variation in non-labor income is lower than in labor income (Figure E.1 in the appendix). Last, note that, to the extent that endogeneity remains an issue for the income-MRS relationship, it suggests that the coefficient of -0.411 (-0.436) *underestimates* the true effect of exogenous income shocks.

We use spending on adverse events to test the effect of preference shocks on MRS. We assume that the occurrence of an adverse event is exogenous, and that expenditure on the event – repairing or replacing an item, paying for healthcare – acts essentially like a negative income shock, by reducing the amount of money available for other consumption.

The first two columns of table 6 show that the occurrence of an adverse event has a significant positive effect on the MRS. This is again supportive of our model with credit constraints over no constraints or narrow bracketing. The next two columns show that the effect remains significant at the 10% level when using event expenditure as the independent variable. These results are echoed in the conditional likelihood estimates in the last two columns. Again, the CL estimates must be rescaled by σ, yielding estimated effects of 0.271 for an adverse event and 0.477 for $100 of adverse event spending.

Similar to the issue of endogenous labor supply above, the amount spent on an adverse event may be correlated with marginal consumption utility through the household's choice of how to respond to the event, attenuating the effect. The household can for example reduce expenditure by doing their own repairs instead of hiring someone (see appendix D.1). We therefore instrument for spending on adverse events with the indicator variable for the occurrence of such an event, in order to estimate the (local) average treatment effect (columns (5) and (6), first stage results in table G.1 in the appendix). The IV approach is valid if the binary variable describing the occurrence of an adverse event satisfies the exclusion restriction, that is, it affects marginal consumption utility only through its effect on what is spent on the event. If the adverse event also increases the marginal value of other consumption independently, the IV coefficient overestimates the effect of adverse event spending. The OLS and IV estimates can therefore be seen as respectively lower and upper bounds on the true effect. The results suggest that the simple OLS substantially underestimates the impact of exogenously imposed adverse event expenditure onto MRS.

Table 6: Preference shocks: effect of spending on an adverse event (in \$100) onto MRS_t. Reduced-form and OLS results in columns (1)-(4), IV estimates in columns (5) and (6), conditional logit estimates in columns (7) and (8).

	OLS	OLS	OLS	OLS	IV	IV	CL	CL
Adv. event (0/1)	0.284 **	0.263 **					0.239 **	
	(0.124)	(0.124)					(0.115)	
Adv. event expense			0.256 *	0.237 *	1.707 **	1.579 **		0.427 **
			(0.147)	(0.141)	(0.789)	(0.791)		(0.194)
1/(sd)$^{(a)}$	-	-	-	-	-	-	0.895 ***	0.895 ***
							(0.042)	(0.042)
Ind FE	yes	yes	yes	yes	yes	yes	yes	yes
Time FE		yes		yes		yes	yes	yes
Observations$^{(b)}$	2547	2547	2543	2543	2467	2467	13560	13552

*Standard errors clustered at the individual level (in parentheses). Significance levels * p<0.10, ** p<0.05, ***p<0.01. (a) Reciprocal of the standard deviation of the error term in the conditional logit model. (b) OLS: max. one observation per week per household. CL: max. eight binary choices per week per household.*

A concern the reader might have at this point is that the observed correlations are not due to the failure of narrow bracketing, but rather that preferences vary over time, and that this changes both experimentally measured MRS, and the households' financial choices. However, if the correlation of MRS and financial variables is due to the household's response to changes in preferences, then it should be strongest for the *endogenous* components of income and spending. This is the opposite of what we see in our data.

3.2 Savings and MRS

The results on income and preference shocks rule out narrow bracketing and the no-constraints version of our model. In this section we test whether there is a relationship between savings and MRS, which distinguishes the partial-constraints from the complete-constraints model. Because our data does not contain a good measure of the savings stock, table 7 reports regressions of change in measured MRS onto linear and squared flow savings terms with different specifications for the intercept (note that the lag in the dependent variable means we have here only two weeks of data).[25] Appendix F provides a detailed justification. All four regressions show a significant negative relationship between flow savings and measured MRS. Individual fixed effects reduce power, but increase the absolute size of the coefficient on linear savings. As we demonstrate in appendix F, this result is consistent with the partial-constraints case, but not the complete-constraints or no-constraints case.

[25]Column (1) is a simple OLS and therefore allows only for a common constant time trend in MRS (restricting all α_i to be equal and $\gamma_t = 0$). Columns (3) and (4) relax the common trend assumption and include individual fixed effects α_i. The constant terms (or individual fixed effects) are significant in each regression, indicating there is a time trend that needs to be accounted for. Allowing the time trend to change over time does not change results significantly (columns (2) and (4)).

Table 7: Savings (flows) and $MRS_t - MRS_{t-1}$.

	OLS	OLS	OLS	OLS
Flow savings Δs	-0.182 **	-0.180 **	-0.392 *	-0.391 *
	(0.0865)	(0.0868)	(0.220)	(0.222)
$0.5\Delta s^2$	-0.00639	-0.00598	0.0247	0.0247
	(0.0268)	(0.0268)	(0.0778)	(0.0778)
Time FE		yes		yes
Ind FE			yes	yes
Observations	1462	1462	1462	1462

*Flow savings measured as income minus expenditure. Standard errors clustered at the individual level (in parentheses). Significance levels * p<0.10, ** p<0.05, ***p<0.01.*

The results in table 7 suggest a significant curvature of R (negative coefficient on Δs) that is constant across the range of possible savings (no effect of $0.5\Delta s^2$). Table G.2 in the appendix shows that the results are similar when including a cubed savings term. This is consistent with "soft" credit constraints, rather than a constant interest rate with a hard credit limit, which would imply that there is only one (minimum) level of savings where the MRS responds strongly to shocks. Note that the coefficients in table 7 may underestimate the average curvature of R, if in addition to common time trends, R is subject to individual- and period-specific shocks and subjects are dissaving (see appendix D.3).

3.3 Spending and MRS

As discussed in section 1.5, the relationship between MRS and spending acts as an additional test of the partial constraints model, and is indicative of the relative importance of income and preference shocks. Columns (1) to (3) of table 8 show that the relationship between spending and MRS is positive and significant in our data: higher current expenditure is related to greater impatience. This is consistent with partial, but not complete constraints. It suggests that high realizations of spending are primarily the result of preference shocks that cannot be smoothed, and are in fact associated with *higher* marginal utility and *lower* levels of "utility-relevant" net consumption.

Table 8: Total and disaggregated spending and MRS_t.

	OLS	OLS	CL	OLS	OLS	CL
Total spending	0.203 **	0.180 **	0.167 *	-	-	-
	(0.089)	(0.092)	(0.090)			
Food and necessities	-	-	-	0.691 ***	0.672 ***	0.621 ***
				(0.157)	(0.158)	(0.240)
Adverse event expenses	-	-	-	0.274 *	0.269 *	0.512 ***
				(0.151)	(0.146)	(0.182)
Large purchases	-	-	-	0.0507	0.0496	0.212
				(0.379)	(0.386)	(0.300)
Bills and rent	-	-	-	-0.104	-0.179	0.093
				(0.407)	(0.412)	(0.437)
Gifts and donations	-	-	-	-0.254	-0.372	0.871
				(0.761)	(0.760)	(0.728)
Personal expenditure	-	-	-	-0.487	-0.754	-1.278
				(0.763)	(0.772)	(0.867)
Social events	-	-	-	-0.919 *	-0.977 **	-0.728
				(0.478)	(0.474)	(0.476)
1/(sd)[(a)]	-	-	0.91 ***	-	-	0.901 ***
			(0.044)			(0.043)
Ind FE	yes	yes	yes	yes	yes	yes
Time FE		yes	yes		yes	yes
Observations[(b)]	2418	2418	12736	2439	2439	12936

*Standard errors clustered at the individual level (in parentheses). Significance levels * p<0.10, ** p<0.05, ***p<0.01. (a) Reciprocal of the standard deviation of the error term in the conditional logit model. (b) OLS: max. one observation per week per household. CL: max. eight binary choices per week per household.*

It is of interest to study which types of spending changes affect measured MRS the most, both through the exposure to shocks and the effect of these shocks on spending, given by the category-specific propensity to consume. The remaining columns of table 8 show the results of such an exercise. Social events are significantly *negatively* correlated with MRS. This suggests that this type of spending is driven by income: households spend in this category when they are relatively well off, pointing to high income elasticity and little "shock-driven" spending. Gifts and personal expenditure – which includes goods such as cigarettes, tea, and phone credit – have insignificant coefficients. Utility bills and rent and large purchases have the lowest correlation with MRS, implying that households are able to smooth this (planned) variation in spending well.

We have seen how spending due to adverse events affects MRS.[26] Remarkably, however, the strongest positive relationship is between MRS and spending on food and household essentials. This indicates that variation in demand for basic household goods is driven by preference shocks, rather than, for example, splurging on a good meal after a successful day at work. Shocks in this category could come from seasonal price fluctuations, though large price changes may be unlikely over the relatively short span of our survey. The size of the coefficient is likely a consequence of the fact that basic consumption needs are unresponsive to income and difficult to delay. An additional reason may lie in the traditional organization of Malian families, where women are expected to cover household needs from their weekly allowance and request additional money as needed. These additional expenses will act like exogenous shocks from the perspective of the household head (whose time preferences we measure). We see the identification of the precise cause of these shocks as an interesting avenue for future research.

[26] As noted earlier, a possible alternative explanation for a positive relationship between MRS and spending is through shocks to expected future income. However, this is hard to reconcile with a positive correlation between MRS and adverse event spending, as these shocks should not be related to positive future income changes.

3.4 Discussion

3.4.1 Robustness

As a test of the robustness of our results, we carry out two additional checks (shown in appendix G). First, we report a set of regressions that include both income shocks and preference shocks, and then all sources of income and spending simultaneously (table G.3). This controls for any covariance in income and preference shocks, which could bias the individual estimated effects of these variables onto MRS. The results are broadly robust to this specification change, and the effect sizes remain the same.

Second, in order to test if there are important nonlinearities, we include quadratic terms in the estimations (table G.4). The coefficient sizes and signs suggest that the main correlations hold as predicted. Although the coefficients on the individual variables are not significant, F-tests show that the income shock variables remain jointly significant in all estimations. Event spending is not significant anymore (note that we do not have enough instruments for both linear and quadratic event spending). F-tests for the inclusion of all the quadratic terms cannot reject that they are jointly insignificant, except in the CL estimates. We interpret these results to mean that the main predictions of the model are robust, but that any potential nonlinear effects are not strong enough to be reliably estimated in this relatively short panel.

3.4.2 Interest Rates and Average MRS

As with many other experimental studies (see Frederick et al. (2002) for a survey), measured MRS in our survey is higher than what can plausibly be explained by external interest rates alone: the mean MRS is 4.7 and the median 4.5. The high average MRS is partially driven by the group of 238 subjects who chose the early payment in every single MPL decision. It could be that this subset of individuals is facing a binding borrowing constraint. Another possibility is that they did not engage with the question, or that they operate under a decision-making heuristic that is not described by our model. If we exclude subjects who make the same choice in all 24 MPL decisions (either the late or the early payment always), the mean MRS falls to 3.45 and the median to 1.75.[27] Yet these numbers still imply an annual interest rate at the higher end of the spectrum reported by Frederick et al. (2002).[28] As suggested by Collins et al. (2009), part of the reason may be one-off (non-monetary) transaction costs of borrowing or saving that drive up the effective interest rate on small payments. A second possible explanation is that our subjects attach a significant probability to future payments not being made, either because the surveyor does not return, or because the subjects themselves become unavailable for interview. A constant hazard rate of survey interruption acts like an additional discount factor and shifts all measured MRS upwards.[29] The model predictions for the effect of financial shocks on changes to the MRS are not affected by the inclusion of such an adjustment. Our key finding remains that external financial changes affect the experimental decisions of at least a proportion of subjects.

4 Conclusion

The above results show, both theoretically and empirically, that the monetary trade-offs our subjects make between time periods have interesting potential uses, but do not relate in a straightforward manner to underlying time preference parameters. What are possible ways forward for the measurement of time preferences from experimental data?

Our results show that individual time preference parameters can only be inferred from a single observation of experimental choices if the individual is a narrow bracketer. By contrast, if our model holds, measured MRS is co-determined by consumption and savings choices, and without information about the marginal utility in this period, the expected marginal utility of consumption next period, and the propensity to consume, time preference parameters are

[27] Repeating the analysis on the effect of shocks on MRS and the correlation between flow savings and MRS for this subsample strengthens the results considerably. We do not report these results, because the exclusion of subjects whose measured MRS is stable through the entire panel biases us towards finding the effects our model predicts.

[28] For example, back-of-the-envelope calculations using the method of Andersen et al. (2008) imply that our average subject has an annual discount rate of 32%, relative to the 10% they find in their study.

[29] Given the political instability of the area and frequent flooding during the rainy season such a hazard rate is not implausible.

not identified (see sections 1.4 and 1.6.1). Moreover, any observed one-off preference reversal in experimental choices for two different periods may be the result of financial shocks and therefore cannot reliably indicate present bias.

The news is slightly better if we have many observations for experimental decision A and B, either for a group of subjects, or for an individual over time. Assuming that the economy is stationary, and that shocks are independent, we have shown that preference reversals towards greater patience from decision A to B can *on average* only occur if $\beta < 1$,[30] (although the converse does not hold: absence of such reversals does not imply time consistency). If individuals are additionally subject to complete credit constraints, it is possible to directly identify δ from decision B and β from the average difference between A and B.

Outside this case, and without non-experimental data, precise individual-level identification of β and δ is not possible, because experimental decisions are determined by the shape of R and savings s. However, in equilibrium the choice of s is itself a function of time preferences. In particular, one may conjecture that an individual with a low discount factor will save less on average, thus creating a relationship between more impatient choices and greater discounting.[31] Indeed, Krusell and Smith (2003) show that in a quasi-hyperbolic model without uncertainty, the set of equilibria and therefore equilibrium realizations of the rate of return on assets depends in monotonic ways on β and δ. Thus, observing long-run average MRS allow some inference on time preference parameters. If a parallel result holds under uncertainty, different time-preference types will exhibit distinct (sets of) stationary equilibrium ergodic distributions and different average $R'(s_t)$, potentially allowing a ranking of individuals by their effective discount factor. Characterizing this connection is a promising direction for future research.

Any further progress can only be made with individual-level information on both experi-

[30] This is also true for naive decision makers, see appendix D.4

[31] As an illustration, in the simplest case of no uncertainty and exponential discounting, a steady state can only occur at $R'(s) = \frac{1}{\delta}$, meaning that measured MRS reveals the inverse discount rate. Note, however, that in general no such steady state exists in the no-constraints model where $R'(s) = 1 + r$ for all s, and therefore $\frac{1}{\delta} = 1 + r$ may not hold.

mental choices and financial variables. One approach would be to use a structural model to identify time preference (semi)parametrically, using expression (3) for decision A and (4) for decision B. This requires measurement of wealth, consumption, and preference shocks, as well as the utility function curvature (for example by measuring risk aversion as suggested by Andersen et al. (2008)). An advantage is that this method works even in the no-constraints case; intuitively, for a *given* MRS and interest rate, a more patient decision maker will have a lower level of consumption today relative to tomorrow. The main disadvantage lies in the very strong data requirements.[32]

A final approach would be to identify experimental subjects for whom experimental choices are informative, either because they are narrow bracketers or because their marginal utility of consumption is constant over time. This is not possible from data that contains only measured MRS. It is also not enough to observe that measured MRS is not correlated with financial shocks (as in Giné et al. (2018)), as this is consistent with a household who is not narrow bracketing, but is able to smooth shocks (as in the no-constraints model). Instead, the researcher would need to be able to estimate the marginal utility of consumption. If it varies but is uncorrelated with MRS, one may conclude that the subject is a narrow bracketer. If both are stable over many periods, one may conclude that the subject is either a narrow bracketer, or an integrated decision maker for whom current and expected consumption utility are the same (either because they are not subject to shocks or because they can smooth these shocks); both cases would then allow the identification of β and δ from experimental choices. Finding methods to identify such subjects could be a promising avenue for future research, because the data requirements for such an exercise may be less stringent than estimating a full structural model (but note that the presence of preference

[32] Andreoni and Sprenger (2012) have suggested that outside consumption and preference parameters could be jointly estimated using a utility function of the form $(c_t + a_t)^\alpha$, where a_t is the payment from the experiment at time t and c_t is consumption at t. Identification is achieved by estimating the curvature α from risk preference experiments, and then examining how measured MRS changes when varying the size of experimental payoffs in order to estimate c_t. If there is no change (including if the subject is always at a corner solution in the convex budget sets of Andreoni and Sprenger (2012)), the parameters are not identified. Their approach requires that experimental payments are large relative to outside consumption, *and* that the subject does not engage in arbitrage after the payments were made.

shocks, which we found to be important in our sample, complicates things because marginal utility may not be monotonic in expenditure).

When a measure of time preferences is needed that does not require repeat measurements and detailed information on consumption utility, the most promising direction is probably to collect alternative experimental measures. Indeed, some authors now replace monetary with primary rewards (see e.g. McClure et al. (2007)) or effort (Augenblick et al. (2015)), which may be harder to arbitrage between different time periods and less affected by preference shocks (although a subject who has to carry out an experimental task or consumes a reward may still choose to reschedule other work or consumption). Another possibility may be to use hypothetical questions, assuming that they are more amenable to narrow bracketing; however, it is worth noting that hypothetical discount rates have been found to be affected by changes in inflation rates, which alter effective interest rates (Krupka and Stephens (2013)). Lastly, our results also support using demand for commitment to identify time-inconsistent preferences, as for example in Ashraf et al. (2006) and Mahajan and Tarozzi (2011).

References

Alan, S., O. Attanasio, and M. Browning (2009). Estimating euler equations with noisy data: two exact gmm estimators. *Journal of Applied Econometrics 24*(2), 309–324.

Ambrus, A., T. L. Asgiersdottir, J. Noor, and L. Sandor (2015). Compensated Discount Functions: An Experiment on the Influence of Expected Income on Time Preferences. Working Paper 168, Economic Research Initiatives at Duke (ERID) Working Paper Series.

Andersen, S., G. Harrison, M. Lau, and E. Rutstrom (2008). Eliciting risk and time preferences. *Econometrica 76*(3), 583–618.

Andreoni, J. and C. Sprenger (2012). Estimating time preferences from convex budgets. *American Economic Review 102*(7), 3333–56.

Ashraf, N., D. Karlan, and W. Yin (2006). Tying Odysseus to the mast: evidence from a commitment savings product in the Philippines. *Quarterly Journal of Economics 121*(2), 635–672.

Attanasio, O. and N. Pavoni (2011). Risk sharing in private information models with asset accumulation: explaining the excess smoothness of consumption. *Econometrica 79*(4), 1027–1068.

Augenblick, N., M. Niederle, and C. Sprenger (2015). Working over time: Dynamic inconsistency in real effort tasks. *The Quarterly Journal of Economics 130*(3), 1067–1115.

Benhabib, J., A. Bisin, and A. Schotter (2010, July). Present-bias, quasi-hyperbolic discounting, and fixed costs. *Games and Economic Behavior 69*(2), 205–223.

Carroll, C. D. (2001). Death to the log-linearized consumption Euler equation! (and very poor health to the second-order approximation). *Advances in Macroeconomics 1*(1).

Carvalho, L. S., S. Meier, and S. W. Wang (2016). Poverty and economic decision-making: Evidence from changes in financial resources at payday. *American Economic Review 106*(2), 260–84.

Cassidy, R. et al. (2018). Are the poor so present-biased? *Institute for Fiscal Studies Working Paper* (18/24).

Chuang, Y. and L. Schechter (2015). Stability of experimental and survey measures of risk, time, and social preferences: A review and some new results. *Journal of Development Economics 117*(151-170).

Collins, D., J. Morduch, S. Rutherford, and O. Ruthven (2009). *Portfolios of the Poor.* Princeton University Press.

Cubitt, R. and D. Read (2007, December). Can intertemporal choice experiments elicit time preferences for consumption? *Experimental Economics 10*(4), 369–389.

Dean, M. and A. Sautmann (2014). Credit constraints and the measurement of time preferences. SSRN Working paper http://ssrn.com/abstract=2423951.

Deaton, A. (1997). *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy.* World Bank.

Deaton, A. and S. Zaidi (2002). *Guidelines for constructing consumption aggregates for* *welfare analysis*, Volume 135. World Bank Publications.

Epper, T. (2017). Income expectations, limited liquidity, and anomalies in intertemporal choice.

Frederick, S., G. Loewenstein, and T. O'Donoghue (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature XL*, 351–401.

Giné, X., J. Goldberg, D. Silverman, and D. Yang (2018). Revising commitments: Field evidence on the adjustment of prior choices. *The Economic Journal 128*(608), 159–188.

Golosov, M., N. R. Kocherlakota, and A. Tsyvinski (2003). Optimal indirect and capital taxation. *Review of Economic Studies 70*, 569–587.

Green, E. J. and S.-N. Oh (1991). Contracts, constraints and consumption. *The Review of Economic Studies 5*, 883–899.

Grosh, M. and P. Glewwe (2000). *Designing household survey questionnaires for developing countries : Lessons from 15 years of the Living Standards Measurement Study*, Volume 1 - 3. World Bank.

Halevy, Y. (2015). Time consistency: Stationarity and time invariance. *Econometrica 83*(1), 335–352.

Hall, R. E. (1978). Stochastic implications of the life-cycle permanent income hypothesis: Theory and evidence. *Journal of Political Economy 86*(6), 971–987.

Harris, C. and D. Laibson (2001, July). Dynamic choices of hyperbolic consumers. *Econometrica 69*(4), 935–57.

Harrison, G. W., M. I. Lau, and E. E. Rutström (2005). Dynamic consistency in Denmark: A longitudinal field experiment. *UCF Economics Working Paper 05-02*.

Karaivanov, A. and R. M. Townsend (2014). Dynamic financial constraints: Distinguishing mechanism design from exogenously incomplete regimes. *Econometrica 82*(3), 887–959.

Karlan, D., R. Osei, I. Osei-Akoto, and C. Udry (2014). Agricultural decisions after relaxing credit and risk constraints. *The Quarterly Journal of Economics 129*(2), 597–652.

Kinnan, C. (2019). Distinguishing barriers to insurance in Thai villages. *Journal of Human*

Resources forthcoming.

Kocherlakota, N. R. and L. Pistaferri (2009). Asset pricing implications of Pareto optimality with private information. *Journal of Political Economy 117*(3).

Krupka, E. L. and M. Stephens (2013). The stability of measured time preferences. *Journal of Economic Behavior & Organization 85*(C), 11–19.

Krusell, P. and A. A. Smith (2003). Consumption-savings decisions with quasi-geometric discounting. *Econometrica 71*(1), 365–375.

Laibson, D. (1997, May). Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics 112*(2), 443–77.

Ligon, E. (1998, October). Risk sharing and information in village economics. *Review of Economic Studies 65*(4), 847–64.

Ludvigson, S. and C. Paxson (2006). Approximation bias in linearized euler equations. *The Review of Economics and Statistics 83*(2), 242–256.

Mace, B. J. (1991). Full insurance in the presence of aggregate uncertainty. *Journal of Political Economy 99*(5), 928–956.

Mahajan, A. and A. Tarozzi (2011, March). Time inconsistency, expectations and technology adoption: The case of insecticide treated nets. *mimeo*.

McClure, S. M., K. M. Ericson, D. I. Laibson, G. Loewenstein, and J. D. Cohen (2007). Time discounting for primary rewards. *The Journal of Neuroscience 27*(21), 5796–5804.

Meier, S. and C. D. Sprenger (2015). Temporal stability of time preferences. *Review of Economics and Statistics 97*(2), 273–286.

Pender, J. L. (1996, August). Discount rates and credit markets: Theory and evidence from rural India. *Journal of Development Economics 50*(2), 257–296.

Rogerson, W. P. (1985). Repeated moral hazard. *Econometrica 53*(1), 69–76.

Runkle, D. E. (1991). Liquidity constraints and the permanent income hypothesis. *Journal of Monetary Economics 27*, 73–98.

Schaner, S. (2015). Do opposites detract? Intrahousehold preference heterogeneity and inefficient strategic savings. *American Economic Journal: Applied Economics 7*(2), 135–74.

Schechter, L. (2007). Risk aversion and expected-utility theory: A calibration exercise. *Journal of Risk and Uncertainty 35*(1), 67–76.

Tanaka, T., C. Camerer, and Q. Nguyen (2010). Risk and time preferences: Linking experimental and household survey data from Vietnam. *American Economic Review 100*(1), 557–571.

Townsend, R. M. (1994). Risk and insurance in village India. *Econometrica 62*(3), 539–591.

Zeldes, S. P. (1989). Consumption and liquidity constraints: An empirical investigation. *Journal of Political Economy 97*(2), 305–346.