Credit Constraints and the Measurement of Time Preferences
PRELIMINARY

Mark Dean* Anja Sautmann†

December 1, 2013

Abstract

Incentivized experiments are commonly used to estimate marginal rate of intertemporal substitution (MRS) in the lab and in the field, and to make inferences about subjects' time preference. This paper considers the implications of an integrated model of behavior in which individuals are subject to financial shocks and credit constraints and take those into account when making experimental choices. The model shows that measured MRS depends on the individual's effective interest rate and her marginal utility of current and future consumption. Experimental responses should therefore be correlated with other variables that describe the subject's financial situation, like savings, income and consumption shocks. We test the model with a panel data set from Mali and find evidence for such effects. We discuss how our model can be combined with repeated time preference measures to identify time preferences and other household characteristics - including credit constraints and the importance of different types of financial shocks.

*Department of Economics, Brown University: mark_dean@brown.edu

†Department of Economics, Brown University: anja_sautmann@brown.edu

We would like to thank Ken Chay, Andrew Foster, Kyle Hyndman, Supreet Kaur, Dilip Mokherjee, Sriniketh Nagavarapu, Nancy Qian, David Weil and participants of seminars at Boston University, Brown University, and Columbia University for their comments and suggestions. All errors are ours. We recognize the support of the Aga Khan Foundation, the Economic and Social Research Council (Development Frontiers Grant ES/K01207X/1), and the Brown University Seed Fund award in the completion of this project.

1 Introduction

Experimental methods have become an important part of the economist's toolkit. Increasingly, applied economists are using incentivised experiments to identify individual-specific preference parameters, and in particular those that govern intertemporal choices (see for example **????**). Oftentimes, such experiments make use of so-called multiple price lists (MPL). A time preference MPL asks subjects in a series of questions to choose between different amounts of money at two given points in time (table 1 depicts the experimental choices used in this paper as a typical example). Effectively, these experiments measure the Marginal Rate of Intertemporal Substitution (MRS) for money between two periods. The underlying assumption in most such studies is generally that the MRS directly reflects "deep" preference parameters, typically discount factors in a temporally separable utility function (see **?** for a comprehensive overview[1]). Recent methodological advances have focused on estimating the curvature of the utility function to more accurately measure the utility changes associated with the different monetary payments [**?**], and on eliciting the point of indifference between different payment times more accurately [**?**].

Outside the experimental realm, however, we know from the standard Euler equation that the MRS is affected not just by time preference parameters, but also by prevailing credit market conditions and the marginal utility derived from consumption at the two different points in time. In the most extreme case of a perfect credit market, the MRS will in fact be fully determined by the market interest rate.[2] At the other extreme, when subjects have no access to outside credit, the MRS will be affected by any income or preference shocks that change the marginal utility of consumption today relative to tomorrow.

In this paper we set out to formally study the relationship between credit constraints, financial shocks and the measurement of time preferences in a model that integrates experimental choices with the subject's broader intertemporal optimization problem. Our model builds on **?**'s analysis of intertemporal decision making when there is hyperbolic discounting and time inconsistency, but allows in addition for imperfect access to credit markets in the face of income and spending shocks. We test the model predictions using a unique panel data set from Mali that contains repeated MPL measurements as well as detailed financial data at the household level. We show that, as predicted by the model, MRS measurements vary systematically with current savings, income, consumption, and spending, and thus cannot be used straightforwardly to back out stable parameters β and δ. We use our model to understand the relative importance of consumption and income shocks in our target population.

Our model starts with the simple assumption that the effective interest rates at which households can borrow and save depends negatively on their current savings; in other words, the conditions for borrowing and saving worsen the larger the amounts borrowed or saved. The extreme cases of no credit constraints (in which the interest rate is constant) and complete credit constraints (in which the interest rate is infinity for any borrowing and negative infinity for any savings) are limit cases of this model. We show that the model implies that experimentally measured MRS in any given time period should be equal to the individual interest rate at the optimal level of savings, which in turn equals the ratio between marginal utility of consumption today and discounted expected marginal utility of consumption tomorrow.[3] Importantly, this conclusion does not rely on the assumption that subjects actively arbitrage experimental payments – only that they adjust optimally to non-experimental shocks, and take this into account when making experimental choices.

[1] In the studies in **?** the reported discount factor is usually simply calculated as the ratio $\frac{x}{y}$ of the payment x today accepted in exchange for payment y in the future.

[2] The relationship between measured MRS and external credit markets has been discussed previously in the literature, for example **?**.

[3] This conclusion holds under the assumption that experimental payments are small relative to income and consumption.

Our model makes predictions for measured MRS that are distinct from existing models, namely that the MRS should be correlated with savings, income and spending shocks. The key elements of the model that lead to these predictions are on the one hand the assumption that experimental choices are affected by subjects' current financial situation, and on the other that our study population has access to imperfect financial instruments. By contrast, the standard approach to modeling choices in time preference experiments is to abstract from choices outside the experiment (see for example **?**), relying on a 'narrow bracketing' assumption which says that subjects make experimental choices with a fixed level of "background consumption" in mind. By definition, narrow bracketing rules out correlations between experimental choices and financial variables. An alternative model that does not assume narrow bracketing but allows for intertemporal consumption smoothing at a fixed market interest rate (i.e. the 'no credit constraints' version of our model) similarly predicts no such correlations, as measured MRS should be equal to the constant interest rate. Finally, a model with spending/income shocks but in which subjects have no ability to borrow and save at all would predict a correlation of MRS with income and consumption, but no correlation with savings; this is because any measured savings are the result of measurement error.

We test the predictions of our model against these alternatives in a unique data set that contains repeated weekly measurements of MRS from 1045 household heads in Bamako, Mali. The data includes detailed weekly measures of spending, income, and whether or not the household has suffered a negative shock (such as damage or loss of property) in that period. We demonstrate that our predictions hold in the data, in particular that MRS is significantly correlated with the level of savings in a given period, a finding that is inconsistent with alternative narrow bracketing models or extreme cases of (full or no) credit constraints. The correlation arises primarily from exogenous shocks to income and spending, demonstrating that the results are robust to endogeneity concerns or reverse causality problems. We use our empirical findings to show that households have particular difficulty smoothing consumption shocks, caused for example by adverse events, while they are better able to respond to income shocks. Our findings open up new avenues of research into the study and measurement of both time preferences and credit constraints.

Only a few recent studies have considered the possibility that measured time preferences depend on subjects' access to financial markets and current financial situation. **?** perform a longitudinal study on a representative sample of 100 members of the Danish population and find that the estimated discount rate tends to fall when a subject reports a worsening of her financial situation since the previous measurement. **?** show that elicited discount rates for hypothetical payoffs respond to macroeconomic events (levels of inflation) as well as individual financial shocks (labor market participation). In contrast, **?** and **?** find no such correlation. We discuss possible explanations for this in section 6.

In applied work and especially in development economics, MPL experiments are increasingly in use to identify individual preferences. At the same time, the quasi-hyperbolic discounting model by **?** has risen to prominence as a potential explanation for the behavioral phenomena observed in the demand for financial products or the uptake of health technologies by the poor. In applied work, this has lead to renewed attention to simple MPL methods that can be used to identify present bias[4] and time inconsistency in large populations in field surveys. The parameters of interest are the present-bias parameter β as well as the discount factor δ, and MPL experiments are used to either estimate these parameters or to classify individual into groups of similar preferences (see e.g. **??**). The results in this paper open up new research questions about the validity of measuring preferences in this manner and ways to improve on the MPL methods. In poor populations who are subject to frequent financial shocks and face substantial credit market constraints it becomes particularly

[4]The tendency to exhibit higher discount rates for choices in which one option is available immediately

Table 1: A Typical Price List Experiment

Decision A		Decision B	
today	in 1 week	in 1 week	in 2 weeks
a_0	a_1	b_1	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

important to ask what these experiments exactly measure and how they relate to other factors that may be relevant for the marginal rate of intertemporal substitution. Our results suggest the intriguing possibility that repeated MPL experiments could be used to estimate time preference parameters in conjunction with the credit constraints that a household faces.

2 Time Preference Experiments and Models of Behavior

Consider the sequence of questions shown in table 1. In Decision A, the subject is asked to make a series of choices between money today and money in one week's time. in Decision B, the subject must make choices between money in one week's time and money in two week's time. These 'multiple price list' (MPL) style experiments have been used in many experimental investigations into time preferences; examples are **?**, **?**, **?**, and **?**.[5]

In this section we discuss two classes of models that can be used to understand behavior in this experiment. The first we refer to as 'narrow bracketing' models, in which decisions in the laboratory are treated in isolation from the outside world. Such models are the basis of analysis in the majority of time preference experiments, even if sometimes implicitly (see for example **???**). Narrow bracketing models imply that time preference experiments can be readily used to understand a subject's time preferences. However, the necessary assumptions are extreme (see **?** for a discussion). We therefore consider a second class of 'integrated' models, in which MPL decisions are taken as part of the subject's broader consumption-savings problem. The model we introduce allows for both time inconsistency (in the manner of **?**) and credit constraints in the form of savings-dependent interest rates. We analyze its implications for the measurement of time preferences with MPL experiments.

2.1 Narrow Bracketing

The term 'narrow bracketing' refers to the idea that behavior in an MPL experiment is divorced from a subjects' day-to-day financial circumstances and decisions, and that it instead reflects deep preference parameters.[6] For example, in an MPL experiment as in table 1, a subject who switches from later to earlier payments at a lower point in decision A than in decision B is assumed to be more "impatient" with regards to immediate payoffs than with regards to future payoffs, and this is interpreted as evidence of present bias.

[5] Typically, one question will be selected at random for payment to avoid wealth/portfolio effects.

[6] See **?** for a description of narrow bracketing, although the idea goes back as least to **?**.

In order to draw this conclusion, experimental subjects must be assumed to ignore both their outside consumption and income, and any opportunity to arbitrage the experimental payments with the help of the financial instruments they use otherwise. To formalize the narrow bracketing model, suppose an individual is making choices between a fixed amount a_1 in $t = 1$ and a varying amount a_0 today in $t = 0$. Utility is time separable and given by the instantaneous utility function $u(c)$ and the discount function δ_t for $t = 1, \ldots, \infty$. Crucially, it is assumed that the individual does not consider his actual current consumption c_t for each point in time, but rather a constant level of "background consumption" $\bar{c}$. Choosing between the two experimental payouts then becomes a choice between $u(\bar{c} + a_0) + \delta_1 u(\bar{c})$ and $u(\bar{c}) + \delta_1 u(\bar{c} + a_1)$. The subject will choose a_0 over a_1 if

$$\begin{aligned} u(\bar{c} + a_0) - u(\bar{c}) &\geq \delta_1 u(\bar{c} + a_1) - \delta_1 u(\bar{c}) \\ \Rightarrow \frac{u(\bar{c} + a_0) - u(\bar{c})}{\delta_1 \left(u(\bar{c} + a_1) - u(\bar{c})\right)} \frac{a_1}{a_0} &\geq \frac{a_0}{a_1}, \end{aligned}$$

and a_1 over a_0 otherwise. If the experimental payments are assumed to be small,[7] the switch point between later and earlier payments in the MPL will approximately occur at

$$\begin{aligned} \frac{u'(\bar{c})}{\delta_1 u'(\bar{c})} &= \frac{a_0}{a_1} \\ \delta_1 &= \frac{a_1}{a_0}. \end{aligned}$$

Similarly, when choosing between payments b_1 at $t = 1$ and b_2 at $t = 2$ in period 0, the switch point will reveal δ_2, at least in approximation. By comparing the two discount factors, we can now for example conclude from $\delta_1 < \delta_2$ that the individual is present biased and time inconsistent.

The crucial assumption in this argument is that the "real world" does not enter subject's decision making in the lab. They neither account for actual variation in their current and future consumption levels, nor do they use the market interest rate to arbitrage (or assess the relative value) of the experimental payments. By contrast, an individual who is subject to financial shocks and takes her current level of outside consumption into account when making experimental decisions may make a "present-biased" choice when she is in dire financial straits, but will make a very "patient" decision on the day of a lottery win.

2.2 An Integrated Model

The assumption that experimental subjects ignore both outside shocks and arbitrage opportunities is clearly extreme. Especially in developing countries, where households are subject to frequent shocks and financial markets are not complete, narrow bracketing may break down on both fronts: households may take their current outside consumption as well as their individual cost of saving and borrowing into account in making experimental choices. We therefore propose an integrated model of time preference measurements in the presence of financial shocks and credit constraints.

Consider the behavior of a decision maker who is infinitely lived. In period 0, their preferences over

[7] **?** introduced a method that removes the need to assume locally linear utility. See section 2.4.2.

consumption streams $\{c\}_{t=0}^{\infty}$ are described as follows:

$$u(c_0) + \beta E_0 \sum_{t=1}^{\infty} \delta^t u(c_t, \rho_t)$$

The parameter $0 < \beta \leq 1$ indexes the degree of present bias. These preferences are quasi-hyperbolic and time inconsistent if $\beta < 1$. The function $u : \mathbb{R}_+ \times \mathbb{R}_+ \to \mathbb{R}$ describes per-period utility, where ρ_t is a stochastic preference parameter that is drawn independently from a distribution f_ρ with support ϱ in each period. The resource constraint is given by

$$\begin{aligned} c_t &= w_t - s_t \\ w_t &= y_t + (1 + r(s_{t-1}))s_{t-1} \\ &= y_t + R(s_{t-1}) \\ w_0 & \text{ given.} \end{aligned}$$

The variable s_t denotes the *stock* of savings at the end of period t, which can be either positive or negative. y_t is income in period t. Income is assumed to be drawn independently from a distribution f_y in each period.

$R(s_t) = (1+r(s_t))s_t$ describes the intertemporal budget constraint. It is modified to allow for the interest rate r to depend on the level of savings. Specifically, we assume that interest rates fall as savings increase (or rise as borrowing increases). This implies that the cost of borrowing rises with the amount of credit, and the return to savings falls as the amount saved increases. We restrict r so that

$$0 \quad < \quad (1 + r(s_t) + r'(s_t)s_t) \equiv R'(s_t)$$

and

$$0 > 2r'(s_t) + r''(s_t)s_t \equiv R''(s_t),$$

The first restriction implies that the resources available in period $t+1$ are increasing in s_t. The second (that R is concave) implies that the marginal rate of return is decreasing in s_t. We further assume that r (and therefore R) is continuously differentiable.

The assumption that R is a concave function is a reduced-form way of modeling credit and savings constraints. We know that households and businesses in developing countries do not have free access to credit. Interest rates vary widely, the rates for borrowing are higher than for saving, and there is evidence for credit rationing, that is, at a given interest rate the total loan size available is restricted below the individual's demand for credit (e.g. **??**). Credit constraints and rationing mean that, unlike in a complete market, the individual level of borrowing affects the interest rate at which the person can borrow. The reasons may lie in the risks of moral hazard and adverse selection for the lenders, and the inability of poor households to put up sufficient collateral. Recent empirical research has also found that there are substantial savings constraints and an unfilled demand for savings instruments. A proportion of savers are in fact willing to accept *negative* interest rates (**??**), indicating that they cannot even have the proverbial "cash under the mattress"; the most basic form of saving money by storing it in the house or on the person until later use. Reasons for these constraints include demands for transfers from others, the risk of theft, loss, or someone

Figure 2.1: An example of the relationship between saving and return.

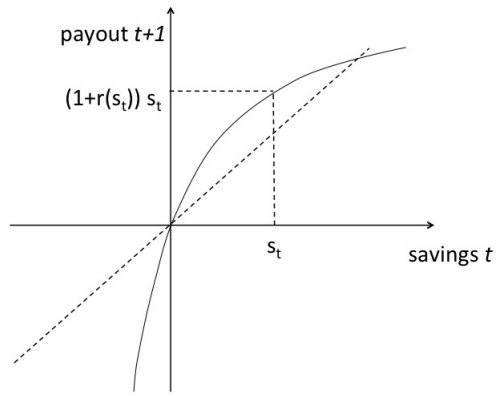

else defaulting on a personal loan, and unexpected increases in inflation (**?**, see **?** for examples).[8] A natural way for these frictions to lead to a concave interest rate function would be for example to assume that households have access to a range of borrowing and savings options at different effective interest rates, and they make use of the most attractive ones first (and pay off loans or dissave from the least attractive options first).

Figure 2.1 illustrates an example of the relationship between s_t and the amount that the decision maker receives in period $t+1$. The degree of credit constraint faced by the decision maker is captured by the curvature of R: The more concave the function, the more the rate of return varies with the amount the decision maker saves or borrows. At one extreme, r is a linear function, in which case the decision maker faces the same interest rate regardless of how much she borrows or saves. We think of such a person as having no credit constraints, and we call this the *no-constraints case*. The no-constraints case represents the standard model of complete financial markets. At the other extreme, as the second derivative $R''(\cdot) \to -\infty$, the cost of borrowing goes to infinity while the rate of return on savings goes to zero.[9] In the limit, the decision maker will effectively face a hard credit constraint, with no borrowing or savings possible. We call this the *full-constraints case*.

There exist several alternative models of credit constraints that manifest themselves in different assumptions about the shape of R. Perhaps the most standard is a model of an asymmetric 'hard' borrowing constraint, in which the agent is free to borrow and save at the market interest rate until their borrowing reaches a hard limit. A hard liquidity constraint is often used in buffer-stock models of wealth and precautionary savings (**??**). A second alternative would be to allow the interest rate to increase with borrowing, but to be invariant (and possibly zero) for savings. Finally, one could imagine a piece-wise linear, "kinked" interest rate function, with one interest rate if the subject is borrowing and another if she is saving.

All of these examples can be approximated arbitrarily well by our model, as in every case the resulting R function is concave (although clearly not continuously differentiable).[10] Thus, the results below all broadly go through for these alternative models, and our focus on a smoothly differentiable return function is largely

[8] Dupas and Robinson also report "unexpected expenditures" for luxury items as reason for women to save less than planned. Constraints that have their roots in the saver's spending behavior are expressions of time-inconsistent preferences, modeled here on the preference side as hyperbolic discounting.

[9] This can be seen by taking the Taylor series expansion $R'(s) \approx R'(0) + R''(0)s$. As $R''(0) \to -\infty$, $R'(s) \to \infty$ for $s < 0$. For $s > 0$, $R'(s)$ falls as $R'(s) \to \infty$, but is bounded below at zero by assumption.

[10] Note also that the piece-wise linear budget constraint becomes indistinguishable from the full-constraints case above if the difference in interest rates for borrowing and saving is so large that the decision-maker never makes use of the financial market.

for analytical convenience. However, we would also argue that our model is more plausible than one with an absolute lower bound on borrowing. It seems likely that, in most cases, people will be able to borrow more as long as they are prepared to pay a high enough interest rate.

One feature that is not commensurate with our model is an interest rate that increases with the level of savings, as suggested by **?**. This would be the case if, for example, a minimum investment was required in a project. We make the assumption that this is not the prevalent situation in our population. One justification is that these lump-sum investment opportunities come along only infrequently, and therefore do not govern day-to-day consumption smoothing for our households.

2.3 The Euler Equation and Marginal Rate of Intertemporal Substitution

We use the results of **?** to derive an Euler equation for a quasi-hyperbolic consumer. Their model analyzes the set of perfect equilibria in stationary Markov strategies of the game between the different 'selves' of this consumer at different points in time. Because shocks are independent over time, the only state variables at time t are cash on hand w_t and the realization of the consumption parameter ρ_t.

? provide a set of conditions under which the equilibrium of the above game can be described by the Strong Hyperbolic Euler Equation (SHEE). We assume these conditions hold in our problem.[11] Essentially, they are standard regularity conditions on the utility function u and the distributions f_ρ and f_y, plus the assumption that the model is parametrized in such a way that the equilibrium consumption function is Lipschitz continuous, which requires β to be not too far from 1 (**?** show that this condition is satisfied for a relatively wide range of parameters).

Definition 1. A consumption function $c : \Re \times \varrho \to \Re_+$ satisfies the Strong Hyperbolic Euler Equation if the following holds for every $w_t \in \Re$:

$$u'(c(w_t, \rho_t), \rho_t) = R'(s_t) E_t \left[\delta \left(1 - (1-\beta) \frac{\partial c(w_{t+1}, \rho_{t+1})}{\partial w_{t+1}} \right) u'(c(w_{t+1}, \rho_{t+1}), \rho_{t+1}) \right] \tag{2.1}$$

The SHEE has the standard Euler Equation interpretation. $u'(c(w_t, \rho_t), \rho_t)$ is the marginal utility of consumption today, while $R'(s_t)$ is the rate at which money today can be converted into money tomorrow. The second term on the right hand side of equation 2.1 is of the form $E_t \delta \beta V'(w_{t+1})$, where $V(w_{t+1})$ is the value of cash on hand in period $t+1$, viewed from the perspective of period t (see equation 2.3).

This version of the SHEE differs from the standard Euler equation in two ways. First, it takes into account the fact that the interest rate changes with savings, so the standard market interest rate term $1+r$ is replaced by the savings-dependent rate of return $R'(s_t)$. Second, since the consumer may be quasi-hyperbolic, the marginal utility of income tomorrow from the perspective of time t is not equal to the expected curvature

[11] 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 ϱ;
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 parametrized such that the equilibrium consumption function is Lipschitz continuous. Effectively this means that β is close to 1.

of the instantaneous utility function; rather, it is modified downward by a factor $(1-\beta)\frac{\partial c(w_{t+1},\rho_{t+1})}{\partial w_{t+1}}$. This adjustment is small either when β is close to 1, or when the marginal propensity to consume is low. If there is any income smoothing, we would generally expect $0 < \frac{\partial c}{\partial w} < 1$. Rewriting $\delta\left(1-(1-\beta)\frac{\partial c(w_{t+1},\rho_{t+1})}{\partial w_{t+1}}\right)$ as $\beta\delta\frac{\partial c_{t+1}}{\partial w_{t+1}} + \delta\left(1-\frac{\partial c_{t+1}}{\partial w_{t+1}}\right)$ shows that the discount factor applied by the decision maker is a weighted average of the short run discount factor $\beta\delta$ and the long-run discount factor δ, where the weight depends on future propensity to consume. If future propensity to consume is high, saving for the future is valued less, as the decision maker anticipates that next period's self will overconsume.

We can rearrange the equation 2.1 to express the marginal rate of intertemporal substitution (MRS) of the consumer as

$$MRS_t \equiv \frac{u'(c(w_t,\rho_t),\rho_t)}{\delta E_t\left[\left(1-(1-\beta)\frac{\partial c(w_{t+1},\rho_{t+1})}{\partial w_{t+1}}\right)u'(c(w_{t+1},\rho_{t+1}),\rho_{t+1})\right]} = R'(s_t). \tag{2.2}$$

Note what this equation implies in the extreme case of no credit constraints, where the interest rate is not dependent on savings: here the MRS of the consumer should not vary, but be constant and equal to the market interest rate. In the limiting case where consumers do not (cannot) borrow or save, consumption is not a choice variable. The identity $w_t = y_t = c_t$ holds in every period, and the MRS will reflect time preference parameters and the relationship between marginal utility of income today and the expected marginal utility of income tomorrow:

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

2.4 Time Preference Measurement

Before discussing the implications of this model further, we need to know what an MPL experiment is actually measuring in this framework. Assume that, at the time the surveyor visits, the subject has optimized her consumption stream given her current period's income level y_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 in one period vs. a_0 immediately. Which will she choose?

Given that optimization has already occurred in this period, one may suspect that the answer to this question depends on if the subject still has access to the financial market and can trade the experimental payments according to $R(s)$, or if she just consumes a_0 and incorporates the payoff of a_1 into her future consumption choices. As it turns out, however, this is not the case. This observation is important, as it demonstrates that our results do not depend on the assumption that experimental subjects will actively arbitrage their experimental payments: only that they have adapted their consumption plan to their interest rate, and take this into account when making their experimental decisions.

If the subject can take the experimental payment to the financial market, then her choice essentially boils down to an arbitrage argument. Assuming a_0 and a_1 are small, the optimal savings rate following receipt of either payment will be close to s_t^*, and so the effective interest rate will be close to $R'(s_t^*)$.

The subject will prefer a_0 over a_1 if the value of a_o at $t+1$ is higher than that of a_1, or

$$\begin{aligned} R(s_t^*+a_0) - R(s_t^*) &\geq a_1 \\ \frac{R(s_t^*+a_0) - R(s_t^*)}{a_0} &\geq \frac{a_1}{a_0} \\ \Rightarrow R'(s_t^*) &\geq \frac{a_1}{a_0}. \end{aligned}$$

If this is the case, then any allocation of a_1 between t and $t+1$ can be bettered by receiving a_0. Conversely, if $\frac{a_1}{a_0} > R'(s_t^*)$, any allocation of a_0 between t and $t+1$ can be improved upon by accepting a_1. Subjects in the experiment should exhibit indifference at $\frac{a_1}{a_0} = R'(s_t^*)$. In effect, the point at which the subject switches from choosing a_1 to a_0 in the experiment measures the slope of the budget constraint at s_t^*.[12]

By contrast, suppose that the experimental income in period 0 is consumed right away, without further opportunity to trade, while period-1 income is added to the subject's stock of wealth in that period. Then the break-even point between a_0 and a_1 is given by

$$\begin{aligned} u(w_t - s_t^* + a_0, \rho_t) + \beta\delta EV(y_{t+1} + R(s_t^*)) &= u(w_t - s_t^*, \rho_t) + \beta\delta EV(y_{t+1} + R(s_t^*) + a_1) \\ \frac{u(w_t - s_t^* + a_0, \rho_t) - u(w_t - s_t^*, \rho_t)}{a_0} &= \frac{a_1}{a_0}\frac{\beta\delta\left[EV(y_{t+1} + R(s_t^*) + a_1) - EV(y_{t+1} + R(s_t^*))\right]}{a_1} \\ \Rightarrow \frac{u'(c_t^*, \rho_t)}{\beta\delta EV'(w_{t+1}^*)} &\approx \frac{a_1}{a_0} \end{aligned}$$

for small experimental payments. But from **?** we know that

$$\beta V'(w_t) = \left(1 - (1-\beta)\frac{\partial c_{t+1}^*}{\partial w_{t+1}}\right) u'(c_{t+1}, \rho_{t+1}) \tag{2.3}$$

and therefore

$$\frac{a_1}{a_0} = \frac{u'(c_t, \rho_t)}{\beta\delta EV'(w_{t+1}^*)} = \frac{u'(c_t^*, \rho_t)}{\delta E_t\left[1 - (1-\beta)\frac{\partial c_{t+1}^*}{\partial w_{t+1}}\right] u'(c_{t+1}, \rho_{t+1})}.$$

This last expression, of course, is just equal to the marginal rate of substitution at t, so in other words, in the case without arbitrage, the experiment measures the slope of the indifference curve at s_t^*. Since in the optimum the slopes of the budget constraint and the indifference curve are equal, the two possible protocols both lead to the same conclusion: the MPL switch point measures MRS_t.

2.4.1 Trade-Offs in Future Time Periods (Decision B)

Before we move to the predictions of the model we discuss what is measured with decision B in the MPL experiment, where the subject chooses between payoffs at two future points in time. In what follows, define

$$\delta\left(1 - \frac{dc_{t+1}}{dw_{t+1}}\right) + \beta\delta\frac{dc_{t+1}}{dw_{t+1}} = d(w_{t+1}),$$

so that

$$\beta\delta V'(w_{t+1}) = d(w_{t+1})u'(c_{t+1})$$

(using equation 2.3). The function $d(\cdot)$ simply describes the modified discount factor as a function of current wealth. Note that d is increasing in w_{t+1} if the marginal propensity to consume is declining in wealth and there is hyperbolic discounting.

The key observation here is that the subject's self at time t anticipates that next period's self knows

[12] Note that this argument must be reexamined if the budget constraint is discontinuous at the optimal consumption level (the household is at a "kink"). In this case, a different interest rate applied to a_0 and a_1 may lead to a convex portion of the budget constraint (see **?**). However, if the payments are small, it can be shown that the subject will apply her 'personal' interest rate given by the slope of her indifference curve at that point.

which payoff she will get from the experiment. Thus, her consumption decision will take this additional income into account, regardless if it accrues in $t+1$ or $t+2$. Since today's self cannot influence future selves' decisions, all that will matter to her is the present value of the two income streams, discounted at the relevant interest rate. The best she can do for her expected utility is to choose the highest level of wealth at time $t+1$.[13] This means that if optimal savings at date $t+1$ are s^*_{t+1}, then b_2 at $t+2$ is equivalent to an increase in wealth Δw at time $t+1$ that solves

$$R(s^*_{t+1} + \Delta w) - R(s^*_{t+1}) = b_{2,}$$

Assuming that these payments are small, we can approximate the interest rate as that faced at s^*_{t+1} and

$$\Delta w = \frac{b_2}{R'(s^*_{t+1})}.$$

This means that the subject is indifferent between b_2 and b_1 if

$$E_t V\left(w_{t+1} + b_1\right) = E_t V\left(w_{t+1} + \frac{b_2}{R'(s_{t+1})}\right),$$

and using a linear approximation of the change in the value function and substituting for $\beta\delta V'(w_{t+1})$, we get

$$\begin{aligned} E_t\left[d(w_{t+1})u'(c_{t+1},\rho_{t+1})\right]b_1 &= E_t\left[\frac{d(w_{t+1})u'(c_{t+1},\rho_{t+1})}{R'(s_{t+1})}\right]b_2 \\ \Rightarrow \frac{b_2}{b_1} &= \frac{E_t\left[d(w_{t+1})u'(c_{t+1},\rho_{t+1})\right]}{E_t\left[d(w_{t+1})\cdot E_{t+1}\left[d(w_{t+2})u'(c_{t+2},\rho_{t+2})\right]\right]} \end{aligned}$$

where we make use of the SHEE for period $t+1$ to substitute for $R'(s_{t+1})$. Note that this is very similar to the SHEE for period t , except for the additional discounting terms and the expectation term in the numerator. Unlike in the narrow bracketing model, however, since today's self realizes that tomorrow's selves will undo her allocation choices, she applies *tomorrow's* effective discounting factor in making her choice, and $\frac{b_2}{b_1}$ does not systematically differ from $\frac{a_2}{a_1}$ due to present bias. Can we then assume that $\frac{b_2}{b_1}$ is in expectation equal to $\frac{a_1}{a_0}$? The answer is in general no. Substituting the Euler equation once more into the above expression gives

$$\begin{aligned} \frac{b_2}{b_1} &= \frac{E_t\left[d(w_{t+1})\cdot E_{t+1}\left[d(w_{t+2})u'(c_{t+2},\rho_{t+2})\right]\cdot R'(s_{t+1})\right]}{E_t\left[d(w_{t+1})\cdot E_{t+1}\left[d(w_{t+2})u'(c_{t+2},\rho_{t+2})\right]\right]}. \\ &= E_t\left[R'(s_{t+1})\right] + \frac{Cov\left(d(w_{t+1})\cdot E_{t+1}\left[d(w_{t+2})u'(c_{t+2},\rho_{t+2})\right], R'(s_{t+1})\right)}{E_t\left[d(w_{t+1})\cdot E_{t+1}\left[d(w_{t+2})u'(c_{t+2},\rho_{t+2})\right]\right]} \end{aligned}$$

The sign on the covariance term will decide how $\frac{b_2}{b_1}$ compares to $E_t R'(s_{t+1})$, which is equal to the expectation of $\frac{a_1}{a_0}$.[14] Consider for example the exponential discounting case, where $d(w_{t+1}) = \delta$, and assume there are no consumption shocks. In this case the MPL measurement is systematically biased upwards, since both u' and R' vary negatively with s_{t+1}. This implies that subjects should appear less patient in future decisions. This continues to hold in non-pathological cases with hyperbolic discounting where the marginal propensity to consume does not respond too strongly to changes in wealth and $d(w)u'(c)$ is decreasing in w. Note that

[13] Assuming that the experimental payments do not affect the interest rate at the optimal savings rate.

[14] Assuming that the economy is stationary, so $E_t R'(s_{t+1}) = E_t R'(s_t)$

the bias becomes small if either there are no credit constraints (R' is constant) or if credit constraints are very high and savings vary little with income (and c_{t+2} becomes independent of s_{t+1}).

2.4.2 Implication for the Measurement of Time Preference Parameters

The argument above has implications for the use of MPL to identify present bias and time discounting factors. Even in expectation $\frac{a_1}{a_0}$ does not give the present discount factor $\beta\delta$. Moreover, it is not possible to compare $\frac{a_2}{a_1}$ and $\frac{b_2}{b_1}$ to learn about β, both because a sophisticated hyperbolic agent will discount future payments at $d(w)$, not at δ, and because there is a systematic difference between the two that is unrelated to time preference parameters.[15] There is an exception, and that is the full-constraints case. Here, self 0 can impose her preferences on future selves by choosing the allocation of payoffs for them, and we have

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

Time preference parameters can therefore be extracted using the average of measured MRS. However, it is important to note that any given observation $\frac{a_1}{a_0}$ will still be affected by consumption and income shocks.

A second case in which the difference between decision A and decision B can teach us something about present bias is if the agent has hyperbolic preferences, but is "naive", that is, she mistakenly believes that future selves will not be present-biased. We do not model this case here, but note that even when the agent is naive, as long as she has access to credit we need information on marginal utility and consumption in addition to MRS measurements in order to identify β and δ.

The experimental literature has innovated in several ways on the original MPL elicitation method in order to recover time preferences under more general assumptions. **?** introduced the "dual MPL method" which estimates risk preferences in addition to time preferences to account for the utility function curvature at $\bar{c}$ in calculating the relative utility derived from a_0 and a_1, in order to avoid using a linear approximation. However, this does not allow for the fact that the slope of the utility function may differ across time periods in the presence of consumption shocks. **?** modify the dual MPL method to allow subjects to choose a point on a continuous budget constraint between payouts at earlier and later points in time. This allows both for utility curvature and discounting to be estimated in a single set of choices, and for point, rather than set identification of the relevant parameters.

Both **?** and **?** suggest that outside consumption and preference parameters could be jointly estimated using a utility function of the form $(c_t + w_t)^\alpha$. Identification is achieved by examining how measured MRS changes with the size of the experimental payoffs. If there is no change (including if the subject is always at a corner solution in the convex budget sets of **?**) then the parameters are not identified. This requires experimental payments to be large relative to outside consumption. While this method can control for shocks in the absence of smoothing (e.g. in the full-constraints case), it breaks down in the case of our model of partial credit constraints. In this case, outside consumption is not a parameter, but the result of optimal intertemporal allocation of resources. Changes in the magnitude of experimental payments will only change

[15] Of course, other authors have commented on the relationship between measured MPL and time preferences once the narrow bracketing assumption has been weakened. See for example, **?** and **?**.

the subject's measured MRS if they change the interest rate, which in turn will change the way resources are allocated between periods. Thus there is no unique outside consumption to estimate. The degree to which this is a problem is related to the degree of credit constraints – the less constrained the subject, the larger the change in allocation to bring about a given change in measured MRS.

In general, measured MRS in our setting could be used to identify time preference parameters, even in the presence of no credit constraints, if we also have an estimate of consumption and marginal utility. Broadly speaking, for a given MRS, a more patient decision maker will have a lower level of consumption today relative to tomorrow.

2.5 Model Predictions: MRS and savings, income, and consumption

We will now use the model to make predictions about the covariance between the MRS, measured with the help of the MPL experiment, and savings, income, and consumption.

First, we note that an increase in s_t will lead to a fall in MRS_t in all but the pathological case of linear R. This follows directly from equation 2.2 and the fact that $R'' < 0$. In fact, the correlation of MRS and savings is entirely driven by R and delivers a first prediction of the model.

Prediction 1: *Savings in period t are negatively correlated with MRS in period t.*

Prediction 1 serves as an important test of a central tenet of the model, that households are subject to a "soft" credit constraint, that is, they can borrow and save to some degree, but the interest rate is not completely fixed. Note that the savings prediction does not hold either under narrow bracketing or in the no-constraints case of complete financial markets. In the full-constraints case where households can neither borrow nor save, there cannot be any (positive or negative) savings.

Next, we consider the impact of an exogenous income shock on the marginal rate of intertemporal substitution in the same period. It is easy to show that savings must increase with cash on hand w. Consider two optimal savings levels s and s' at wealth $w_t > w_t'$. It must be the case that

$$\begin{aligned} u(w_t - s, \rho) + \beta\delta EV(R(s) + y_{t+1}) &\geq u(w_t - s', \rho) + \beta\delta EV(R(s') + y_{t+1}) \Rightarrow \\ u(w_t - s, \rho) - u(w_t - s', \rho) &\geq \beta\delta EV(R(s') + y_{t+1}) - \beta\delta EV(R(s) + y_{t+1}) \end{aligned}$$

If $s \leq s'$, the left hand side is positive: it is the increase in instantaneous utility from reducing savings from s' to s. Moreover, by the concavity of u we have

$$\begin{aligned} u(w_t' - s, \rho) - u(w_t' - s', \rho) &> u(w_t - s, \rho) - u(w_t - s', \rho) \\ \Rightarrow u(w_t' - s, \rho) + \beta\delta EV(R(s) + y_{t+1}) &> u(w_t' - s', \rho) + \beta\delta EV(R(s') + y_{t+1}) \end{aligned}$$

contradicting the optimality of s' for w_t'. Thus, an increase in cash on hand will increase savings and reduce measured MRS.

Third, we consider the effect of preference shocks. In general, measured consumption spending and true consumption "value" do not always perfectly line up. If we think of the c in the utility function as all expenditure on consumption goods in a given time period, we have to account for consumption spending that does not translate into immediate utility gains. An example would be an event which requires an unexpected and possibly urgent expenditure, such as the theft of a productive asset, illness of a family member, or damage to one's house; the adverse event and the subsequent spending to "undo" the event do not actually increase the decision maker's utility in the same way as, say, a meal would. Another example is a large exogenous

lump-sum expense that has to be made in full, like paying one's rent. A household that is subject to such an adverse event or lump-sum cost has a higher marginal utility to additional consumption spending than a household with the same level of c but without this type of expenditure. We model this as a spending shock that increases the marginal utility of consumption everywhere. Note that these types of shocks could also be seen as negative income shocks. Importantly, such shocks will typically *raise* consumption expenditure c while at the same time increasing marginal utility from consumption.

We can show formally that a shock of the type described above will lead to an increase in consumption and a reduction in savings. Assume that preference parameters ρ and ρ' are such that $\frac{\partial u(c,\rho)}{\partial c} < \frac{\partial u(c,\rho')}{\partial c}$ for all c. Let s and s' be the optimal savings rates for the two parameters respectively (for some fixed w_t). If $s \leq s'$ (i.e. the consumer saves more when current-period marginal utility is higher) it must be the case that

$$\begin{aligned} u(w_t - s, \rho') - u(w_t - s', \rho') &> u(w_t - s, \rho) - u(w_t - s', \rho) \\ &\geq \beta\delta EV(R(s') + y_{t+1}) - \beta\delta EV(R(s) + y_{t+1}) \\ \Rightarrow u(w_t - s, \rho') + \beta\delta EV(R(s) + y_{t+1}) &> u(w_t - s', \rho') + \beta\delta EV(R(s') + y_{t+1}), \end{aligned}$$

a contradiction.

These two observations give rise to prediction 2.

Prediction 2: *Exogenous preference shocks that monotonically increase current-period marginal utility increase consumption spending and MRS. Exogenous income shocks cause MRS to fall.*

Outside of exogenous shocks to consumption and income, the relationship between these variables and MRS is less clear. Consider for example the relationship between consumption and MRS generated by an income shock. In general, consumption may be increasing or decreasing in wealth, as illustrated by **?**. But in the case of exponential discounting (i.e. $\beta = 1$), an increase in income will always lead to an increase in period t consumption.[16] In the case of hyperbolic discounting, where $\beta < 1$, the marginal propensity to consume $\frac{\partial c(w_{t+1},\rho_{t+1})}{\partial w_{t+1}}$ may be decreasing in wealth, and this can lead to pathological cases in which an increase in income at time t makes the consumer reduce consumption in order to bind their $t+1$ self to a lower consumption rate. However, **?** show that, if β is sufficiently close to 1 and f_y and u are three times continuously differentiable, then the consumption function will be strictly increasing. Thus, an increase in income will generally lead to both an increase in consumption and a fall in MRS, causing a negative correlation. Taken together, the presence of income and consumption shocks means that the correlation of consumption spending and MRS is indeterminate. If households are subject to important consumption shocks, then MRS and c will be correlated positively. On the other hand, if variation in consumption is driven mostly by income shocks, then MRS and c with be correlated negatively. This will later on provide us with a test which types of shocks are most important for these households, in that they cannot be smoothed well through the financial market.

Our model also implies that income and consumption shocks that affect savings will have some persistence, because a change in savings s_t will lead to a corresponding change in w_{t+1} for every realization of y_{t+1} (as

[16] We can see this by considering the limit of the optimal policy for the finite T period problem, which we can solve by backward induction. In period T it is clear that consumption is increasing in w_T, as everything will be consumed. Assume that consumption is an increasing function of income in period $t+1$, then it must be the case that it is an increasing function in period t. Assume not, and assume that for some w_t equation 2.1 holds. If $\frac{\partial c(w,\rho_t)}{\partial w} \leq 0$ then an increase in w will lead to an (weak) increase in $u'(c(w,\rho_t),\rho_t)$, a strict increase in savings (which would lead to a fall in $R'(s_t)$) and a strict increase in w_{t+1} for every possible realization of y_{t+1}. By assumption, this would increase $c(w_{t+1},\rho_{t+1})$ in every state and so reduce $u'(c(w_{t+1},\rho_{t+1}),\rho_{t+1})$. Thus the left hand side of equation 2.1 would increase while the right hand side would fall, a contradiction.

$R' > 0$). Thus, a shock in period t acts as an income shock of the same sign in period $t+1$. Similarly, a consumption shock that increases the marginal utility of consumption in period t will lead to a decrease in savings in period t, and acts like a negative income shock in period $t+1$. However, the relationship will be attenuated if credit constraints are high and savings therefore respond little to shocks, or if financial shocks are large, so that changes in asset holdings matter little compared to current income and spending shocks. These observations provide our third prediction:

Prediction 3: *Exogenous income shocks in period t increase savings and decrease MRS in period $t+1$. Preference shocks that monotonically increase current-period marginal utility decrease period $t+1$ savings and increase MRS in $t+1$. The effect is smaller if credit constraints are high and shocks are large.*

3 Data

The MPL experiments were carried out as part of a larger panel survey in Sikoroni, a peri-urban area at the outskirts of Bamako, Mali. The survey was the baseline of a randomized control trial for a health care program for children. It took place over seven weeks in the Fall of 2012 after Ramadan. Households were visited weekly by a team of trained surveyors and answered questions on income and spending. The first visit included a set of detailed demographic questions, and a survey module on access to credit and borrowing and lending was part of the last visit. In weeks 1 to 4, household heads participated in a multiple price list time preference experiment.

Table 2 shows summary statistics for the population of subjects. There are in total 1017 households for which we have some information on financial variables and time preference measurements. Household heads who were working long hours or traveling and could not be reached or found did not participate in the MPL experiments. Our analysis includes all household heads in the sample that were present at least for some of the survey visits.

Our participants are overwhelmingly male and in prime working age, although only a small proportion hold a salaried job. About half are fully literate. Households have on average more than six people, with a large number of children and, due to polygamous households, slightly more than one wife/mother per household. The sample is fairly characteristic for the area, but there is some degree of selection: survey participants were chosen according to the criteria of the NGO providing the health care program based on a proxy-means test for income. This means that the household had to have children under five; moreover, part of the proxy test were questions about formal savings instruments and steady employment. Household heads who said that they had a savings account or that one household member was holding a salaried job at the time of the proxy test were not included in the program and therefore not part of the survey.

Although life in Sikoro is shaped by the rural traditions and institutions of much of Mali's population (there are for example traditional healers, story tellers, and "neighborhood chiefs" much like the village chiefs outside the city), labor market and income structures are predominantly urban. While households do own the occasional farm animal or cultivate a small plot, domestic production is relatively unimportant, and the household heads typically report income earning activities like sale and resale, motorcycle or car repair, taxi driving, or tailoring. Only a minority of households own any business assets, so the reported income from these activities is largely the return to labor.

The time preference experiment consists of one or more multiple price list experiments over payoffs at different points in time. We use three weeks of data from weeks 2 to 4 of the survey on the two multiple price lists shown in table 1. These MPL measure trade-offs between the current week and the next, and next

Table 2: Characteristics of experimental subjects.

	N	Mean	SD
Salaried employment	1008	12.2%	0.327
Male	1009	87.2%	0.334
$<$ 25 years	1009	4.96%	0.217
$>$ 44 years	1009	26.26%	0.440
Can read and write	1009	50.05%	0.500
Household size	1013	6.290	3.146
Children 0-4	1013	1.567	0.794
Mothers	1013	1.190	0.459

week and one week after.[17] Each decision in the MPL is a choice between a payment of CFA 300 (about US$ 0.60) made 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 design 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 lists of the current visit was randomly selected for payout, using a bag of numbered dry beans from which the subject drew one at the end of the experiment. Subjects then either received their monetary payout, if a decision was selected that prescribed an amount of money paid today, 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 any 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 time preference experiment in week 1 consisted only of choices over payouts in either one week or two weeks from the visit to demonstrate to subjects that the surveyors would actually return next week. The choices from this experiment are not used here.

Table 3 shows a summary of the three weeks of MPL choices. Decision A refers to the trade off between today and a payment in one week's time, decision B to trade-offs between payoffs one and two weeks away. The top row in the table shows the first decision in the list in which the earlier payoff was chosen. For example, if the subject chooses CFA 300 in one week over CFA 150 today, but CFA 200 today over CFA 300 next week, the number used is CFA 200. The lowest possible MPL value is therefore CFA 50, the highest value was set to 450 (for individuals who chose the later payment always). Note that reporting exact "indifference points" between earlier and later payments is not possible given the discrete nature of the experimental choices and the problem of classifying individuals who either chose the earlier payment or the later payment always. We will discuss this in some more detail below.

The table shows that subjects' decisions in the different MPL experiments are somewhat persistent: a sizable proportion choose the same switch point in both decision A and decision B, and across weeks. That said, the correlation is far from perfect; between 30 and 40% of subjects in each week choose different switch points in A and B in the same week, and the correlation of choices between weeks is only around 0.61-0.67. In 13-18% of cases subjects make a more "patient" choice in decision A than decision B, and around 6-10%

[17] In week 3, an additional MPL experiment was carried out which is not used here, concerning choices between payouts two and three weeks away.

Table 3: Experimental choices.

	week 1		week 2		week 3	
	A	B	A	B	A	B
avg. switch at or below (CFA)	157.0	155.6	153.5	152.4	158.4	154.6
correlation A	weeks 1 and 2: 0.61			weeks 2 and 3: 0.67		
correlation B	weeks 1 and 2: 0.62			weeks 2 and 3: 0.64		
A=B	64.40%		65.39%		69.82%	
more patient in A	18.47%		16.17%		13.32%	
more patient in B	17.13%		18.45%		16.86%	
pay neg. interest	9.66 %	8.15%	7.38%	5.52%	7.37%	6.86%
inconsistent	14.76%	13.93%	10.16%	11.71%	11.13%	10.51%
N	969		965		961	

of subjects are willing to pay a weakly negative interest rate, choosing CFA 300 in one week over CFA 300 or more today. Both patterns cannot be explained in the standard "narrow bracketing" framework with discounting, but they are possible in the presence of financial shocks. As is fairly typical in these types of experiments, a proportion of subjects were recorded as making inconsistent decisions. These are choices in which the surveyor recorded repeated switches between earlier and later payoffs. As is standard in the literature, these observations are excluded from the regressions below.

The lack of a clear "present bias", that is, a systematically higher level of "impatience" in decision A over decision B, is notable in our data, and perhaps surprising given earlier results. However, our findings are similar to recent work by **?**, in which repeated time preferences are taken as in our data with the experimenter visiting the subjects on each occasion, reducing both transaction costs and implicit risk of the later payment choice.

In addition to the weekly time preference measures, our data also contains weekly income and spending data. To avoid problems of asymmetric information, this information was collected from the household head and the wife/wives separately. Income data was collected by source and can be broadly categorized into labor income (income from sales, contracting work, or salaried work) and non-labor income (rent collected as a landlord, government and military pensions and other transfers, transfers from charitable organizations, loan repayment, etc.). Spending was equally split into different categories, like food expenditure, household items, durables, rent payments, debt service, etc. Of particularly importance for our analysis is the expenditure category of "adverse events": here we asked for the occurrence of and amount spent on events like the loss of an animal, damage to a building, or illness of a family member.

Table 4 reports summary statistics of the financial variables of interest. Income and spending are shown both for the household head only and for the whole household (using separately reported amounts by the women in the family). We additionally break out income from labor and non-labor sources. The idea is to separate income from "exogenous" sources over which the household has little influence from income which can be affected by the household's choice of labor supply (see below). Earnings from selling household possessions were not included in either category, both because the decision to sell an item may be an endogenous response to financial shocks and because the sales price does not constitute new earnings, but rather an exchange of an illiquid for a more liquid asset.[18] Spending includes any monetary outlays of the household, purchases of

[18]Including or excluding large sales in either income category has no effects on our results because these sales are rare in our sample.

Table 4: Weekly Savings, Income, and Consumption.

	min	med	mean	max	sd
Income	0	30	58.66	1309	92.38
Income (household head)	0	20	38.69	1110	58.65
Labor income	0	20	32.53	260	41.40
Nonlabor income	0	0	7.95	1001	41.54
Spending	0	62.00	93.51	1210	102.59
Spending (household head)	0	49.10	71.91	1210	76.08
Adverse event	0	0	33.2%	1	0.471
Event spending	0	0	5.80	600	23.57
Savings (flow)	-1176	-24.50	-34.85	1147	80.88
Savings (flow) (household head)	-1176	-21.70	-33.21	955	75.62
No formal credit	0	0	46.6%	1	0.499

All amounts converted to US$.

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 RoSCA[19] or to pay off a debt. We also break out expenses for adverse events, again with the idea that these events constitute primarily spending shocks. Savings in the table are given by the difference between income and spending.

Some notes on data quality and the match with the model variables are in order. First, savings as reported here are a flow variable, whereas savings in the model s_t are a stock variable. The true stock of savings is unobserved. We discuss this problem further in the next section. Second, spending does not directly correspond to consumption, but rather represents the outflow of cash. Separating consumption from spending is very difficult, as true consumption is unobserved. We partly deal with this in the model by allowing for the presence of spending shocks which do not directly contribute to consumption utility (see above). Empirically, we test the effects of spending shocks by using reports on adverse events. However, the imperfect link between consumption and spending opens another possibility, namely that some variation in spending is an endogenous response to shocks and not related to changes in consumption utility. In particular, consider spending on durable consumption items. This could be true consumer durables like a TV, or a bag of rice that is eaten up over many days. Just like an adverse event, the purchase of a durable does not translate fully into an increase in consumption utility, and it will raise the marginal utility of consumption. However, unlike a true shock, this type of spending can be timed to coincide with positive income shocks. This type of endogeneity, which arises from changes in financial variables in *response* to exogenous shocks, rather than the shocks themselves, also exists on the income side. Specifically, in response to (exogenous) consumption or income shocks, households can adjust their labor supply in the market to achieve partial smoothing. This means that, in order to test the effect of income and consumption shocks on MRS, we have to use variation in those variables that is plausibly exogenous rather than endogenous.

As a third point, we may be concerned about selection bias; it is possible 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. We can only indirectly test this by comparing financial variables for households that have some weeks of missing data and households that do not. It appears that there is some selection bias: households with missing experimental data in one or two weeks

[19] Rotating savings and credit association, a form of savings club, called "tontine" in Mali.

(411 out of 2547 observations) have on average lower spending, income, and savings. The difference is significant at the 10% level for income; households with missing data report on average \$52 compared to \$60 weekly income. Spending on adverse events is nearly identical for both types of households, so they seem to be subject to similar shocks.

Lastly, the financial data is likely measured with error. In particular, using our information on consumption and income, it appears that there are on average negative savings. While it is possible that our sample of households as a whole is dis-saving, this is typical for household surveys and commonly interpreted as a sign of under-reported income (see e.g. **?**) It is therefore possible that our savings variable is subject to substantial measurement error. We will address this again when discussing individual empirical tests.

4 Testing the Predictions of the Model

In this section we test if predictions 1-3 hold in our data. The basic empirical model is

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

where X_{it} is the financial variable of interest (converted to US\$ 100) and ϵ_{it} is an error term clustered at the person level. The coefficient α_i is a person fixed effect, and in some specifications we also allow for period fixed effects γ_t.

MRS_t represents the marginal rate of substitution measured by the MPL experiment. Since we only have discrete brackets given by the nine possible switch points in the list, we approximate the subject's MRS by taking the ratios of payments tomorrow and payments today for each decision, and using the midpoint between those ratios. The MRS for individuals who never switch decisions within an MPL cannot be inferred from the experiment, so we chose several plausible endpoint values and checked the robustness of our estimates to these changes. The results reported here use 0.708 as the lowest and 8 as the highest MRS. Later on we will deal with the censoring of these observations more carefully.

4.1 Prediction 1: Savings and MRS

The defining feature of the model we propose is the correlation between savings and MRS. This correlation is within the model framework almost mechanical: it is simply given by the curvature of the individual's budget constraint. At the same time, neither the narrow bracketing model nor the two extreme cases of no constraints or full constraints generate a systematic correlation between savings and the marginal rate of intertemporal substitution. In the narrow bracketing model, savings may actually vary as a consequence of financial shocks; however, the subjects do not take these outside factors into account in their experimental choices. In a complete financial market, the interest rate and therefore the MRS should not vary even without narrow bracketing. Finally, if subjects were fully credit constrained (or at a kink of their budget constraint), then savings would never be different from zero. In this case our measure of savings would be entirely a product of measurement error, and again there should be no correlation with MRS.

Table 5 shows the estimation results using head's savings and household savings as the independent variable X_{it}. The effect of savings onto MRS and therefore the interest rate is clearly and significantly negative, both for the individual and the household as a whole. Note that our savings variable corresponds not to s_t in the model, but to $y_t - c_t$, or the *change* in savings at t. This is justified as long as we assume that the stock of savings $R(s_{t-1})$ at period entry is identically distributed over time and is therefore captured by

Table 5: Savings and MRS_t.

	(1) MRS (A)	(2) MRS (A)	(3) MRS (A)	(4) MRS (A)
Savings (I-E) Head	-0.239**	-0.226**		
	(0.0802)	(0.0822)		
Savings (I-E)			-0.277***	-0.267***
			(0.0742)	(0.0760)
Constant	4.603***	4.690***	4.585***	4.672***
	(0.0267)	(0.0646)	(0.0262)	(0.0641)
Ind. fixed effects	yes	yes	yes	yes
Time fixed effects		yes		yes
Observations	2547	2547	2547	2547

Standard errors in parentheses

+ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

the individual fixed effect and the error term ϵ_{it}. We do not observe the stock of savings in the first week, and in the following weeks, we only observe s_t, but not $R(s_t)$. Moreover, any measure of cumulative savings constructed from the individual s_t would amplify the measurement error. We therefore report regressions that use only the individual variation coming from changes in savings for each household for identification. When including lagged savings as a proxy for $R(s_{t-1})$ in these regressions, the estimated coefficients on current savings do not significantly change.[20]

One concern the reader might have is that the correlation between MRS and savings is not due to credit constraints after all. The model postulates that the correlation is driven by the curvature of the budget constraint and that the level of savings changes the MRS; however, one might suspect that the causal relationship is actually reversed and that the MRS, as a measure of an individual's preferences, drives changes in savings behavior, completely independently of credit constraints. Specifically, consider the possibility that β and δ are not constant, but subject to preference shocks. In this case, a low $\beta\delta$ will lead the individual to be less patient in this period and dis-save, while at a high $\beta\delta$ she will accumulate savings. As is straightforward to see, within our model the presence of shocks to β and δ does not change the fundamental fact that the MRS equals the slope of the budget constraint. Thus, as long as subjects are taking the outside world into account in their experimental choices, the identity $MRS_t = R'(s_t)$ remains unchanged. If there are no credit constraints, preference shocks can only cause a simultaneous change in savings and measured MRS if the subject engages in a very strong form of narrow bracketing: she adjusts her financial choices to the shock (changing savings), and she adjusts her choices in the experiment to the shock (choosing a different $\frac{a_2}{a_1} = \frac{1}{\beta\delta}$), but in her experimental choice, she *ignores* the change in her marginal utilities from consumption that results from her changed savings behavior.

Nonetheless, we would like to address this "reverse causality" concern, and we do so in the next section where we examine income and consumption shocks.

[20]The coefficients on past savings are all positive, but not significant (the coefficients on current savings are virtually unchanged). This is likely due to both measurement error and the imperfect mapping of s_{t-1} into w_t, as well as the lower number of observations when including lags.

Table 6: Consumption shocks and MRS_t.

	(1) MRS (A)	(2) MRS (A)	(3) MRS (A)	(4) MRS (A)
Adv. event (0/1)	0.284*	0.263*		
	(0.124)	(0.124)		
Adv. event expenses			0.256+	0.237+
			(0.147)	(0.141)
Constant	4.588***	4.678***	4.665***	4.755***
	(0.0411)	(0.0737)	(0.00869)	(0.0586)
Ind. fixed effects	yes	yes	yes	yes
Time fixed effects		yes		yes
Observations	2547	2547	2543	2543

Standard errors in parentheses
+ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

4.2 Prediction 2: Income and Consumption shocks and MRS

Prediction 2 states that consumption shocks should be positively correlated with MRS, and income shocks negatively. As was alluded to further above, a problem in testing this prediction is that especially labor income and durables spending may be endogenous rather than exogenous; see for example **?** on "income smoothing" strategies that households use in order to achieve consumption smoothing. Moreover, depending on the relative importance of spending and income shocks, the relationship between consumption spending and MRS is indeterminate and cannot be used as a test of the model.

We overcome this problem by testing the effect of those types of spending and income onto MRS that are plausibly exogenous shocks. For spending shocks, we use what we call "adverse events"; illness, damage or destruction of a household possession, and other unexpected calamities. We have information on both the occurrence of such events and any spending on them.

The results are shown in table 6. Clearly, the occurrence of an adverse event has a significant and positive effect on the MRS. Moreover, even though expenditure on adverse events may be partly endogenous – households may delay repairs or choose less expensive temporary solutions if income is tight – the effect remains significant at the 10% level when using event expenditure as the independent variable.

Deciding which portion of income variation is exogenous and which is endogenous is less obvious. For example, it is likely that households can affect the timing of many sources of income depending on their current need. We approach this by breaking out labor income and non-labor income as described above. Labor income includes all income-earning activities that require labor input from the household (work in their own business, work for a piece-rate or time dependent pay, salaried work). Non-labor income includes important income sources like rent paid to the household, government and military pensions and transfers etc. Variation in non-labor income is likely not entirely exogenous, but the household has relatively less power over non-labor income than over labor income. Indeed, table 7 supports this notion: whereas the effect of labor income onto MRS is small and insignificant, non-labor earnings have a large and significant negative effect, confirming prediction 2.

So far we have not directly addressed the reverse causality concern noted earlier. One might argue that the effects we see are due to preference shocks that change the MRS, which then affects households' financial choices, rather than financial shocks that affect the effective interest rate and MRS of the household. However,

Table 7: Income and MRS_t.

	(1) MRS (A)	(2) MRS (A)	(3) MRS (A)	(4) MRS (A)
Nonlabor earnings	-0.364**	-0.361**		
	(0.112)	(0.112)		
Labor earnings			0.0690	0.0537
			(0.155)	(0.156)
Constant	4.714***	4.806***	4.658***	4.754***
	(0.0377)	(0.0665)	(0.0623)	(0.0860)
Ind. fixed effects	yes	yes	yes	yes
Time fixed effects		yes		yes
Observations	2542	2542	2498	2498

Standard errors in parentheses
+ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

note that the evidence on consumption and income shocks actually points in the opposite direction. If we think that the correlation of MRS and financial variables is due to the household's response to changes in preferences, then the correlation should be strongest for the endogenous components of income and spending. At the same time, exogenous shocks, which operate through the effect on marginal utility from consumption and on $R'(s_t)$, should not change the experimental choices of a subject who uses narrow bracketing. Of course, this is exactly the opposite of what we see in tables 6 and 7. In fact, the portion of income variation that can be more plausibly affected by the household is barely correlated with MRS, consistent with the idea that the household uses its discretionary labor supply to smoothe income shocks, whereas the MRS responds most strongly to the exogenous portions of income and spending, indicating that the relationship is determined by the difficulty households have in smoothing these shocks via the financial market.

4.3 Prediction 3: Financial variables and expected MRS (Decision B)

In order to test prediction 3, we repeat the regressions in tables 5 to 7 using the MRS elicited in decision B. Table 8 shows the results for individual fixed effects (the regressions with week fixed effects are in the appendix). While the coefficients generally show the correct signs, most of them are not significant. Of course, the relationship of current shocks with future MRS is expected to be less strong; first, because the effect of a given shock on savings in future periods is weakened over time, and second, because decision B measures an expectation over possible future shocks; the larger and less persistent individual period shocks, the lower is the correlation between current financial variables and future MRS.

5 Discussion

Summarizing the findings above, the evidence suggests that our model captures an important feature of how subjects make their choices in the MPL experiment. This affects the interpretation of heterogeneity in experimental choices. While households with a high or low current MRS can still be considered – and will behave – more or less "impatient", our results suggest that interpersonal heterogeneity as well as intrapersonal changes over time may be caused by the subjects' changing financial situation, rather than fundamental differences in their time preferences. It also has the somewhat pessimistic implication that without further

Table 8: Savings, consumption shocks, and income shocks and MRS_{t+1}.

	(1) MRS (B)	(2) MRS (B)	(3) MRS (B)	(4) MRS (B)	(5) MRS (B)	(6) MRS (B)
Savings (I-E)	-0.148+					
	(0.0853)					
Savings (I-E) Head		-0.135				
		(0.0908)				
Adv. event expenses			0.210			
			(0.162)			
Adv. event (0/1)				0.191+		
				(0.114)		
Labor earnings					0.155	
					(0.143)	
Nonlabor earnings						-0.151
						(0.154)
Constant	4.589***	4.597***	4.626***	4.578***	4.588***	4.653***
	(0.0296)	(0.0297)	(0.00945)	(0.0375)	(0.0464)	(0.0123)
Observations	2546	2546	2542	2546	2497	2541

Standard errors in parentheses; all regressions include individual fixed effects.
+ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

information on the household's utility function, consumption levels, and marginal propensity to consume, we will not be able to directly uncover subjects' "deep" time preference parameters β and δ from experimental choices, unless subjects are naive hyperbolic discounters and completely credit constrained.

At the same time, our results do suggest that we can use the model in conjunction with our data to learn more about the constraints that the households in our sample are under. For example, we can study which types of shocks are most important for these households, in the sense that they cannot be smoothed in the financial market and directly affect the household's marginal consumption utility. It is possible that the degree to which households are credit and savings constrained depends on the nature of the shock they are subject to; for instance, it may be more difficult to obtain credit to pay for healthcare expenses than for a temporary revenue shortfall in business, or the reverse. As argued above, we would expect spending shocks and income shocks to have exactly opposite effects on the relationship between spending and MRS. If this relationship is positive, then the variation in spending is dominated by spending shocks, whereas if it is negative, the variation in spending is dominated by income shocks.

As the OLS regressions (1)-(4) in table 9 show, the relationship between spending and MRS is clearly and significantly positive. This tells us not only that consumption spending is subject to significant shocks that cannot be smoothed in the financial market, but also that these shocks are quantitatively more important for households' utility streams than the effects of income shocks. Note that the underlying relationship that is measured here is that between the marginal utility of consumption at c_t and MRS_t. This suggests that high realizations of spending are primarily the result of negative shocks and are in fact associated with *lower* levels of "utility-relevant" consumption and *higher* marginal utility.

After we established that the correlation between spending and MRS is governed by spending shocks, we

Table 9: Income vs. Spending Shocks and Credit Constraints

	(1) MRS (A) OLS	(2) MRS (A) OLS	(3) MRS (A) OLS	(4) MRS (A) OLS	(5) MRS (A) IV	(6) MRS (A) IV	(7) MRS (A) IV	(8) MRS (A) IV
Expenses	0.221*	0.200*			1.148*	1.134*		
	(0.0909)	(0.0929)			(0.457)	(0.488)		
Expenses Head			0.241*	0.220*			1.304*	1.276*
			(0.103)	(0.105)			(0.518)	(0.548)
Ind. fixed effects	yes	yes	yes	yes	yes	yes	yes	yes
Time fixed effects		yes		yes		yes		yes
Constant	4.475***	4.569***	4.509***	4.601***	3.602***	3.613***	3.744***	3.769***
	(0.0856)	(0.113)	(0.0743)	(0.101)	(0.432)	(0.505)	(0.375)	(0.437)
Observations	2547	2547	2547	2547	2547	2547	2547	2547

Standard errors in parentheses
+ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

can ask what the true impact of exogenous shocks to spending on the marginal rate of substitution (and the effective interest rate) actually is. As noted earlier, the fact that some spending variation may be endogenous, for example due to the timing of durables purchases, means that we are potentially underestimating the true size of the effect. In order to answer this question, we re-estimate the spending-MRS regression, but use the occurrence of an adverse event as an instrument for spending shocks. The results are reported in regressions (5)-(8) in table 9, and they indicate that the simple OLS regressions underestimate the size of the effect by a factor of at least 4. Specifically, the IV estimates show that a $100 increase in spending as part of an exogenous shock (equivalent to a $100 income shortfall) increases the MRS from 3.60-3.77 to 4.75-5.05. In other words, these households, extremely "impatient" to begin with, become even more focused on current over future consumption when they face unexpected expenditures that they cannot adequately smoothe through the market.

6 Literature Review

To our knowledge, there are four other papers that make use of repeated time measures from the same individuals and relate them to financial variables.[21] Two of these offer tentative support to our model.

? perform a longitudinal study of time preferences from a representative sample of 100 members of the Danish population. While they do not collect detailed financial information, they do ask subjects whether they consider their current financial situation to have improved or worsened since the previous survey. They find that those who reported an improvement tended also to report a fall in their estimated discount rate. More recently, **?** report data from the Seattle and Denver Income Maintenance Experiments which were run in the 1970s. These experiments constructed a panel data set of residents in these two areas, collecting detailed income information and asking hypothetical questions designed to elicit the participants discount rate. In this sample of 1194 subjects, **?** find that changes in income are negatively correlated with changes in discount rates, in line with the predictions of our model.

In contrast, **?** do not find a correlation between changes in income and changes in measured discount rate in a sample of 250 individuals surveyed at a tax filing center in Massachusetts. One potential explanation is

[21] **?** takes repeated time series measurements in a laboratory study in Vancouver, Canada, and studies how consistent individual choices are over time. However, financial information is not collected from the subjects.

the fact that **?** use annual income data, rather than the much higher frequency data in our study, and that the variation in income that they observe is not closely enough related with income shocks to be detectable in the relatively small sample. Our measures may better reflect the financial situation of the individual at the point at which time preferences were elicited. Another possible reason is that the Malian participants in our experiment suffer proportionally larger unexpected income shocks, or are on average more credit constrained, than individuals in the US.

? perform an experiment in rural Malawi to examine whether people tend to revise their intertemporal choices. Subjects were first asked to allocate money across an intertemporal budget (in the manner of **?**). In an unannounced visit at a later date, they were then offered the chance to change their initial allocation. The authors find no significant effect of measured shocks – a death in the family, "unexpected income shortfalls" – on the subject's choice to revise earlier decisions. This is surprising given our findings, although the authors do point out that the signs of the coefficients are negative, as predicted by our model. They also state that the results do not necessarily generalize to other populations, because there were few deaths (2% of households) and income shortfalls were generally small. This is due to the fact that their sample is rural and the experiment was carried out in the growing season, a time when these households have virtually no income. Another possible reason may be the sample size of 661 households (compared to a total of over 2500 decisions in our experiment).

7 Conclusions

In this paper we have developed a model of flexible credit constraints and applied it to choices in a multiple price list experiment of the type that are often used to measure time preferences. We show that our model predicts a number of correlations between measured time preferences and other economic variables. We test these predictions using repeated time preference measures in a poor population in Bamako, Mali. Broadly speaking, our predictions hold.

In future work, we intend to use our model to extract estimates of underlying parameters from our data. At present, our results on the identification of time preference parameters are somewhat negative, in that they suggest that there is no identification from experimental data alone. However, in principle our data could be used to identify these parameters. Essentially, for a given measured MRS, a higher current consumption relative to future consumption implies a higher discount rate. Using this insight, we can potentially estimate time preference parameters at an individual level, and compare these to the parameter estimates obtained from the 'narrow bracketing' approach.

Our results also suggest the intriguing possibility that we can use the repeated measurement of MRS to learn more about the type and severity of the credit constraints households are under. Broadly speaking, a household that is more credit constrained should exhibit bigger MRS responses to shocks than one that is not constrained. In extremis, a household that has no credit constraints should have an MRS that is invariant to spending and income shocks, while one that is completely constrained should see the full impact of these shocks on MRS. More generally we may be able to use measures of MRS to determine the nature of credit constraints – for example whether it is easier to insure certain types of shock, whether formal or informal insurance is the more important, and the 'shape' of the R curve.

References

Appendix

Table 10: Savings, consumption shocks, and income shocks and MRS_{t+1} with week fixed effects.

	(1) MRS (B)	(2) MRS (B)	(3) MRS (B)	(4) MRS (B)	(5) MRS (B)	(6) MRS (B)
Savings (I-E)	-0.142+					
	(0.0853)					
Savings (I-E) Head		-0.128				
		(0.0909)				
Adv. event expenses			0.201			
			(0.161)			
Adv. event (0/1)				0.181		
				(0.113)		
Labor earnings					0.147	
					(0.144)	
Nonlabor earnings						-0.149
						(0.153)
Constant	4.633***	4.640***	4.671***	4.622***	4.629***	4.699***
	(0.0659)	(0.0658)	(0.0594)	(0.0698)	(0.0810)	(0.0591)
Observations	2546	2546	2542	2546	2497	2541

Standard errors in parentheses

+ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$