

CREDIT CONSTRAINTS AND THE MEASUREMENT OF TIME PREFERENCES*

Mark Dean[†] Anja Sautmann[‡]

April 4, 2014

Abstract

Incentivized experiments are commonly used to estimate marginal rates of intertemporal substitution (MRS) in the lab and in the field in order to make inferences about individual time preferences. This paper considers 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 which is equal to the relative 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 and shocks to income and consumption. We test the model using a new a panel data set from Mali and find evidence for such effects. Our results imply that the relationship between experimentally elicited MRS and time preferences is not straightforward. However, measured MRS can be useful in determining the importance of different types of financial shocks to the household.

JEL: D90, D03, D14, C90, C81, B41

*Word Count: 14327. We would like to thank David Atkin, Ken Chay, Andrew Foster, Yoram Halevy, Kyle Hyndman, Supreet Kaur, Dilip Mokherjee, Sriniketh Nagavarapu, Nancy Qian, David Weil and seminar participants at Boston University, Brown University, Columbia University, and Yale University for their comments and suggestions. All errors are 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.

[†]Department of Economics, Brown University: mark_dean@brown.edu

[‡]Corresponding author: Department of Economics, Brown University Box B Providence, Rhode Island 02912 USA. Tel +1 401 863 6972. Fax +1 401 863 1970. Email anja_sautmann@brown.edu

1 Introduction

Experimental methods have become an important part of the economist’s toolkit. Increasingly, applied economists are using incentivized experiments to identify individual-specific preference parameters, and in particular those that govern intertemporal choices (see for example Ashraf et al. [2006], Mahajan and Tarozzi [2011], Tanaka et al. [2010], Schaner [2013]). Effectively, these experiments measure the Marginal Rate of Intertemporal Substitution (MRS) for money between two periods. The underlying assumption in most such studies is that the MRS directly reflects “deep” preference parameters, typically discount factors in a temporally separable utility function; see Frederick et al. [2002] for a comprehensive overview.¹ Recent methodological advances have focused on controlling for the curvature of the utility function at a given point in time [Andersen et al., 2008], and on eliciting the point of indifference between payments at different times more accurately [Andreoni and Sprenger, 2012].

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 be fully determined by the market interest rate. 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.² The identification of time preferences directly from measured MRS therefore requires the assumption of ‘narrow bracketing’ – that experimental choices are divorced from decisions and conditions outside the laboratory.

In this paper we set out to formally study the relationship between credit constraints, financial shocks and measured MRS in a model that integrates experimental choices with the subject’s broader intertemporal optimization problem. Our model allows for time inconsistency, and 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 MRS measurements as well as detailed financial data at the household level. We show that, as predicted by the model, MRS measurements vary systematically with income, consumption, spending shocks, and savings. Because narrow bracketing fails, MRS measurements cannot be easily used to recover time preference parameters. However, the model implies that repeated MRS measures can be used to understand the relative importance of different types of financial shocks for a population. In our sample, we find that expenditure shocks play an important role in determining a household’s MRS. Surprisingly, expenditure on basic needs like food and household items seems to be an important source of these shocks.

Our model builds on Harris and Laibson [2001]’s analysis of intertemporal decision making when there is hyperbolic discounting and time inconsistency. We augment their model 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. The model implies that experimentally measured MRS equals 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 discounted expected marginal utility of consumption tomorrow.³

¹In the studies in Frederick et al. 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.

²The relationship between measured MRS and external credit markets has been discussed previously in the literature, for example Cubitt and Read [2007] and Farmer and Geanakoplos [2009].

³Under the assumption that experimental payments are small relative to income and consumption.

Importantly, this conclusion does not rely on the assumption that subjects actively arbitrage experimental payments – only that they adjust consumption optimally to non-experimental shocks, and that they take the resulting change in their “real world” MRS into account when making experimental choices.

The model predicts that measured MRS should be correlated with financial variables, including income shocks, spending shocks, and savings. By contrast, the standard approach to modeling behavior in time preference experiments is to abstract from choices outside the experiment (see for example Andersen et al. [2008]), 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 (the ‘complete credit constraints’ case) would predict a correlation of MRS with income and consumption, but no correlation with savings (which are zero by assumption).

We test the predictions of our model against these alternatives in a unique panel data set of 1017 household heads in Bamako, Mali. Along with repeated multiple price list experiments designed to elicit MRS, the data includes detailed weekly measures of spending, income, and whether or not the household has suffered a negative expenditure shock (such as damage or loss of property) in that period. We demonstrate that MRS is significantly correlated with exogenous shocks to income and expenditure, as well as with the flow of savings in any given period. Together, these findings rule out narrow bracketing and the pathological cases of complete credit constraints and no credit constraints.

We further use our model to examine what can be learned from experimental measurement of MRS. Our first result is negative: when narrow bracketing fails, there is no simple way to use these measures to estimate time preference parameters, except in the extreme case of perfect credit constraints. The difference between measured trade-offs over earlier and later time periods cannot be used to measure time-inconsistency, first, because sophisticated subjects will anticipate the discount factor their ‘future self’ will apply, and second, because this comparison is subject to an independent statistical bias. Point measures of time preference may be confounded by other factors that affect MRS; for example, a household that has just suffered a negative financial shock may appear present-biased, even if their choices are perfectly time-consistent. This result has relevance for the growing applied literature that has made use of experimental methods to measure individual specific discount factors and to identify time inconsistency (see e.g. Ashraf et al., 2006, Mahajan and Tarozzi, 2011).

More positively, we show that repeated MRS measurements can be used to measure the relative importance of different types of shock for a household. In our sample, we find the presence of significant spending shocks which are quantitatively important relative to income shocks, demonstrated by the positive correlation between expenditure and MRS. A surprising finding is that intertemporal differences in food expenditure are the result of shocks with significant effects on the household’s MRS. 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. Harrison et al. [2005] 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. Krupka and Stephens [2013] show that elicited discount rates for hypothetical payoffs

Table I: A Price List Experiment

Set A		Set B	
in 1 week	today	in 2 weeks	in 1 week
a_1	a_0	b_2	b_1
CFA 300	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

respond to macroeconomic events (levels of inflation) as well as individual financial shocks (labor market participation). Carvalho et al. [2014] find that low-income households exhibit more present bias just before payday than just after payday. Halevy [2012] shows in a student population that violations of exponential discounting are largely due to instability of measured MRS over time, in line with our findings. In contrast, Meier and Sprenger [2010] and Giné et al. [2013] find no correlation between financial shocks and measured MRS. We discuss possible explanations for this in section 6.

Section 2 describes our model and its predictions. Section 3 describes our data. Section 4 tests the predictions of the model. Section 5 discusses the implications for the measurement of time preferences and the relative importance of financial shocks. Section 6 relates our work to the existing literature and section 7 concludes.

2 Time Preference Experiments and Models of Behavior

Consider the sequences of decisions shown in table I. 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. In Set B, the subject makes choices between money in one week’s time (b_1) and money in two week’s time (b_2). ‘Multiple price list’ (MPL) experiments of this kind have been used in many experimental investigations into time preferences; examples are Coller and Williams [1999], Andersen et al. [2008], Benhabib et al. [2010], and Meier and Sprenger [2010].⁴ Going from the top to the bottom of each list, the earlier payout becomes more and more attractive, and the parameter of interest is the point at which the subject switches to choosing the early over the late payment.

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 Ashraf et al. [2006], Andersen et al. [2008], Andreoni and Sprenger [2012]). 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 Cubitt and Read [2007] 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 Laibson [1997]) and credit constraints in the form of

⁴The questions are either hypothetical or one question is selected at random for payment to avoid wealth/portfolio effects.

savings-dependent interest rates.

2.1 Narrow Bracketing

The term ‘narrow bracketing’ refers to the idea that behavior in an MPL experiment is divorced from a subject’s day-to-day financial circumstances, and that it instead reflects deep preference parameters.⁵ For example, in an MPL experiment as in table I, a subject who switches from later to earlier payments at an earlier point in set A than in set 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.

In order to draw this conclusion, experimental subjects are 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 $u(c)$ together with time-preference parameters δ and β . Thus, in period 0, expected utility is

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

where c_t denotes consumption in $t = 0, \dots, \infty$, δ is the discount factor, and the parameter $0 < \beta \leq 1$ indexes the degree of present bias. These preferences are quasi-hyperbolic and time inconsistent if $\beta < 1$.⁶ Despite this utility function, it is assumed that during the experiment, 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 $u(\bar{c} + a_0) + \beta\delta u(\bar{c})$ and $u(\bar{c}) + \beta\delta u(\bar{c} + a_1)$. The subject will choose a_0 over a_1 if

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

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

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

Similarly, when choosing between payments b_1 at $t = 1$ and b_2 at $t = 2$ in period 0, the switch point will in approximation reveal δ alone. By comparing the two, we can now for example conclude from $\beta\delta \approx \frac{a_1}{a_0} < \frac{b_2}{b_1} \approx \delta$ that the individual is present-biased and time-inconsistent.

The crucial assumption in this argument is that the “real world” does not affect decisions made in the experiment. Subjects neither account for actual variation in their current and future consumption levels, nor

⁵See Rabin and Weizsacker [2009] for a description of narrow bracketing, although the idea goes back as least to Tversky and Kahneman [1981].

⁶The assumption that the consumer is infinitely lived eases analysis but is not fundamental to the results.

⁷Andersen et al. [2008] introduced a method that removes the need to assume locally linear utility, see section 5.1.

do they use the market interest rate to arbitrage (or assess the relative value) of the experimental payments. By contrast, an exponential discounter 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 experimental choices in the presence of financial shocks and credit constraints.

Consider the behavior of a decision maker whose preferences are described by

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

The only modification is that the function $u : \mathbb{R}_+ \times \mathbb{R} \rightarrow \mathbb{R}$ denotes per-period utility, where ρ_t is a stochastic preference parameter that is drawn independently from a distribution F_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 represents 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 further assume that r (and therefore R) is continuously differentiable and obeys

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

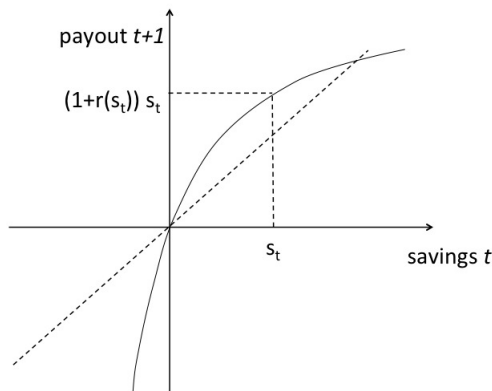
and

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

The first restriction implies that the resources available in period $t + 1$ are increasing in s_t , the second that the marginal rate of return to savings is decreasing in s_t .

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

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



credit. Interest rates vary widely, the rates for borrowing are higher than for saving, and there is evidence for credit rationing: at a given interest rate the total loan size available is restricted below the individual’s demand for credit (e.g. Banerjee and Duflo [2010, 2012, first version 2004]). 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 (Ashraf et al. [2006], Dupas and Robinson [2013a]), 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 else defaulting on a personal loan, and unexpected increases in inflation (Dupas and Robinson [2013b], see Collins et al. [2009] 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 $t + 1$. One can interpret the curvature of R as indexing the degree to which the consumer is credit constrained: 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*. At the other extreme, as the second derivative $R''(0) \rightarrow -\infty$, 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 *complete-constraints case*.

There exist several other models of credit constraints that manifest themselves in different assumptions

⁸Dupas 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.

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

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 (Deaton [1991], Carroll [2001]). 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, our focus on a smoothly differentiable return function is largely for analytical convenience. However, we would also argue that it is quite plausible to assume that people will always be able to borrow or save more as long as they are prepared to pay a high enough interest rate.

One feature that can not be accommodated by our model is an interest rate that increases with the level of savings, as suggested by Kaboski and Townsend [2011]. This would be the case if, for example, a minimum investment was required in a project. We must assume 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 Harris and Laibson [2001] to derive an Euler equation for the quasi-hyperbolic consumer of our model. 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 shock ρ_t .

Harris and Laibson [2001] provide a set of conditions under which the equilibrium of the above game can be described by what they call 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 (Harris and Laibson show that this condition is satisfied for a relatively wide range of parameters).

Definition 1. A consumption function $c : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}_+$ satisfies the Strong Hyperbolic Euler Equation if

¹⁰Note 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.

¹¹These 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. $0 < \underline{x} \leq \frac{-cu''(c,\rho)}{u'(c,\rho)} \leq \bar{x} < \infty$, on $[0, \infty)$ for all ρ in the support of g ;
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.

the following holds for every $w_t, \rho_t \in \mathbb{R}$:

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

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 expectation term on the right hand side of equation 2.1 is the expected marginal value of cash on hand in period $t + 1$, viewed from the perspective of period t ,

$$E_t \beta \delta V'(w_{t+1}) = \left(\beta \delta \frac{\partial c_{t+1}}{\partial w_{t+1}} + \delta \left(1 - \frac{\partial c_{t+1}}{\partial w_{t+1}} \right) \right) u'(c_{t+1}, \rho_{t+1}). \quad (2.2)$$

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, the marginal utility of income tomorrow from the perspective of time t is modified by a factor

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

Essentially, 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 (time-variant) weight is given by the future propensity to consume. If 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

$$MRS_t \equiv \frac{u'(c(w_t, \rho_t), \rho_t)}{E_t [d_{t+1} u'(c(w_{t+1}, \rho_{t+1}), \rho_{t+1})]} = R'(s_t). \quad (2.3)$$

Note that, in the extreme case of no credit constraints, where the interest rate is not dependent on savings, this equation implies that the MRS of the consumer should not vary, but be constant and equal to one plus the market interest rate. In the limiting case where consumers do not (cannot) borrow or save, the propensity to consume equals one, $d_{t+1} = \beta \delta$, and the identity $w_t = y_t = c_t$ holds in every period. 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 of the experiment, 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 whether the subject still has access to the financial market and can trade the exper-

imental 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 boils down to an arbitrage argument. If 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 $R'(s_t^*)$. As subjects can move money between periods at this rate, they will prefer a_0 over a_1 if the value of a_0 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 $\frac{a_1}{a_0} > R'(s_t^*)$ and the experimental payments are small enough, 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 approximately $\frac{a_1}{a_0} = R'(s_t^*)$. 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 point of indifference between a_0 and a_1 is given by

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

for small experimental payments. Using equation (2.2), we have

$$\frac{a_1}{a_0} \approx \frac{u'(c_t^*, \rho_t)}{E_t [d_{t+1} u'(c(w_{t+1}, \rho_{t+1}), \rho_{t+1})]}.$$

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 provides an approximation for 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 conclusion that the MPL switch point measures MRS_t .

Note that the simple all-or-nothing choices of the experiment mean that we only obtain a range for the MRS, given by the last choice of a_1 over a_0 and the first choice of a_0 over a_1 . Moreover, even for very

¹²This 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 Cubitt and Read [2007]). 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.

small experimental payments, some subjects may be so close to one of the cut-offs given by the experimental decisions that the linear approximation of the curvature of R or u becomes inaccurate. The size and direction of this inaccuracy is indeterminate.

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. The key observation here is that the subject's self at time t anticipates that next period's self knows which payoff she will get from the experiment. Thus, her consumption decision at $t + 1$ will take this additional income into account, regardless of whether 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 option that gives 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

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

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

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

where we make use of the SHEE for period $t + 1$ to substitute for $R'(s_{t+1}^*)$. This expression is similar to the SHEE for period t , except for the additional discounting terms and the expectation term in the numerator. Note that d depends on the propensity to consume, which may itself be a function of future shocks, so the future applied discount factor is uncertain from the perspective of today's self, but its expectation is affected by w_t and ρ_t . Similarly, the expected future marginal utilities in $t + 1$ and $t + 2$ are functions of today's realizations of w_t and ρ_t .

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.

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

First, we note that an increase in s_t will lead to a fall in MRS_t . This follows directly from equation 2.3 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 (Savings and MRS): *Savings in period t are negatively correlated with MRS in period t .*

This prediction will serve as a 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 where R is linear. In the complete-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.

Prediction (Income shocks and MRS): *Exogenous income shocks cause MRS to fall.*

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 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. A closely related case is a large exogenous lump-sum expense that has to be made in full, like buying a bulk food item.¹⁴ A household that is subject to such an adverse event 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 an increase in the marginal utility of consumption. We can show formally that a shock of this type 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 $s \leq s'$

¹⁴Unlike an adverse event, a large purchase is typically not unexpected, and the household may save up towards such an expense. If there are credit constraints, however, the effects will be qualitatively similar to an unanticipated shock.

(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.

Prediction (Spending shocks and MRS): *Exogenous preference shocks that monotonically increase current-period marginal utility increase MRS.*

In the presence of income and consumption shocks, our model implies that the correlation between consumption spending and MRS depends on the relative importance of the two types of shocks. On the one hand, Harris and Laibson [2001] show that the consumption function will be strictly increasing in wealth (if β is sufficiently close to 1 and f_y and u are three times continuously differentiable). Thus, an income shock will generally lead to an increase in consumption expenditure and a fall in MRS. On the other, consumption shocks lead to a rise in consumption expenditure as well as a rise in measured MRS. Thus, if households are subject to large consumption shocks, then MRS and c will be correlated positively (subjects who spend more are more impatient), while if they are subject to large income shocks they will be negatively correlated (those who spend more will be less impatient). This will later on provide us with a test which types of shocks are most important for these households, in the sense 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 $R' > 0$). Thus, a shock in period t acts as a wealth 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 wealth shock in period $t + 1$.

Prediction (Shocks and future MRS): *Exogenous income shocks in period t increase savings and decrease expected MRS in period $t + 1$. Spending shocks that monotonically increase current-period marginal utility decrease period $t + 1$ savings and increase expected MRS in $t + 1$.*

2.6 Summary

We have shown that our model of ‘soft’ credit constraints makes four predictions about the relationship between MRS and other financial variables. These predictions allow us to disentangle this model from the alternatives of narrow bracketing, no credit constraints, and complete credit constraints.

The predictions regarding spending and income shocks differentiate between narrow bracketing and no credit constraints on one hand, and complete and soft credit constraints on the other. The first two models predict no relationship between shocks and MRS, in the first case by assumption and in the second because such shocks do not affect the effective interest rate faced by the household. In contrast, both the complete and soft credit constraint models predict that income shocks should lower MRS, while expenditure shocks should raise them.

The predictions regarding savings and future MRS help us to further differentiate between the complete- and soft-credit constraints case. Under complete constraints, savings are zero, and any measured difference between income and spending is due to measurement error and therefore uncorrelated with MRS. Similarly, if the household cannot save, shocks in one period do not affect the financial situation in the next. Thus,

under the (admittedly strong) assumption that shocks are independent across time, shocks in period t should have no effect on expected MRS in period $t + 1$.

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 participated in multiple price list time preference experiments in the first four visits. In the first visit (week 0), demographic information was collected, and in the following weeks, household heads answered detailed questions on income and spending.

Table II 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 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. There are on average more than six people in a household, with a large number of children and, due to polygamous households, slightly more than one wife/mother. 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 eligible for the program and did therefore not participate in 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 data from weeks 1 to 3 of the survey on the two multiple price lists shown in table I. These MPL measure trade-offs between the current week and the next (A), and next week and one week after (B).¹⁵ 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 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).

¹⁵In week 3, an additional MPL experiment was carried out which is not used here, concerning choices between payouts two and three weeks away.

Table II: 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

Table III: 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	

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 0 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. These choices are not used here.

Table III shows a summary of the three weeks of MPL choices. The top row in the table shows the average 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). Due to the discrete experimental choices, we cannot report exact "indifference points" between earlier and later payments. We will discuss this issue and the problem of classifying individuals who either always chose the earlier payment or always the later payment in some more detail below.

The next rows in the table show the week-to-week correlations in decisions in A and in B, and the

Table IV: 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 (yes/no)	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\$.

proportions of subjects who made the same, more patient, or less patient decisions in A compared to B. The results show 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. However, 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. The table also shows that around 6-10% of subjects are willing to pay a weakly negative interest rate, that is, they choose CFA 300 one week later over CFA 300 or more earlier. Both patterns cannot be explained in the standard “narrow bracketing” framework with quasi-hyperbolic 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 (second-to-last row). These are choices in which the surveyor recorded repeated switches between earlier and later payoffs. These observations are excluded from the analysis.

The lack of present bias is notable in our data, and perhaps surprising given the results of other studies. However, our findings are consistent with recent work by Halevy [2012]. As in our survey, in Halevy the experimenter visits the subjects on each occasion, which lowers transaction costs and implicit risk for the subjects. Moreover, it reduces self-selection into experiment participation based e.g. on current financial need. Similarly, Andreoni and Sprenger estimate present-bias parameters β that are never significantly less than 1.

In addition to the weekly time preference measures, our data also contains weekly income and spending data. Both the household head and mothers in the household were interviewed. 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, loan repayment, etc.). Spending was equally split into different categories, like food expenditure, personal expenditure, rent payments, debt service, etc. Of particular 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 IV 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 also break out labor and non-labor income in order to separate income

from “exogenous” sources and income which can be affected by the household’s choice of labor supply (see below).¹⁶ Spending includes any monetary outlays of the household, 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 RoSCA¹⁷ or to pay off a debt. We also break out expenses for adverse events, again with the idea that these events constitute primarily exogenous 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, as in our model, 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 spending shocks which do not directly contribute to consumption utility. Empirically, we test the effects of spending shocks by using reports on adverse events. As a third point, 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 errors, which appear as 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 or inconsistent data and households that do not. It appears that there is some selection bias: households with missing experimental data in one or two weeks (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, using our information on consumption and income, it appears that there are on average negative flow savings. While it is possible that our sample of households as a whole is dissaving, this discrepancy is typical for household surveys and commonly interpreted as a sign of under-reported income (see e.g. Deaton [1997]) It is therefore possible that our savings variable is subject to measurement error. We will address this again when discussing individual empirical tests.

4 Testing the Predictions of the Model

In this section we test the four predictions from the theory section in turn. 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 individual level. The coefficient α_i is an individual 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 use the midpoint between those ratios. The MRS for individuals who never switch decisions within an MPL cannot be inferred from

¹⁶Earnings 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. This has no effects on our results because these sales are rare in our sample.

¹⁷Rotating savings and credit association, a form of savings club, called “tontine” in Mali.

Table V: Consumption shocks and MRS_t .

	(1)	(2)	(3)	(4)	(5)	(6)
	MRS(A)	MRS (A)	MRS(A)	MRS (A)	MRS (A)	MRS (A)
	OLS	OLS	OLS	OLS	IV	IV
Adv. event (0/1)	0.284*	0.263*				
	(0.124)	(0.124)				
Adv. event expenses			0.256+	0.237+	1.707*	1.579*
			(0.147)	(0.141)	(0.686)	(0.690)
Constant	4.588***	4.678***	4.665***	4.755***	4.579***	4.663***
	(0.0411)	(0.0737)	(0.00869)	(0.0586)	(0.0550)	(0.0817)
Ind. fixed effects	yes	yes	yes	yes	yes	yes
Time fixed effects		yes		yes		yes
Observations	2547	2547	2543	2543	2543	2543

Standard errors in parentheses

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

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.

We begin by discussing the relationship between consumption and income shocks and MRS. These predictions are relatively easy to test and allow us to rule out the full-constraint and narrow bracketing cases. We then discuss the prediction relating savings to MRS, which allows us to differentiate between the complete- and soft-constraints case. Testing this prediction is more difficult, as we do not directly observe the household’s stock of savings. We present supporting evidence based on flow savings. Finally we discuss the relation of financial variables in period t with expected MRS (decisions in B) in period $t + 1$.

4.1 Income and Consumption shocks and MRS

Our predictions state that consumption shocks should be positively correlated with MRS, and income shocks negatively. A problem in testing these predictions is that especially labor income and durables spending may be endogenous rather than exogenous; see for example Murdoch [1995] on “income smoothing” strategies that households use in order to achieve consumption smoothing. We have already seen that 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 subcomponents 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 V. Clearly, the occurrence of an adverse event has a significant and positive effect on the MRS, and moreover, the effect remains significant at the 10% level when using event expenditure as the independent variable. Since the actual amount spent on an adverse event may be partly endogenous – households may delay repairs or choose less expensive temporary solutions if money is tight – we also carry out an IV regression in which we instrument for spending on adverse events with the indicator variable for the occurrence of one or more such events. The estimates suggest that the simple OLS underestimates the impact of the shock onto MRS by an order of magnitude; an exogenous increase in spending due to an adverse event by \$100 raises MRS by 1.58-1.71 points.

Table VI: Income and MRS_t .

	(1)	(2)	(3)	(4)
	MRS (A)	MRS (A)	MRS (A)	MRS (A)
Labor earnings	0.0690 (0.156)	0.0537 (0.155)		
Nonlabor earnings			-0.364** (0.121)	-0.361** (0.126)
Constant	4.658*** (0.0501)	4.754*** (0.0806)	4.714*** (0.0101)	4.806*** (0.0585)
Ind. fixed effects	yes	yes	yes	yes
Time fixed effects		yes		yes
Observations	2498	2498	2542	2542

Standard errors in parentheses

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

Deciding which portion of income variation is exogenous 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). The majority of labor income comes from sources which can be endogenously adjusted, such as making and selling items, driving a taxi etc. Non-labor income includes important income sources like rent paid to the household, government and military pensions and transfers. The timing and variation of 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 VI 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, in line with our prediction for exogenous income shocks.

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 the result of omitted variable bias. Specifically, one might argue that the effects we see are due to preference shocks, which on the one hand change experimentally measured MRS, and on the other affect households' financial choices. However, the evidence on consumption and income shocks 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. 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. This is exactly the opposite of what we see in tables V and VI. 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.

4.2 Savings and MRS

In this section we test if there is a significant correlation between savings and MRS. In the model this correlation occurs if individuals are subject to some, but not complete credit constraints, and it is determined

by the curvature of the individual’s budget constraint. 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), 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.

The basic relationship underlying our empirical test is

$$\begin{aligned}
MRS_{it} &= R'(s_{it}) \\
&= 1 + r(s_{it}) + r'(s_{it})s_{it} \\
&= 1 + r(s_{it}) + r'(s_{it})R(s_{it-1}) + r'(s_{it})(y_{it} - c_{it})
\end{aligned} \tag{4.1}$$

Note that we only observe the flow of savings $y_{it} - c_{it}$: the total stock of savings s_{it} is unobserved. Now suppose we estimate

$$MRS_{it} = \alpha_i + \beta(y_{it} - c_{it}) + \gamma_t + \epsilon_{it}, \tag{4.2}$$

The individual fixed effect α_i will pick up the subject’s average interest rate and stock of savings over the three periods. A time period fixed effect γ_t may be added to account for unobserved population shocks to credit market conditions or preferences (e.g. religious events, weather changes etc.).

Under the alternative models we consider, we would expect the estimated coefficient β to be zero. In the no-constraints case, measured flow savings are independent from past savings (by the permanent income hypothesis), so the error term will be independent of y_{it} and c_{it} , and $r'(s_{it}) = 0$, because r is unaffected by past or current savings. In the complete-constraints case, $y_{it} - c_{it}$ represents only measurement error uncorrelated with MRS .¹⁸ If subject makes experimental decisions with narrow bracketing then measured MRS should be unrelated to any outside variables.

We see in table VII that the coefficient on savings is significantly negative, both for household head and total household savings. This is just what is expected if r is decreasing in s_{it} . The results of this regression therefore support the partial credit constrains model over the alternatives of complete and no credit constraints, as well as narrow bracketing. It may be tempting to interpret β as the slope of r ; however, even if r'' is close to 0, under the alternative hypothesis $r(s_{it})$ will be negatively correlated with $y_{it} - c_{it}$. Thus, if $r' \neq 0$, then $\beta \leq r'(s_{it})$, i.e. the coefficient overestimates the true slope of r in absolute terms.

¹⁸There is a possibility that individuals systematically misreport $y_{it} - c_{it}$ depending on the shocks they were subject to, or that one measure is reported less accurately than the other. This could lead to a spurious correlation between flow savings and MRS in the case of complete credit constraints if y_{it} is reported relatively accurately, but c_{it} is estimated by the subject and reported closer to average consumption than is really the case. In this situation, c_{it} will be greater than y_{it} in periods with a negative income shock and smaller than y_{it} under a positive shock. Unfortunately, we cannot rule this possibility out with certainty. However, it is generally assumed that consumption is reported more accurately than savings. In this case, measured savings would be predominantly driven by the negative of consumption. Thus measured savings would be high when consumption and income are low, meaning MRS and savings should be positively correlated, the opposite of what we find. Further evidence against the perfect constraints case is provided by the positive relationship between consumption and MRS demonstrated in section 5.2.

Table VII: Savings and MRS_t .

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

Standard errors in parentheses

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

4.3 Financial variables and expected MRS

In order to test our last prediction, we repeat the regressions in tables V to VII using the MRS elicited with the set of choices B. The argument in section 2.4.1 shows that the subject's expected future MRS should be affected by current financial conditions as well, as long as today's shocks are transmitted through their effect on the savings stock. Table VIII 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 individual period shocks, the lower is the correlation between current financial variables and future MRS.

5 Applications

In this section we consider the implication of our model for what can be learned from experiments that measure a subject's MRS. We first consider the circumstances under which it is possible to extract information about time preferences. Second, we show how we can use repeated MRS measures to understand the relative importance of income and spending shocks in changes in consumption. Finally we use it as a tool to determine the source of these expenditure shocks.

5.1 Implication for the Measurement of Time Preference Parameters

Our model has implications for the use of MPL to identify present bias and time discounting factors. Note first that $\frac{a_1}{a_0}$ by itself does not identify the present discount factor $\beta\delta$ even in expectation, because the subject applies the weighted discount factor d_{t+1} . Moreover, unlike in the narrow bracketing model, where the difference between $\frac{a_1}{a_0}$ and $\frac{b_2}{b_1}$ identifies β , $\frac{b_2}{b_1}$ does not provide a second moment that allows us to disentangle the time preference parameters in a straightforward manner. 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 (d_{t+2}), rather than the discount factor today's self would like to see applied (δ). Therefore, $\frac{b_2}{b_1}$ does

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

	(1)	(2)	(3)	(4)	(5)	(6)
	MRS (B)	MRS (B)	MRS (B)	MRS (B)	MRS (B)	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.

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

not systematically differ from $\frac{a_1}{a_0}$ due to present bias.

However, $\frac{b_2}{b_1}$ is separately subject to a bias relative to $\frac{a_1}{a_0}$ which is unrelated to β and δ , but will lead to a systematic difference between choices in A and in B. Substituting the Euler equation into our expression for $\frac{b_2}{b_1}$ gives

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

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 $\frac{a_1}{a_0} = R'(s_t)$, assuming that the economy is stationary. Consider for example the exponential discounting case, where $d(w_{t+1}) = \delta$, and assume there are no consumption shocks. In this case, $\frac{b_2}{b_1}$ is biased upwards relative to $\frac{a_1}{a_0}$, since both u' and R' vary negatively with s_{t+1} ; this implies that subjects should on average appear less patient in future decisions. The argument continues to hold in non-pathological cases with quasi-hyperbolic discounting where the marginal propensity to consume does not respond too strongly to financial shocks. Note that 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 (c_{t+2} becomes independent of s_{t+1}).

All this implies that we can make inferences about time-preference parameters only in the special case of complete credit constraints. Here, self 0 can impose her preferences on future selves by choosing the

allocation of payoffs for them. We have

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

as long as spending and income shocks are iid. However, it is important to note that any given observation $\frac{a_1}{a_0}$ will still be affected by consumption and income shocks. Therefore we can compare average MRS in decisions A and B on the population level, but in order to identify individual preference parameters we need repeated MPL experiments.

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, we need information on marginal utility and consumption in addition to MRS measurements in order to identify β and δ when the agent has access to credit.

Andersen et al. [2008] and 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 measuring risk preferences in addition to time preferences to estimate the curvature parameter α ¹⁹, and then examining how measured MRS changes when the size of experimental payoffs at two points in time is varied 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.

This 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, but consumes the additional income at the time it is received. Identification breaks down if the subject reallocates resources optimally in response to the realized experimental payoff: in this case, the experimental payments will only affect the subject’s MRS if the interest rate changes, which in turn reflects an optimal reallocation of resources that depends itself on the size of the payments. This means there is no unique level of outside consumption to estimate. The degree to which this is a problem is related to the level of credit constraints – the less constrained the subject, the larger the change in allocation to bring about a given change in measured MRS.

This argument highlights the need for measuring consumption along with MRS to identify time preferences. In our setting, time preference parameters could be estimated with data on MRS, consumption c_t , shocks ρ_t , propensity to consume at current wealth w_t , and the curvature of u . This is possible even in the no-constraints case; intuitively, time preferences are identified because for a *given* MRS, a more patient decision maker will have a lower level of consumption today relative to tomorrow. One complication with this approach is the separation of spending shocks from ‘true’ consumption in order to accurately identify marginal utility.

¹⁹This is the “dual MPL method” suggested first by Andersen et al. [2008].

5.2 The effect of income and spending shocks on consumption

While the above result is somewhat pessimistic about identifying time preferences from experimental measures of MRS, our model suggests that these measures can instead be of use in understanding the financial shocks and constraints that affect a household. One example is the role of shocks for consumption expenditure. As argued earlier, spending shocks and income shocks have exactly opposite effects on the relationship between spending and MRS. We can therefore use the sign of this relationship to understand the relative size of these shocks.

Consider a special case of the model above in which consumption shocks are unexpected expenses that do not affect utility, and over which the household has no control. Examples are repairs or maintenance to business assets, or health expenses for a relative. In this case we can write $u(c_t, \rho_t) = u(n_t)$ where $n_t = c_t - \rho_t$: utility depends on “net consumption” n_t , and consumption shocks act essentially as negative income shocks. The individual maximizes

$$E \sum_{t=0}^{\infty} u(n_t)$$

with wealth at $t = 0$ given and the constraint

$$s_t = w_t - n_t - \rho_t.$$

The variable $c_t = n_t + \rho_t$ represents total spending, and we can assume that the decision maker chooses n_t conditional on wealth w_t and expenditure shock ρ_t .

In order to understand the relative importance of shocks to wealth and consumption for the intertemporal optimization problem, we examine the relationship between total spending c_t and $R'(s_t)$. More precisely, suppose we regress MRS_t onto spending c_t . The sign of the regression coefficient β is determined by the sign of the covariance of the two variables, and, since $R'' < 0$, by the inverse sign of the covariance of savings and spending. Note that, in the canonical consumption-savings model without spending shocks, $c_t = n_t$, and we would expect $\beta < 0$, as any change in spending is caused by changes in wealth. But if consumption shocks are important, increases in spending are *negatively* associated with net consumption n_t , leading to a *positive* correlation with the *MRS*.

In order to understand the precise relationship between wealth and consumption shocks and the correlation between MRS and spending, we take a Taylor series expansion of s around \bar{w} and $\bar{\rho}$, the average levels of the wealth and spending shocks:

$$s_t(w_t, \rho_t) \approx \bar{w} - \bar{\rho} - n(\bar{w}_t, \bar{\rho}_t) + \left(1 - \frac{\partial n_t}{\partial w_t}\right) (w_t - \bar{w}_t) - \left(1 + \frac{\partial n_t}{\partial \rho_t}\right) (\rho_t - \bar{\rho}_t)$$

where $\frac{\partial n_t}{\partial w_t}$ is the marginal propensity to consume out of wealth and $\frac{\partial n_t}{\partial \rho_t}$ the effect of spending shocks on true consumption. Note that $(w_t - \bar{w}_t)$ incorporates both income shocks and the current stock of savings relative to \bar{w} .

We can similarly approximate spending as

$$c_t(w_t, \rho_t) = c(\bar{w}, \bar{\rho}) + \frac{\partial n_t}{\partial w_t} (w_t - \bar{w}_t) + \left(1 + \frac{\partial n_t}{\partial \rho_t}\right) (\rho_t - \bar{\rho}_t).$$

Assuming that spending shocks are independently distributed, and so uncorrelated with wealth shocks, this

implies

$$Cov(s_t, c_t) \approx \left(1 - \frac{\partial n_t}{\partial w_t}\right) \frac{\partial n_t}{\partial w_t} Var(w_t) - \left(1 + \frac{\partial n_t}{\partial \rho_t}\right)^2 Var(\rho_t).$$

Since shocks to w_t are equivalent to negative shocks to ρ_t , we have $-\frac{\partial n_t}{\partial w_t} = \frac{\partial n_t}{\partial \rho_t}$, and therefore

$$Cov(s_t, c_t) \approx \left(1 - \frac{\partial n_t}{\partial w_t}\right) \left[\frac{\partial n_t}{\partial w_t} Var(w_t) - \left(1 - \frac{\partial n_t}{\partial w_t}\right) Var(\rho_t) \right].$$

This equation says that the covariance of savings and spending is the weighted difference between the variance of income and consumption shocks, scaled by the marginal propensity to save. The weights depend on the marginal propensity to consume additional income (or negative consumption shocks), $\frac{\partial n_t}{\partial w_t}$. If $\frac{\partial n_t}{\partial w_t}$ is high, net consumption and therefore total spending closely follows income, but since spending shocks are almost entirely compensated by consumption changes, total spending is relatively unaffected by consumption shocks. On the other hand, if the marginal propensity to consume is low, income shocks have little effect on net consumption, but consumption shocks translate almost entirely into spending changes. Since the propensity to consume reflects the household's optimal response to shocks, the covariance between savings and spending not only depends on the relative size of shocks, but also on the willingness of the household to change (net) consumption in response to these shocks.

Taken together, this means that the regression coefficient β will be positive only if

$$\frac{\partial n_t}{\partial w_t} Var(w_t) < \left(1 - \frac{\partial n_t}{\partial w_t}\right) Var(\rho_t). \quad (5.1)$$

Note that, in principle, savings could replace MRS in this exercise. However, the stock of savings is notoriously difficult to measure in populations such as ours, and spending is a component of (flow) savings, so that measurement error can lead to spurious correlations. By contrast, the MRS is an easy to execute and independent measure of the subject's interest rate. Note also that, while the marginal propensity to save scales the relationship between shocks and savings, it does not scale the relationship between MRS and shocks. For example, a household that faces high credit constraints may have a low propensity to save, but a very strong relationship between saving and MRS – in fact the latter is the cause of the former. This observation means that our results is robust to heterogeneity in credit constraints or the marginal propensity to consume.

The OLS regressions (1)-(4) in table IX show that the relationship between spending and MRS is clearly and significantly positive in our data: higher current expenditure is related to more impatience on the part of the household. A crude estimate of the marginal propensity to consume in our population²⁰ is 0.38, implying that expenditure shocks are at least half the size of income shocks. This tells us not only that spending is subject to significant shocks that cannot be smoothed in the financial market, but also that high realizations of spending are primarily the result of negative shocks and are in fact associated with *lower* levels of “utility-relevant” net consumption and *higher* marginal utility. At the same time, the positive relationship between expenditure and MRS provides further evidence that our households have some ability to smooth consumption: in the pathological case of complete credit constraints, $\frac{\partial n_t}{\partial w_t} = 1$ and consumption spending is determined entirely by income. Thus, assuming that consumption and income shocks are uncorrelated, higher expenditure in any given period should be associated with lower MRS.

²⁰Calculated by regressing changes in spending on changes in income, controlling for week fixed effects.

Table IX: Income vs. Spending Shocks in the Relationship of Spending and MRS

	(1)	(2)	(3)	(4)
	MRS (A)	MRS (A)	MRS (A)	MRS (A)
Expenses	0.221*	0.200*		
	(0.0909)	(0.0929)		
Expenses Head			0.241*	0.220*
			(0.103)	(0.105)
Constant	4.475***	4.569***	4.509***	4.601***
	(0.0856)	(0.113)	(0.0743)	(0.101)
Ind. fixed effects	yes	yes	yes	yes
Time fixed effects		yes		yes
Observations	2547	2547	2547	2547

Standard errors in parentheses

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

5.3 What are the sources of expenditure shocks?

A second use of MRS measurement suggested by our model is to understand which shocks affect measured MRS and so consumption utility the most. Different types of expenditure may differ in the size of shocks they are subject to, the change in spending in response to these shocks (given by the category-specific propensity to consume), and the degree to which the household can smoothe the shocks through the financial market. By examining the relationship of the MRS with various different components of expenditure, we can draw conclusions about the role of income and consumption shocks for different types of expenditure. Note that equation (5.1) continues to hold if we replace ρ with category-specific shocks and $\frac{\partial n_t}{\partial w_t}$ with the propensity to adjust category spending in response to wealth changes.

Table X shows the results of such an exercise. As we have seen earlier, spending due to adverse events has a significant positive coefficient in the regression. Although mostly insignificant, the negative coefficients on social events, gifts and donations, and personal expenditure (this includes “luxury goods” like cigarettes or tea and phone credit) suggest that variation in these types of spending is likely driven by income: the propensity to consume is high or there are simply no significant “shocks” to these types of spending. Utility bills and rent spending have almost no correlation with MRS, reflecting that their payment schedule is independent from income shocks, but also that households are able to smoothe this (anticipated) variation in spending well.

Interestingly, expenditure on food and other essentials for the household by the household head is significantly positively correlated with MRS. This indicates that variations in demand for food and basic household goods are driven by negative expenditure shocks, rather than, for example, splurging on a good meal after a successful day at work. This is likely partly a consequence of the fact that basic consumption needs are unresponsive to income, and that unexpectedly higher demand in this category is less predictable, and less easy to delay, than other shocks. Lastly, one reason may lie in the traditional household organization of Malian families, where women receive a weekly allowance from husbands and request additional money when the need arises. This makes food and household expenditures appear much more like an exogenous shock from the perspective of the household head.

Table X: MRS and Disaggregated Spending

	(1)	(2)
	MRS (A)	MRS (A)
Food Expenditure	0.717*** (0.165)	0.699*** (0.164)
Personal Expenditure	-0.388 (0.760)	-0.634 (0.769)
Adv. event expenses	0.274+ (0.152)	0.270+ (0.148)
Bills and Rent	-0.0272 (0.388)	-0.103 (0.392)
Social Events	-0.830+ (0.471)	-0.881+ (0.468)
Gifts and Donations	-0.298 (0.741)	-0.408 (0.742)
Constant	4.663*** (0.0641)	4.768*** (0.0927)
Observations	2530	2530

Standard errors in parentheses

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

6 Literature Review

To our knowledge, there are six other papers that make use of repeated time measures from the same individuals. Four of these offer support to our model.

Harrison et al. [2005] 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, Krupka and Stephens [2013] 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, Krupka and Stephens find that changes in income are negatively correlated with changes in income, in line with the predictions of our model.

In another recent study, Carvalho et al. [2014] use the natural variation in liquidity constraints caused by the pay cycle to examine the effect of these constraints on a number of cognitive and behavioral measures in a sample of low income US households. They show that before-payday participants behave as if they are more present-biased when making choices about monetary rewards. However, this difference is not observed for intertemporal choices about non-monetary real effort tasks, supporting a liquidity-constraint interpretation.

Halevy [2012] reports the results of a laboratory study which at time t elicited preferences between payments at t and $t + 1$ and between payments at $t + 1$ and $t + 2$, and at time $t + 1$ measured preferences

over payments at $t + 1$ and $t + 2$. This allows him to categorize subjects who exhibit non-stationary time preferences into those that violate time invariance (i.e. preferences over the relative value of immediate and delayed payments change) and time consistency (i.e the relative valuations of payments that occur in $t+1$ and $t+2$ change). While time inconsistency could be a consequence of present-bias in a narrow-bracketing setting, a quasi-hyperbolic narrow bracketer would still be expected to make time invariant choices. The author reports significant violations of time invariance, particularly among subjects who are ‘liquidity constrained’ in one period, that is, they exhibit an unusually strong preference for the earlier payment.

In contrast, Meier and Sprenger [2010] 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. This may be due to the fact that the variation in annual, retrospective income data from tax returns is not closely enough related with current income shocks (rather than past shocks or permanent income changes) 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.

Giné et al. [2013] 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 Andreoni and Sprenger [2012]). In an unannounced visit at a later date, they were then offered the chance to change their initial allocation. While the observed changes go in the predicted direction, 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 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 the experiment was carried out in the growing season, a time when the rural households in their sample 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. Further possibilities include using repeated MRS measurements as a diagnostic to measure the welfare effects of different types of shock, and the effectiveness of policies designed to improve households’ ability to smoothe income.

Mark Dean: Department of Economics, Brown University

Anja Sautmann: Department of Economics, Brown University

Appendix

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

	(1)	(2)	(3)	(4)	(5)	(6)
	MRS (B)	MRS (B)	MRS (B)	MRS (B)	MRS (B)	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

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

References

- S. Andersen, G.W. Harrison, M.I. Lau, and E.E. Rutstrom. Eliciting risk and time preferences. *Econometrica*, 76(3):583–618, 2008. ISSN 1468-0262.
- James Andreoni and Charles Sprenger. Estimating time preferences from convex budgets. *American Economic Review*, 102(7):3333–56, 2012. doi: 10.1257/aer.102.7.3333. URL <http://www.aeaweb.org/articles.php?doi=10.1257/aer.102.7.3333>.
- Nava Ashraf, Dean Karlan, and Wesley Yin. Tying Odysseus to the mast: evidence from a commitment savings product in the Philippines. *Quarterly Journal of Economics*, 121(2):635–672, 2006.
- Abhijit Banerjee and Esther Duflo. Giving credit where it is due. *Journal of Economic Perspectives*, 24(3): 61–79, 2010.
- Abhijit Banerjee and Esther Duflo. Do firms want to borrow more? testing credit constraints using a directed lending program. *mimeo*, 2012, first version 2004.
- Jess Benhabib, Alberto Bisin, and Andrew Schotter. Present-bias, quasi-hyperbolic discounting, and fixed costs. *Games and Economic Behavior*, 69(2):205–223, July 2010. URL <http://ideas.repec.org/a/eee/gamebe/v69y2010i2p205-223.html>.
- Christopher D. Carroll. A theory of the consumption function, with and without liquidity constraints. *Journal of Economic Perspectives*, 15(3):23–45, 2001.
- Leandro S. Carvalho, Stephan Meier, and Stephanie W. Wang. Poverty and economic decision-making: Evidence from changes in financial resources at payday. *Mimeo*, 2014.
- Maribeth Coller and Melonie Williams. Eliciting individual discount rates. *Experimental Economics*, 2(2): 107–127, December 1999. URL <http://ideas.repec.org/a/kap/expeco/v2y1999i2p107-127.html>.
- Daryl Collins, Jonathan Morduch, Stuart Rutherford, and Olivia Ruthven. *Portfolios of the Poor*. Princeton University Press, 2009.
- Robin Cubitt and Daniel Read. Can intertemporal choice experiments elicit time preferences for consumption? *Experimental Economics*, 10(4):369–389, December 2007. URL <http://ideas.repec.org/a/kap/expeco/v10y2007i4p369-389.html>.
- Angus Deaton. Savings and liquidity constraints. *Econometrica*, 59(5):1221–1248, 1991.
- Angus Deaton. *The Analysis of Household Surveys: A Microeconomic Approach to Development Policy*. World Bank, 1997.
- Pascaline Dupas and Jonathan Robinson. Savings constraints and microenterprise development: Evidence from a field experiment in kenya. *American Economic Journal: Applied Economics*, 5(1):163–192, 2013a.
- Pascaline Dupas and Jonathan Robinson. Why don’t the poor save more: evidence from health savings experiments. *American Economic Review*, 103(4):1138–1171, 2013b.
- J. Doyne Farmer and John Geanakoplos. Hyperbolic discounting is rational: Valuing the far future with uncertain discount rates. Cowles Foundation Discussion Papers 1719, Cowles Foundation for Research in Economics, Yale University, August 2009. URL <http://ideas.repec.org/p/cwl/cwldpp/1719.html>.

- Shane Frederick, George Loewenstein, and Ted O'Donoghue. Time discounting and time preference: A critical review. *Journal of Economic Literature*, XL:351–401, 2002.
- Xavier Giné, Jessica Goldberg, Dan Silverman, and Dean Yang. Revising commitments: Field evidence on the adjustment of prior choices. *mimeo*, 2013.
- Yoram Halevy. Time consistency: Stationarity and time invariance. Micro Theory Working Papers 2012-19, Microeconomics.ca Website, June 2012.
- Christopher Harris and David Laibson. Dynamic choices of hyperbolic consumers. *Econometrica*, 69(4): 935–57, July 2001. URL <http://ideas.repec.org/a/econ/emetrp/v69y2001i4p935-57.html>.
- Glenn W. Harrison, Morten I. Lau, and E. Elisabet Rutström. Dynamic consistency in denmark: A longitudinal field experiment. *UCF Economics Working Paper*, 05-02, 2005.
- Joseph P. Kaboski and Robert M. Townsend. A structural evaluation of a large-scale quasi-experimental microfinance initiative. *Econometrica*, 79(5):1357–1406, 09 2011. URL <http://ideas.repec.org/a/econ/emetrp/v79y2011i5p1357-1406.html>.
- Erin L. Krupka and Melvin Stephens. The stability of measured time preferences. *Journal of Economic Behavior & Organization*, 85(C):11–19, 2013. URL <http://ideas.repec.org/a/eee/jeborg/v85y2013icp11-19.html>.
- David Laibson. Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 112(2): 443–77, May 1997. URL <http://ideas.repec.org/a/tpr/qjecon/v112y1997i2p443-77.html>.
- Aprajit Mahajan and Alessandro Tarozzi. Time inconsistency, expectations and technology adoption: The case of insecticide treated nets. *mimeo*, March 2011.
- Stephan Meier and Charles Sprenger. Stability of time preferences. IZA Discussion Papers 4756, Institute for the Study of Labor (IZA), February 2010. URL <http://ideas.repec.org/p/iza/izadps/dp4756.html>.
- Jonathan Murdoch. Income smoothing and consumption smoothing. *Journal of Economic Perspectives*, 9 (3):103–114, 1995.
- Matthew Rabin and Georg Weizsacker. Narrow bracketing and dominated choices. *American Economic Review*, 99(4):1508–43, September 2009. URL <http://ideas.repec.org/a/aea/aecrev/v99y2009i4p1508-43.html>.
- Simone Schaner. Do opposites detract? intrahousehold preference heterogeneity and inefficient strategic savings. *mimeo*, 2013.
- Tomomi Tanaka, Colin Camerer, and Quang Nguyen. Risk and time preferences: Linking experimental and household survey data from vietnam. *American Economic Review*, 100(1):557–571, 2010.
- Amos Tversky and Daniel Kahneman. The framing of decisions and the psychology of choice. *Science*, 211(4481):453–458, 1981. doi: 10.1126/science.7455683. URL <http://www.sciencemag.org/content/211/4481/453.abstract>.