Credit Constraints and the Measurement of Time Preferences

Mark Dean* Anja Sautmann†

January 6, 2020

Appendix for Online Publication

A Credit and Savings Constraints and the Shape of R

Foundations for Concavity. A key assumption of the paper is that the marginal returns to saving $R'(s_t)$ are decreasing in the (possibly negative) stock of savings s_t. What motivates a concave R? To begin with, note that we assume that "credit and savings" in an incomplete market may encompass a range of technologies used to shift consumption between periods: formal and informal borrowing, physical storage of money or assets like gold, or investment in an enterprise or capital good that yields returns later. The function R describes the effective return earned by the portfolio of instruments that the household uses at each level of saving. R is linear only if the household borrows and saves all funds at the same constant rate of return.

There are several possible underlying reasons for a concave R. First, various kinds of market imperfections – from lenders with market power to information asymmetries – can

*Department of Economics, Columbia University. Email mark.dean@columbia.edu.
†Corresponding author: Massachusetts Institute of Technology, 400 Main Street, Cambridge, MA 02141, USA. Tel +1-617-258-0667, email: sautmann@mit.edu.

mean that individual demand for financial instruments affects the interest rate (price) paid. Second, if there are large frictions in the capital market, households may save by investing in a productive asset or enterprise (e.g. Rosenzweig and Wolpin, 1993). Diminishing returns to capital in the production function then imply decreasing returns to savings. Indeed, De Mel et al. (2008) have shown empirically that small businesses in Sri Lanka that receive a cash or in-kind grant exhibit lower returns if their owners report higher household wealth to begin with.

Last, concavity of R also arises if the household has access to a range of different borrowing and savings tools with different rates of return and with limits to the amounts that can be borrowed or saved through each. Empirically, limits on the amounts that can be borrowed through one instrument are common, especially in informal markets and when borrowers cannot provide collateral. Poor households are known to use complex and varied portfolios of financial instruments (see Collins et al. (2009) for examples). There is also evidence for substantial savings constraints and an unfilled demand for savings instruments, to the point that some savers are willing to accept *negative* interest rates (Ashraf et al. (2006); Dupas and Robinson (2013a)). This indicates that they cannot even store the proverbial "cash under the mattress", possibly due to demands for transfers from others, the risk of theft or loss, and inflation (Dupas and Robinson (2013b)).

We assume that households will use these different instruments in order of attractiveness. For example, when reducing its debt, the household will first pay off the loan with the highest interest rates, or when short of money, it will first exhaust its store credit before taking a loan with interest. This means that the household uses more expensive forms of credit the more it borrows, and earns lower interest rates the more it saves. As long as the available instruments individually exhibit constant or decreasing returns, the resulting returns function is continuous and concave, and can be approximated arbitrarily well by a differentiable R. An extreme version of such limits is a hard credit constraint, as in classical liquidity constraint models, where the agent is free to borrow at the market interest rate down to a hard limit B (as in buffer-stock models of wealth and precautionary savings, Zeldes (1989); Deaton (1991)). Adverse selection, moral hazard, and contract enforcement problems can all lead to this form of credit rationing, because restricting loan sizes helps lenders mitigate the risks of intentional and involuntary default (see e.g. Ghosh et al. (2000) for an overview). This model can be recast as a linear budget constraint with a "kink" at B, which is another limit case of our model.

Non-differentiable R. key predictions from section 1.5 are robust to a concave and piecewise linear (non-differentiable) return function, as well as the possibility of infinite slope implied by a hard constraint. As an illustration, consider the simple case in which the decision maker faces one interest rate on savings, and another, higher rate for borrowing, leading to a budget curve that is concave and piecewise linear with a kink at zero.

MRS as measured by our experiment equals the ratio of marginal utility today to the discounted expected value of money tomorrow. For savings not equal to zero, the SHEE will still hold with equality, so measured MRS is equal to the prevailing interest rate. At zero savings, measured MRS can take many values, but is bounded between the interest rates for borrowing and saving. Thus, measured MRS would now be a set-valued, but still weakly decreasing function of savings, meaning that our prediction regarding MRS and savings continues to hold.

Shocks will also have the same impact on measured MRS in the piecewise linear case as they do in the partial constraints case. For example, our current proof in Appendix B shows that savings must weakly decrease in response to a negative income shock. If savings strictly decrease, then the negative relationship between savings and marginal returns is enough to guarantee that measured MRS must (weakly) rise. If savings do not decrease (as they might not at the kink), current consumption must decrease, while future expected marginal value of income remains the same, which will also increase the MRS. These arguments readily extend to the case in which there are multiple kinks in the return function, and to the case of a single, hard borrowing constraint.

Convexity in the returns function. It may in principle be possible that R is convex rather than concave in some places. This can for example arise when there are lump-sum investment or borrowing opportunities, such as an indivisible asset, or a formal loan for a larger amount that offers a lower effective interest rate. We do not believe these instances to be important in our data. Our population has very little access to formal banking. Moreover, in weekly data and considering week-to-week intertemporal trade-offs, it is likely that an opportunity as the one just described is a rare event. Lastly, since these instruments also have in common that they are often illiquid, and oftentimes are used to transform one form of asset holdings into another (rather than for saving and dissaving for consumption smoothing), they may be better described by a shock to the interest rate curve, rather than a convex portion to that curve.

B Proof of Propositions in Section 1.4

Proof of Proposition 1: Case 1: Suppose that a_0 must be consumed right away, while a_1 is added to the stock of wealth in the next period. Then the decision maker strictly prefers a_0 over a_1 if

$$u(w_t - s_t^* + a_0, \rho_t) + \beta\delta E_t V(y_{t+1} + R(s_t^*)) > u(w_t - s_t^*, \rho_t) + \beta\delta E V(y_{t+1} + R(s_t^*) + a_1)$$

or

$$\frac{u(w_t - s_t^* + a_0, \rho_t) - u(w_t - s_t^*, \rho_t)}{a_0} > \frac{a_1}{a_0}\frac{\beta\delta\left[E_t V(y_{t+1} + R(s_t^*) + a_1) - E_t V(y_{t+1} + R(s_t^*))\right]}{a_1} \tag{1}$$

The fractions on the left and right converge to the derivatives of u and $\beta\delta E_t V$, respectively, as a_0 and a_1 converge to zero. Now take any sequence of payments such that both transfers converge to zero, while their ratio stays constant at $\hat{R}$. Since equation 3 holds in the limit, and using the fact that $\beta\delta V'(w_{t+1}^*) = d_{t+1}u'(c(w_{t+1}, \rho_{t+1}), \rho_{t+1})$ (see Harris and Laibson (2001)), there must be some nonzero lower bound a_0 (and associated a_1) where inequality

(1) also holds. Finally, using equation (2) to equate MRS with the derivative of R gives the desired result.

A similar argument shows that a_1 is preferred if $\frac{u'(c_t^*, \rho_t)}{\beta\delta E_t V'(w_{t+1}^*)} < \frac{a_1}{a_0}$.

Case 2: Suppose the decision maker can re-optimize their savings choice after receiving a_0. Preferences between monetary payments can then be established via an arbitrage argument. Consider a decision maker that would like to change their consumption in t by $c(a_1)$ if promised a_1 in $t+1$. Wealth in period $t+1$ will then be $R(s_t^* - c(a_1)) + a_1$. Now if the decision maker received a_0 instead, one possible strategy would be to change consumption by the same amount $c(a_1)$ and have $R(s_t^* - c(a_1) + a_0)$ in $t+1$. These two strategies now provide exactly the same consumption in period t, and differ only in the amount of money available in period $t+1$. Thus, receiving a_0 must be strictly preferred to a_1 if

$$R(s_t^* - c(a_1) + a_0) > R(s_t^* - c(a_1)) + a_1$$

or

$$\frac{R(s_t^* - c(a_1) + a_0) - R(s_t^* - c(a_1))}{a_0} > \frac{a_1}{a_0} \tag{2}$$

Constructing a sequence as above and noting both that the fraction on the left converges to the derivative of R and that equation (3) holds in the limit, there must be a non-zero a_0 below which inequality (2) holds and the earlier payment is preferred. Parallel arguments show that the later payment is strictly preferred if $R'(s_t) < \frac{a_1}{a_0}$.

Proof of Proposition 2: Consider a decision maker in period t who is given a surprise choice between b_1 in $t+1$ and b_2 in $t+2$. We will proceed by first approximating the value to the decision maker of receiving b_2 in period $t+2$ with the value of receiving an appropriately discounted amount in period $t+1$ upon learning w_{t+1}, which we call $t(b_2 : w_{t+1})$. We will then examine the "approximation error" that results from replacing b_2 with this transfer. As we will show, this "error" is equal to zero unless the decision maker is time inconsistent and the interest rate varies with savings. In other cases we can determine the direction of this error as long as the decision maker is not too present biased.

Let $s(b_2 : w)$ be the savings rate in period $t+1$ if the decision maker has wealth w in that period and anticipates receiving a transfer of b_2 in period $t+2$. We define the equivalent transfer to b_2 as

$$t(b_2 : w) = R^{-1}(R(s(b_2 : w)) + b_2) - s(b_2 : w)$$

Note that this is exactly the discounted value of b_2 in period $t+1$ if the interest rate is constant.

The decision maker at time t will prefer the earlier payment b_1 if

$$\beta\delta E_t V(w_{t+1} + b_1) > \beta\delta E_t V(w_{t+1} + t(b_2 : w_{t+1})) + E_t O(b_2 : w_{t+1}), \tag{3}$$

where O is the approximation error. If we let $W(\cdot)$ equal the value of wealth in period $t+2$ from the perspective of the period t self, we can write

$$\begin{aligned} O(b_2 : w_{t+1}) = {} & \beta\delta u\left(w_{t+1} - s(b_2 : w_{t+1}), \rho_{t+1}\right) \\ & + \beta\delta^2 E_{t+1} W\left(y_{t+2} + R\left(s(b_2 : w_{t+1})\right) + b_2\right) \\ & - \beta\delta u\left(w_{t+1} + t(b_2 : w_{t+1}) - s(0 : t(b_2 : w_{t+1}) + w_{t+1}), \rho_{t+1}\right) \\ & - \beta\delta^2 E_{t+1} W\left(y_{t+2} + R\left(s(0 : t(b_2 : w_{t+1}) + w_{t+1})\right)\right) \end{aligned}$$

where $s(0 : t(b_2 : w_{t+1}) + w_{t+1})$ is savings in period $t+1$ with wealth w_{t+1} and transfer $t(b_2 : w_{t+1})$ in $t+1$ but no transfer in period $t+2$. We will return to the nature of O below. For now, note that this implies that the earlier payment will be preferred to the latter payment if

$$\beta\delta E_t \frac{V(w_{t+1}+b_1) - V(w_{t+1})}{b_1} > \frac{b_2}{b_1}\left[\beta\delta E_t \frac{V(w_{t+1}+t(b_2 : w_{t+1})) - V(w_{t+1})}{b_2} + E_t \frac{O(b_2 : w_{t+1})}{b_2}\right].$$

Note that

$$\lim_{b_2 \to 0} \frac{t(b_2 : w)}{b_2} = \lim_{b_2 \to 0} \frac{R^{-1}(R(s(b_2 : w)) + b_2) - s(b_2 : w)}{b_2}$$
$$= \frac{1}{R'(s(0 : w))}$$

and so, holding the ratio of payments $\frac{b_2}{b_1}$ constant, letting $b_1 \to 0$ gives:

$$\beta\delta E_t V'(w_{t+1}) > \frac{b_2}{b_1}\left[\beta\delta E_t\left[\frac{V'(w_{t+1})}{R'(s(w_{t+1}))}\right] + E_t O_b(b_2 : w_{t+1})\right]$$

where $O_b(b_2 : w_{t+1}) = \lim_{b_2 \to 0} \frac{O(b_2 : w_{t+1})}{b_2}$ and we use $s(w_{t+1})$ as a shorthand for s$(0 : w_{t+1})$. Using that $\beta\delta V'(w_{t+1}) = d_{t+1}u'(c_{t+1})$ and $R'(s(w_{t+1})) = MRS_{t+1} = \frac{u'(c_{t+1},\rho_{t+1})}{d_{t+2}u'(c_{t+2},\rho_{t+2})}$ gives

$$\frac{E_t\left[d_{t+1}u'(c_{t+1},\rho_{t+1})\right]}{E_t\left[d_{t+1}d_{t+2}u'(c_{t+2},\rho_{t+2}) + O_b(b_2 : w_{t+1})\right]} > \frac{b_2}{b_1}. \tag{4}$$

Using equivalent arguments to those deployed in the proof of Proposition 1, this is enough to conclude the existence of a lower bound $\bar{b_1}$ such that for all pairs of payments $b_1 < \bar{b_1}$, $b_2 = b_1\hat{R}$**,** the earlier payment will be preferred if $\hat{R}$ is below the left hand side of the above inequality. An equivalent argument shows that the decision maker prefers b_2 if the opposite inequality holds.

We now discuss the nature of $O_b(b_2 : w_{t+1})$ (for convenience we will suppress the preference parameter ρ). In order to do so we add and subtract to $O(b_2 : w)$ the value of the status quo, i.e. receiving no transfers:

$$\beta\delta u\left(w - s(0 : w)\right) + \beta\delta^2 E_{t+1}W\left(y + R\left(s(0 : w)\right)\right).$$

Doing so allows us to write

$$\frac{O(b_2:w)}{b_2} = \beta\delta\frac{u\left(w-s(b_2:w)\right)-u\left(w-s(0:w)\right)}{b_2}$$
$$-\beta\delta\frac{u\left(w+t(b_2:w)-s(0:w+t(b_2:w))\right)-u\left(w-s(0:w)\right)}{b_2}$$
$$+\beta\delta^2\frac{E_{t+1}W\left(y_{t+2}+R\left(s(b_2:w)\right)+b_2\right)-E_{t+1}W\left(y+R\left(s(0:w)\right)\right)}{b_2}$$
$$-\beta\delta^2\frac{E_{t+1}W\left(y_{t+2}+R\left(s(0:w+t(b_2:w))\right)\right)-E_{t+1}W\left(y+R\left(s(0:w)\right)\right)}{b_2}$$

Taking the limit as b_2 goes to zero gives

$$\beta\delta\frac{u'(c(w))}{R'(s(0:w))}\left[-1-R'(s(0:w))s_b(0:w)+s_w(0:w)\right]$$
$$+\beta\delta^2E_{t+1}W'\left(y_{t+2}+R(s(0:w))\right)\left[1+R'(s(0:w)s_b(0:w)-s_w(0:w)\right]$$

where s_w and s_b are respectively the marginal propensity to save from current income and the future transfer b. This means that

$$O_b(b_2:w_{t+1}) =$$
$$\left[\beta\delta^2E_{t+1}W'\left(y+R(s(0:w))\right)-\frac{\beta\delta u'(c(w))}{R'(s(0:w))}\right]\left[1+R'(s(0:w))s_b(0:w)-s_w(0:w)\right]$$

The approximation error will therefore be zero if either of the two bracketed terms is equal to zero. If $\beta=1$, the agent is *dynamically consistent*, and W will be equal to the value of money in period $t+2$ from the perspective of $t+1$. Thus, the Euler equation for the $t+1$ agent ensures that the first term is equal to zero.

If $\beta<1$, we can explore the nature of the error term by deriving expressions for s_b and s_w using the SHEE. Note that the SHEE must continue to hold in response to infinitesimal changes in income in the current period. The derivatives of the left and right hand side of the SHEE in response to a change in w must be equal, and so (suppressing the transfer term in savings $s(0 : w) = s(w))$

$$(1 - s_w(w))\, u''(c(w)) = s_w(w) R''(s(w)) E\left[d\left(w'\right) u'(c(w'))\right] \\ + s_w(w) \left[R'(s(w))\right]^2 E\left[d'(w') u'(c(w')) + d(w') c'(w') u''(c(w))\right]$$

where $w' = y + R\left(s(w)\right)$. This implies

$$s_w(w) = \frac{u''(c(w))}{u''(c(w)) + R''(s(w)) E[d(w') u'(c(w'))] + [R'(s(w))]^2 E[d'(w') u'(c(w')) + d(w') c'(w') u''(c(w'))]}$$

Similarly, taking derivatives with respect to an anticipated increase in y (or equivalently, b_2) in all states of the world gives

$$-s_b(w) u''(c(w)) = s_b(w) R''(s(w)) E\left[d\left(w'\right) u'(c(w'))\right] \\ + \left(R'(s(w)) s_b(w) + 1\right) R'(s(w)) E\left[d'(w') u'(c(w')) + d(w') c'(w') u''(c(w'))\right]$$

meaning

$$s_b(w) = \frac{-R'(s(w)) E[d'(w') u'(c(w')) + d(w') c'(w') u''(c(w'))]}{u''(c(w)) + R''(s(w)) E[d(w') u'(c(w'))] + [R'(s(w))]^2 E[d'(w') u'(c(w')) + d(w') c'(w') u''(c(w'))]}$$

The key term for the evaluation of $O_b(b_2 : w_{t+1})$ is $1 + R'(s(w)) s_b(w) - s_w(w)$. Plugging in the above expressions gives

$$1 - \frac{[R'(s(w))]^2 E[d'(w') u'(c(w')) + d(w') c'(w') u''(c(w'))] + u''(c(w))}{R''(s(w)) E[d(w') u'(c(w'))] + [R'(s(w))]^2 E[d'(w') u'(c(w')) + d(w') c'(w') u''(c(w'))] + u''(c(w))} \tag{5}$$

Notice first that $R''(s(w)) = 0$ when there are no credit constraints, so that $O_b(b_2 : w_{t+1}) = 0$. Thus, the approximation error is equal to zero if either there is no present bias, or if interest rates are not a function of savings s.

Outside these cases, it is possible to sign the term $O_b(b_2 : w_{t+1})$ under mild conditions. Note that $R''(s(w))$ and $u''(c(w'))$ are negative, and $d(w')$, $c'(w')$, and $u'(c(w'))$ are positive, so as long as $d'(w')u'(c(w'))$ is small in magnitude,[1] the numerator and denominator will both be negative, and the whole fraction will be less than 1. Thus the expression (5) is positive and less than 1.

Next, consider the term $\beta\delta^2 E_{t+1}W'(y+R(s(w))) - \frac{\beta\delta u'(c(w))}{R'(s(w))}$. Recall that if the agent were time consistent, this would equal zero. A present-biased decision maker will consume more in period $t+1$ than is desirable from the perspective of the agent in period t, meaning that the marginal utility of consumption will be lower, and so this expression will be positive. Thus it must be the case that $O_b(b_2 : w_{t+1}) > 0$. Looking at the expression in 4, the presence of a positive $O_b(b_2 : w_{t+1})$ term will reduce the left-hand side, meaning that it will reduce the break-even value of $\frac{b_2}{b_1}$ – in other words, it will make Decision B look more patient than without the bias. □

C Proofs for the Claims in Section 1.5

We begin by proving the two predictions in section 1.5 that are not proved in the text.

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

Proof. We show that savings s_t must increase with with y_t and thus reduce $R'(s_t)$. 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

[1] Specifically if $|\left[R'(s(w))\right]^2 E\left[d'(w')u'(c(w'))\right]| < |\left[R'(s(w))\right]^2 E\left[d(w')c'(w')u''(c(w'))\right] + u''(c(w))|$. This will be the case if β is close to 1 – i.e. the decision maker is not too present biased – as in the limit where $\beta = 1$, d becomes constant.

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, $s > s'$ and $R'(s) < R'(s')$. □

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

Proof. 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 given w_t). If $s \leq s'$ (i.e. the consumer saves more when current-period marginal utility is higher) it must be the case that

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

a contradiction. Thus, $s > s'$ and $R'(s) < R'(s')$. □

We next discuss the conditions under which these predictions mean that MRS will exhibit covariance with the financial variable of interest of the appropriate sign. In each case, the object of interest is the expectation of a sample covariance between a variable x (which could be income shocks, preference shocks, spending, etc.) and MRS, calculated from a T length sample. That is, we consider

$$E\left(T^{-1}\sum_{t=1}^{T}\left(MRS_t - \overline{MRS}\right)(x_t - \bar{x})\right),$$

where $\overline{MRS} = T^{-1} \sum_{t=1}^{T} MRS_t$ and $\bar{x} = T^{-1} \sum_{t=1}^{T} x_t$.

We will show the following: if x_t is positively (negatively) related to MRS_t for any (fixed) realization of the other state variables,[2] as in the propositions above, we can conclude that the related sample covariance will be weakly positive (negative) in expectation, provided that

1. the distribution for the initial savings stock s_0 and the parameterization of the model described in section 1.2 is such that, at any subsequent time t the distributions of MRS and x conditional on the remaining state variables have finite second moments; and

2. the magnitude of the covariance between MRS_t and x_t is larger than between MRS_t and x_s for any $s \neq t$, that is

$$|E\left[\left(MRS_t - E(MRS)\right)\left(x_t - E(x)\right)\right]| \geq$$
$$|E\left(\left(MRS_t - E(MRS)\right)\left(x_s - E(x)\right)\right)|$$

   for any $t>0$ and $s \neq t$.

Whether or not this second condition holds will of course depend on the underlying parameters of the model. However, this formulation has an intuitive interpretation - financial variables must be more strongly related to MRS in the concurrent period than in future or lagged periods. This holds true for the simulations we run on various parameterized versions of the model in appendix D.

Note that this covariance is zero for $s > t$ if income and preference shocks are iid, so this assumption requires only that MRS is more strongly related to current than to lagged values of these variables.

[2]The state variables are the preference shock, the income shock, and last-period savings.

First, note that

$$E\left(T^{-1}\sum_{t=1}^{T}\left(MRS_t-\overline{MRS}\right)(x_t-\bar{x})\right)=$$

$$E\left(T^{-1}\sum_{t=1}^{T}MRS_t\left(x_t-\bar{x}\right)\right)-E\left(T^{-1}\sum_{t=1}^{T}\overline{MRS}\left(x_t-\bar{x}\right)\right)=$$

$$E\left(T^{-1}\sum_{t=1}^{T}MRS_t\left(x_t-\bar{x}\right)\right),$$

so we can focus on the expression in which only x is centered.

We will assume that MRS and x are negatively related and show that the sample covariance is negative in expectation. The argument for a positive relation and positive covariance is entirely analogous.

Fix some $1 \leq \tau \leq T$. In order to prove the result, it is sufficient to show that, for any such τ,

$$E\left[MRS_\tau\left(x_\tau-\bar{x}\right)\right]\leq 0$$

as the sample covariance is simply the sum of such expectations across all τ.

Note that

$$E\left[MRS_\tau\left(x_\tau-\bar{x}\right)\right]=E\left[MRS_\tau\left(x_\tau-T^{-1}\sum_{t=1}^{T}x_t\right)\right]$$

$$=\frac{1}{T}\sum_{t=1}^{T}E\left[MRS_\tau\left(x_\tau-x_t\right)\right]$$

It is enough therefore to show that $E\left[MRS_\tau\left(x_\tau-x_t\right)\right]$ is weakly negative for every τ. If $t=\tau$ then we have $E\left[MRS_\tau\left(y_\tau-y_\tau\right)\right]=0$. Otherwise, add and subtract $E(x)$:

$$E\left[MRS_\tau\left(x_\tau-x_t\right)\right]=$$

$$E\left[MRS_\tau\left(x_\tau-E(x)\right)\right]-E\left[MRS_\tau\left(x_t-E(x)\right)\right]$$

The first term is the covariance of MRS_τ and x_τ. Using the law of iterated expectations we have

$$E\left[MRS_\tau\left(x_\tau - E(x)\right)\right] = E\left[E\left(MRS_\tau\left(x_\tau - E(x)\right)|z_\tau\right)\right],$$

where z_τ is the value of the remaining state variables at time τ. By assumption, we know that for any z_τ, MRS_τ is a monotonically decreasing function of x_τ. This is sufficient to guarantee that the conditional covariance is weakly negative under assumption (1) of finite second moments for MRS and x (see for example Schmidt (2003)). The unconditional covariance, as the expectation of weakly negative numbers, must also be weakly negative. The second term is the covariance between MRS_τ and x_t for $t \neq \tau$ which, by assumption (2), is smaller in magnitude than the first term. Thus the difference between the two must also be weakly negative.

C.1 Relationship between Spending and MRS

In the partial constraints model with both income and preference shocks the relationship between spending and measured MRS is indeterminate: on the one hand, a positive income shock (such as an unexpected payment) will lead to an increase in consumption expenditure,[3] accompanied by a *fall* in MRS. On the other hand, a preference shock which increases the marginal utility of consumption (such as the illness of a child) will lead to a rise in consumption spending as well as a *rise* in measured MRS. Thus, the two different types of shock can lead to either a positive or negative relationship between consumption spending and measured MRS. The average relationship can therefore tell us something about the relative importance of each type of shock.

We illustrate this claim using a simple example with a particular type of preference shock, where the household spends money that does not generate utility in the way that regular consumption does. Consider for example a household that usually spends all its income on food. In one period, suppose an asset used for household production breaks and has to be

[3] Harris and Laibson (2001) show that the consumption function is strictly increasing in wealth provided that β is sufficiently close to 1 and f_y and u are three times continuously differentiable.

replaced at the cost of \$10. A total expenditure of x (on food and the asset) then leads to the same marginal utility of consumption that would be associated with the expenditure level $x - \$10$ absent the need to replace the asset. Such a "spending shock" therefore causes variation in the marginal utility of consumption at given observed consumption levels. We can write this as $u(c_t, \rho_t) = u(n_t)$ where $n_t = c_t - \rho_t$: utility depends on "net consumption" n_t. The individual maximizes

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

with wealth at $t = 0$ given and under the constraint $s_t + n_t = w_t - \rho_t$. The decision maker chooses n_t conditional on wealth w_t and preference shock ρ_t. Note that preference shocks act essentially as negative income shocks; they reduce the funds available for consumption and savings $s_t + n_t$. We assume that the researcher observes consumption expenditure c_t but not n_t or ρ_t.

In order to examine the relative importance of shocks to w_t and ρ_t, one can examine the relationship between spending c_t and MRS_t. Since $R'' < 0$, this is determined by the negative of the covariance of savings and spending. In the canonical model without preference shocks, we have $c_t = n_t$, and any increase in spending is caused by an increase in income. An increase in income also leads to higher savings and so to a lower MRS. By contrast, a preference shock – i.e. an increase in ρ_t – increases c_t, but reduces savings, leading to a higher MRS.

Taking a Taylor series expansion of s around the average levels of the wealth and preference shocks $\bar{w}$ and $\bar{\rho}$, we get

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

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 marginal effect of preference shocks on consumption. Note that $(w_t - \bar{w})$ incorporates both income shocks and differences in the current stock of savings relative to $\bar{w}$. We can similarly approximate spending as

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

Assuming that preference shocks are distributed independently of wealth, this implies[4]

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

For the purpose of consumption choice, changes to w_t are equivalent to negative changes to ρ_t, so we have $-\frac{\partial n_t}{\partial w_t} = \frac{\partial n_t}{\partial \rho_t}$, and therefore

$$Cov(s_t, c_t) \approx \left(1 - \frac{\partial n_t}{\partial w_t}\right)\left[\frac{\partial n_t}{\partial w_t} Var(w_t) - \left(1 - \frac{\partial n_t}{\partial w_t}\right) Var(\rho_t)\right]. \tag{6}$$

Equation (6) says that the covariance of savings and spending is the difference between the variance of income and preference shocks, weighted by $\frac{\partial n_t}{\partial w_t}$, and scaled by the marginal propensity to save. If the marginal propensity to consume is high, net consumption and therefore total spending closely follows income, but total spending is relatively unaffected by preference shocks, since those are almost entirely compensated by (net) consumption changes n_t. If $\frac{\partial n_t}{\partial w_t}$ is low, income shocks have little effect on net consumption, but preference shocks translate almost entirely into spending changes. Taken together, this means that MRS and consumption will be positively related if

$$\frac{\partial n_t}{\partial w_t} Var(w_t) < \left(1 - \frac{\partial n_t}{\partial w_t}\right) Var(\rho_t) \tag{7}$$

and negative otherwise.[5] Thus, with an estimate of $\frac{\partial n_t}{\partial w_t}$ we can bound the relative variance of the two types of shock.

[4]Here we can think of expectations being taken with regard to the distribution of ρ_t and w_t conditional on w_{t-1}, though the following approximation holds for any distribution in which wealth and preference shocks are independent.

[5]Following the logic of the proof in appendix C, under mild conditions this ensures that the expectation of the sample covariance calculated from a T-length sample will have the same sign.

D Model Extensions

D.1 Endogenous Income Sources

While some income variation is exogenous and beyond the control of the household, the household can likely also endogenously change their income in response to shocks. Most obviously they can adjust the time or effort spent on working, but in our population informal transfers are also likely to be an important endogenous income source (gifts and remittances are about 2% of total average income). Households may also request RoSCA payments or sell items they own.

Here we analyze a model which allows for this possibility. We show that the results of section 1.5 go through under mild conditions: MRS is negatively related to exogenous income shocks, positively related to preference shocks, and negatively related to savings. We also show that endogenous increases in income will be *positively* related to MRS.[6] For example, a preference shock that increases the marginal utility from consumption or an exogenous drop in income will lead the household to increase its endogenous (labor) income and decrease its savings, and measured MRS will be higher. This implies that sources of income which are more under the household's control will be less negatively or even positively related to MRS. Moreover, the endogenous income response has an attenuating effect not only on the relationship between MRS_t and total income, but also between MRS_t and preference shocks.

In order to capture endogenous income, we extend our model to allow for endogenous labor supply. Other endogenous income sources such as gifts would work in a similar way, assuming they are subject to a utility cost that is increasing and convex in the size of the gift (for example because households have altruistic preferences, because they suffer some shame, or because they will have to reciprocate at some point in the future). In addition to non-labor income y_t, in each period the consumer has one unit of leisure, any fraction of which she can either consume (as leisure l_t) or sell at wage rate W. Utility is now given by

[6] Assuming that changes are driven by shocks to income and preferences, and not shocks to the wage rate.

$$u(c_0, \rho_0, l_0) + \beta E_0 \sum_{t=1}^{\infty} \delta^t u(c_t, \rho_t, l_t),$$

and current-period resources are $y_t + W(1 - l_t) + R(s_{t-1})$.

Define the instantaneous function $v(e_t, \rho_t)$ to be the utility of net spending e_t, defined as any expenditure above labor earnings $W(1 - l_t)$ in period t, assuming that the leisure-labor trade-off has been made optimally, i.e.

$$v(e_t, \rho_t) = \max_{c_t, l_t} u(c_t, \rho_t, l_t)$$

subject to

$$c_t - W(1 - l_t) = e_t.$$

We can think of this as the instantaneous utility obtained if the decision maker adjusts their savings by the amount e_t in period t, and therefore rewrite the dynamic problem purely in terms of choice of e_t:

$$v(e_0, \rho_0) + \beta E_0 \sum_{t=1}^{\infty} \delta^t v(e_t, \rho_t)$$

such that

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

We showed in section 1.5 that exogenous income shocks (i.e. changes in y_t) are negatively related to MRS_t. The argument holds equivalently if the function v is concave in its first argument. Conditions on the utility function u which guarantee that this is the case are well known: for example that u is concave in both c and l, and that consumption and leisure are both normal goods (see Bordley (1995)). Similarly, the result pertaining to the relationship between preference shocks and MRS goes through if we replace the original assumption, that an increase in ρ_t everywhere increases the derivative of u with respect to c, with the assumption that it everywhere increases the derivative of v with respect to e. Finally, savings remain negatively related to MRS by identity. By contrast, endogenous increases in income (i.e. in $W(1 - l_t)$) will be *positively* related to MRS.[7]

Lastly, note that some types of *spending* may also respond to shocks in similar ways as labor supply. Specifically, consider spending that does not directly enter the utility function, but is a response to a preference shock; say, the repair of a motorbike. The amount spent can be adjusted to some degree at the cost of a reduction in utility (e.g. spending time searching for a used motorbike part and doing the repair oneself). An argument similar to the above applies; instead of increasing income by increasing effort, the household can reduce spending by increasing effort, and this endogenous response may attenuate the effect of the relationship between MRS and spending on shocks as well as MRS and income shocks. We address this by instrumenting for spending on shocks with the occurrence of the shock (see section 3).

D.2 Serially Correlated Income and Preference Shocks

Our baseline model assumes that income and preference shocks are independent over time. There are, however, circumstances in which this may be unrealistic, especially for income – for example, poor business conditions in one week may be predictive of bad conditions in the following week, lowering income today as well as expected income tomorrow. Do our results hold up in the case of correlated shocks? For the relationship between savings and MRS the answer is yes. The relationship between savings and the interest rate is mechanical, rather than behavioral, and so the two remain negatively related by assumption. Moreover, it is still the case that measured MRS reflects the interest rate and so the third prediction of section 1.5 holds.

For the relationship between shocks and MRS, the answer depends on the degree of cor-

[7] Assuming that changes are driven by shocks to income and preferences, and not shocks to the wage rate.

relation. Consider first whether income shocks are negatively related to MRS. In the case of the partial constraints model, this boils down to the question of whether the marginal propensity to consume from the shock is less than one: if so, then a positive income shock will lead to increased savings, and so a fall in MRS. If not, then the shock will lead to a *decrease* in savings and so an *increase* in MRS, the opposite of our prediction.

The question when correlated income shocks lead to an increase in savings has been studied in macroeconomics (see for example Uribe and Schmitt-Grohé (2017)). Broadly speaking, the answer is that if a shock raises current income more than permanent income, savings will increase, because the decision maker wishes to move resources into the future. Thus, if an AR(1) income process has a parameter of less than one, a positive shock will increase savings, as current income increases more than lifetime income. If the parameter is greater than one, however, a positive shock can decrease savings and the results of section 1.5 will not hold. In the case of exponential discounting, quadratic utility, and constant interest rates, this result can be readily established analytically (see Uribe and Schmitt-Grohé (2017) chapter 2). Maliar and Maliar (2004) show that the same result holds in the special case of quasi-hyperbolic discounting, exponential utility, and no credit constraints.

Outside these simple cases, and in particular in the presence of partial credit constraints, analytical results can no longer be obtained.[8] We therefore numerically simulate our more complex model and show that for autoregressive parameters less than one, the marginal propensity to save is strictly positive for a wide range of parameters. Thus, we know our results are not robust for autoregressive parameters greater than 1, but we have not been able to find counterexamples for parameters less than 1.

These simulation results are based on the following parameterization of the model. Utility is assumed to be of the Constant Relative Risk Aversion (CRRA) form. We abstract from preference shocks:

$$u(c, \rho) = u(c) = \frac{c^{1-\gamma} - 1}{1 - \gamma}$$

[8] Uribe and Schmitt-Grohé (2017) (chapter 4) use numerical simulations to derive similar results for the case of upward-sloping interest rates and capital accumulation, for example.

Log income is assumed to follow an autoregressive process

$$\ln(y_t) = \mu \ln(y_{t-1}) + \epsilon_t,$$

where ϵ_t is drawn independently from a normal distribution with standard deviation σ. We use an AR(1) process in log income in order to ensure that income is positive.

The interest rate function is given by

$$R(s) = a - ae^{-bs}.$$

$R(s)$ is concave, with $R(0) = 0$ and a savings ceiling a (note that $R(s) \to a$ as $s \to \infty$). b determines the curvature of R.

This means that the model has seven parameters: γ (curvature of the utility function), μ (persistence of income shocks), σ (standard deviation of income shocks), a and b (curvature of the interest rate function), β (present bias) and δ (discount rate).

This model can be solved numerically using Euler equation iteration. There are two state variables: income y and savings s. The numeric approach approximates the optimal policy (i.e. consumption) on a discrete grid of values for the two state variables. This in turn allows us to approximate the marginal propensity to consume (MPC) from income shocks at each point on this grid.

In table D.1, we report the maximal MPC across all grid points for various different parameterizations of the model. The key question is whether the MPC is always less then one: if so, then savings will rise with increases in income, as they do in the iid case, meaning that income and MRS are negatively related. Each row varies the degree of income autocorrelation, given by μ. Each column reports results for different configurations for the other six variables, as described below:

- Baseline: $\gamma = 2$, $\sigma = 0.01$, $a = 1.05$, $b = 0.98$, $\beta = 0.97$, δ=0.99

- High risk aversion: As Baseline, apart from $\gamma = 4$
- High income variation: As Baseline, apart from $\sigma = 0.04$
- High credit constraints: As Baseline, apart from $a = 0.94$ and $b = 1.06$
- High present bias: As Baseline, apart from $\beta = 0.95$
- High discounting: As Baseline, apart from $\delta = 0.9$

The final row reports the average correlation between income and savings across 100 repetitions of 20-period samples, drawn from the stationary distribution of the model under each configuration.

Table D.1: Maximal MPC and simulated correlation of income and savings from parameterized model.

Persistence of the income shock μ	Baseline	High: Risk aversion γ=4	High: Variance σ=0.04	High: Fin. constraints a=0.94, b=1.06	High: Present bias β=0.95	High: Discounting δ=0.9
	Maximal MPC in simulations					
0.65	0.78	0.67	0.78	0.79	0.80	0.85
0.70	0.81	0.70	0.80	0.82	0.83	0.88
0.75	0.84	0.74	0.83	0.84	0.85	0.91
0.80	0.87	0.78	0.86	0.87	0.88	0.94
0.85	0.90	0.83	0.89	0.91	0.92	0.98
0.90	0.94	0.88	0.93	0.94	0.95	0.95
0.95	0.98	0.95	0.97	0.98	0.99	0.99
	Average estimated correlation between income and savings (100 simulations)					
	0.55	0.58	0.58	0.50	0.48	0.55

Baseline: γ=2, σ=0.01, a=1.05, b=0.98, β=0.97, δ=0.99. Each column represents a different parameterization of the model as described in the text. Each row uses a different μ, the persistence of the income shock.

Our results show that, for a wide range of parameterizations, the marginal propensity to consume is less than one (meaning that savings increase as a result of income shocks) and the average sample correlation between income and saving is positive.

The same reasoning holds true when considering correlated preference shocks: as long as shocks are not 'too' correlated, our results go through. This can be seen directly in the special case of appendix C.1, where preference shocks work as negative income shocks.

D.3 Shocks to the Returns Function R_i

So far we have assumed that the interest rate faced by the household is determined only by the level of savings through R. However, interest rates may vary for other reasons, for example if investment opportunities available to the household change over time.

Shocks to R that are observed by the household introduce a new force that can lead to a relationship between savings and MRS. To see this, consider the no-constraints version of the model with a period-specific interest rate, in which

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

Here, (for example) a positive shock to the interest rate in period t will lead to an increase in MRS. In order to determine the resulting change in savings, we can take the derivative of $(1+r_t)E_t\left[d_{t+1}u'(c(w_{t+1},\rho_{t+1}),\rho_{t+1})\right]$ – i.e. the marginal value of savings – with respect to a change in r (holding s_t constant). If this derivative is positive (negative), then $u'(c(w_t,\rho_t),\rho_t)$ must rise (fall) in order to maintain the Euler equation, meaning that consumption must fall (rise) and savings rise (fall) as a result of the change. The relevant derivative is

$$E_t\left[d_{t+1}u'(c_{t+1},\rho_{t+1})\right] + (1+r_t)s_tE_t\left[d_{t+1}u''(c_{t+1},\rho_{t+1})\frac{\partial c_{t+1}}{\partial w_{t+1}} + \frac{\delta d_{t+1}}{\delta w_{t+1}}u'(c_{t+1},\rho_{t+1})\right].$$

The first expectations term captures the mechanical effect of the increase in the interest rate on the marginal value of savings (the substitution effect of making saving more attractive than immediate consumption), and it is always positive. The second term captures the impact of the change in wealth caused by higher interest rates, first through the impact on future marginal utility (the income effect), and second through the change in the discount factor in the quasi-hyperbolic model (where $\beta < 1$). With exponential discounting, the second expectation will be negative, as $\frac{\delta d_{t+1}}{\delta w_{t+1}}$ equals zero, the discount rate and marginal propensity to consume are positive, and the utility function is concave. Thus, if the household is in debt and savings are negative, savings and measured MRS will be *positively* related; both the substitution and income effect drive the household to reduce current consumption. If savings are positive, then the relation may be positive or negative, depending on whether the income or substitution effect dominates. Higher savings tend to make a negative relationship more likely. With quasi-hyperbolic discounting, the wealth effect depends in addition on the effect of wealth on the marginal propensity to consume and therefore the discount factor.

These results mean that shocks to the interest rate would introduce noise in the estimations we perform in section 3. If savings are negative or the substitution effect dominates the interest rate effect on savings, then the relationship of MRS with savings studied in section 3.2 constitutes a test of the relative importance of exogenous shocks to R, and endogenous changes to the interest rate due to changes in s. The fact that we find a negative relationship implies that the latter are more important.

If instead the household is a net saver and the income effect of the interest change dominates, then the negative relationship between MRS and savings could be due to either endogenous or exogenous changes in interest rates.Could the no-constraints model be correct, yet financial shocks are correlated with the interest rate due to this effect? Note first that the results we document in section 3 are robust to the inclusion of time dummies, which control for aggregate changes in the interest rate, e.g. due to macroeconomic conditions. Note further that shocks to the interest rate would lead to a relationship of MRS with endogenous (labor) income and consumption spending, but not with financial shocks. Thus, some third event would have to cause both a change in in the individual-specific interest rate and individual-specific financial shocks, at a high enough frequency to be picked up in our three-week panel. Given the nature of the shocks we examine (for example the sickness of a family member or payment received from the government), we believe that the channel highlighted in our model is the most plausible explanation.

D.4 Naive Decision Makers

So far we have assumed that subjects with time-inconsistent preferences are sophisticated: they realize that their future selves will make choices in order to maximize utility functions which are different from their own. Another possibility is that subjects are naive, and believe that in the future they will behave in a manner consistent with their current preferences. There is evidence that a significant fraction of the population are at least partially naive (see for example DellaVigna (2009)).

A naive decision maker will choose consumption in period 0 in order to maximize

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

under the assumption that for periods $s > 0$, they will make choices that maximize

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

In this case the subject believes that their decisions for all future periods will be governed by the standard exponential Euler equation, and so

$$\hat{E}_0 u'(c(w_s, \rho_s), \rho_s) = \hat{E}_0 \left[R'(s_s)\delta u'(c(w_{s+1}, \rho_{s+1}), \rho_{s+1})\right]$$

where $\hat{E}_0$ indicates the (incorrect) expectations of the naive period 0 agent and $c(w_{s+1}, \rho_{s+1})$ is their (incorrect) prediction of the consumption function for all periods $s > 0$. Moreover, because the time 0 agent assumes that the time 1 agent maximizes the same utility function as the time 0 agent, a standard envelope theorem argument implies that

$$V'(w_1) = \hat{E}_0 u'(c(w_1, \rho_1), \rho_1)$$

where V is the perceived value of wealth in period 1 from the perspective of the period 0 agent. The first-order condition for consumption in period 0 therefore implies

$$u'(c_0, \rho_0) = \beta\delta\hat{E}_0 \left[R'(s_0)u'(c(w_1, \rho_1), \rho_1)\right].$$

Using these results, we can first establish that an equivalent result to Proposition 1 holds using the same arguments as in appendix B, replacing inequality 3 with

$$R'(s_t^*) = \frac{u'(c_t^*, \rho_t)}{\hat{E}_t\left[\beta\delta u'(c(w_{t+1}, \rho_{t+1}), \rho_{t+1})\right]} > \hat{R}.$$

One can also derive an equivalent of Proposition 2, replacing inequality 4 with

$$\frac{\hat{E}_t\left[u'(c_{t+1}, \rho_{t+1})\right]}{\hat{E}_t\left[\delta u'(c_{t+2}, \rho_{t+2})\right]} > \hat{R}.$$

Note that the approximation error $O(b_2 : w_{t+1})$ is zero because the subject assumes that they will be time-consistent.

The three predictions from section 1.5 hold in the case of the naive model. The relationship between measured MRS and the savings rate comes directly from equation (2). For the results regarding income and preference shocks, an examination of the proofs from appendix C show that the key ingredients are (1) the additive separability of the utility function between consumption in period t and wealth in period $t+1$; (2) the concavity of the per-period utility function; and (3) the fact that utility is increasing in wealth in period $t+1$, all of which are true in the naive model.

One key difference between the naive model and either the exponential or the sophisticated hyperbolic model is that it introduces a new channel which can lead to a gap between $\frac{a_2}{a_1}$ and $\frac{b_2}{b_1}$. Because the naive decision maker believes that they will be more patient in period $t+1$ than they are in period t, they believe that their consumption rules will also differ, and that they will save more at any given level of wealth. Thus, even if the economy is in steady state, the household believes that it will save more on average in the future period, and so

Table E.1: Characteristics of experimental subjects.

		Observations	Mean	Std. Dev.	Minimum	Maximum
Respondent:	Holds salaried employment	1008	12.20%	32.70%	0	1
	Male	1009	87.20%	33.40%	0	1
	Under 25 years	1009	4.96%	21.70%	0	1
	45 years and over	1009	26.26%	44.00%	0	1
	At least four years of school	1009	34.49%	47.56%	0	1
	Can read and write	1009	50.05%	50.00%	0	1
Household:	Number of members	1013	6.29	3.15	2	22
	Children under 15	1013	3.31	2.03	0	13

Sample: 1017 respondents with at least one completed MPL. Four subjects without demographic information.

will face a lower interest rate. This in turn means that, ceteris paribus, $\frac{b_2}{b_1}$ will be lower than $\frac{a_2}{a_1}$ and the household will look more patient for choices involving future payments.

E Data Summary

E.1 Experimental Data

\

Table E.2: Weekly distribution of experimentally elicited MRS in decision A.

Implied MRS at	week 1		week 2		week 3	
interval midpoint	Number	Percent	Number	Percent	Number	Percent
0.71	53	6.4%	38	4.4%	43	5.0%
0.80	5	0.6%	1	0.1%	3	0.4%
0.93	22	2.7%	25	2.9%	17	2.0%
1.1	122	14.7%	122	14.0%	113	13.2%
1.35	61	7.4%	74	8.5%	95	11.1%
1.75	58	7.0%	93	10.7%	74	8.6%
2.5	53	6.4%	56	6.4%	84	9.8%
4.5	54	6.5%	47	5.4%	54	6.3%
8	402	48.4%	415	47.7%	375	43.7%
Total	830	100	871	100	858	100

Distribution of MRS in decision A, ratio of earlier to later payment at interval midpoint. Assuming value 0.708 for individuals who always choose the late payment and 8 for those who always choose the earlier payment. Inconsistent choices excluded.

E.2 Income by Source

Figure E.1 shows the distribution of non-zero income realizations. We break out labor income, where we include all income-earning activities that require labor input from the household; working in one's own business (in activities such as making and selling items, or driving a taxi), working for a piece-rate or time-dependent pay, or working for a regular salary. Non-labor income includes important income sources like rent paid to the household or transfers from other households. Since the timing and variation of non-labor income is likely not entirely exogenous to the household's consumption needs either, we further categorized government and military payments, rent paid to the household, and loans paid back to the household by others as most likely to be "exogenous", that is, less likely to be determined by a choice that the household made. By comparison, payments received from a RoSCA (tontine), sales revenues, and transfers and gifts from other households might be actively requested or generated by the household according to their consumption smoothing needs, and are thus to some degree "endogenous". We decided on these categories before running any data analysis of effects by income category.

Figure E.1: Distribution of different types of income by source, excluding person-weeks with zero income in each category. Labor income has the most non-zero observations, followed by exogenous non-labor income. The overall amounts of non-labor income tend to be smaller.

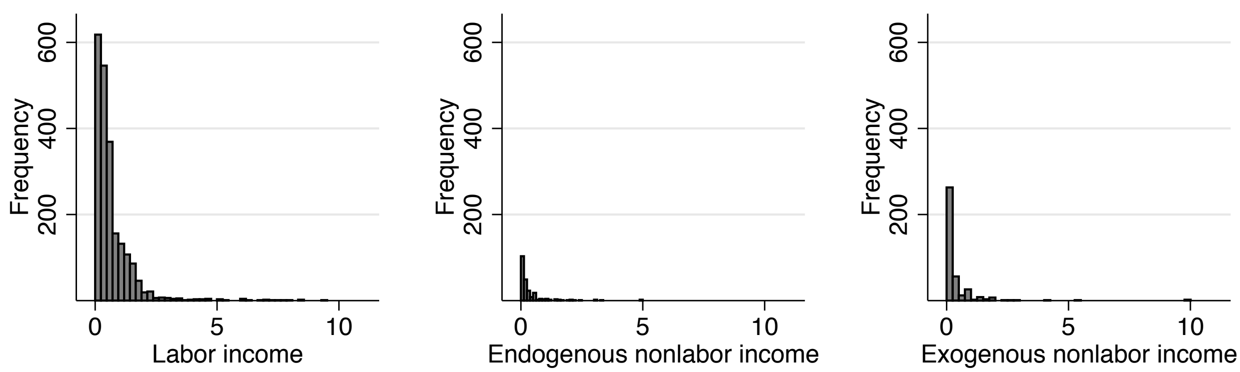

Figure E.1 shows that the income distribution is very skewed. While our outcome variable (MRS) can only take a few fixed values, the reader may be concerned that results are driven by extreme points in the independent variables. We re-estimated all main regressions after transforming income variables, spending variables, and flow savings using the inverse hyperbolic sine function to reduce the influence of such extreme values, and the results are qualitatively unchanged (regression tables available on request).

F Justification for Savings Regression

The theory relates the stock of savings s_{it} to MRS through the function $R(s_{it})$. Our data does not contain good measures of the savings stock of the household, but we can construct a test of this relationship based on flow savings (the difference between income and consumption). If our model is correct, then a Taylor expansion of R' around s_{it-1} links last-period MRS, the change in savings, and current MRS:

$$\begin{aligned} R'_{it}(s_{it}) &= R'_{it}(s_{it-1}) + R''_{it}(s_{it-1})(s_{it} - s_{it-1}) + \frac{1}{2}R'''_{it}(s_{it-1})(s_{it} - s_{it-1})^2 + o_{it} \\ MRS_{it} &= MRS_{it-1} + R''_{it}(s_{it-1})(\Delta s_{it}) + R'''_{it}(s_{it-1})\frac{1}{2}(\Delta s_{it})^2 + \eta_{it} \end{aligned} \tag{8}$$

where $\Delta s_{it} = y_{it} - c_{it}$ represents the flow of savings in t, and the person and time indices on R account for variation in financial conditions between periods and individuals (see section D.3). The term η_{it} includes the approximation error o_{it}; if R_{it} differs from week $t-1$ to t, for example due to a change in credit market conditions, then $MRS_{it-1} = R'_{it-1}(s_{it-1})$ and η_{it} also contains the difference $R'_{it}(s_{it-1}) - R'_{it-1}(s_{it-1})$ (see section D.3).

Now consider the regression equation

$$\Delta\widetilde{MRS}_{it} = \lambda_1\widetilde{\Delta s}_{it} + \lambda_2\left[0.5\widetilde{\Delta s}_{it}^2\right] + \alpha_i + \gamma_t + \epsilon_{it}. \tag{9}$$

$\Delta\widetilde{MRS}_{it}$ and $\widetilde{\Delta s}_{it}$ denote values we observe in our data. $\Delta\widetilde{MRS}_{it}$ is the difference in observed experimental trade-offs between period $t-1$ and t, and is measured with error r_{it}, which will enter ϵ_{it} along with η_{it} above. The fixed effects α_i and γ_t capture any period-specific or individual-specific trends in R'. If there are no such trends, these coefficients will be zero.[9] Measured flow savings $\widetilde{\Delta s}_{it} = \tilde{y}_{it} - \tilde{c}_{it} = \Delta s_{it} + z^y_{it} - z^c_{it}$ may contain measurement

[9] Because measured flow savings are on average negative, estimating the equation without any constant will bias the estimate of λ_1 if such trends are present. Note also that an individual fixed effect can capture changes in experimental decisions over time that are not related to the financial market, due for example to changing levels of trust in the surveyors by our sample subjects. This corresponds to a time trend in (non-classical) measurement error.

error in income and consumption, z_{it}^y and z_{it}^c, respectively.

We first use this regression to test our model against the null hypothesis that the households in our sample are subject to complete credit constraints. In the complete-constraints case, $MRS_{it} = \frac{u'(y_{it})}{E[\beta\delta u'(y_{it+1})]}$ in each period, there are no savings, and (true) flow savings equal zero. Measured flow savings $\widetilde{\Delta s}_{it} = z_{it}^y - z_{it}^c = z_{it}$ therefore represent only measurement error. We would expect λ_1 to be zero under this model.[10] By contrast, if our model holds, we expect $\lambda_1 < 0$, since R is concave.[11]

Under the complete-constraints model, the estimate of λ_1 will not be significantly different from zero, under the assumption that $z_{it} \perp \widetilde{MRS}_{it}$. We do however expect MRS_{it} to vary with (true) income, so if z_{it} is non-classical measurement error and correlated with true income or spending, this assumption may be violated. Could such a correlation be responsible for our negative estimate for λ_1? Our survey asked individuals to report income and spending separately. A plausible source of a correlation between income or spending and error z_{it} is therefore differential under- or over-reporting; and since reported flow savings are on average negative, and reported income has a lower variance than reported spending, this would imply that true income is negatively correlated with flow savings (for example, $\tilde{y}_{it} = \alpha_{it}^y y_{it}$ and $\tilde{c}_{it} = \alpha_{it}^c c_{it}$ so that $z_{it} = (\alpha_{it}^y - \alpha_{it}^c) y_{it}$ and $\alpha_{it}^y < \alpha_{it}^c$ on average). Since MRS is negatively correlated with income in the full constraints model, this error structure would introduce a positive bias and predict $\hat{\lambda}_1 > 0$ if $\lambda_1 = 0$.

Another type of bias could occur if the measurement error in MRS, r_{it}, were directly correlated with z_{it}, in other words, whenever subjects underreport income relative to consumption, they also systematically overreport their MRS (and vice versa), perhaps due to a mistaken perception of being more constrained.[12] However, we judge this source of bias to be fairly

[10]Under the no-constraint and narrow bracketing models $\lambda_1 = 0$ as well: with narrow bracketing, experimental MRS is unrelated to any outside financial variables by assumption, and in the no-constraints case, the interest rate is unaffected by savings, and $R'' = 0$ and $R''' = 0$.

[11]The exception would be if flow savings $\widetilde{\Delta s}_{it}$ and the unobserved variation in MRS, ϵ_{it}, are strongly positively correlated. See below for an additional discussion on potential error correlations.

[12]But note that this *cannot* be driven by perceived consumption $\hat{c}_{it}$. For example, a plausible effect could be that subjects who perceive the household's consumption to be lower than what it really was will also overestimate the marginal utility of additional consumption. This would again lead to a positive correlation

Table G.1: Instrumental variables regression with individual fixed effects (left) and individual and time fixed effects (right). First and third columns reproduce results from table 6, second and fourth columns show first stage regressions.

	Individual FE		Individual and time FE	
	MRS (A)	Adv. event expenses	MRS (A)	Adv. event expenses
Adv. event (0/1)		0.161 ***		0.160 ***
		(0.014)		(0.015)
Adv. event expenses	1.707 **		1.579 **	
	(0.789)		(0.791)	
Observations	2467	2467	2467	2467

Standard errors clustered at the individual level (in parentheses).
*Significance levels * p<0.10, ** p<0.05, *** p<0.01.*

implausible under complete credit constraints, because subjects with no ability to save or borrow should realize that it is not actually possible that they consumed an amount different from their income. Moreover, the way we collected our income and consumption data – by asking household members individually for their income from different sources, as well as their spending on different types of items – makes it unlikely that the difference between total reported consumption and income (i.e. z_{it}) is even salient to the subject.

G Additional Regression Results and Robustness Checks

G.1 Additional Results

Table G.1 reports the first stage of the instrumental variables regression in table 6.

Table G.2 repeats the regressions from Table 7, but includes a cubed savings term. The first-order effect of savings is slightly stronger than with only squared savings included.

G.2 Robustness Checks

We conduct three different robustness checks.

Table G.3 shows MRS regressed on all income and spending components. The two columns titled "IV" again instrument for adverse event spending.Exogenous income is significantly and negatively correlated with the MRS elicited from experimental decisions (MRS_t) in every

of flow savings and MRS.

Table G.2: MRS change over time as a function of (flow) savings, third-order Taylor expansion.

	OLS	OLS	OLS	OLS
Flow savings Δs	-0.231 *	-0.228 *	-0.593 *	-0.592 *
	(0.134)	(0.135)	(0.328)	(0.330)
$0.5\Delta s^2$	-0.0240	-0.0234	-0.0861	-0.0860
	(0.0341)	(0.0344)	(0.134)	(0.134)
$(1/3)\Delta s^3$	0.00705	0.00695	0.0317	0.0317
	(0.00819)	(0.00822)	(0.0286)	(0.0287)
Time FE		yes		yes
Ind FE			yes	yes
Observations	1462	1462	1462	1462

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

specification. Adverse event spending is positively and significantly correlated with MRS, except in the OLS regressions where 'other spending' is included. This specification does not change the coefficient sizes, but does lead to somewhat larger standard errors, possibly due to collinearity. Event spending remains significantly related to MRS in the IV specification, even when other spending is included.

Table G.4 shows the estimation results for exogenous shocks when including quadratic terms. This preserves the expected coefficient signs, but renders the coefficients individually insignificant. However, we can carry out F-tests (likelihood ratio tests) for the joint significance of the terms relating to each type of shocks, and these show that the coefficients on non-labor exogenous income are reliably jointly significant at the 1% level in all specifications. By contrast, adverse event spending is not significant in the OLS specifications and exhibits some concavity. The coefficients suggest that the effect on MRS is strongest at about $322 weekly expenses and declines after. This may be a result of the endogeneity of the amount spent and an implicit selection effect: only households that are not liquidity constrained can afford to spend several hundred dollars out of pocket on an unexpected event in a given week. Unfortunately we do not have two independent instruments for the linear and the squared terms and therefore cannot carry out IV regressions. The coefficient sizes

Table G.3: Income and spending effects on MRS_t.

	OLS	OLS	OLS	OLS	IV	IV	CL
Labor income			-0.185	-0.189	-0.153	-0.159	-0.262 *
			(0.142)	(0.143)	(0.163)	(0.142)	(0.136)
Nonlabor income "endogenous"			-0.330	-0.321	-0.268	-0.265	-0.316
			(0.251)	(0.258)	(0.261)	(0.270)	(0.282)
Nonlabor income "exogenous"	**-0.409 *****	**-0.409 *****	**-0.382 *****	**-0.384 *****	**-0.378 *****	**-0.380 *****	**-0.379 ****
	(0.142)	**(0.149)**	**(0.125)**	**(0.133)**	**(0.128)**	**(0.134)**	**(0.171)**
Other spending			0.268 **	0.245 *	0.192	0.177	0.215 *
			(0.128)	(0.131)	(0.141)	(0.132)	(0.119)
Adv. event expense	**0.252 ***	**0.233 ***	**0.251**	**0.222**	**1.683 ***	**1.562 ****	**0.390 ****
	(0.145)	**(0.139)**	**(0.182)**	**(0.183)**	**(0.865)**	**(0.873)**	**(0.199)**
1/(sd)[(a)]	-	-	-	-	-	-	0.916 ***
							(0.044)
Ind FE	yes	yes	yes	yes	yes	yes	yes
Time FE		yes		yes		yes	yes
Observations[(b)]	2540	2540	2390	2390	2289	2289	12608

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

and signs in the conditional likelihood specification echo those of the OLS estimates but the shock variables are not significant. However, unlike in the OLS estimates, likelihood ratio tests for the coefficients on the quadratic terms show that they jointly contribute significantly (at the 10% level) to improving the conditional likelihood model fit.

Table G.4: Exogenous shocks with quadratic terms of all variables.

	OLS	OLS	OLS	OLS	CL
Nonlabor income	**-0.190**		**-0.195**	**-0.152**	**-0.096**
"exogenous"	**(0.281)**		**(0.268)**	**(0.285)**	**(0.356)**
(Nonlabor income	**-0.030**		**-0.028**	**-0.029**	**-0.030**
"exogenous")^2	**(0.039)**		**(0.038)**	**(0.037)**	**(0.043)**
Adv. event expenses		**0.580**	**0.534**	**0.328**	**0.426**
		(0.358)	**(0.356)**	**(0.424)**	**(0.456)**
(Adv. event expenses)^2		**-0.090**	**-0.079**	**-0.027**	**-0.019**
		(0.072)	**(0.072)**	**(0.126)**	**(0.151)**
Labor Income	-0.055			-0.043	-0.090
	(0.184)			(0.204)	(0.202)
(Labor Income)^2	0.023			-0.034	-0.058
	(0.044)			(0.044)	(0.064)
Nonlabor income	0.27			0.248	0.388
"endogenous"	(0.462)			(0.474)	(0.435)
(Nonlabor income	-0.181			-0.175	-0.251 *
"endogenous")^2	(0.147)			(0.147)	(0.132)
Other Spending				-0.086	-0.321
				(0.169)	(0.214)
(Other Spending)^2				0.0554 ***	0.107 **
				(0.018)	(0.045)
1/(sd)[(a)]				-	0.921 ***
					(0.045)
Ind FE	yes	yes	yes	yes	yes
Observations[(b)]	2484	2543	2540	2390	12616
Joint significance of coefficients (p-value):					
Nonlabor ex. income (all)	0.02 ***	-	0.02 ***	0.02 ***	0.14
Adverse event spending (all)	-	0.16	0.17	0.38	0.18
All quadratic terms	0.46	0.21	0.40	0.72	0.09 *

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

Overall, these results suggest that the sample is not quite large enough to deliver the required power for including higher order terms. Nonetheless, the effects are qualitatively robust and the simple linear estimates do not appear to mask important nonlinear effects.

Lastly, one may be worried about using as a proxy of MRS the midpoint of the interval that contains it according to the subject's decisions. The conditional logit approach resolves this by explicitly modeling the subject's binary decisions and assigning the (unobserved) MRS value that fits these observed choices best. As an additional test of the robustness of the OLS results, we carry out 200 repeated regressions of MRS on financial variables, where we

Table G.5: Regression coefficients from 200 repeated OLS estimates using random draws for the MRS variables from the interval in which they must lie according to experimental choices.

	(Estimation approach)	Coefficient distribution on financial variable			
		1st percentile	Mean	99th percentile	Prop. $\|t\| \geqq 1.69$
Adverse event (0/1)	(fixed effects)	0.181	0.283	0.399	88%
Spending on adverse event	(fixed effects)	0.060	0.264	0.487	49%
Nonlabor income, "exogenous"	(fixed effects)	-0.566	-0.424	-0.313	100%
Savings	(constant)	-0.261	-0.181	-0.097	77%
Savings	(fixed effects)	-0.570	-0.382	-0.164	56%

Mean, 1st, and 99th percentile of 200 coefficient estimates when estimating effect of each financial variable (left) on MRS, using in each regression a new set of uniform random draws from the interval which contains the unobserved MRS (intervals 0.66-0.75 and 6-10 for lowest and highest possible values). Last row shows proportion of coefficient estimates that are significant at 5% level.

draw the MRS from a uniform distribution over the interval of possible values defined by the subject's experimental choices. For the lowest and highest intervals, we use as endpoints 300/450=0.66 and 10 respectively, so that the averages over the lowest and highest interval equal the values we assigned in the regressions reported in the main text.

Summary statistics on the set of regression estimates obtained from this procedure are reported in table G.5. The mean coefficients on an adverse event dummy and on adverse event spending are 0.283 and 0.264, respectively, and the coefficient on exogenous non-labor income is -0.424 on average. The coefficient on savings is on average -0.181 with a constant term and -0.382 with fixed effects. 100% of these estimates have the predicted signs, and between 49% and 100% are themselves significant at the 5% level.

H Conditional Likelihood Estimation

Recall that the decision maker chooses a_1 over a_0 if

$$\frac{u'(c_t,\rho_t)}{E_t\left[d(w_{t+1},\rho_{t+1})u'(c_{t+1},\rho_{t+1})\right]} \leq \frac{a_1}{a_0}$$

and a_0 over a_1 otherwise. The left-hand side is the MRS.

Consider decision set A. For each decision maker i in period t we observe 8 decisions $k = 1, \ldots, 8$. Let $a_0^k = 50k$ and $a_1 = 300$, so that low k corresponds with a high MRS. We model the set of decisions A in a given week as a latent discrete choice problem, separately for each binary choice

$$Prob\left(i \text{ chooses } a_1 \text{ over } a_0^k \text{ in } t\right) = P\left(MRS_{it} + \epsilon_{itk} \leq \frac{a_1}{a_0^k}\right).$$

Assuming that the error term has a logistic distribution with scale parameter σ, the probability is

$$P\left(\epsilon_{itk} \leq \frac{a_1}{a_0^k} - MRS_{it}\right) = \frac{exp\left[\frac{1}{\sigma}\left(\frac{a_1}{a_0^k} - MRS_{it}\right)\right]}{1 + exp\left[\frac{1}{\sigma}\left(\frac{a_1}{a_0^k} - MRS_{it}\right)\right]}.$$

Without an error term, we would by necessity expect decisions to be monotonic, that is, there is a single switch point within a set of choices, and each set of choices pins down MRS_t up to an interval (of varying size across k). We would like to estimate the MRS_{it} as a function of an individual fixed parameter and the effect of financial variables between different periods using conditional likelihood estimation (Chamberlain 1984).

Denote the choice of a_1 over a_o^k in t by $y_{itk} = 1$ (0 otherwise). Assume for the moment that there are only two levels of k and two periods. Now consider for example the probability that $y_{i11} = y_{i12} = 1$ but $y_{i21} = 0$ and $y_{i22} = 0$, given by

$$\frac{exp\left[\frac{1}{\sigma}\left(\frac{a_1}{a_0^1} - MRS_{i1}\right)\right] exp\left[\frac{1}{\sigma}\left(\frac{a_1}{a_0^2} - MRS_{i1}\right)\right]}{\prod_{t=1}^{2}\prod_{k=1}^{2}\left\{1 + exp\left[\frac{1}{\sigma}\left(\frac{a_1}{a_0^k} - MRS_{it}\right)\right]\right\}}$$

Then the probability that $y_{i11} = y_{i12} = 1$ conditional on $\sum_{t,k} y_{itk} = 2$ is

$$\frac{exp\left[\frac{1}{\sigma}\left(\frac{a_1}{a_0^1} - MRS_{i1}\right)\right] exp\left[\frac{1}{\sigma}\left(\frac{a_1}{a_0^2} - MRS_{i1}\right)\right]}{\sum_{t,k|\sum_{t,k} y_{itk}=2}\left(exp\left[\frac{1}{\sigma}\left(\frac{a_1}{a_0^{k_1}} - MRS_{it_1}\right)\right] exp\left[\frac{1}{\sigma}\left(\frac{a_1}{a_0^{k_2}} - MRS_{it_2}\right)\right]\right)},$$

where $MRS_{it} = \alpha_i + \lambda X_{it} + \gamma_t$. Since $\frac{\alpha_i}{\sigma}$ is constant, this term cancels from all the exponents. This construction works equally with three periods and eight choices k, and all possible choice combinations. Denoting $m_{itk} = \frac{1}{\sigma}\frac{a_1}{a_0^k} - \beta X_{it} - \tilde{\gamma}_t$ where $\beta = \frac{\lambda}{\sigma}$ and $\tilde{\gamma}_t = \frac{\gamma_t}{\sigma}$, the log conditional likelihood is

$$\ln L = \sum_{i=1}^{N} \left\{ \sum_{t,k} y_{itk} m_{itk} - \ln \left(\sum_{\{d_{tk}\} \in B_{itk}} exp \left[\sum_{t,k} d_{tk} m_{itk} \right] \right) \right\},$$

where B_{itk} contains all indicator sets such that $\sum d_{tk} = \sum y_{itk}$. This is equivalent to estimating a conditional logistic regression with person fixed effects and period dummies, where we can in addition cluster the error terms at the person level. For computational reasons and given the short panel, it is assumed that the error variance σ is constant across the population. We parameterize the MRS as a linear or quadratic function of the financial variables as in the main text. We exclude any subjects who in any set of decisions switch between early and late payments more than three times. This does not occur frequently, but we assume that these participants have not understood the experimental decisions well enough to contribute meaningfully to the MRS estimation.

I Relationship to the Literature

From a theory standpoint, the closest paper to ours is Pender (1996). He considers an integrated model of choices over experimental payments in two periods in different credit environments, including where interest rates increase with borrowing. Our model extends this work in a number of ways, for example by allowing for preference shocks, time inconsistency and endogenous labor responses, and by showing that the results do not rely on active arbitrage. From an empirical perspective, Pender's identification strategy, which is based on experimental data only, requires that (large) experimental payments themselves affect the household's credit market position, whereas our model and empirics consider the role of exogenous financial shocks on experimental choices.

Concurrent to our paper, Epper (2017) considers theoretically a model in which experimental subjects have a (hard) liquidity constraint, positive income expectations and integrate their experimental choices into their broader consumption plans. While the set-up of this paper shares many features with our own, the aims of the two papers are different: Epper

(2017) shows how such a set up can lead naturally to many observed 'anomalies' in intertemporal choice experiments, even if subjects have standard exponential preferences. In contrast we focus on what can and cannot be learned from experimental choices in the face of soft credit constraints and income and preference shocks.

Another theoretical contribution is Cubitt and Read (2007), who consider the implications of a model of arbitrage with two distinct interest rates for borrowing and saving. Similar to Pender, their predictions derive from the assumption that subjects apply different interest rates to experimental payments at different points in time, but they do not consider the impact of financial shocks. As discussed in section 1.6.1, Andreoni and Sprenger (2012) estimate a parameter that captures background consumption, but identification in their model of experimental choices requires changes to the MRS that cannot be arbitraged. None of these papers explore what can be learned about consumption smoothing, credit constraints or financial shocks from MRS experiments.

On the empirical side, several papers investigate the variation of individual's time preference measures over time (see Chuang and Schechter (2015) for an overview). On balance, papers that study the relationship of experimental time preference measures with external financial conditions offer support for the general idea that the latter can influence the former. None of these papers make predictions about, or separately study, the relationship of MRS with individual-level variation in the inputs into the household's financial decision problem, i.e. financial shocks, spending and savings. They therefore cannot perform the range of tests we use to learn about the household's financial decision-making.

Five studies correlate time preferences with some measure of subjects' outside financial situation in a range of populations and find a relationship.[13] Harrison et al. (2005) correlate measured time preferences with a variety of questions on perceived changes in their financial situation. The coefficients on these survey measures are jointly insignificant, although one

[13] In related work, Meier and Sprenger (2010) study whether credit constraints and future liquidity are correlated with time preference in cross sectional, rather than panel data. They find no evidence for such a relationship. Haushofer et al. (2013) find that experimenter induced income shocks increase impatience. However, their design precludes the credit constraints channel we consider in this paper.

particular question regarding the general economic situation in the country is significantly related to the measured discount rate.[14] Krupka and Stephens (2013) report data collected in the Seattle and Denver Income Maintenance Experiments from the 1970. They find that changes in weekly income and monthly hours worked (as well as inflation rates) are correlated with changes in the discount rate, which they interpret as a response to real interest changes in line with our model predictions. Concurrent with our own study, Carvalho et al. (2016) and Ambrus et al. (2015) report a relationship between financial variables and measured MRS. Carvalho et al. use the natural variation in liquidity caused by the pay cycle to examine the effect 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 in choices over monetary rewards but not over non-monetary real effort tasks, supporting a liquidity-constraint interpretation. Ambrus et al. find that only subjects without a stable background income exhibit present bias. Also contemporaneous with our paper, Cassidy et al. (2018) find a causal link between liquidity and measured present bias using randomized timing of experimental payments.

Three other studies do not find a correlation between financial variables and measured discount rates. Meier and Sprenger (2015) survey a sample of 250 individuals at a tax filing center in Massachusetts and find no relationship between time preferences and demographic or income changes. Similarly, in Chuang and Schechter's sample of 49 households from Paraguay, income changes between 2007 and 2009 have no significant impact on (hypothetical) time preference measures in the two survey rounds. The findings in these papers may be different from ours due to the smaller sample size and the fact that events from the past year have already entered subjects' intertemporal optimization, or do not reflect income shocks .Giné et al. (2018) perform an experiment in rural Malawi to examine the revision of intertemporal choices. Subjects were first asked to allocate money across an intertemporal budget, and were then offered to change their initial allocation later. The authors

[14]"Turning to the economic conditions in the country as a whole, would you say that at the present time economic conditions are better or worse than they were X months ago?"

find no significant effect of measured shocks – a death in the family or "unexpected income shortfalls" – on the subject's choice to revise. This is surprising given our findings, but the authors also state that there were few deaths (2% of households) and income shortfalls were generally small, so that the presence of undetected large effects cannot be rejected given the wide standard errors.

Lastly, some indirect support for our model comes from Halevy (2015b). He reports time t preferences between payments at t and $t+1$ and between payments at $t+1$ and $t+2$, and time $t+1$ preferences over payments at $t+1$ and $t+2$. This allows the categorization of non-stationary time preferences into those that violate time invariance (changing relative value of immediate and delayed payments) and time consistency (changing relative valuations of payments that occur in $t+1$ and $t+2$). While time inconsistency could be a consequence of present bias, a quasi-hyperbolic narrow bracketer should still make time-invariant choices. The author reports significant violations of time invariance, consistent with temporary liquidity constraints.

Bibliography

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

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

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.

Robert F. Bordley. Possible convexity of the indirect utility function due to nonlinear budget constraints. *Economics Letters*, 49(1):19–25, July 1995.

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

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

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

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.

Suresh De Mel, David McKenzie, and Christopher Woodruff. Returns to capital in microenterprises: evidence from a field experiment. *Quarterly Journal of Economics*, 123(4): 1329–1372, 2008.

Angus Deaton. Savings and liquidity constraints. *Econometrica*, 59(5):1221–1248, 1991.

Stefano DellaVigna. Psychology and economics: Evidence from the field. *Journal of Economic literature*, 47(2):315–72, 2009.

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.

Thomas Epper. Income expectations, limited liquidity, and anomalies in intertemporal choice. 2017.

Parikshit Ghosh, Dilip Mookherjee, and Debraj Ray. Credit rationing in developing countries: an overview of the theory. *Readings in the theory of economic development*, pages 383–401, 2000.

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

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

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

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.

Johannes Haushofer, Daniel Schunk, and Ernst Fehr. Negative income shocks increase discount rates. Technical report, 2013.

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

Lilia Maliar and Serguei Maliar. Quasi-geometric discounting: A closed-form solution under the exponential utility function. *Bulletin of Economic Research*, 56(2):201–206, 2004.

Stephan Meier and Charles Sprenger. Present-Biased Preferences and Credit Card Borrowing. *American Economic Journal: Applied Economics*, 2(1):193–210, January 2010.

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

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

Mark R. Rosenzweig and Kenneth I. Wolpin. Credit market constraints, consumption smoothing, and the accumulation of durable production assets in low-income countries: Investments in bullocks in India. *Journal of Political Economy*, 101(2):223–244, 1993.

Klaus D Schmidt. On the covariance of monotone functions of a random variable. 2003.

M. Uribe and S. Schmitt-Grohé. *Open Economy Macroeconomics*. Princeton University

Press, 2017. ISBN 9781400885305.

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