

Payday Loans, Uncertainty, and Discounting: Explaining Patterns of Borrowing, Repayment, and Default*

Paige Marta Skiba
Vanderbilt Law School

Jeremy Tobacman
University of Pennsylvania and NBER

August 21, 2008

Abstract

Ten million American households borrowed on payday loans in 2002. Typically, to receive two weeks of liquidity from these loans households paid annualized (compounded) interest rates over 7000%. Using an administrative dataset from a payday lender, we seek to explain demand-side behavior in the payday loan market. We estimate a structural dynamic programming model that includes standard features like liquidity constraints and stochastic income, and we also incorporate institutionally realistic payday loans, default opportunities, and generalizations of the discount function. Method of Simulated Moments estimates of the key parameters are identified by two novel pieces of evidence. First, over half of payday borrowers default on a payday loan within one year of their first loans. Second, defaulting borrowers have on average already repaid or serviced five payday loans, making interest payments of 90% of their original loan's principal. Such costly delay of default, we find, is most consistent with partially naive quasi-hyperbolic discounting, and we statistically reject nested benchmark alternatives.

*We would like to thank Daniel Benjamin, Steve Bond, Martin Browning, Stefano DellaVigna, Ed Glaeser, David Laibson, Ulrike Malmendier, Markus Mobius, John Muellbauer, Sendhil Mullainathan, Matthew Rabin, Jesse Shapiro, and seminar audiences at Berkeley, Penn, Wharton, Illinois, Cornell, Minnesota, the Federal Reserve Board of Governors, the OCC, the FDIC, Harvard, Oxford, the Society of Computational Economics, Copenhagen, and Yale for valuable feedback. Matthew Smith and Owen Ozier provided excellent research assistance.

JEL Codes: D14 (Personal Finance), D91 (Intertemporal Consumer Choice; Life Cycle Models and Saving)

1 Introduction

Payday loans are one of the most expensive forms of credit in the world. Borrowers typically pay non-annualized finance charges of 18% for loans lasting two weeks. These terms imply an annualized cost of payday loan liquidity of $1.18^{26} - 1 = 7295\%$. Truth-in-Lending regulations result in posted Annual Percentage Rates (APRs) for two-week-long loans of $18\% \cdot 26 = 468\%$. Since finance charges generally do not depend on loan length, month-long and week-long payday loans respectively carry annualized liquidity costs of $1.18^{12} - 1 = 629\%$ and $1.18^{52} - 1 = 546745\%$.

Despite these high interest rates, ten million distinct American households borrowed on payday loans in 2002 and the industry's growth rate exceeds 15% annually (Robinson and Wheeler 2003). Consumption models offer three main complementary explanations for such a phenomenon. First, consumers may have very high discount rates, particularly in the short term (Phelps and Pollak 1968, Laibson 1997b, Frederick, Loewenstein and O'Donoghue 2002). Second, consumers may experience shocks that cause large, unanticipated variation in the marginal utility of consumption (Deaton 1991, Carroll 1992). Third, consumers may have overoptimistically rosy forecasts of the future, in regard to either their own time preferences (Akerlof 1991, O'Donoghue and Rabin 1999a) or the probability of favorable shocks (Brunnermeier and Parker 2005, Browning and Tobacman 2007).¹

This paper evaluates the contributions of these candidate explanations by nesting them in a single structural model, and estimating that model using detailed measurements from a unique administrative panel dataset of 100,000 borrowers and 800,000 payday loans. Provided by a financial

¹None of these three explanations can suffice if consumers have lower-cost credit alternatives. However, payday borrowers seem to have limited outside options. A representative survey of one thousand payday loan customers found that 56.5% had bank cards, but 61% of those with cards did not use them in the past year because they would have exceeded their credit limits (Elliehausen and Lawrence 2001, Table 5-16). This survey also reported that 15.4% of payday borrowers had declared bankruptcy in the previous 5 years (Table 5-18). Also, IoData (2002) report the results of surveys of 2600 payday borrowers in six states. In that sample 55% had credit cards, but only 34% "almost always" or "sometimes" paid off their balances at the end of the month. Two-thirds of respondents were deterred from applying for credit at some point during the previous five years by the expectation they would be rejected.

services firm that offers payday loans, the data include complete histories of loan initiations, repayments, and defaults from the time this firm began offering payday loans in January 2000 through August 2004. Rich demographic information is also available. Means of borrowing probabilities, loan sizes, and default probabilities, conditional on the amount of time elapsed since each customer's first loan, identify the structural model's parameters.

Results of estimation using the Method of Simulated Moments indicate that uncertainty and time-consistent discounting go partway toward accounting for the observed phenomena. In this benchmark case we estimate two-week discount rates of 21% and default costs of about \$300. At the estimated parameter values borrowing occurs, and the average borrowing rate over a year roughly matches the empirical observations. However, average empirical loan sizes exceed simulated loan sizes by more than 10%: the simulated consumers smooth consumption as they ramp up their loan sizes too gradually in anticipation of default. In addition, simulated default rates at the estimated parameter values are 50% higher than the empirical rates. Simulated sophisticated quasi-hyperbolic discounters have higher short-term discount rates and hence borrow more aggressively, providing a better match of loan magnitude data. However, sophisticated quasi-hyperbolic discounters also exhibit preferences for commitment. In this setting default acts as a form of commitment, as defaulters are excluded from future access to payday loans. Real consumers' *failure* to take rapid advantage of this form of commitment results in evidence against sophisticated quasi-hyperbolic discounting. At the estimated parameter values, sophisticated quasi-hyperbolics have default rates twice as high as the empirical values. The last model we study, naive quasi-hyperbolic discounting, helps to account simultaneously for initial borrowing, moderate default rates, and delayed defaults. Naive quasi-hyperbolics incorrectly predict that future selves will absorb the default costs, depressing overall default rates toward the empirical values. In some specifications we cannot reject the hypothesis of perfect naivete, in which quasi-hyperbolic consumers believe that future selves will exactly implement the current self's preferences.

This paper complements a rapidly growing literature on payday loans. Distinguished contributions by Caskey (1991, 1994, 2001, 2005) drew attention to the topic and studied "fringe banking" more generally. Flannery and Samolyk (2005) used store-level data from two payday lenders to study profitability of the payday lending industry; and Skiba and Tobacman (2007b) complement that profitability work using asset-return and micro-level data. Surveys of payday borrowers have been conducted by Eliehausen and Lawrence (2001) and IoData (2002). State Departments of Finance

have analyzed the industry as well. Stegman and Faris (2003) and Stegman (2007) review regulatory considerations and policy proposals. Morse (2006), Melzer (2007), and Skiba and Tobacman (2007a) estimate causal effects of access to payday loans. Consumer advocates and the industry lobby have also produced numerous (separate) studies.

Section 2 introduces the data and presents the key empirical facts. We devote Section 3 to the model and its predictions. Section 4 reviews the Method of Simulated Moments estimation procedure. We report and interpret the estimation results in Section 5. Section 6 discusses possible extensions, and in Section 7 we conclude.

2 Borrowing, Repayment, and Default Data

2.1 Source and Features

We conduct our analysis using a new, proprietary dataset from a provider of financial services that offers payday loans (Skiba and Tobacman 2007a). Employment, banking, and demographic information are collected during the loan application process,² and detailed administrative records about payday borrowing, repayment, and default at this firm have also been provided. The dataset also includes information on how frequently borrowers are paid. This is a crucial variable, since the dynamic programming model we describe in the next section must specify the timing of decision-making.

In typical payday borrowing scenarios customers immediately receive their loans in cash, in exchange for personal checks for the loan principal plus finance charges. The personal checks are dated on the borrowers' next paydays, which are also the due dates of the loans. If a loan isn't first renewed or repaid in person, the store can deposit the customer's collateralizing check.

We study the population of adults who borrowed at least once from a Texas outlet of this company between September 2000 and August 2004; who did not change pay frequency; and who did not get "second chances" to borrow again after a default.³ We replace with missing the top and

²Defaults—and positive loan outcomes—are reported to a subprime credit rating agency called Teletrack. Scores computed by Teletrack are used by most payday lenders, so default at this company decreases future payday borrowing opportunities elsewhere. See below for additional discussion.

³About 1% of people who default on loans are allowed to borrow again after their defaults. Lacking information about demand- and supply-side selection of these individuals, we omit them and assume in the model that default leads to exclusion from future payday

bottom 0.1 percent of checking account balance and income (this is done for everyone, not by pay frequency). In addition, to address the occasional occurrence of multiple observations within a time period in the data, within a pay period we take the maximum value of the loan amount, the default indicator, the checking account balance, and income.

These procedures result in a sample encompassing 776,667 loans for 101,377 borrowers. We focus further on the 51,636 individuals who are paid bi-weekly and their 335,376 loans. Table 1 summarizes the characteristics of borrowers in the dataset, which resemble those in the IoData (2002), Wiles and Immergluck (1999), and Elliehausen and Lawrence (2001) survey samples. A majority of borrowers are female, a large share are Black or Hispanic, and the typical borrower is in her mid-thirties. Renting is twice as common in this population as in the general US population. Typical checking account balances at the time of loan application are very low: the mean balance is \$283,⁴ and the median is about \$100.⁵ The average income of individual borrowers in our dataset is about \$1700 per month. Since payday loan approval requires a steady income source, it is not surprising that these data imply annual household incomes generally above poverty levels. The borrowers who are paid biweekly, who represent half our population, are evidently similar to the full population. Information about borrowing patterns is presented in Table 2.

Moments

Using the data described in the previous subsection, we compute moments used to identify parameters of the discount function. The moments we use are conditional expectations, where we condition on the amount of time elapsed since an individual's first loan.

The moments we consider are shown, with two-standard-error bands, in green in Figures 1-3. In each graph the horizontal axis represents the number of pay periods since an individual's first payday loan at this company. The vertical axis plots, in turn, the fraction of the population borrowing, average debt at time t , and the average default rate conditional on borrowing. We

credit.

⁴All dollar amounts reported in the paper have been deflated to January, 2002, dollars using the CPI-U.

⁵We believe these checking account balances reflect little measurement error, since borrowers show their current bank statements to the lender at the time of their loan applications. However, borrowers may have strategic motives to keep these balances low and/or to maintain assets in other accounts.

define “default” as occurring when a borrower bounces a check collateralizing his or her payday loan, since that is when he or she immediately faces the main costs of default. These costs include bounced check fees imposed by the payday lender and the borrower’s bank, the annoyance caused by the lender’s collection efforts, and stigma or shame (Gross and Souleles 2002).

Summarizing key features of the moments, first, many borrowers turn to payday loans on a regular basis for liquidity. On average customers borrowed 5.5 times per year that we observe them; 25 percent borrowed 10 or more times and 10 percent borrowed 20 or more times. Renewing loans is common, rather than paying the loan in full, resulting in significant durations of indebtedness. Almost half of the loans in our sample were renewed. Conditional on renewing once, loans were renewed three times on average, yielding a two-month period of indebtedness for consumers who receive biweekly paychecks. Second, loans average roughly \$300. This amount rises slightly as time elapses from a borrower’s first loan. Third, default rates tend to be highest on early loans. First-loan default rates approach 12%, while someone borrowing one year after her first loan is half as likely to default.

3 Dynamic Programming Model

In order to study the phenomena described in the previous section, we adopt a dynamic programming model of consumption, saving, borrowing, and default. In the tradition of Carroll (1992, 1997), Deaton (1991), and Zeldes (1989) the model includes income uncertainty and liquidity constraints, and we also incorporate the option to borrow on institutionally realistic payday loans. Research by economists and psychologists has suggested two extensions to this benchmark model, “quasi-hyperbolic discounting” and “overoptimism” (Akerlof 1991, Laibson 1997a, O’Donoghue and Rabin 1999a, O’Donoghue and Rabin 1999b, O’Donoghue and Rabin 2001, Angeletos, Laibson, Repetto, Tobacman and Weinberg 2001). We permit these generalizations of preferences and beliefs in a way that nests the benchmark alternatives.

3.1 Timing and Demographics

We solve the model separately for consumers who receive their paychecks at weekly, biweekly, semimonthly, and monthly intervals. Denote these four groups of consumers as having pay-cycle duration $d \in \{W, BW, SM, M\}$. When we solve the model for a consumer with pay cycle duration d , we

interpret a period in the model as having length corresponding to d . For each d , the model concerns the behavior of the *typical* individual with that pay-cycle duration. Because the periods in the model are very short and we estimate the model’s parameters using one year of data on each individual, we ignore lifecycle considerations and solve an infinite-horizon version of the model.

3.2 The State Space

Consumers begin period t by receiving after-tax real income Y_{it} . We assume that $y_{it} = \log(Y_{it})$ equals a constant plus an AR(1) component and white noise. Specifically,

$$\begin{aligned} y_{it} &= \mu^y + u_{it}^y + \eta_{it}^y, \text{ where} & (1) \\ u_{it}^y &= \alpha^y u_{it-1}^y + \varepsilon_{it}^y, \\ \eta_{it}^y &\sim N(0, \sigma_{\eta^y}^2), \quad \varepsilon_{it}^y \sim N(0, \sigma_{\varepsilon^y}^2) \end{aligned}$$

The autoregressive component is approximated with a five state Markov process; we call the persistent Markov state v_t^y .⁶

Cash-on-hand at time t is $X_t \geq 0$, which includes Y_t . We assume that savings in the liquid asset earn a gross after-tax real return R^X derived from the average yield of Moody’s AAA municipal bond yields and the CPI-U from 2000 to 2004 (Gourinchas and Parker 2002, Laibson, Repetto and Tobacman 2007). The annual average yield is 3%. We take this interest rate to the $1/12$, $1/24$, $1/26$, and $1/52$ powers when we insert it into the simulation model for the M , SM , BW , and W pay-cycle durations, respectively (so the value of R^X depends on d). Let I_t^X equal net investment in X at time t , so $X_{t+1} = (X_t + I_t^X) R^X + Y_{t+1}$.

We denote outstanding payday loan debt at the start of period t by $D_t \geq 0$. In accord with the institutional rules at the company that provided our data, consumers can’t take out loans larger than $D^{\max} = \min\{\$500, \lambda^d \exp(v_t^y)\}$, where the credit limit λ^d depends on the pay-cycle duration. Specifically, $\lambda^M = 0.295$, $\lambda^{SM} = 0.59$, $\lambda^{BW} = 0.639$, and $\lambda^W = 0.639$. Let I_t^D equal net investment in D at time t , so $D_{t+1} = (D_t - I_t^D) R^D$. Though we focus on the case where consumers can renew loans indefinitely, we also consider the possibility that the number of renewals is capped. If ψ_t is the number of

⁶Because of the curse of dimensionality, we estimate the income process separately from the rest of the model. See the Appendix, where we also describe our procedures in more detail.

times an outstanding loan has been renewed, then $\psi_t \in \{0, 1, \dots, \psi^{\max}\}$.⁷ In addition, consumers are permitted to default on their loans. We assume that default has one benefit and two costs: Outstanding payday loans are erased, but there is an immediate pecuniary cost of default k ,⁸ and consumers who default are excluded from the opportunity to borrow on payday loans in the future.⁹ In order to implement the opportunity to default, we add the state variable ϕ_t as an indicator for “ever defaulted;” $\phi_t = 1 \Rightarrow D_\tau = 0 \quad \forall \tau \geq t$. When consumers borrow on payday loans, they face a gross per-period nominal interest rate of 1.18, regardless of d . We adjust this for inflation (which does depend on d) to obtain the interest rate on payday loans R^D .

Summarizing the description of the model so far, the state variables at time t include the liquid-asset level, the amount of outstanding debt, the Markov state of the income process, an indicator for “has already defaulted,” and a counter for the number of times a loan has been renewed. We write the vector of state variables as $\Lambda_t = \{X_t, D_t, v_t^y, \phi_t, \psi_t\}$.

3.3 The Choice Space

The choice variables at time t are I_t^D , I_t^X , and whether or not to default ($n_t \in \{0, 1\}$). Consumption is calculated as a residual:

$$C_t = -I_t^X - I_t^D (1 - n_t) - n_t k$$

The allowable values of the choice variables depend on the state as follows. First, consumers who have already defaulted cannot borrow again or default again: $\phi_t = 1 \Rightarrow I_t^D = 0$, $n_t = 0$, and $I_t^X \geq -X_t$ (i.e., $C_t \leq X_t$). Second, consumers who are not in debt and have not defaulted can consume available cash-on-hand and borrow up to the credit limit: for them,

⁷After any period in which $D_t = 0$, ψ is reset to 0. We interpret “renewal” to mean that some debt is carried over from the previous period.

⁸Consumers default on payday loans by allowing the checks they’ve written to collateralize the loans to bounce. Assuming pecuniary costs of default may permit a priori estimates of their magnitude, since the costs of bounced checks are known. However, k is also meant to capture psychic costs of default and the costs of being harassed to repay, which can’t be measured directly.

In addition, note that if the pecuniary cost of default is greater than the cost of debt service, then impatient consumers who receive extremely bad shocks will generally service their debt and postpone paying the costs of default.

⁹This approximates the impact of the actual credit-scoring procedures used in the payday-loan industry (Skiba and Tobacman 2007a). Alternatively, one might believe that default would restrict access in the future to even higher-priced credit than payday loans. See Section 6 on Extensions for further discussion.

$I_t^X \geq -X_t$ and $I_t^D \geq -D^{\max}$ (i.e., $C_t \leq X_t + D^{\max}$). Default is assumed to be impossible unless a consumer is borrowing.

Third, a consumer who is borrowing can choose to default. We assume $I_t^D = D_t$ if default is chosen, and then $I_t^X \geq -(X_t - k)$ (i.e., $C_t \leq X_t - k$). If default is not chosen and the renewal limit has been reached (i.e., $\psi_t = \psi^{\max}$), then $I_t^D = D_t$ and $I_t^X \geq -(X_t - D_t)$ (ie, $C_t \leq X_t - D_t$). If default is not chosen and the renewal limit has *not* been reached, then $I_t^D \in [D_t - D^{\max}, D_t]$ and $I_t^X \geq -(X_t - I_t^D)$ (i.e., $C_t \leq X_t - I_t^D$).

If the marginal utility of wealth at all t is positive, then the optimal I_t^X and I_t^D will imply that $\min[X_t - Y_t, D_t] = 0$. As a result, the choice variables $\{I_t^X, I_t^D, n_t\}$ can be replaced by $\{C_t, n_t\}$ without loss of generality. Consequently we focus only on choices of consumption and default.

3.4 Preferences

We assume that instantaneous utility is of the constant relative risk aversion (CRRA) form:

$$u(C) = \frac{C^{1-\rho}}{1-\rho}$$

Total utility is the discounted sum of these instantaneous utilities,

$$U_t(\{C_\tau\}_{\tau=t}^\infty) = u(C_t) + \beta \sum_{\tau=t+1}^\infty \delta^\tau u(C_\tau^E),$$

where we permit the discounting to be quasi-hyperbolic. When $\beta < 1$, from the time t perspective, the discount factor between periods t and $t + 1$ is lower than the discount rate between subsequent pairs of adjacent periods. Note that when $\beta = 1$ this reduces to exponential discounting.

We consider the possibility that consumers are “partially naive” about the degree to which their current time preferences will be respected by future selves. Future selves will apply the same quasi-hyperbolic discount function as the current self, but the current self believes that future selves will instead apply the discount function $\{1, \beta^E \delta, \beta^E \delta^2, \beta^E \delta^3, \dots\}$, $\beta \leq \beta^E \leq 1$, generating consumption realizations C_τ^E which are expected erroneously when $\beta \neq \beta^E$.

3.5 Equilibrium

We can write the consumers' optimization problem as follows:

$$\begin{aligned} \max_{C_t, n_t} \quad & u(c_t) + \beta \delta E_t [V_{t,t+1}(\Lambda_{t+1}) | (C_t, n_t)] \\ \text{s.t.} \quad & \text{Dynamic budget constraints,} \end{aligned} \tag{2}$$

where the time t state variables are $\Lambda_t = \{X_t, D_t, v_t^y, \phi_t, \psi_t\}$, and the value function is given by the following functional equation:

$$V_{t-1,t}(\Lambda_t) = u(C_t^*) + \delta E_t [V_{t,t+1}(\Lambda_{t+1}) | (C_t^*, n_t^*)].$$

We iterate the functional equation numerically until it converges in order to find the Markov Perfect Equilibrium decision rules.¹⁰ Hyperbolic discount functions can produce value functions that are discontinuous and policy functions that are discontinuous and non-monotonic but, as in our application, sufficient noise ensures smoothness (Harris and Laibson 2001).

A shortcoming of this model is its assumption of partial equilibrium. Though we find the question of how payday loan companies choose their repayment terms, interest rates, and other contract rules extremely interesting (DellaVigna and Malmendier 2004, Skiba and Tobacman 2007b), we focus here on the necessary first step of explaining the behavior of consumers in response to the observed supply-side choices.

3.6 Simulation

Uncertainty prevents direct calculation of the theory's predictions. As a result, once we have solved for the equilibrium policy rules, we generate $J_s = 20000$ paths of consumption and income shocks according to the stochastic shock processes described above, and we simulate the behavior of a population of consumers that behave according to the policy rules. From the resulting simulated panel dataset we calculate analogues to the empirical moments presented in Section 2. We do this over the space of parameter vectors $\theta \in \mathfrak{R}^l$, where θ potentially includes $\beta, \delta, \beta^E, \rho, k$, and the income shock parameters. Then, as described in the next section, we formally compare the moments of the empirical and simulated data in order to estimate θ .

¹⁰Uniqueness and convergence of value functions are not guaranteed. Due to concerns about the stability of the iteration, we don't use hill-climbing techniques to find solutions to the Euler Equation for each state in each period. Instead, we rely on efficient local grid searches and Matlab's built-in Nelder-Mead simplex-based optimization routine.

The computational problem entails solving for the equilibrium policy rules, creating the simulated data, and computing the descriptive moments from the simulated data. This requires about 5000 lines of Matlab code excluding comments and 4 minutes on a 3.2GHz Pentium 4. Because of the high dimensionality of the θ space, we have to perform these computations roughly 10,000 times in order to implement the baseline estimation strategy described in the next section. Some robustness checks are an order of magnitude more computationally intensive.

3.7 Predictions

The model described in the past section has nine free parameters:

$$(\delta, \beta, \beta^E, \rho, k, \mu^y, \alpha^y, \sigma_{\eta^y}^2, \sigma_{\varepsilon^y}^2).$$

We identify the income process parameters (the last four) separately (see Appendix A). Intuition for identification of other parameters follows.

First, note that we observe individuals borrowing at extremely high interest rates. The mere fact of borrowing can in principle be explained by high (exponential or hyperbolic) discount rates or by large, temporary negative shocks. We also observe people borrowing *repeatedly* on payday loans. The median number of loans taken in the year after an individual's first loan (including the first loan) is 3, and the mean is about 7. Ordinarily, a precautionary savings intuition would apply: anticipating states of the world that could make this repeated borrowing optimal, consumers would accumulate a buffer stock of precautionary wealth. The absence of such buffer stocks could, in some cases, identify non-exponential discounting. However, this intuition doesn't apply in this case because our budget constraint has a kink. Exponential discount rates between the (very low) return on liquid savings and the (very high) payday loan interest rate can predict a lack of precautionary savings when the shock is normal and borrowing on payday loans when the shock is low.¹¹

Since the fact of borrowing has little power to identify parameters on its own, we also consider data on defaults. Recall from the model and the institutional details described above that default entails two costs: first, there are immediate pecuniary and non-pecuniary costs, as the lender and the borrower's bank impose fees and the lender attempts to collect, and second, there are delayed costs in the form of exclusion from future access to

¹¹The exponential discount rates necessary to account for repeated borrowing at 18% per two weeks, however, are unusually large.

payday loans. Sophisticated hyperbolic consumers would choose to default immediately on payday loans, both to avoid the costs of repayment *and* to commit future selves to behave patiently and not borrow. Naive hyperbolic discounting potentially provides a solution, since it predicts procrastination on default. We estimate the degree to which this can account for the data below.

A remaining question is whether we can identify $\beta < 1$, if consumers are naive.¹² Consumption models remain well-behaved when $\beta^E > 1$, and delays in default rise as β^E grows above β . However, the maximum amount of overoptimism about β comes when $\beta^E = 1$, since then the consumer believes future selves will exactly implement the current self's preferences. In addition, as β^E rises above 1, the current self anticipates less and less borrowing in the future, which decreases the current costs of default (since the option value to borrow is believed to be low). Calibrationally, in order to have enough naivete from β^E to cause delays in default, but not so much naivete that the option value bought by repayment falls, we conjecture we must have $\beta < 1$. After these other adjustments, the immediate cost of default k adjusts to match average default rates.

4 Estimation Procedure

Our estimation strategy applies the method of simulated moments (MSM), as developed by McFadden (1989), Pakes and Pollard (1989), and Duffie and Singleton (1993) (for a review, see Stern 1997). MSM enables us to formally test the nested hypotheses of exponential discounting ($\beta = \beta^E = 1$), perfect sophistication ($\beta^E = \beta$), and perfect naivete ($\beta^E = 1$), and to perform specification tests on the model.

Denote the vector of individual-level empirical data for individual i by m_i^e and the vector of corresponding data for simulated individual j by $m_j^s(\cdot)$.¹³ From these data, we can formulate the MSM moment conditions in the following five ways.

¹²Naivete is consistent with behavior in two other credit markets. First, pawn lenders offer potential borrowers the opportunity to sell their items outright for the same amount as the principal of the possible loan. This implies the pawn loan is effectively a free, exclusive option to repurchase the good (although the repurchase price equals the loan principal plus an interest payment). In addition, the dollar-valued elasticity of credit card takeup is three times higher with respect to teaser rates than post-teaser rates (Ausubel 1999).

¹³Here and in all that follows, $m_j^s(\cdot)$ is the simulation approximation to the true theoretical value $m_j(\cdot)$.

The first way ignores heterogeneity in the population entirely: We construct $m^e = E(m_i^e)$ and $m^s(\theta) = E(m_j^s(\theta))$, and the moment conditions are $g^s(\theta) = m^s(\theta) - m^e = 0$. Thus we estimate the model once to find a single θ for the whole population. Uncertainty in the estimate of θ is then partially attributable to heterogeneity in the population.

The second approach “controls” for observed heterogeneity in the population, and again estimates a single θ . Specifically, let Z_i^e be a vector of discrete and continuous demographic characteristics including for example race, gender, homeownership status, age, and income. We could write $m_i^e = \alpha + Z_i^e \Gamma^e + \epsilon_i$; regress m_i^e on Z_i^e to obtain $(\hat{\alpha}, \hat{\Gamma}^e)$; choose as “typical” demographics Z^e the means of the continuous variables in Z and the medians of the discrete variables; and compute $\hat{m}^e = \hat{\alpha} + Z^e \hat{\Gamma}^e$. In this case, the moment conditions become $\hat{g}^s(\theta) = m^s(\theta) - \hat{m}^e = 0$, and the model is interpreted as characterizing the behavior of a typical individual in the population. This approach has the advantage that it uses more of the available information than the first method, while still requiring that the model be solved and estimated only once. The disadvantages are the restriction that demographic characteristics affect the moments linearly and homogeneously,¹⁴ and the possibility that the Z_i^e are correlated with θ_i and hence controlling for Z_i^e will bias down the amount of uncertainty attributed to θ .

Third, we could partition the population on some demographic characteristics and estimate the model separately for the resulting subpopulations. Formally, let Π be the set of all partitions on demographic characteristics of the population. Consider some $\pi \in \Pi$. For example, the elements $\omega \in \pi$ might be the set of female homeowners, the set of female renters, the set of male homeowners, and the set of male renters. For each ω we could construct $m_\omega^e = E(m_i^e \mid i \in \omega)$ and $m^s(\theta_\omega) = E(m_j^s(\theta_\omega))$ and adopt the moment conditions $g_\omega^s(\theta_\omega) = m^s(\theta_\omega) - m_\omega^e = 0$. This approach has the disadvantage that it requires the model to be solved and estimated separately for each ω , but each such estimation could be performed independently.

Hybrids between the second and third approaches are possible. Thus, fourth, we could control for some demographic variables and solve the model separately after partitioning the sample on other demographics. A final and final alternative would be to assume a functional form for the effect of Z on θ and estimate the parameters of that functional form. In full generality, the simulated data could be given by $m_j^s(\theta, \Gamma^s, Z_j^s)$; in practice we might

¹⁴Though this could be generalized: instead we could assume $m_i^e = g(Z_i^e) + \epsilon_i$ for an arbitrary function g .

assume $m_j^s = m_j^s(\theta + Z_j^s \Gamma^s)$ and focus on the Z 's in some partition π . Then for each $\omega \in \pi$ we could construct $m_\omega^s(\theta, \Gamma^s) = E\left(m_j^s(\theta + Z_j^s \Gamma^s) \mid j \in \omega\right)$, and the moment conditions become $g_\omega^s(\theta, \Gamma^s) = m_\omega^s(\theta, \Gamma^s) - m_\omega^e = 0$. In this case all of the g_ω^s depend on the same parameters, so these moment conditions must be stacked and used to *simultaneously* estimate (θ, Γ^s) .

Under the first formulation above, suppose θ_0 is the true parameter vector. (Under the third formulation above, we would begin by assuming $\theta_{\omega 0}$ is the true parameter vector.) Assume that m^e (or m_ω^e) has N elements, and its asymptotic variance-covariance matrix is Ω (or Ω_ω). Then m^s (or m_ω^s) correspondingly has N elements. Let W be an $N \times N$ weighting matrix. Let $q(\theta) \equiv g(\theta) \cdot W^{-1} \cdot g(\theta)'$ be a scalar-valued loss function, equal to the weighted sum of squared deviations of simulated moments from their corresponding empirical values.

Our procedure is to minimize the loss function $q(\theta)$ and define the MSM estimator as,

$$\hat{\theta} = \arg \min_{\theta} q(\theta). \quad (3)$$

Pakes and Pollard (1989) demonstrate that, under certain regularity conditions satisfied here,¹⁵ $\hat{\theta}$ is a consistent estimator of θ_0 , and $\hat{\theta}$ is asymptotically normally distributed. For $W = \Omega$, $\hat{\theta} \rightarrow N(\theta_0, \Sigma)$ asymptotically, where

$$\Sigma = \left[\left(\frac{\partial g_i(\hat{\theta})}{\partial \theta_j} \right)' \Omega^{-1} \left(\frac{\partial g_i(\hat{\theta})}{\partial \theta_j} \right) \right]^{-1} \quad (4)$$

Note that the derivatives in this expression are evaluated at $\hat{\theta}$. The intuition for this equation is most easily seen by analogy to the case of estimation of one parameter by one moment with the familiar method of moments. First, observe that the standard error of the parameter estimate is increasing in the standard error of the empirical moment. If the moment is imprecisely estimated, we can attach little certainty to the parameter estimate. Also, suppose in this simple one-parameter, one-moment case that $m^s(\theta)$ is very flat near the optimum. Then large changes in the parameter would have only a small effect on the simulated moment. Consequently, the location of the true minimum of the loss function will be relatively uncertain. Conversely, if $m^s(\theta)$ is steeply sloped very close to the optimum, so small changes in the parameter have a dramatic effect on the moment, one can be

¹⁵Principally, the moment functions must be continuous in the parameters at $\hat{\theta}$.

relatively confident about the parameter estimate: the standard error of $\hat{\theta}$ will be small.

Returning to the general case of MSM, the expression for Σ above likewise indicates that large derivatives of the moments as functions of the parameters result in small standard errors for the estimated parameters. In other words, if the moments are very sensitive to the parameters, the parameters are more likely to be precisely estimated. The Ω^{-1} term in the center captures the notion that redundant or imprecisely estimated empirical moments in general do not tightly constrain the MSM parameter estimates.

MSM also allows us to perform specification tests. If the model is correct, $q(\theta)$ is distributed $\chi^2(N - p)$.

We implement a two-stage variant of the first approach above, and apply the relevant simulation correction in the variance formula. This approach closely follows Gourinchas and Parker (2002) and Laibson et al. (2007).

5 Results

5.1 Parameter Estimates

To find estimates of θ , we adopt the income process parameters estimated in Appendix A, and assume those estimates are independent of the other parameters (though see the Extensions below). For a weighting matrix, we adopt the inverse of the diagonal of the VCV of the empirical moments. We focus just on the conditional expectations of the borrowing probability, the default probability conditional on borrowing, and the loan size conditional on borrowing; and we perform the estimation using the first approach in Section 4, where uncertainty in the parameter estimates incorporates heterogeneity in the population.

We find the results reported in Table 4. First, in Column 1, we consider the exponential discounting case, with the default cost fixed at $k = 200$ and the coefficient of relative risk aversion fixed at $\rho = 2$. We obtain a very precise estimate of the exponential discount factor, of 0.8161. Note that this is a two-week discount factor; it would correspond to an annual discount factor of $0.8161^{26} = 0.0051$. Evidently, such sharp exponential discounting implies almost no concern for future years. However, not surprisingly, very high short-term discount rates are required to predict any borrowing on payday loans. The quantitative findings change considerably in Columns 2 and 3, when we allow k and ρ to vary, but the qualitatively high exponential discount rates persist.

In Columns 4 and 5 we report the perfectly sophisticated quasi-hyperbolic estimates. Here we find estimates of β significantly below 1, and the goodness-of-fit measures improve somewhat over the exponential measures. However, we continue to find very low estimates of δ , perhaps because they are required in order to predict delayed default.

Columns 6 and 7 of the table present the perfectly naive quasi-hyperbolic estimates. In these cases, δ is estimated to be much closer to 1, and we find highly significant measures of naivete: $\hat{\beta} \approx 0.5$, with standard errors of about 0.002. The goodness-of-fit measures tell a mixed story: q falls in this case to lower values than in the exponential or sophisticated cases, but ξ (which incorporates the uncertainty in the income process estimates) remains high.

In Columns 2, 5, and 7, (the exponential, perfectly sophisticated, and perfectly naive cases, respectively), we estimate default costs of $k \approx 200$. Bounced check fees imposed by the lender and the borrower’s bank typically total about \$50, implying the non-pecuniary costs of stigma, sense of irresponsibility, and pressure to repay amount to roughly \$150. We find this figure reasonable, and note also that default costs are bounded above by the cost of repaying a loan, which equals the outstanding balance.

In Column 3, when the discount function is exponential and ρ is also free, we find a better fit for a lower value of ρ and a higher value of k . Intuitively, lower ρ helps to induce larger loan sizes, improving the fit. Lower ρ also decreases the value of the option to borrow, and hence requires larger immediate default costs to deter immediate default.

These results are illustrated in Figures 1-3, which plot the simulated moments, calculated at the parameter estimates, for the exponential, perfectly sophisticated, perfectly naive, and partially naive cases, along with their empirical analogues.

5.2 Discussion

These estimates provide suggestive evidence that the naive and sophisticated quasi-hyperbolic models perform better than the exponential model at explaining payday borrowing, repayment, and default. One source of additional, out of sample evidence that naivete might be more relevant than sophistication for payday borrowers is that naivete predicts more procrastination, and the fact that payday loan borrowers don’t have other available sources of credit (e.g., credit cards) might be because of procrastination in applying for them.

Two complementary perspectives may also aid interpretation. First,

we have described the cost of payday loans as interest, in the language of intertemporal choice. Alternatively, we could have described the finance charges as (approximately \$54) convenience charges. Instead of the dynamic programming model of intertemporal choice, we could conduct a careful calibration of the small fixed costs that consumers would have to pay to sustain alternatives. For example, to the extent that consumers could engage in precautionary saving in order to avoid needing to borrow on payday loans when bad shocks arrive, we could model the costs of undertaking that precautionary saving. Second, the quasi-hyperbolic model does not distinguish between self-control problems it induces and spousal control problems (Ashraf 2005). Issues of spousal control could be modelled explicitly, but the model studied here may remain a useful reduced-form representation.

Our estimates above are for a discrete-time discount function with two weeks between periods. Thus, estimates of $\beta < 1$ imply payday borrowers treat the present qualitatively differently from dates two weeks in the future (McClure, Laibson, Loewenstein and Cohen 2004, McClure, Ericson, Laibson, Loewenstein and Cohen 2006). We find annualized discount rates ranging from over 500% for the exponential and sophisticated cases to 150% for the naive case.

6 Extensions

The estimates in the previous section are reasonably suggestive, particularly combined with the intuition for identification, but a number of extensions would help to clarify the reach and the limitations of these results.

6.1 Identifying Moments

First, we could modify the set of moments used in the MSM procedure to estimate the parameters.

6.1.1 Income Moments

To mitigate the curse of dimensionality we estimated the parameters of the income process separately from the preference parameters of primary interest (c.f. Appendix 1). Instead, moments characterizing income could potentially be combined with moments on borrowing, repayment, and default to simultaneously identify income process parameters and preference parameters.

6.1.2 Other Moments

More generally, the procedure for optimally choosing the moments for GMM or MSM estimation is not obvious.¹⁶ Section 2 presented a large collection of possible moments. Our selection here used the three most central pieces of information about loans – whether they occur, how much they are for, and whether they are repaid – and attempts to retain transparency about how the moments identify the parameters.¹⁷

6.2 Consumption Shocks

Uncertainty about income is incorporated into the model, but consumers may face other sources of risk. Specifically, stochastic shocks to consumption needs like car repairs, funeral expenses, or health-care costs (Hubbard, Skinner and Zeldes 1995, Palumbo 1999) could motivate people to take out short-term loans by temporarily raising the marginal utility of consumption. Increasing the variance of income shocks approximates the effect of including consumption shocks, and causes no qualitative differences in our results.¹⁸

6.3 Overoptimism about Shocks

A growing literature studies consumer overoptimism. We could formally consider this possibility by assuming that beliefs about the mean of the log-income process may be incorrect, ie, that consumers believe Equation 1 has some μ^E which differs from μ .¹⁹

¹⁶Gallant and Tauchen (1996) propose a criterion based on a score computed from an auxiliary model.

¹⁷Covariance moments seem interesting here, but for our (unbalanced) panel empirical covariances are defined for only few observations. Small-sample biases described by Altonji and Segal (1996) deter our use of these moments.

¹⁸Consumption shocks could be studied in more detail. First, shocks could be included in our model as an additional stochastic process. For example, the log consumption shock could equal a constant plus an AR(1) plus white noise, with the AR(1) approximated with a Markov process, like our treatment of the income process. The consumption shock parameters could be estimated simultaneously with the other parameters of the model, or separately, for observably-similar individuals, using data from the Consumer Expenditure Survey. Rather than using consumption shocks in dollar terms, it would alternatively be possible to model them as taste shocks that cause proportional shifts in utility.

¹⁹We are agnostic about why beliefs may be incorrect in this way. One possible explanation is that overoptimism adaptively offsets risk aversion. Another possibility is that consumers have a preference for believing the future will be rosy, and willingly sacrifice actual outcomes in order to act consistently with those beliefs (Brunnermeier and Parker 2005). We consider overoptimism about income to be another hypothesis potentially testable in this context.

6.4 Default Costs

Default on a payday loan entails immediate costs of bounced check fees and pressure to repay, and delayed costs of exclusion from future access to credit. Rather than estimating a fixed pecuniary cost for the immediate costs as in the benchmark model, we could consider several alternative assumptions. First, we could assume default incurs an immediate *utility* cost of k_u , in addition to loss of the option to borrow. This may be more natural for capturing the psychological costs of default, but would be less easily interpretable. Second, defaulting could cause only some probability of exclusion from credit in the future. Since the estimated immediate costs of default must adjust to match observed default rates, this change would have the simple effect of increasing \hat{k} .

A third possibility is that the costs of default depend on the length of the borrower's history with the company. Bounced check fees and the company's collections procedures do not depend on histories, but psychological costs of default may. Gross and Souleles (2002) argue that stigma associated with bankruptcy fell during the 1990s. Our results would require payday loan default costs to fall over the course of a single year.

Underlying our estimates of default costs is a counterintuitive prediction. When payday borrowers delay default, they eventually default when their marginal utility of consumption is relatively low. Thus defaults might be positively correlated with contemporaneous income, and this could be tested out-of-sample.

6.5 Heterogeneity

Alternative hypotheses might emphasize possible heterogeneity in the population. Specifically, one could suppose there are three groups. People in one group have unobservable characteristics disposing them to become frequent payday loan borrowers (they could have high hedonic tastes for payday borrowing). People in the second group are responsible, standard consumers. They experience bad shocks which are genuinely temporary; they each borrow once on payday loans; and then they repay and disappear from our sample. People in the third group are poor financial risks. They discover payday loans; at some point they achieve credit scores that qualify them for loans, but barely; and then they repeat a pattern of irresponsibility by quickly defaulting and disappearing.

This story requires heterogeneity in two dimensions – hedonic taste for payday loans, and unobservable creditworthiness. Heterogeneity could be

introduced along other dimensions as well. It is unclear which dimensions of heterogeneity to prioritize. We think it would nonetheless be fruitful to attempt to estimate heterogeneity, improving the model’s fit and further illustrating the financial decision-making of the payday borrowing population. However, we doubt this would affect the main results reported above, since typical payday borrowers default after paying interest on 6.5 payday loans. These borrowers drive the parameter estimates, so a large share of a heterogeneous population would be found to pay large costs to delay default, exhibiting the signature behavior of naivete.

6.6 Variation in Pay Frequency

Our benchmark model pertains to consumers who are paid biweekly. This is the largest group in our sample, but we could also estimate the model using data on consumers paid weekly, semimonthly, or monthly. Empirical moments are similar across all these groups, so we believe estimates of discounting parameters would differ little across them. However, the *lack* of difference in parameter estimates would be surprising because interest rates on payday loans do not depend on the loan duration. Thus the annualized cost of liquidity ranges from $1.18^{12} - 1 = 629\%$ for monthly payday loans to $1.18^{52} - 1 = 546,750\%$ for weekly payday loans.

6.7 Additional Credit Instruments

We assume that consumers don’t have access to credit instruments other than payday loans. A search model of loan choice (Hortacsu and Syverson 2004) would be necessary to study payday loan adoption in the context of the consumer portfolio (Musto and Souleles 2006). However, survey evidence indicates that most payday borrowers do not have available liquidity on credit cards (supra note 3). In addition, since few loan products in the US carry interest rates higher than payday loans, introducing common alternatives would make it more difficult to account for payday loan use.

6.8 Policy Implications

Structural estimation has the advantage of facilitating out-of-sample predictions, particularly about policy counterfactuals. Consumer advocates and policymakers support a variety of restrictions on payday lending. Most prominently, payday loans are already prohibited in 12 states (Consumer Federation of America 2006); the Nelson-Talent Amendment restricts APRs extended to military members (or their dependents) to no more than 36%

(SA 4331, Amendment to S2766/HR5221, the 2007 Defense Authorization Bill); and the FDIC requires that no individual receive payday loans covering more than 90 days out of every year (FDIC 2005). The payday lending industry, by contrast, favors removal of interest rate and finance charge caps that currently exist in many states. Our estimated model can be used to simulate behavior and firm profitability under these policy alternatives and measure welfare implications.²⁰

7 Conclusion

This paper has studied payday borrowing, repayment, and default behavior, and explanatory models of uncertainty and discounting. We have estimated a structural dynamic programming model with income uncertainty, institutionally realistic payday loans, and the option to default. Though the mere fact people borrow at payday loans' high interest rates does not identify the discount function, delays in default on payday loans push our estimates toward a model of naive quasi-hyperbolic discounting. These results allow simulation analysis of policy alternatives, and contribute to the literature on high-frequency consumption behavior.

²⁰Under time-inconsistent preferences, three main possibilities have been studied as welfare criteria. Laibson (1997a) considers the Pareto criterion. Laibson, Repetto and Tobacman (1998) adopt the perspective of the $t = 0$ self. O'Donoghue and Rabin (2001) argue that the right approach is to treat each temporal self identically and discount their consumption exponentially.

A Income Appendix

A.1 Data and Procedures

Though a huge literature has studied how income fluctuates at frequencies of a year or more, little is known about patterns of *high-frequency* income fluctuations. A full analysis of high-frequency income processes is beyond the scope of the current paper. Here we use income data acquired from the payday lender to estimate a standard income process that has both persistent and transitory shocks, but where the period equals the pay cycle duration.²¹

To perform the estimation, we restrict the sample of individuals in the same way we restrict to calculate the non-income moments in Section 2, dropping people who never borrowed from a Texas outlet of the company, people who were allowed to borrow after defaulting, and people whose pay frequency changed. For the individuals who remain, we estimate separate income processes for those paid weekly, bi-weekly, semi-monthly, and monthly.

In all cases, we assume that empirical $\log(\text{income})$ equals an individual fixed effect, plus an AR(1), plus white noise.²² The parameters of the theoretical income process (the mean, the autocorrelation coefficient, the variance of the persistent shock, and the variance of the transitory shock) are estimated by GMM, where as moments we use the mean of the fixed effects and the first year of autocovariances (e.g., we use the first 12 autocovariances for people paid monthly, and the first 24 autocovariances for people paid semi-monthly). This procedure assumes that simulated consumers have the average income fixed effect and interprets empirical heterogeneity in individual mean income as uncertainty in the theoretical mean.

We compute the autocovariances observation-by-observation. Since this is a panel dataset with a great deal of missing data, we are unable to restrict to a subset of individuals whose observations constitute a balanced panel. This could introduce bias, since income is more likely to be observed when individuals are seeking payday loans, and people seeking loans may have recently received bad shocks.²³ We attempt to partially address this concern, and test the robustness of our estimates, by considering three different sets

²¹In one other difference from many standard income process estimates, data limitations force us to estimate individual, not household, income.

²²Our data refer to take-home (aftertax) pay.

²³However, payday loan applications are only approved for individuals with steady income sources: applications require submission of a recent pay stub.

of observations. First, we include all income observations.²⁴ Second, we examine only the income observed on dates people applied for payday loans. Third, we examine only the income observed on dates people received payday loans. (Roughly ten percent of payday loan applications are rejected.) In each case, we bootstrap a full empirical variance matrix of mean income and of the first year’s worth of income autocovariances.

A.2 Estimates and Discussion

Our estimates are reported in Table 3. For borrowers paid biweekly, we find an autocorrelation coefficient of 0.194, a variance of transitory shocks of 0.073, a variance of persistent shocks of 0.109, and a (log) mean of 6.634. These numbers are estimated with very high precision. Restricting the sample of observations we use, still focusing on borrowers paid biweekly, implies estimates of smaller but more persistent shocks. For borrowers paid at other frequencies, the picture is less clear. In some cases the standard errors become anomalously large or anomalously small, and in other cases we find that all the variance in income is attributed by the GMM estimates to either the persistent shocks or the transitory shocks. These results may arise because so much data is missing.

Compared to estimates of annual income processes, we generally find shocks to be larger (standard deviations generally exceed 25%) but less persistent. Our findings could be cross-validated by simulating annual income processes from the estimated (weekly, bi-weekly, semi-monthly, and monthly) processes, and comparing to estimates from standard, lower-frequency sources with less missing data. Propensity score matching could also be used with data from standard sources to estimate annual income processes for individuals similar to our sample of payday borrowers. However, the approach we’ve adopted directly provides sample-relevant estimates for use in the paper’s high-frequency dynamic programming model.

²⁴In addition to the income data the payday lender collects directly, at the time of loan application, the lender receives some additional data in regular updates from its credit scorer, Teletrack. About 1/3 of all our income observations are from these updates.

References

- Akerlof, George**, “Procrastination and Obedience,” *American Economic Review*, May 1991, *81* (2), 1–19.
- Altonji, Joseph and Lewis Segal**, “Small-Sample Bias in GMM Estimation of Covariance Structures,” *Journal of Business and Economic Statistics*, 1996, *14* (3), 353–66.
- Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg**, “The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation,” *Journal of Economic Perspectives*, Summer 2001, *15* (3), 47–68.
- Ashraf, Nava**, “Spousal Control and Intra-Household Decision Making: An Experimental Study in the Philippines,” 2005. HBS mimeo.
- Ausubel, Lawrence**, “Adverse Selection in the Credit Card Market,” 1999. University of Maryland mimeo.
- Browning, Martin and Jeremy Tobacman**, “Discounting and Optimism Equivalences,” 2007. University of Oxford mimeo.
- Brunnermeier, Markus and Jonathan Parker**, “Optimal Expectations,” *American Economic Review*, September 2005, *94* (5), 1092–1118.
- Carroll, Christopher**, “The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence,” *Brookings Papers on Economic Activity*, 1992, *1992* (2), 61–156.
- , “Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis,” *Quarterly Journal of Economics*, February 1997, *112* (1), 1–55.
- Caskey, John P.**, “Pawnbroking in America: The Economics of a Forgotten Credit Market,” *Journal of Money, Credit, and Banking*, February 1991, *23* (1), 85–99.
- , *Fringe Banking: Check-Cashing Outlets, Pawnshops, and the Poor*, New York: Russell Sage Foundation, 1994.
- , “Payday Lending,” *Association for Financial Counseling and Planning Education*, 2001, *12* (2).

- , “Fringe Banking and the Rise of Payday Lending,” in Patrick and Howard Rosenthal, eds., *Credit Markets for the Poor*, Russell Sage Foundation, 2005.
- Consumer Federation of America**, “Legal Status of Payday Lending by State,” 2006.
- Deaton, Angus**, “Saving and Liquidity Constraints,” *Econometrica*, 1991, 59 (5), 1221–1248.
- DellaVigna, Stefano and Ulrike Malmendier**, “Contract Design and Self-Control: Theory and Evidence,” *Quarterly Journal of Economics*, May 2004, 119 (2), 353–402.
- Duffie, Darrell and Ken Singleton**, “Simulated Moments Estimation of Markov Models of Asset Prices,” *Econometrica*, 1993, 61 (4), 929–952.
- Elliehausen, Gregory and Edward C. Lawrence**, *Payday Advance Credit in America: An Analysis of Customer Demand*, Credit Research Center, Georgetown University, 2001.
- FDIC**, “Guidelines for Payday Lending,” 2005. Available at <http://www.fdic.gov/regulations/safety/payday/>.
- Flannery, Mark and Katherine Samolyk**, “Payday Lending: Do the Costs Justify the Price?,” *Federal Reserve Bank of Chicago Proceedings*, April 2005.
- Frederick, Shane, George Loewenstein, and Ted O’Donoghue**, “Time Discounting and Time Preference: A Critical Review,” *Journal of Economic Literature*, 2002, 40 (2), 351–401.
- Gallant, Ronald and George Tauchen**, “Which Moments to Match?,” *Econometric Theory*, October 1996, 12 (4), 657–81.
- Gourinchas, Pierre-Olivier and Jonathan Parker**, “Consumption over the Life Cycle,” *Econometrica*, January 2002, 70 (1), 47–89.
- Gross, David and Nicholas Souleles**, “An Empirical Analysis of Personal Bankruptcy and Delinquency,” *Review of Financial Studies*, Spring 2002, 15 (1), 319–347.
- Harris, Christopher and David Laibson**, “Dynamic Choices of Hyperbolic Consumers,” *Econometrica*, July 2001, 69 (4), 935–957.

- Hortacsu, Ali and Chad Syverson**, “Product Differentiation, Search Costs, and Competition in the Mutual Fund Industry: A Case Study of the SP500 Index Funds,” *Quarterly Journal of Economics*, 2004, 119 (2), 403–456.
- Hubbard, Glenn, Jonathan Skinner, and Stephen Zeldes**, “Precautionary Saving and Social Insurance,” *Journal of Political Economy*, 1995, 103 (2), 360–399.
- IoData**, “Payday Advance Customer Research: Cumulative State Research Report,” September 2002.
- Laibson, David**, “Golden Eggs and Hyperbolic Discounting,” *Quarterly Journal of Economics*, 1997, 62 (2), 443–478.
- , “Hyperbolic Discount Functions and Time Preference Heterogeneity,” 1997. Harvard mimeo.
- , **Andrea Repetto, and Jeremy Tobacman**, “Self-Control and Saving for Retirement,” *Brookings Papers on Economic Activity*, 1998, (1), 91–196.
- , —, and —, “Estimating Discount Functions with Consumption Choices over the Lifecycle,” Working Paper 13314, National Bureau of Economic Research August 2007.
- McClure, Samuel, Keith Ericson, David Laibson, George Loewenstein, and Jonathan Cohen**, “Time Discounting for Primary Rewards,” 2006.
- McClure, Samuel M., David I. Laibson, George Loewenstein, and Jonathan D. Cohen**, “Separate Neural Systems Value Immediate and Delayed Monetary Rewards,” *Science*, 2004, 306 (5695), 503–507.
- McFadden, Daniel**, “A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration,” *Econometrica*, 1989, 57 (5), 995–1026.
- Melzer, Brian**, “The Real Costs of Credit Access: Evidence from the Payday Lending Market,” 2007. Available at <http://www.chicagogsb.edu/research/workshops/finance/docs/Melzer-CreditAccess.pdf>.

- Morse, Adair**, “Payday Lenders: Heroes or Villains?,” 2006. University of Michigan mimeo.
- Musto, David K. and Nicholas S. Souleles**, “A Portfolio View of Consumer Credit,” *Journal of Monetary Economics*, January 2006, *53* (1), 59–84.
- O’Donoghue, Ted and Matthew Rabin**, “Doing It Now or Later,” *American Economic Review*, March 1999, *89* (1), 103–124.
- and —, “Incentives for Procrastinators,” *Quarterly Journal of Economics*, August 1999, *114* (3), 769–816.
- and —, “Choice and Procrastination,” *Quarterly Journal of Economics*, February 2001, *116* (1), 121–160.
- Pakes, Ariel and David Pollard**, “Simulation and the Asymptotics of Optimization Estimators,” *Econometrica*, 1989, *57* (5), 1027–1057.
- Palumbo, Michael**, “Uncertain Medical Expenses and Precautionary Saving Near the End of the Life Cycle,” *Review of Economic Studies*, April 1999, *66* (2), 395–421.
- Phelps, Edmund and Robert Pollak**, “On Second-Best National Saving and Game-Equilibrium Growth,” *Review of Economic Studies*, 1968, *35*, 185–199.
- Robinson, Jerry and John Wheeler**, “Update on the Payday Loan Industry: Observations on Recent Industry Developments,” Technical Report, Stephens, Inc. September 2003.
- Skiba, Paige Marta and Jeremy Tobacman**, “Do Payday Loans Cause Bankruptcy?,” 2007. University of Oxford mimeo.
- and —, “The Profitability of Payday Lending,” 2007. University of Oxford mimeo.
- Stegman, Michael**, “Payday Lending,” *Journal of Economic Perspectives*, Winter 2007, *21* (1), 169–190.
- Stegman, Michael A. and Robert Faris**, “Payday Lending: A Business Model That Encourages Chronic Borrowing,” *Economic Development Quarterly*, February 2003, *17* (1), 8–32.

Stern, Steven, “Simulation-Based Estimation,” *Journal of Economic Literature*, December 1997, 35 (4), 2006–2039.

Wiles, Marti and Daniel Immergluck, “Short Term Lending,” Illinois Department of Financial Institutions 1999.

Zeldes, Stephen, “Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence,” *Quarterly Journal of Economics*, 1989, 104 (2), 275–298.

Figure 1: Fraction Borrowing




Figure 1: This figure plots the fraction of the population borrowing on a payday loan, at each period in the one year after a borrower's first loan. Empirical values and their two-standard-error bands are shown in green, and the other lines are simulation model predictions at estimated parameter values.

Figure 2: Loan Amount Conditional on Borrowing

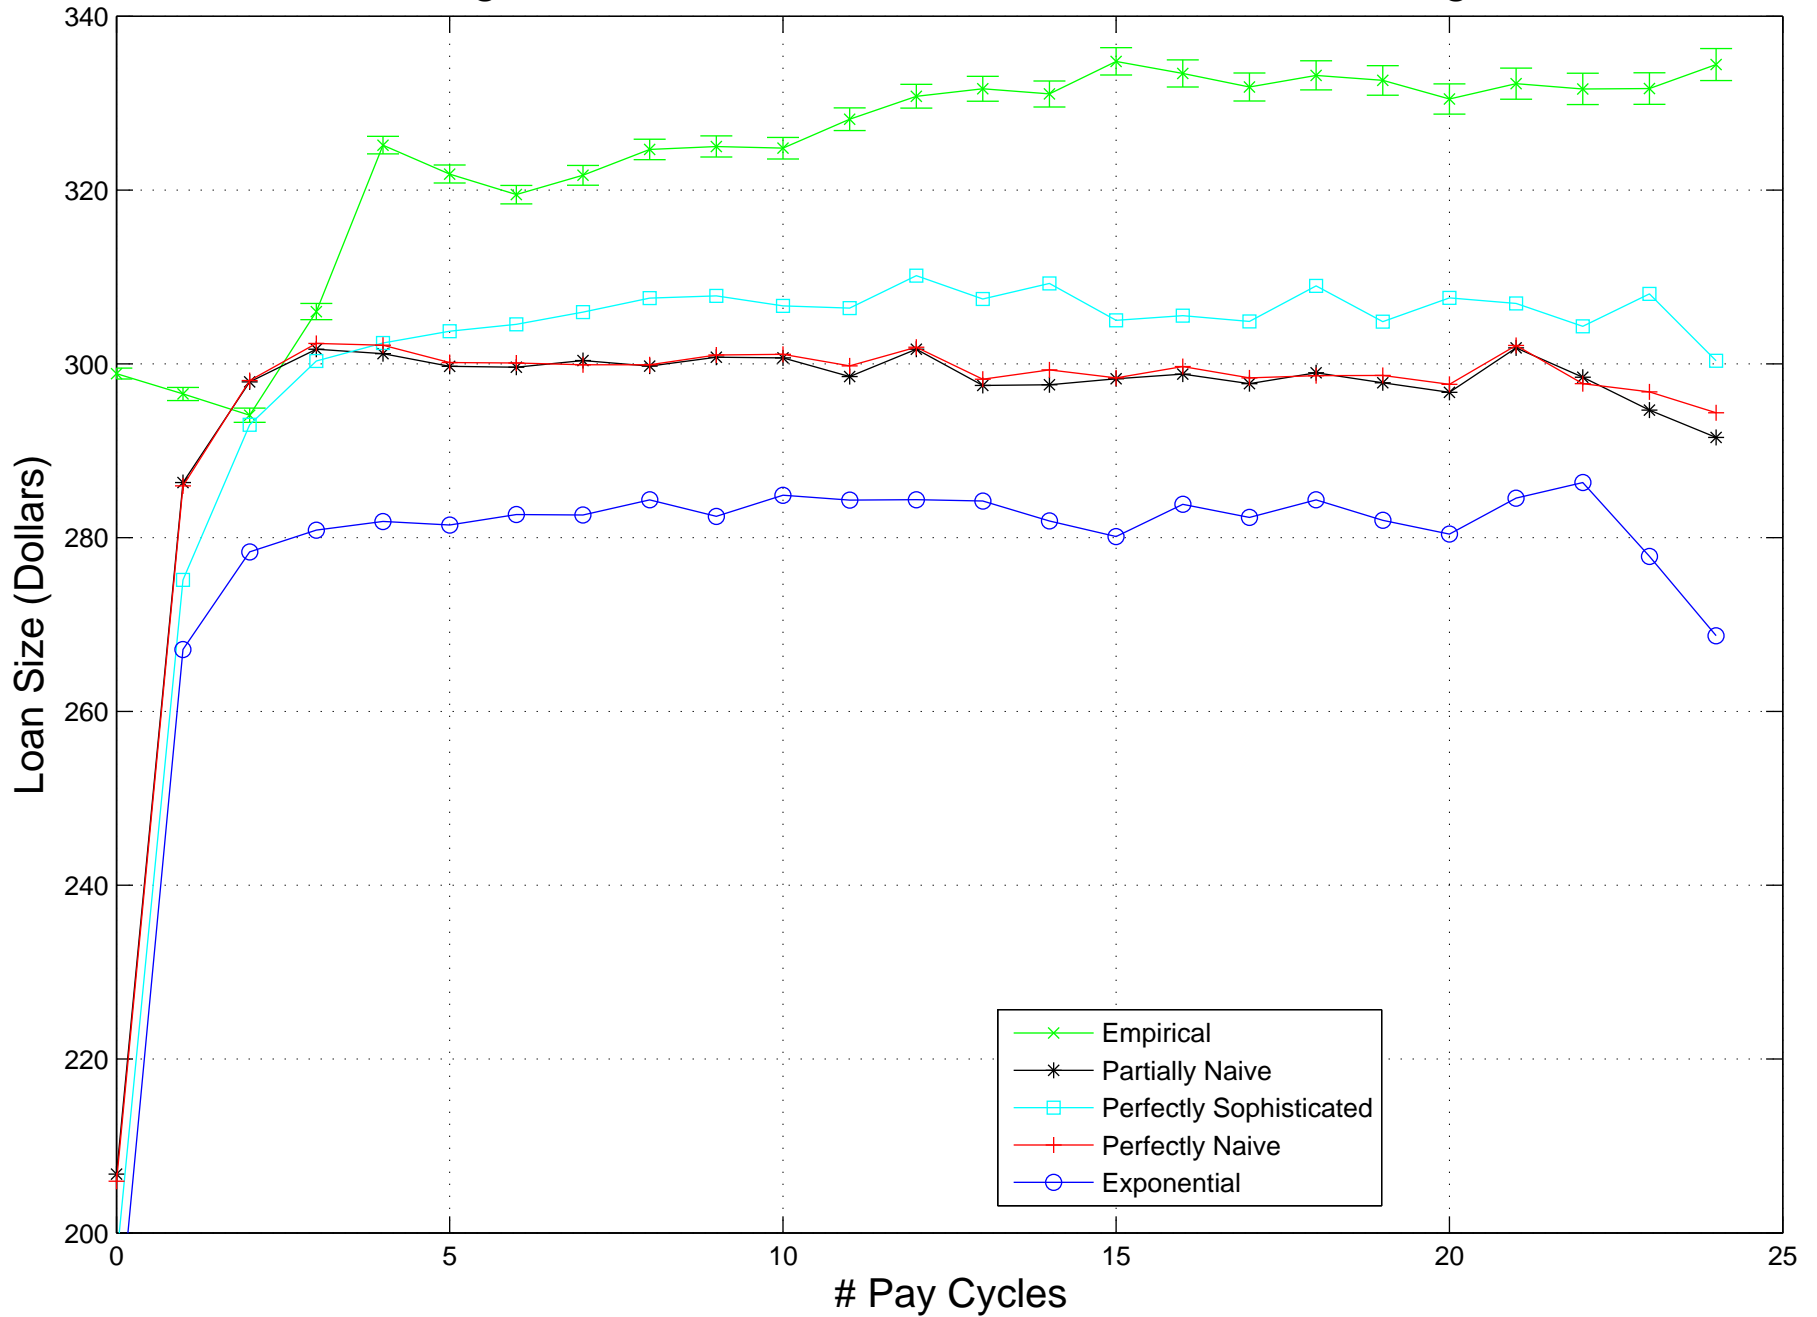


Figure 2: This figure plots loan amounts conditional on borrowing on a payday loan, for each period in the year after a borrower's first loan. Empirical values and their two-standard-error bands are shown in green, and the other lines are simulation model predictions at estimated parameter values.

Figure 3: Default Rate Conditional on Borrowing

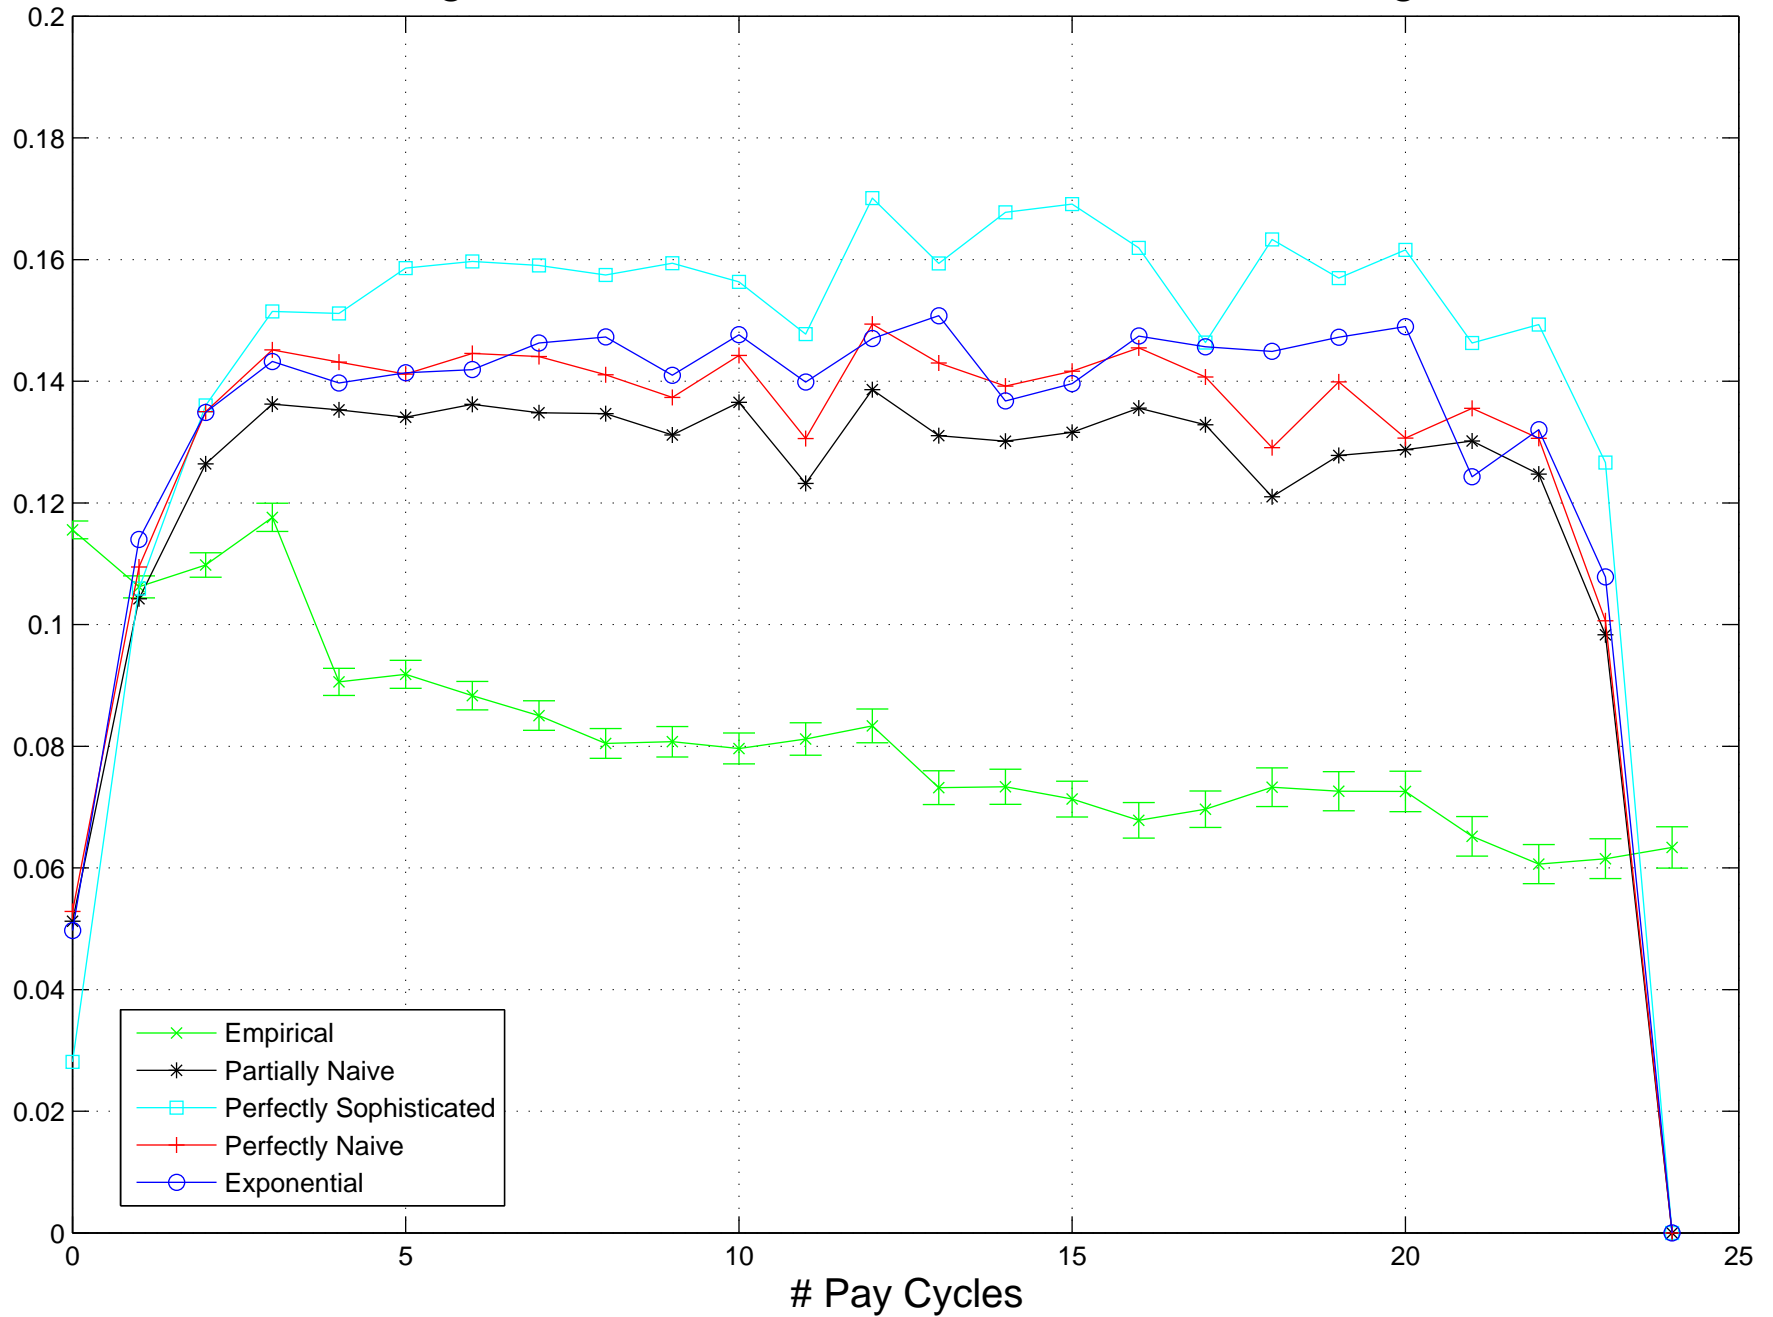


Figure 3: This figure plots default rates conditional on borrowing on a payday loan, for each period in the year after a borrower's first loan. Empirical values and their two-standard-error bands are shown in green, and the other lines are simulation model predictions at estimated parameter values.

TABLE 1: BORROWER DEMOGRAPHICS

Variable	All		Biweekly	
	Mean	N	Mean	N
Age	37.30 (11.4)	101,374	35.60 (10.2)	51,634
Female	0.62	46,828	0.64	23,493
Black	0.41	46,640	0.41	23,387
Hispanic	0.36	46,640	0.36	23,387
Owns Home	0.37	45,617	0.35	23,387
Months at Current Residence	71.30 (97.4)	101,377	65.40 (87.7)	51,636
Checking Account Balance (\$)	283 (577)	99,347	271 (546)	50,732
NSF's on Bank Statement	0.83 (2.42)	101,377	0.82 (2.38)	51,636
Direct Deposit	0.71	65,054	0.71	33,368
Wages Garnished	0.03	45,617	0.03	23,387
Monthly Pay (\$)	1732 (1050)	64,984	1742 (971)	33,334
Months at Current Job	4.89 (7.67)	65,053	4.44 (7.28)	33,368
Paid Weekly	0.12	101,377	0	
Paid Biweekly	0.51	101,377	1	
Paid Semimonthly	0.19	101,377	0	
Paid Monthly	0.18	101,377	0	

Notes: Data provided by a company that makes payday loans. Included are all available demographics for payday borrowers at this company in Texas between 9/2000 and 8/2004. We restrict to individuals who did not change pay frequency and were not allowed to borrow after a default. Quantities are calculated from information provided at the time of each individual's first payday loan. Some (fixed) demographics are only collected when customers seek pawn loans at this company, causing the number of observations to vary. "NSF's" are "Not Sufficient Funds" events like bounced checks.

TABLE 2: BORROWING, REPAYMENT, AND DEFAULT

	(1)	(2)	(3)
	Full Sample	Biweekly	Biweekly, one year
<i>Loan-Level Statistics</i>			
Loan Size (\$)	307.94 (140.29)	323.23 (136.16)	317.55 (133.96)
Interest Payments (\$)	54.78 (25.26)	57.52 (24.55)	56.40 (24.16)
Pr(Default--Bounced Check)	0.089 (0.284)	0.089 (0.284)	0.097 (0.296)
Pr(Default--Write-off)	0.041 (0.197)	0.04 (0.196)	0.047 (0.211)
N	776,667	424,233	335,376
<i>Individual-Level Statistics</i>			
# of Loans Per Person	7.55 (9.19)	8.22 (9.96)	6.50 (6.07)
\$ Loans Per Person	2341.72 (3216.62)	2655.62 (3587.30)	2062.48 (2213.03)
Pr(Default--Bounced Check)	0.51 (0.50)	0.54 (0.50)	0.51 (0.50)
Pr(Default--Write-off)	0.295 (0.46)	0.322 (0.47)	0.297 (0.46)
N	101,377	51,636	51,636

Source: Authors' calculations using data from a financial services provider that offers payday loans. The top panel of the table reports loan-level information, while the bottom panel reports individual-level information. Column 1 pertains to all individuals in the baseline sample, ie, people who borrowed in Texas between September 2000 and August 2004, did not change pay frequency, and were not given the opportunity to borrow after a written-off default. Column 2 restricts to borrowers who were paid biweekly. Column 3 continues to examine the biweekly group, but only examines their experiences in the first year after their first loan.

TABLE 3: INCOME PROCESS CHARACTERISTICS AND ESTIMATION RESULTS

<i>Income process:</i>	$y_{it} = \log(Y_{it}) = \mu_t^y + u_{it}^y + \eta_{it}^y$		$u_{it}^y = \alpha^y u_{it-1}^y + \varepsilon_{it}^y$	
	α^y	σ_{η}^2	σ_{ε}^2	μ^y
Biweekly Pay				
All Observations	0.194 (1.996E-03)	0.073 (6.853E-20)	0.109 (3.385E-03)	6.634 (4.851E-19)
Loan Applications	0.507 (3.373E-03)	0.020 (2.961E-19)	0.080 (1.381E-03)	6.574 (4.352E-18)
Loan Approvals	0.568 (3.301E-03)	0.023 (2.798E-19)	0.071 (1.540E-03)	6.552 (4.050E-18)
Semimonthly Pay				
All Observations	0.044 (3.215E-03)	0.000 (2.168E-17)	0.124 (4.189E-03)	6.771 (1.301E-06)
Loan Applications	0.612 (5.326E-03)	0.052 (3.701E-18)	0.045 (6.165E-03)	6.702 (3.744E-17)
Loan Approvals	0.594 (5.985E-03)	0.045 (4.872E-18)	0.049 (5.065E-03)	6.710 (5.523E-17)
Weekly Pay				
All Observations	0.634 (2.819E-03)	0.147 (1.049E-38)	0.062 (1.557E-02)	5.994 (1.761E-38)
Loan Applications	0.815 (6.687E-03)	0.072 (1.484E-38)	0.052 (8.321E-03)	5.993 (4.708E-38)
Loan Approvals	0.786 (6.138E-03)	0.069 (1.343E-38)	0.056 (5.986E-03)	6.006 (4.397E-38)
Monthly Pay				
All Observations	0.513 (3.988E-03)	0.142 (3.358E-09)	0.000 (6.609E-02)	7.201 (3.541E-08)
Loan Applications	0.340 (8.130E-03)	0.000 (4.007E-08)	0.085 (3.791E-03)	7.128 (1.055E+06)
Loan Approvals	0.375 (7.380E-03)	0.000 (3.739E-08)	0.090 (4.026E-03)	7.142 (5.100E+03)

Source: Authors' estimation using data on take-home pay, as observed by the payday lender. Standard errors in parentheses. The dynamics of income estimation includes a household fixed effect. This table only reports standard errors, but the full covariance matrix is used in the second-stage estimation. In estimates of the consumption-savings-borrowing model in the paper, we use the first case reported in this table. All estimates pertain to income processes with the stated time periods and are not annualized.

TABLE 4: PARAMETER ESTIMATES

Parameter Estimates	Exponential			Perfectly Sophisticated Hyperbolic		Perfectly Naive Hyperbolic		Partially Naive Hyperbolic		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
$\hat{\theta}$										
$\hat{\beta}$	1	1	1	0.83799	0.70048	0.4949	0.49574	0.49529	0.49752	0.53205
	-	-	-	(0.00386)	(0.00370)	(0.00227)	(0.00193)	(0.00491)	(0.00496)	(0.00380)
$\hat{\beta}^E$	1	1	1	$\hat{\beta}$	$\hat{\beta}$	1	1	1.0152	1.0204	0.90042
	-	-	-	-	-	-	-	(0.05321)	(0.04235)	(0.04668)
$\hat{\delta}$	0.81605	0.7225	0.84772	0.79022	0.82016	0.96956	0.96995	0.96937	0.96938	0.96989
	(0.00196)	(0.00329)	(0.00051)	(0.00129)	(0.00024)	(0.00007)	(0.00041)	(0.00023)	(0.00068)	(0.00063)
$\hat{\rho}$	2	2	0.50499	2	2	2	2	2	2	1.7655
	-	-	(0.00398)	-	-	-	-	-	-	(0.01549)
\hat{k}	200	174.38	308.72	200	230	200	197.46	200	198.68	207.22
	-	(0.279)	(0.614)	-	(0.48)	-	(2.1239)	-	(3.1475)	(3.3501)
Goodness-of-fit										
$\chi(\hat{\theta}, \hat{\chi})$	2.55E+05	1.92E+05	1.52E+05	1.71E+05	1.69E+05	1.27E+05	1.27E+05	1.27E+05	1.27E+05	1.25E+05
$\xi(\hat{\theta}, \hat{\chi})$	1.00E+05	63828	46209	62214	34398	74068	76344	50937	50754	50080

Source: Authors' estimation based on the simulation model in the paper and data from a payday lender. Standard errors in parentheses. All estimates use a two-stage Method of Simulated Moments procedure, for the subpopulation of borrowers who are paid biweekly. In all cases, the moments used to estimate the parameters are conditional probability of borrowing, the average amount borrowed conditional on borrowing, and the default rate conditional on borrowing. From the time of each individual's first loan, one year of data is used, implying we have $3 \times 26 = 78$ moments. In every case the specification test is rejected. If the model were correct, csi would have a chi-squared distribution with degrees of freedom equal to 78 minus the number of estimated parameters. Discount factor estimates are for biweekly periods and are not annualized.