Revealed Preference Discount Rates*

Joseph P. Hall†

Abstract

This paper measures household discount rates using default decisions rather than asset prices or choices in laboratories. When a borrower decides whether to default, the relative sensitivity of that decision to short-run payment burdens versus long-run debt obligations reveals the borrower's rate of time preference. I formalize this insight in an endogenous default model and show that two comparative statics—the effect on default of a change in monthly payments and of a change in total repayment—are sufficient statistics for the discount factor. Applying this method to over six million U.S. auto loans, I estimate an average annual discount factor of $\beta = 0.79$, implying a 27% annual discount rate—far steeper than the 4% typically assumed in macroeconomic calibrations. Discount rates decline monotonically with income and with credit score, and both gradients are modest relative to the level, indicating that high default rates among low-score borrowers arise mostly from income risk and costs of default. Extending the method across consumer debt contracts at maturities from two to seventeen years reveals a steeply declining term structure of discount rates—including within borrower—evidence of a discount rate schedule rather than a single rate. These findings provide the first large-scale, revealed-preference distributional estimates of time preferences for households who do not participate in asset markets.

JEL: D14, D15, E21, E43, G10, G40, G51, H23, H55

*This version: August 25, 2026. I thank Amit Seru, Benjamin Hébert, Greg Buchak, Robin Greenwood, Jeffrey Zweibel, Hanno Lustig, Sudheer Chava, Xindi He, Indrajit Mitra, and the Georgia Tech Financial Services Innovation Lab. Excellent research assistance was provided by Victor Huang. This paper subsumes and expands upon the content of a prior paper circulated with the title "Liquidity Constraint and Strategic Default: a Sufficient Statistics Approach".

†Georgia Institute of Technology. Email: jhall390@gatech.edu. Address: 800 W Peachtree St NW, Atlanta, GA 30308. Phone: +1 (781) 715-5862

1 Introduction

A borrower deciding whether to default on a loan makes an intertemporal choice: forgo current payments to consume more today, at the cost of penalties that reduce future utility. The central contribution of this paper is to show that the comparative statics of this decision—how default responds to changes in short-run versus long-run repayment obligations—are sufficient statistics for the household's subjective discount rate. In an endogenous default model with concave, time-separable utility, the ratio of the sensitivity of default to total debt ($\partial p/\partial L$) relative to the sensitivity to monthly payments ($\partial p/\partial m$) identifies the discount factor β without observing consumption, income paths, or the utility function. This replaces "marginal willingness to save" with "marginal willingness to avoid default" as the identifying variation for time preferences—a shift that opens time-preference measurement to the vast majority of households who hold debt but no meaningful financial assets.

Applying this method to over six million U.S. auto loans from administrative credit bureau records, I estimate an average annual discount factor of $\beta = 0.79$ (s.e. 0.003), implying that borrowers value a dollar next year at only 79 cents relative to a dollar today. This corresponds to an annual discount rate of approximately 27% (unchanged by the marginal-utility correction at the state persistence measured in the panel, and no lower than 25% under literature persistence values)—far exceeding the $\beta \approx 0.96$ (4% annual rate) routinely used in macroeconomic calibrations. Reinterpreting the quasi-experimental estimates from Ganong and Noel (2020) on mortgage borrowers through the same sufficient-statistics lens yields a similar magnitude: $\beta = 0.76$ per year. Across different populations, debt types, and identification strategies, the evidence consistently points to steep impatience among borrowing households.

The heterogeneity in estimated discount factors is at least as informative as the level. Revealed discount rates decline monotonically with income and with credit score, and both gradients are modest—roughly four percentage points from the bottom to the top of each distribution, against a level near 27 percent. The score gradient carries a direct implication: because it is an order of magnitude too small to account for the enormous score–default gradient, the welldocumented fact that low-credit-score individuals default more frequently is mostly attributable to income risk, thin financial buffers, and lower costs of default. This distinction matters for credit policy, welfare analysis, and the structural modeling of household default.

Why must we look beyond asset prices? In the United States, only 12% of households participate in any form of risk-free financial investment, and the top 1% holds 75% of risk-free assets (Board of Governors of the Federal Reserve System (2022)). The Euler equation prices risk-free bonds using the discount factor of the marginal investor, who is wealthy and patient. The time preferences of the remaining 88% of households—the constrained, the hand-to-mouth, the borrowers—are not identified by these prices. Laboratory experiments offer an alternative, but the relevance of lab-elicited discount rates to real-world financial decisions is unclear, in part because behavior under small stakes may differ substantially from behavior under large stakes (Coller and Williams (1999), Harrison et al. (2002), Rabin (2000)). The approach developed here uses real-stakes decisions at economically meaningful time horizons, observed at population scale.

I complement the sufficient-statistics approach with a second, independent strategy. Using the 2022 Survey of Consumer Finances, I observe the highest interest rate at which each household currently borrows. Because credit is supplied along an upward-sloping schedule, this maximum borrowing rate provides a revealed-preference lower bound on the household's discount rate. After adjusting for consumption-smoothing demand using a calibrated life-cycle income model (Guvenen et al. (2021)), the median implied upper bound on β is 0.91—consistent with at least 9% annual discounting even before accounting for the discreteness of credit products that makes these bounds conservative. Consumption-smoothing demand itself, as measured by the life-cycle model, explains almost none of the cross-sectional variation in borrowing rates ($R^2 < 2\%$), leaving heterogeneity in impatience as the dominant source of variation.

These estimates have direct implications for public finance, credit regulation, and welfare analysis. The U.S. federal government applies a social discount rate near 2–3% per year (House (2023)). If a large share of the population has discount factors near 0.80, then cost-benefit analyses of policies with long-run payoffs—early childhood education, climate mitigation, infrastructure—will systematically overweight future benefits relative to the preferences of the citizens these policies aim to serve. Hendren and Sprung-Keyser (2020) show that the relative ranking of early childhood versus adult education programs flips at discount rates in this range. For consumer credit policy, the results imply that short-term payment relief will be far more effective at preventing default than principal reduction or long-term debt forgiveness, consistent with the quasi-experimental evidence in Ganong and Noel (2020).

The remainder of the paper proceeds as follows. Section 2 reviews related literature. Section 3 develops the sufficient-statistics framework. Section 4 applies this framework to auto loan data and presents the main estimates. Section 5 documents heterogeneity by income, credit score, and geography. Section 7 places these estimates in the context of prior work. Section 8 develops the complementary lower-bound approach using borrowing rates and a life-cycle income model. Section 9 discusses interpretation, including what the estimated β captures and extensions to non-separable utility. Section 10 concludes.

2 Related Literature

This paper connects to several literatures.

Discount rates from asset markets. The standard approach to measuring time preferences uses the Euler equation to link the price of risk-free bonds to the discount factor of the marginal investor (Lustig and Chien (2005); Bansal and Yaron (2004)). Calibrations of heterogeneous-agent models routinely set $\beta \approx 0.96$ to match the capital-output ratio (Guvenen (2006)) or government bond yields. Di Tella et al. (2024) construct a "zero-beta rate" from equity portfolios, estimating an average of 8.3% per year. These approaches identify the preferences of the marginal investor—necessarily a wealthy, patient household—which may differ substantially from those of the typical household. My approach instead identifies discount factors for the borrowing population, which constitutes the majority of U.S. households.

Laboratory and field experiments. A large experimental literature elicits discount rates using monetary tradeoffs over time. Coller and Williams (1999) and Harrison et al. (2002) are foundational contributions. Cohen et al. (2020) survey 222 empirical studies of time preferences published through 2016 and provide a comprehensive catalog of methods and findings. Of these 222 studies, 131 are lab-based, 56 use field data, and only 24 use publicly available datasets. The 112 studies that report point estimates of the subjective discount rate have a median of 22% per year and a mean of 26%—considerably higher than standard macro calibrations, though primarily from small samples with small stakes. The relevance of lab-elicited discount rates to real-world financial behavior is unclear, as preferences over small-stakes gambles may differ substantially from those governing large-stakes intertemporal decisions (Rabin (2000)). My approach complements this literature by using revealed preference from large-stakes decisions at population scale.

Structural models of default and household debt. Ganong and Noel (2020) use a regression discontinuity in the HAMP mortgage modification program to show that short-term payment reductions prevent default far more effectively than equivalent-value principal reductions. As I show, applying the sufficient-statistics formula to their reported estimates yields $\beta = 0.76$, consistent with steep impatience. Dobbie and Song (2020) conduct a randomized experiment on credit card borrowers and find that delayed interest write-downs improve outcomes while immediate payment reductions do not; however, the payment reduction effect is not statistically significant, so the formula does not produce a well-identified β in their setting. Ganong and Noel (2019) estimate that approximately one-third of consumers in their data are extremely impatient "low-β" types, but do not produce full distributional estimates. The structural default model in my paper shares the basic tradeoff in Indarte (2023), where borrowers weigh current consumption against default penalties.

Sufficient statistics methods. The sufficient-statistics approach to welfare and policy analysis (Chetty (2009)) emphasizes that key policy-relevant parameters can be recovered from a small number of empirically estimable moments, without specifying the full structure of the model. My application to default behavior is, to my knowledge, the first to use this methodology to identify discount rates. The identifying logic is analogous to recovering an MRS from behavioral responses: the ratio of two default sensitivities reveals the intertemporal tradeoff.

Discount rates in public finance. Discount rates are central to cost-benefit analysis. The United Kingdom and France publish declining discount rate schedules for horizons up to 350 years (Arrow et al. (2014)). The U.S. government recently lowered its recommended social discount rate from 3% to 2% (House (2023)). Hendren and Sprung-Keyser (2020) show that the ranking of public investments is sensitive to the discount rate, with adult education dominating at high rates and early childhood education at low rates. DeMarzo et al. (2023) show that patient citizens are harmed when their government borrows excessively. My estimates suggest that the actual discount rates of the median household are an order of magnitude higher than those used in policy analyses.

3 A Sufficient Statistics Approach to Measuring Discount Rates

What can we learn about households' impatience from their default decisions? Interest rates on observed debt provide lower bounds on discount rates (Section 8), but they cannot point-identify the discount factor because credit products are discrete and rationed. Default decisions, however, encode a richer intertemporal tradeoff. I show that in a simple endogenous default model with separable utility, the causal effects on default of changes in repayment obligations at two different time horizons are sufficient to point-identify the discount factor β between those two horizons, up to a single marginal-utility ratio whose key input—the persistence of the marginal borrower's economic state—I measure directly in the credit-bureau panel; at the measured persistence the implied correction is negligible, and under substantially lower literature values of persistence it moves the annual estimate by less than two percentage points (Appendix D). No functional form for the utility function is required beyond strict concavity and monotonicity for the identifying identity itself.

The economic intuition is simple. An impatient borrower is very sensitive to short-run changes in payment obligations—a higher monthly payment today sharply increases the temptation to default. A patient borrower, by contrast, responds primarily to changes in total debt, which affect long-run consumption. The ratio of these two responses reveals the weight the borrower places on the present versus the future.

3.1 Two-Period Endogenous Default Model

The model is set in two discrete time periods $t \in \{1, 2\}$, representing the dates at which payments are due. The starting balance of the loan is L, and the payment in each period is m (monthly payment), with the second-period payment equal to $L - m$. In the base case, $L = 2m$. The two-period structure is a clean special case of a general model rather than a knife-edge simplification: Appendix B extends the model to T payments with default as an optimal stopping problem and shows (Proposition 1) that every default probability remains strictly interior for any horizon T, provided the punishment has a flow component and staying in good standing is valued as one period of standing plus the option to default later. The same appendix characterizes the two modeling choices that instead produce degenerate corner probabilities—comparing default against committed repayment of all remaining installments, and letting the punishment vanish at maturity—so the sensitivities derived below are the $T = 2$ case of well-behaved general formulas.

Period 1 income $y_1 > 0$ is uncertain and distributed according to $F(y_1)$. No specific distributional assumption is needed. The punishment for default is σ, encompassing credit-market exclusion, social stigma, and garnishment (Indarte (2023)).

The borrower has a strictly increasing, strictly concave utility function over consumption, additive across periods with discount factor β:

$$
\begin{aligned}
V_2^P &= u(C_2^P) \\
V_2^N &= u(C_2^N) - \sigma \\
V_1^P &= u(C_1^P) + \beta \left[\int_0^{y_2^*} V_2^N dF(y_2|y_1) + \int_{y_2^*}^{\infty} V_2^P dF(y_2|y_1) \right] \\
V_1^N &= u(C_1^N) + \beta V_2^N
\end{aligned}
\tag{1}
$$

where y_2^* is the endogenous default threshold in period 2. The budget constraints are:

$$
\begin{aligned}
C_1^P &= y_1 - m \\
C_1^N &= y_1 \\
C_2^P &= y_2 - (L - m) \\
C_2^N &= y_2
\end{aligned}
\tag{2}
$$

The borrower defaults when the utility gain from avoiding payment exceeds the punishment σ. Let y_2^* satisfy:

$$
u(C_2^P(y_2^*)) = u(C_2^N(y_2^*)) - \sigma. \tag{3}
$$

The period-1 default threshold y_1^* is determined by:

$$
u(C_1^{P*}) + \beta \mathbb{E}[V_2^P] \;=\; u(C_1^{N*}) + \beta \mathbb{E}[V_2^N]. \tag{4}
$$

I do not require closed-form solutions for y_1^* or y_2^*. Instead, I focus on how small changes in m and L perturb the default probability $p_1 = F(y_1^*)$ and thereby reveal β.

3.2 Default Sensitivities as Sufficient Statistics

Differentiating $p_1 = F(y_1^*)$ with respect to m and L yields:

$$\frac{\partial p_1}{\partial m} = f(y_1^*)\frac{\partial y_1^*}{\partial m}, \qquad \frac{\partial p_1}{\partial L} = f(y_1^*)\frac{\partial y_1^*}{\partial L}.$$

Applying implicit differentiation to the threshold condition (4) (full derivation in Appendix A) yields:

$$\begin{aligned}\frac{\partial p_1}{\partial L} &= f(y_1^*)\,\beta(1-p_2)\,\frac{\mathbb{E}\left[u'(C_2^P(y_2)) \mid y_2 > y_2^*\right]}{u'(C_1^{P*}) - u'(C_1^{N*})},\\ \frac{\partial p_1}{\partial m} &= f(y_1^*)\frac{u'(C_1^P)}{u'(C_1^{P*}) - u'(C_1^{N*})} - \frac{\partial p_1}{\partial L}.\end{aligned} \tag{5}$$

The first expression captures the effect of raising long-run obligations: higher L reduces future consumption for borrowers who repay, lowering their continuation value and making default more attractive. The second captures the effect of raising current payments: higher m reduces current consumption, increasing the marginal utility of consuming today and thus increasing the temptation to default by not paying. Crucially, the ratio of these effects eliminates the default-margin density $f(y_1^*)$ and the scale of utility, leaving β multiplied by a single ratio of marginal utilities—the object λ defined in equation (6) below, whose calibrated value is close to one.

Rearranging:

$$\frac{\frac{\partial p_1}{\partial L}}{(1-p_2)\frac{\partial p_1}{\partial m} + \frac{\partial p_1}{\partial L}} = \beta\frac{E\left[u'\left(c_2^P\right)|y_2 > y_2^*\right]}{u'\left(c_1^P\right)}$$

Two further steps take this exact identity to an estimator. The first is the empirically supported simplification that late default is negligible ($p_2 \approx 0$, verified in Appendix J). The second deserves scrutiny rather than a passing assumption: defining the marginal-utility ratio

$$\lambda \equiv \frac{\mathbb{E}\left[u'(C_2^P) \mid y_2 > y_2^*, y_1 = y_1^*\right]}{u'(C_1^{P*})}, \tag{6}$$

the exact identity reads $A/(A+B) = \beta\lambda$, and the simplified formula below sets $\lambda = 1$. The numerator of λ is the average marginal utility of future repayers conditional on standing at today's default margin, while the denominator is the marginal utility of the marginal defaulter herself; with concave utility these need not coincide. Appendix D disciplines λ with a measurement rather than an assumption: the ratio's key input is the persistence of the marginal borrower's economic state, and the credit-bureau panel measures it directly. Across 21.4 million delinquency-spell transitions for auto borrowers, the distress state that puts a borrower at the default margin persists with annualized probability 0.992—only 0.3 percent of delinquent loan-months cure by the next snapshot. At this measured persistence the two marginal utilities cancel almost exactly ($\lambda = 1.00$ at the preferred insurance pass-through), so the simplified estimator is measured, not assumed, to be nearly exact. Under substantially lower persistence values from the earnings literature the correction remains small: at $\rho = 0.95$ (Storesletten et al., 2004), $\lambda = 0.95$ and the corrected estimate is $\beta = 0.80$ (a 25% annual rate) against the baseline 0.79 (27%). Second, the data themselves bound the correction: identification requires $\lambda \geq (\hat{\beta})^T = 0.47$, which refutes exactly the low-persistence, high-insurance calibrations that would generate large corrections. Third, because a common λ rescales all subgroup estimates by the same monotone transformation, the cross-sectional comparisons at the heart of Section 5 – across credit scores, incomes, and geographies – are unaffected by any uniform value of the correction. Together with the borrowing-rate evidence of Section 8, which bounds discount rates from below, the corrected estimates are bracketed by two independent methodologies with oppositely-signed biases. Under the simplification $\lambda = 1$:

$$\frac{\frac{\partial p_1}{\partial L}}{\frac{\partial p_1}{\partial m} + \frac{\partial p_1}{\partial L}} = \beta \tag{7}$$

This maps neatly to intuition. If default decisions are entirely forward-looking and strategic

($\partial p / \partial L$ dominates), then $\beta = 1$: perfect patience. If default decisions are entirely myopic and driven only by current payment burdens ($\partial p / \partial m$ dominates), then $\beta = 0$. The formula smoothly nests all intermediate cases.

3.3 Recovering the Annual Discount Factor

The formula produces the discount factor over the horizon between typical default and loan maturity. To convert to an annual rate under the assumption of time-consistent discounting:

$$\beta_{annual} = \left(\frac{\frac{\partial p_1}{\partial L}}{\frac{\partial p_1}{\partial m} + \frac{\partial p_1}{\partial L}} \right)^{1/T} \tag{8}$$

where T is the average time (in years) between default and loan maturity.[1] Standard errors are computed via the delta method (Appendix K).

It is worth being precise about the estimand, because Section 6 will show that the measurement differs across horizons within the same borrower. Let $D_i(t)$ denote borrower i's discount function and define the revealed annualized yield at horizon T as $y_i(T) = D_i(T)^{-1/T} - 1$, the zero-coupon rate of the borrower's discount curve. Each contract family prices exactly one point on this curve—the default margin trades consumption today against the stream over the family's default-to-maturity horizon—so equation (8) recovers $y_i(\tau_f)$ at the family's horizon τ_f. Every cross-sectional statement in this paper is therefore a statement about the distribution of a horizon-labeled rate: the auto-loan heterogeneity of Section 5 concerns the three-year revealed discount rate, an object that is well-defined without any assumption on the curve's shape because the horizon is common across the groups compared (a decomposition reported there verifies this empirically). Statements comparing rates across horizons are statements about the curve, and are treated separately in Section 6.

The assumption of time-consistent discounting deserves comment, because it is an interpretation rather than a testable restriction: a two-period contract offers exactly one intertemporal margin, so the sufficient statistic delivers one point on the borrower's discount curve, and con-

[1] In the auto loan sample, $T = 3.2$ years unconditionally.

verting it to an annual rate requires taking a stand on the curve's shape. Appendix B makes this precise and shows it is not a limitation of the framework but of the contract. In the T-period extension, the response of initial default to a perturbation of the payment due at horizon s is $\partial p_1/\partial m_s = f(y_1^*)\, D(s-1)\, Q_s\, \mathbb{E}[u'(c_s)]/\Delta u_1'$ (Proposition 3), where $D(\cdot)$ is an arbitrary discount function and Q_s is the observable probability of remaining in good standing through s. Ratios of these responses across horizons therefore identify the entire discount function nonparametrically from a menu of per-horizon payment experiments (Proposition 4); quasi-hyperbolic β–δ preferences are point-identified from perturbations at just two distinct future horizons (Corollary 5); and with three or more horizons the exponential null itself becomes testable. The required variation—randomized, ceteris paribus perturbations of individual payments within a fixed borrower population—is not delivered by comparing borrowers who selected into different contract lengths, and to my knowledge no such experiment has been run. The theory for measuring the full discount curve is thus worked out and waiting on the experiment; with the data at hand, the honest object is the single point that equation (8) reports.

3.4 Application in Experimental and Quasi-Experimental Settings

Before turning to the auto loan data, I demonstrate the framework's power by reinterpreting three prominent quasi-experimental studies through this lens. Table 1 summarizes the results. Each application uses the simplified ($\lambda = 1$) formula of equation (7); the marginal-utility correction of Appendix D applies to these settings symmetrically, so a common correction leaves the cross-study comparison intact.

Ganong and Noel (2020) study mortgage modifications using a regression discontinuity in HAMP. They find that short-term payment reductions substantially reduce default, while long-term principal reductions have almost no effect. Converting both estimates to semi-elasticities—using the 11 percentage point LTV reduction and a pre-modification LTV of approximately 120% to express the principal treatment as a percent change in total balance—and applying the sufficient-statistics formula with a five-year horizon yields:

$$\beta_{\mathrm{GN}} = 0.76 \quad \text{(s.e. 0.30)}.$$

The large standard error reflects the imprecision of the principal reduction estimate, which is not statistically significant in the original study. Nevertheless, the point estimate is consistent with steep impatience—similar to the auto loan estimate.

Indarte (2023) decomposes the bankruptcy filing response into a moral hazard channel (response to discharge value, analogous to A) and a liquidity channel (response to payment burden, analogous to B), using two independent identification strategies: a regression kink design exploiting homestead exemption limits for the moral hazard estimate, and an instrumental variable using adjustable-rate mortgage resets for the liquidity estimate. She finds that 83% of the filing response operates through the liquidity channel and only 17% through moral hazard. Both estimates are converted to comparable per-dollar units via an NPV adjustment, and the two identification strategies use independent samples ($N \approx 99$ million for the RKD, $N \approx 1$ million for the ARM IV), so the covariance is zero. Applying the sufficient-statistics formula with a seven-year horizon (the median remaining mortgage duration in her sample) yields:

$$\beta_{\mathrm{Indarte}} = 0.78 \quad \text{(s.e. 0.03)}.$$

This is the most precisely estimated external application of the formula, with a 95% confidence interval of $[0.72, 0.83]$. The point estimate is remarkably close to the auto loan estimate of $\beta = 0.79$, despite coming from an entirely different population (homeowners at risk of bankruptcy), debt type (mortgages), and identification strategy.

Dobbie and Song (2020) study credit card borrowers using a randomized experiment and find that delayed interest write-downs improve repayment outcomes while immediate payment reductions do not. Because the payment reduction effect is not statistically significant in any specification, the formula does not produce a well-identified β in their setting. This is itself informative: the absence of a measurable short-run payment effect means the identifying variation required by the sufficient-statistics approach is not present in the Dobbie–Song experiment.

Table 1: External Validation: Applying the Sufficient-Statistics Formula to Other Studies

Study	Setting	Identification	T	β	SE	Annual Rate
Auto loans (this paper)	U.S. auto loans	Logit	3.2	0.79	0.003	27%
Ganong and Noel (2020)	HAMP mortgages	Two RDs	5	0.76	0.30	31%
Indarte (2023)	Household bankruptcy	RKD + ARM IV	7	0.78	0.03	29%
Dobbie and Song (2020)	Credit card debt	RCT	4	0.69	0.40	—

Notes: Each row applies the formula $\beta = (A/(A+B))^{1/T}$ to estimates from the indicated study, with standard errors via the delta method. T is the horizon in years. The Indarte (2023) estimates use the 83/17 moral-hazard/liquidity decomposition from her Tables 1 and 3, with the NPV-adjusted coefficients ensuring comparable per-dollar units. The Dobbie and Song (2020) estimate is excluded from the comparison because the short-run payment coefficient is not statistically significant ($t = 0.33$), producing an uninformative confidence interval.

4 Application to Auto Loan Data

4.1 Data

Auto loans offer an attractive empirical setting: maturities vary substantially, payments are highly salient, and auto loans are one of the most common forms of installment credit, with wide participation across the income distribution. I use a 10% random sample of U.S. credit files containing over six million auto loans originated between 2005 and 2024. Table 2 summarizes the key variables.

Table 2: Summary of Auto Loan Sample

Variable	Non-null Obs.	Mean	Median	25th Pctl	75th Pctl
Loan Amount	6313128	22800.77	20000	13209	29122
Loan Balance	6404640	20440.13	17500	9877	27216
Monthly Payment	6197142	2241.88	395	292	533
Maturity (Months)	6408812	60.24	61	48	72
Credit Score	6408812	771.76	700	623	774
Age	6392929	45.48	45	33	56
Estimated Income/Mo	6316100	3885.30	3417	2500	4667

The typical repayment horizon is five years (60 months), and the three most popular maturities—36, 60, and 72 months—account for over half of observations. Figure 1 shows steady origination volumes since 2010.

Figure 1: Auto Loan Originations Over Time

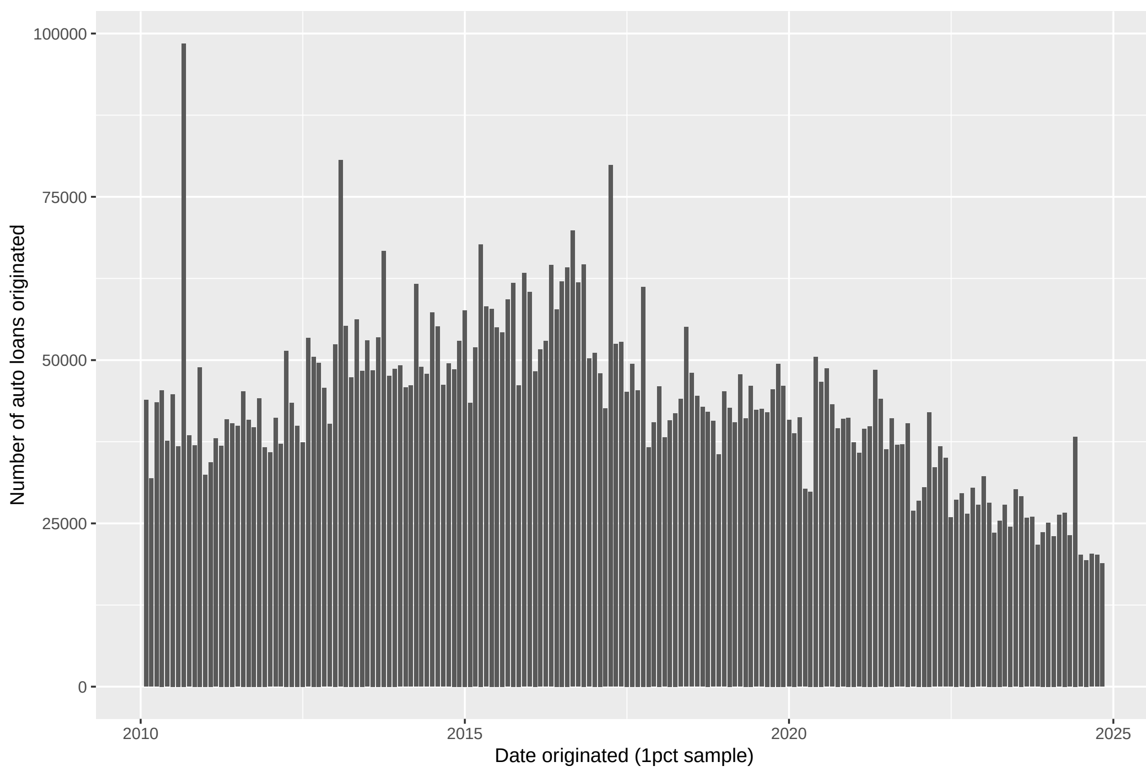

Notes: Number of auto loans originated in the 10% credit bureau sample, plotted by origination year.

4.2 Identification

To estimate $\partial p_1/\partial m$ and $\partial p_1/\partial L$, I exploit variation in contractual maturity and loan size across borrowers with the same credit score. The identifying assumption is:

> Conditional on credit score, maturity and loan size are as-if randomly assigned with respect to pre-existing determinants of default.

This assumption is plausible because lenders rely on standardized underwriting algorithms that map credit scores to allowable loan structures. Conditional on score, residual variation in contract terms reflects idiosyncratic factors—dealer negotiations, vehicle choice, down payment size—rather than private information about future default risk. The model explicitly calculates the discount rate under the assumption that the correlation between loan terms and default is entirely due to moral hazard, in which larger loans cause more default, with adverse selection, in which worse borrowers select into larger loans, set to zero.

Several pieces of evidence support this assumption. First, Figure 2 shows that the unconditional relationship between maturity and default rate is highly idiosyncratic, with no systematic monotonic pattern. This is consistent with lenders adjusting terms (e.g., requiring larger down payments) to equalize default risk across maturities. Second, even if adverse selection is present, a directional argument favors the identifying assumption. If lenders could privately observe worse credit risks conditional on score, they would offer smaller loans, biasing against the observed pattern. Third, if the degree of adverse selection is equal across all maturities, it cancels out in the ratio of Equation (7).

An instrumental-variables program supplements these arguments, and I report its state completely. Instruments built from the level of market interest rates at origination fail, in three separate designs, for a structural reason worth recording. A rate level varies only at the origination-cohort level, and cohort composition co-moves with the rate cycle, so no fixed-effect structure recovers exclusion. Instruments built from lender policy menus perform differently. The staggered expansion of 72- and 84-month auto terms across lenders, and cross-lender differences in loan-size menus, yield two leave-one-out instruments with first-stage F statistics of 221 and above 47,000 when clustered by lender. A control function using these instruments without lender fixed effects fails an exclusion diagnostic—lenders with long-maturity menus serve observably riskier clienteles—and the within-lender design that identifies from menu changes over time is the appropriate specification. The estimates from that specification are the subject of ongoing work; the identification of the estimates reported in this paper rests on the fixed-effects design above, on the diagnostics of Section 4.4, and on the agreement documented in Table 3: two external designs whose identification is quasi-experimental—the HAMP regression discontinuity and the bankruptcy regression kink—recover discount factors within one to four points of the fixed-effects specifications when fed through the same formula, across three FE structures whose sensitivity ratio is additionally invariant to the logit, average-marginal-effect, and linear-probability metrics under geographic-cohort cells.

Table 3: Agreement between quasi-experimental designs and fixed-effects specifications

	population and horizon	identification	$\widehat{\beta}$	(s.e.)
Panel A: well-identified external designs, through the formula				
HAMP modifications	distressed mortgagors, 5 yrs	regression discontinuity	0.76	[0.72, 0
Bankruptcy filing	at-risk homeowners, 7 yrs	RKD + ARM instrument	0.78	–
Panel B: fixed-effects specifications, auto sample				
baseline	all auto borrowers, 3.2 yrs	score-ventile $\times$ zip-month cells	0.787	(0.003
geographic-cohort	all auto borrowers, 3.0 yrs	zip3 $\times$ origination-year cells	0.803	–
low-score subsample	bottom score decile, 3.2 yrs	varying-coefficient, evaluated at d1	0.789	(0.004

Notes: Panel A applies the sufficient-statistics formula to the reported comparative statics of Ganong and Noel (2020) (bracket: 95 percent confidence interval) and Indarte (2023), as in Section 3.4. Panel B reports this paper's estimator under three specifications; the geographic-cohort row uses the extract sample and its ratio is metric-invariant (logit 0.523, average marginal effects 0.523, linear probability 0.511; `results/metric_diagnosis.csv`). The low-score row evaluates the varying-coefficient surface at the bottom score decile, the population segment closest to the HAMP sample. Replication: `Code/pace/68_mortgage_fe_logits.R`, `Code/make_paper_exhibits.py`.

Figure 2: Unconditional Default Rate by Maturity

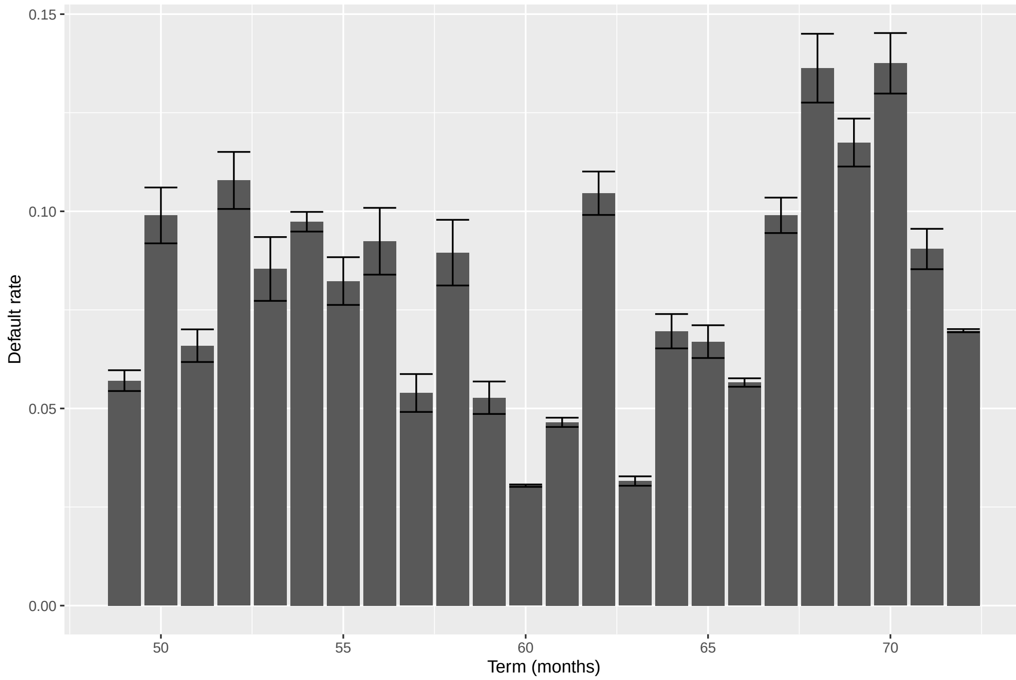

Notes: Each bar shows the unconditional default rate for auto loans of the given contract maturity. The absence of a monotonic pattern is consistent with lenders adjusting terms to equalize default risk.

4.3 Main Results

Table 4 reports estimates from two logit models of default. Each includes credit-score ventile fixed effects and zipcode-by-month fixed effects, which flexibly absorb borrower risk classification and local economic conditions.

$$\text{Default}_i = \beta_S \log(\text{MonthlyPayment}_i) + \gamma_{v(i)} + \delta_{z(i),t(i)} + \varepsilon_i \tag{9}$$

$$\text{Default}_i = \beta_L \log(\text{TotalPayments}_i) + \gamma_{v(i)} + \delta_{z(i),t(i)} + \varepsilon_i \tag{10}$$

The short-run coefficient is $\beta_S = 0.50$ (s.e. 0.008) and the long-run coefficient is $\beta_L = 0.43$ (s.e. 0.006). Both are large and precisely estimated, and their ratio is the finding: the future component of obligations receives less than half the weight of an equally sized change in the current payment, which is the signature of steep discounting that equation (8) converts into β.[2]

Table 4: Default Sensitivity to Short-Run and Long-Run Payment Obligations

Dependent Variable:	Default	
Model:	(1)	(2)
Variables		
log(Monthly Payment)	0.5000***	
	(0.0083)	
log(Total Payments)		0.4283***
		(0.0058)
Fixed-effects		
Score Ventile	Yes	Yes
Zip Code × Month	Yes	Yes
Fit statistics		
Observations	1,434,210	1,434,208
Squared Correlation	0.37087	0.37278
Pseudo R^2	0.37593	0.37766
BIC	3,423,797.9	3,421,505.8

IID standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Applying the sufficient-statistics formula (8):

Table 5: Sufficient-Statistics Estimate of the Discount Factor β

Variable	Estimate	Std. Error	Description
$\frac{\partial p_1}{\partial m}$	0.500	0.008	Sensitivity of default to monthly payment
$\frac{\partial p_1}{\partial L}$	0.428	0.006	Sensitivity of default to total payments
T	3.223	0.004	Horizon: years between average default and maturity
β	0.787	0.003	Annual discount factor: $\left(\frac{\partial p_1/\partial L}{\partial p_1/\partial L + \partial p_1/\partial m}\right)^{1/T}$

[2] The logit estimation sample is 1.43 million loans: observations in fixed-effects cells with no variation in default contribute no likelihood and drop from the estimation, which removes most of the 6.3 million-loan sample's singleton zip-month cells. The defaulter timing moments use the full sample.

The estimated annual discount factor is $\beta = 0.787$ (s.e. 0.003), implying that the average auto loan borrower values a dollar next year at only 79 cents relative to a dollar today. The equivalent annual discount rate is approximately 27%. The marginal-utility correction of Appendix D leaves this estimate essentially unchanged at the persistence measured in the panel ($\rho = 0.992$; $\lambda = 1.00$); under the lower persistence values estimated in the earnings literature the corrected rate falls no further than 25% ($\beta = 0.80$). Because any common value of the correction rescales every subgroup identically, the cross-sectional comparisons that follow are invariant to it.

Two decomposition results establish that the estimate measures default behavior rather than the annualization clock. First, the level verdict does not depend on the horizon estimate T: the measured sensitivity ratio is $R = 0.465$ over the default-to-maturity horizon, and consistency with the macroeconomic benchmark $\beta = 0.96$ would require $T = \ln R / \ln 0.96 = 18.8$ years—on loans that last five. A generous 20% error in T moves the implied annual rate from 27% only to 22%. Second, decomposing the cross-sectional variation in Section 5 into the contribution of the sensitivity ratio R and the contribution of the horizon T, the sensitivity ratio accounts for 84% of the variance in log discount factors across the 142 geographic clusters (89% across state-cycle cells) and reproduces the cross-group ranking with rank correlation 0.97; the horizon contributes 2%. The heterogeneity this paper reports is variation in how default responds to payments and obligations, not variation in loan timing.

4.4 Is It Liquidity?

The model cannot rule out liquidity constraints or high expected income growth as drivers of low measured discount factors;[3] rather, the evidence that the estimates reflect impatient preferences is quantitative, and this subsection assembles it. For mortgages, Ganong and Noel (2020) find

[3] Nor can the estimator itself: Appendix F reports a Monte Carlo exercise in which economies with entirely liquidity-driven default produce estimated discount factors near the empirical value rather than near zero, because total obligations load mechanically on the payment that liquidity-driven default does respond to. The same exercise shows the converse discipline: across simulated economies spanning true discount factors from 0.55 to 0.99 and liquidity shares up to one half—each calibrated to the empirical default rate—no combination of patient preferences and admissible contamination produces an estimate as low as the one measured.

that most defaults follow cash-flow shocks; if binding liquidity constraints drove default here, the sensitivity ratio would measure desperation rather than discounting. Liquidity constraints fail to predict three features of the data, expectations of income growth fail as an explanation for three reasons, and the quantitative adjustments both channels admit are small; the four diagnostics below carry the argument, and the closing paragraph states the magnitudes.

Table 6 collects the two diagnostics that carry the most weight.

Table 6: Liquidity diagnostics

Panel A: other tradelines of auto defaulters (measured at six months after chargeoff)				
Defaulters with other open, clean trades at default				2,462,105
Mean number of other tradelines				10.7
Share still paying at least one other trade				65.1%
Share with all other trades current				13.6%
Share with any other chargeoff				22.7%
Share defaulting on every other trade				1.6%
Panel B: the estimate across local labor-market shocks				
	mean Δu (pp)	$\widehat{\beta}$	(s.e.)	defaults
Unemployment-change tercile 1	-1.39	0.774	(0.005)	93,528
Unemployment-change tercile 2	-0.46	0.786	(0.005)	93,972
Unemployment-change tercile 3	+0.88	0.789	(0.006)	92,499
Expansion Cohort	-0.33	0.792	(0.003)	231,371
Recession Cohort	-0.25	0.731	(0.011)	48,628

Notes: Panel A follows every auto defaulter holding at least one other open tradeline with no chargeoff at the default event, and measures the status of those other tradelines at the archive snapshot nearest six months after the chargeoff; a tradeline is paying if its last-payment date falls within three months of the snapshot. Panel B estimates the sufficient statistic within terciles of the change in state unemployment over the loan's first eighteen months, terciles taken within origination year, and separately for recession-year origination cohorts (2007–2009, 2020). Replication: `Code/pace/41_cross_loan_events.R`, `Code/pace/44_shock_splits.R`.

First, auto default is overwhelmingly a portfolio choice. Among 2.46 million auto defaulters holding other open, clean tradelines at the time of default, 65 percent are still actively making payments on at least one other account six months after the auto chargeoff (they hold 10.7 other tradelines on average), and only 1.6 percent default on every account. Money continued to flow to other creditors through the event. The only form of the objection that undermines revealed preference—total cash-flow collapse—is bounded between 1.6 and roughly 35 percent of defaults.

Second, the estimates do not track liquidity shocks. Sorting loans within origination year into terciles of the change in state unemployment over the loan's first eighteen months, the estimated discount factor is flat: 0.774, 0.786, and 0.789 from the mildest to the most severe exposure. If the estimates repackaged liquidity shocks, this gradient would be steep. (Loans originated in recession years do show lower estimates, 0.731 against 0.792; that comparison confounds borrower composition and is not assigned an interpretation.)

Third, the within-person schedule is inconsistent with a binding budget constraint as the driver. A liquidity constraint binds at the person-month—one budget—while Section 6 shows the same borrower simultaneously revealing an annual rate near 29 percent on the personal-loan margin, 12 percent on the auto margin, and 1 percent on the mortgage margin. Inability to pay pushes every obligation toward default together; margin-specific, horizon-ordered revealed rates are the structure of choice.

Fourth, the measured persistence of distress closes the remaining theoretical channel. Delinquency spells at the default margin cure at a rate of 0.26 percent per quarter (annualized persistence 0.992; Appendix D.2). Liquidity narratives built on transitory shocks contradict this permanence directly, and permanent-state narratives operate through marginal utility—precisely the correction λ, which at the measured persistence equals one.

Quantitatively, the two channels admit only small adjustments. Correcting the sufficient-statistics formula for the income process measured in the panel changes the average β by less than 0.01 (the annual rate moves from 27 percent to no less than 25 percent under any admissible calibration; Appendix D), and the borrowing-rate lower bounds of Section 8—whose bias runs in the opposite direction—independently indicate double-digit annual discounting for the borrowing population. What survives is the narrow claim that some share of defaults are true cash-flow collapses, bounded above by roughly 35 percent. The simulation harness shows that this residual cannot manufacture the headline estimate: across preference calibrations spanning $\beta = 0.55$ to 0.99, contamination at the ceiling share depresses the estimated discount factor by at most five points, and no simulated combination of patience and contamination—including contamination of half of all defaults—produces an estimate as impatient as the one measured, because the simulated estimator attenuates toward patience in every world it is given. Within the examined model class, the measured impatience requires genuine preference-side impatience.

5 Heterogeneity in Estimated Discount Factors

The estimates in this section vary the sufficient statistic across observable subgroups. A preliminary decomposition disciplines their interpretation: across geographic clusters and state-cycle cells, 84–89% of the cross-group variance in log discount factors is attributable to the default-sensitivity ratio R alone, with rankings reproduced at rank correlation 0.97 when the annualization horizon is held fixed at its common mean. The maps and gradients below therefore reflect differences in default behavior itself.

A varying-coefficient implementation sharpens the score and income gradients. Interacting the two logits' payment and obligation terms with standardized score, standardized log income, and census region lets every default in the sample inform every cell, and the resulting three-year revealed rates carry standard errors of approximately 0.4 percentage points where cell-splitting yields multiples of that. The rates decline monotonically in credit score, from 26.7 percent in the lowest decile to 23.2 percent in the highest, and decline in income, from 25.8 percent in the bottom quartile to 21.8 percent in the top (Figure 3; Table 7). Both gradients are modest relative to the level of the estimates.

Table 7: Varying-coefficient margins

Panel A: three-year revealed annual rate (%) by credit-score decile										
	d1	d2	d3	d4	d5	d6	d7	d8	d9	d10
rate	26.7	25.8	25.3	25.0	24.6	24.3	24.0	23.7	23.4	23.2
	(0.4)	(0.4)	(0.4)	(0.4)	(0.4)	(0.4)	(0.4)	(0.4)	(0.4)	(0.4)
Panel B: by income quartile										
	q1	q2	q3	q4						
rate	25.8	24.6	23.5	21.8						
	(0.4)	(0.4)	(0.4)	(0.6)						

Figure 3: Three-year revealed discount rates by credit score and income

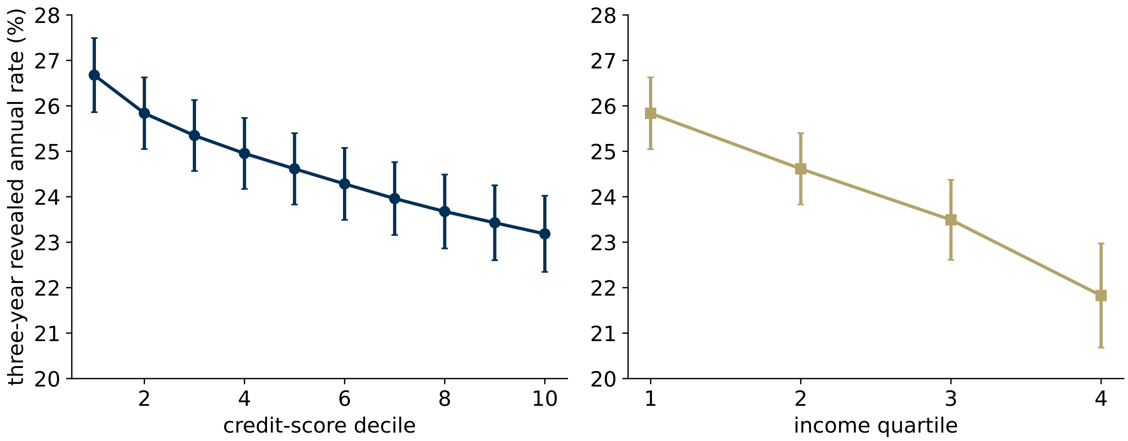

Notes: Rates from the varying-coefficient specification, evaluated at the median of the dimensions held fixed; bars are 95 percent parametric-bootstrap intervals. Replication: `Code/pace/49_interacted_logits.R` and `Code/make_paper_exhibits.py` from `results/vc_r_surface.csv`.

The preceding section establishes steep average impatience. But the distribution of β is at least as important as its mean. This section estimates β separately by credit score ventile, income ventile, and U.S. state. For each subgroup, I re-estimate the two logit models, recompute the sensitivity ratio, and obtain β with delta-method standard errors.

Before turning to the subgroup estimates, it is worth stating exactly what the pooled estimate of the previous section measures when the population is heterogeneous. Appendix C shows (Proposition 8) that the pooled statistic is, exactly, a weighted average of type-level statistics, $\hat{\beta} = \sum_i \omega_i s_i \beta_i^{stat} / \sum_i \omega_i s_i$, with weights equal to each type's total default sensitivity $s_i = f_i(y_i^*)\, u_i'(C_1^{P*}) / [u_i'(C_1^{P*}) - u_i'(C_1^{N*})]$: the mass of the type's members at its default margin, scaled by the responsiveness of its marginal defaulter. The pooled $\hat{\beta}$ is therefore always bounded by the smallest and largest type-level values—aggregation cannot manufacture a number outside the convex hull of group discount factors—and it upweights groups that sit near the default margin, in the same way that instrumental-variable estimates weight compliers. Numerical verification and magnitudes are in Appendix C; in two-type calibrations the wedge between the pooled statistic and the population-weighted mean is at most about 0.01. Two implications follow for this section. First, the subgroup estimates below are the primitive objects, and the full-sample number is their sensitivity-weighted mean rather than a population mean.

Second, because the weights favor households at the margin of default, the pooled estimate is the discount factor of precisely the population whose behavior default-related policy operates on.

Figure 4: Estimated β by Credit Score Ventile

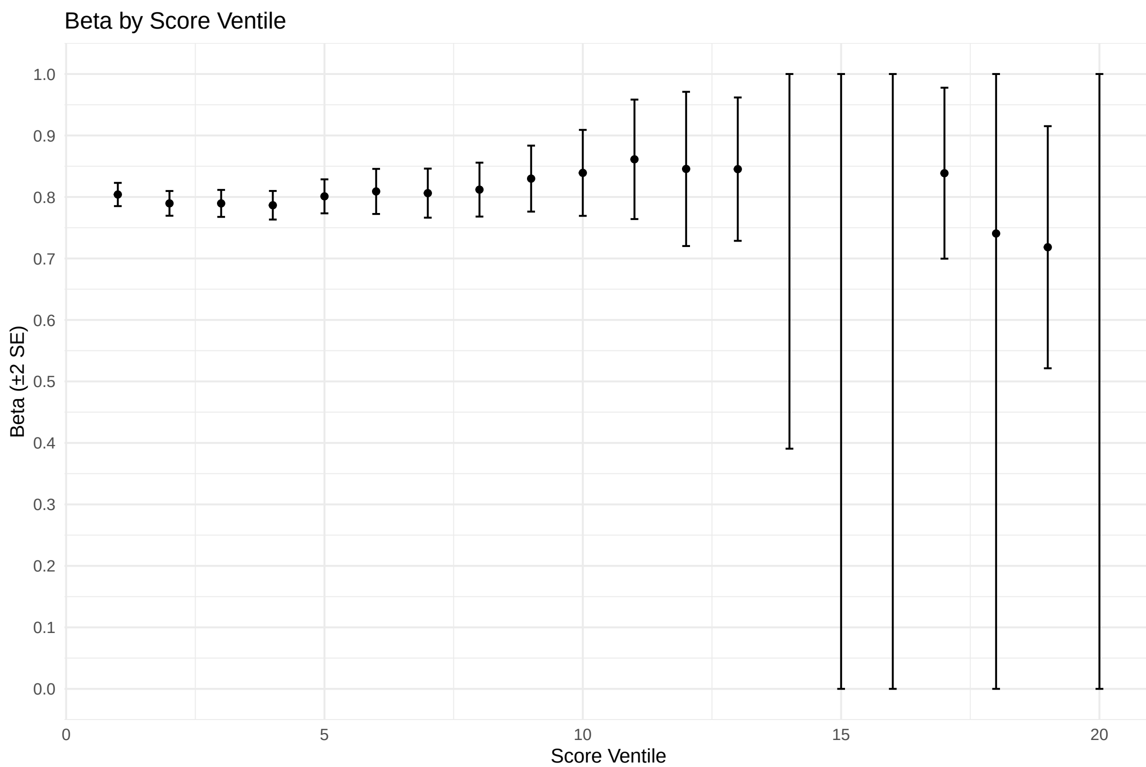

Notes: Error bars show ± 2 standard errors. Wide bands in the upper ventiles reflect few defaults; the varying-coefficient estimates of Figure 3 resolve the gradient these bands cannot.

Figure 5: Estimated β by Income Ventile

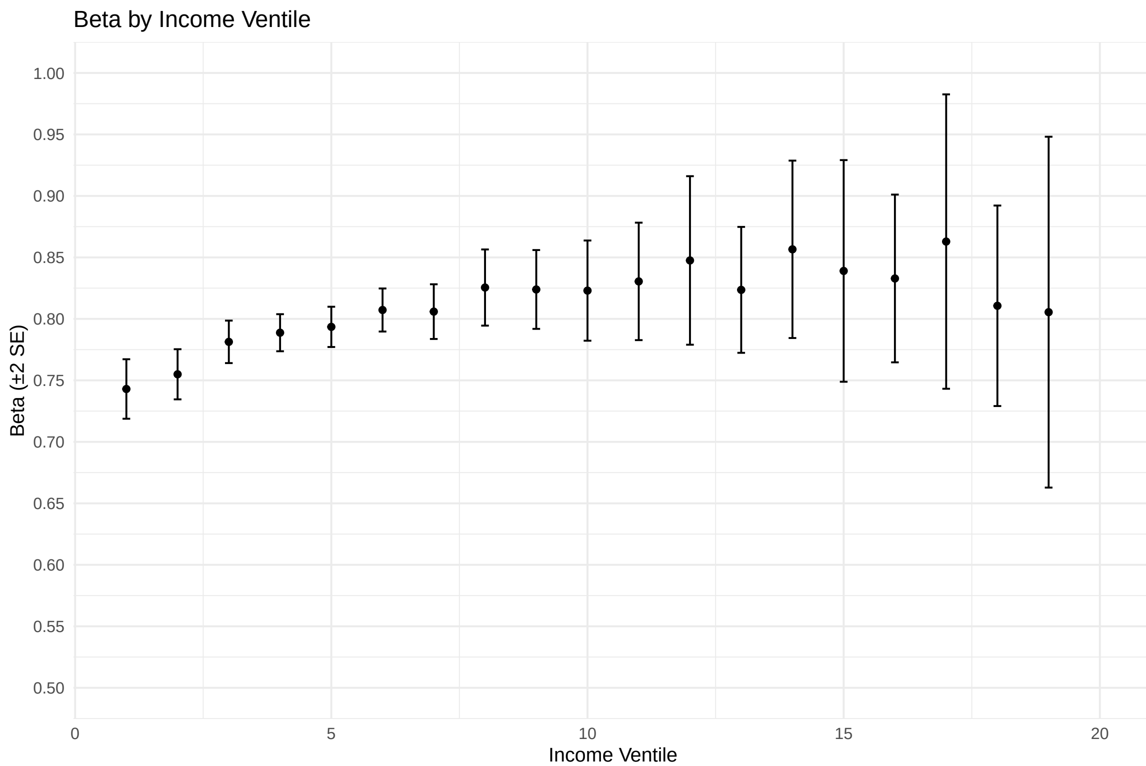

Notes: Error bars show ± 2 standard errors. Low-income borrowers are substantially more impatient.

Figure 6: Estimated β by State

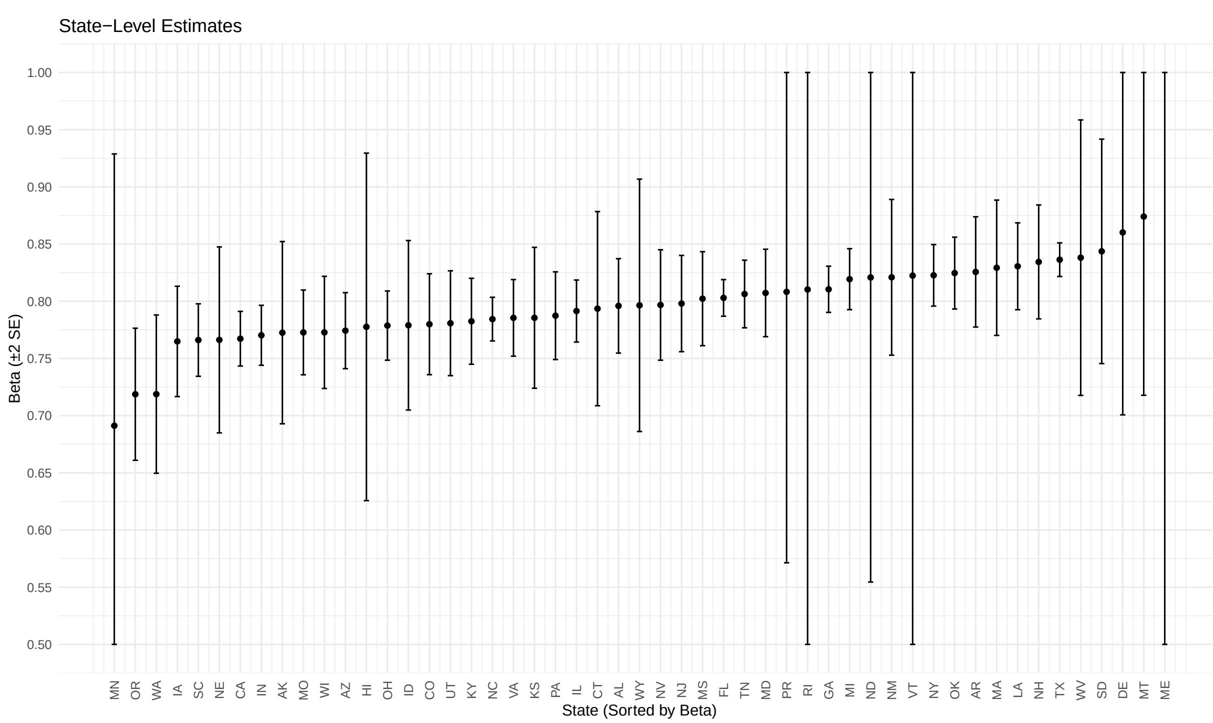

Table 8: State-Level Discount Factor Estimates

State	Beta	SE	N
MN	0.6912	0.1188	106299
OR	0.7187	0.0289	76697
WA	0.7188	0.0346	128864
IA	0.7649	0.0241	72523
SC	0.7661	0.0158	103495
NE	0.7662	0.0406	39562
CA	0.7673	0.0119	623313
IN	0.7702	0.0131	139600
AK	0.7725	0.0398	12478
MO	0.7728	0.0185	127183
WI	0.7728	0.0245	107439
AZ	0.7743	0.0166	130653
HI	0.7776	0.0760	20135
OH	0.7787	0.0151	237281
ID	0.7790	0.0370	36909
CO	0.7799	0.0221	107594
UT	0.7808	0.0229	70515
KY	0.7825	0.0188	83162
NC	0.7844	0.0095	214638
VA	0.7855	0.0167	164137
KS	0.7855	0.0308	57470
PA	0.7874	0.0192	250457
IL	0.7914	0.0136	221882
CT	0.7935	0.0424	58263
AL	0.7960	0.0206	106409
WY	0.7964	0.0552	13067
NV	0.7968	0.0241	56195
NJ	0.7980	0.0210	160247
MS	0.8023	0.0205	58887
FL	0.8029	0.0080	454937
TN	0.8063	0.0148	137186
MD	0.8073	0.0191	117419
PR	0.8082	0.1184	41139
RI	0.8103	0.1650	16993
GA	0.8105	0.0101	203167
MI	0.8193	0.0133	205098
ND	0.8208	0.1331	17054
NM	0.8209	0.0340	42475
VT	0.8224	0.4009	16030
NY	0.8227	0.0134	306696
OK	0.8246	0.0157	84997
AR	0.8256	0.0241	63014
MA	0.8293	0.0296	107804
LA	0.8306	0.0190	89300
NH	0.8344	0.0249	34025
TX	0.8363	0.0073	605128
WV	0.8381	0.0602	41744
SD	0.8436	0.0491	20326
DE	0.8602	0.0798	18860
MT	0.8740	0.0781	20799
ME	1.1407	0.4449	31992

Three patterns emerge. First, the score gradient is small relative to the score–default gradient (Figure 4). The binned equal-default splits detect no gradient within their wide bands; the varying-coefficient design, which pools power across cells, resolves a monotone decline of 3.5 percentage points from the bottom to the top decile. Under either reading, the variation in the preference parameter is an order of magnitude too small to account for the variation in default rates across scores, so income volatility, thin financial buffers, and differing costs of default must drive the score–default relationship.

Second, β rises with income (Figure 5). The varying-coefficient quartile gradient is four percentage points of the annual rate; the binned ventile estimates place the lowest-income ventiles substantially below the rest, and the difference between the first and third income ventiles is rejected at conventional levels. The pattern is consistent with tighter liquidity constraints, higher exposure to income shocks, and higher marginal utility of current consumption among low-income households.

Third, there is meaningful geographic variation (Figure 6, Table 8), with state-level estimates spanning a range of roughly 0.10 in β.

In the upper credit-score ventiles, the error bands on β are extremely wide. This is not a defect of the method but a feature: high-score borrowers rarely default, so the sensitivity of default to payment structure is imprecisely estimated. For these populations, default-based identification is unnecessary—their credit is not rationed, and their marginal borrowing rates directly reveal their discount rates.

6 A Term Structure of Revealed Discount Rates

The framework so far reports one point on each borrower's discount curve: the annual rate implied by the single horizon a contract offers. Appendix B shows that responses at multiple horizons identify the discount function itself (Proposition 4), and notes that the ideal per-horizon payment experiment has not been run. Different loan products, however, supply different horizons: extending the estimator to the major fixed-term consumer debt families in the same credit-bureau panel prices the default margin at maturities from under two years to nearly two decades.

Applying the identical estimator with a uniformly-defined default cost to each family yields a steeply declining term structure of revealed annual discount rates: approximately 31% for personal installment loans (mean default-to-maturity horizon 1.7 years), 12% for auto loans (3.0 years), and 1–1.5% for home-equity loans and first mortgages (17 years).[4] Selection across products does not drive the pattern: among borrowers who hold both contract types, the same

[4] These family-level estimates use the aggregate sensitivity-ratio variant of the estimator with the default cost defined identically (the outstanding balance at default) in every family, so that cross-family comparisons hold the cost measurement fixed. The auto-loan value under this variant (12%) sits below the baseline logit estimate (27%); the two variants differ in weighting and are not interchangeable across tables.

person's revealed annual rate is 28.9% on the personal-loan margin against 12.0% on the auto margin, and 11.4% on the auto margin against 1.0% on the mortgage margin.

Because two measured horizons point-identify quasi-hyperbolic preferences (Corollary 5), the dual-holder estimates also deliver the schedule's two-parameter summary. Writing $D(t) = \beta_{pb}\,\delta^t$ and inverting the pair of horizon factors yields $\beta_{pb} = 0.71$ with $\delta = 1.01$ from the auto–mortgage margins, and $\beta_{pb} = 0.55$ with $\delta = 1.09$ from the auto–personal-loan margins: a large present-bias wedge set against long-run discounting statistically indistinguishable from zero (δ estimates marginally above one reflect the measurement burden in long-horizon T and are not interpreted as negative long-run rates). Under this summary, the population object that heterogeneity analysis should target is the joint distribution of (β_{pb}, δ); estimating it at the group level from two family margins per cell is the natural extension of Section 5 and is left for the full panel.

Table 9: Within-person estimates across contract families

Dual-holder sample	Margin	Horizon (yrs)	$\widehat{\beta}$	annual rate	defaults
Auto × mortgage	Auto	2.9	0.898	11.4%	1,062,436
Auto × mortgage	First mortgage	16.8	0.990	1.0%	45,745
Auto × pers inst	Auto	3.0	0.893	12.0%	2,088,628
Auto × pers inst	Personal installment	1.7	0.776	28.9%	2,196,405

Notes: Each pair of rows estimates the sufficient statistic separately on the two contract families of borrowers who hold both, with the default cost defined as the outstanding balance at default in every family. Consumer identity is matched on hashed consumer keys. Replication: `Code/pace/46_within_person.R`.

A declining term structure is the signature of non-exponential discounting (Figure 7 and Table 9): short-horizon tradeoffs are priced at rates an order of magnitude steeper than long-horizon ones, within the same borrower, at market stakes. This is the pattern that quasi-hyperbolic and related present-biased specifications were built to describe (Laibson et al., 2024), measured here from field behavior rather than elicitation. Two caveats discipline the interpretation. First, horizon and stakes co-move across contract families—mortgages are both the longest and the largest obligations—so the declining schedule is also consistent with attention that scales with consequence; both readings reject a single constant discount rate, and separating them requires within-family variation in horizon at fixed loan size, which the auto data

Figure 7: The revealed discount curve, traced across loan families

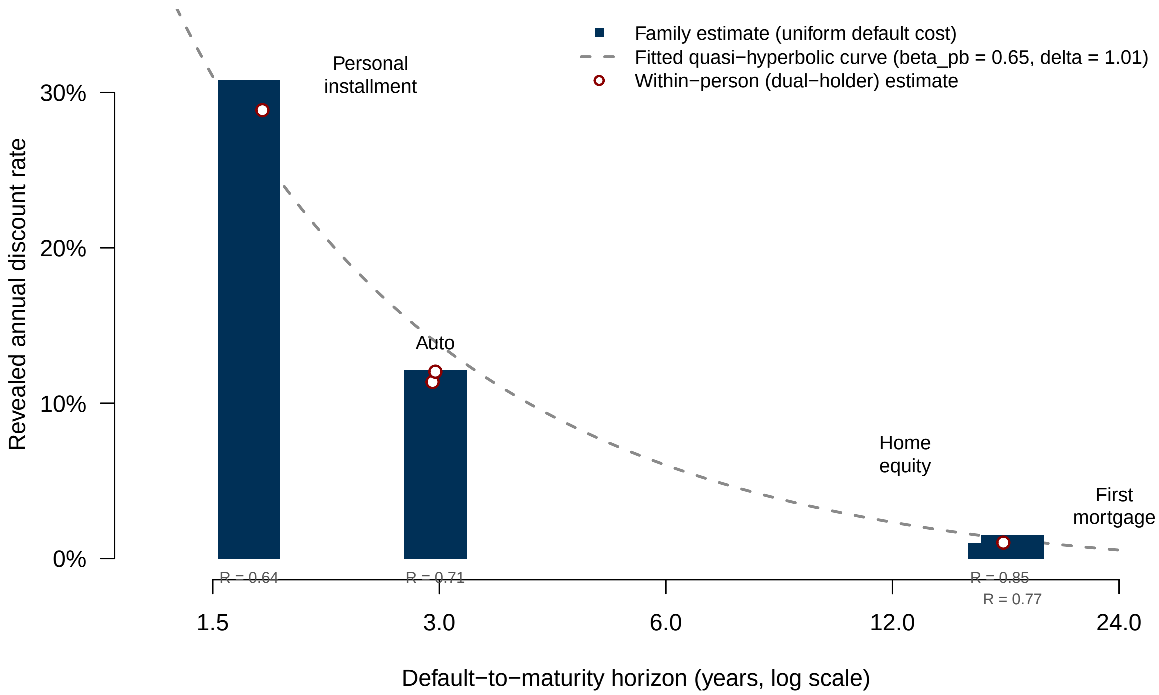

Notes: Each bar is the annualized revealed discount rate from applying the sufficient-statistics estimator to one fixed-term loan family, plotted at the family's mean default-to-maturity horizon (log scale); the default cost is defined identically (outstanding balance at default) in every family. The measured sensitivity ratio R is printed beneath each bar. Hollow markers are the within-person estimates from borrowers holding both contract types. The dashed line is the quasi-hyperbolic discount function $D(t) = \beta_{pb}\delta^t$ fitted to the four family points ($\beta_{pb} = 0.66$, $\delta = 1.01$); the estimates trace it point by point. Replication: `Code/fig_term_structure.R` from `results/consequence_norm_beta.csv` and `results/within_person_beta.csv`.

can support and is left for the estimation on the full panel. Second, the long-horizon estimates inherit the heaviest measurement burden in the default-to-maturity horizon, and the sensitivity ratio R is reported alongside every annualized figure so the measurement and the clock can be read separately.

7 Comparison with Prior Estimates

Table 10 places the estimates from this paper alongside prior discount rate estimates from the literature. Despite using fundamentally different methods, populations, and identification strategies, a consistent picture emerges: households who borrow and default—as opposed to the wealthy households who price financial assets—exhibit discount rates far steeper than standard calibrations assume.

Table 10: Discount Rate Estimates Across Studies

Study	Method	Population	β	Annual Rate
This paper				
Auto loans (full sample)	Sufficient statistics	U.S. auto borrowers	0.79	27%
Auto loans (bottom income quartile)	Varying-coefficient FE	Low-income borrowers	0.79	26%
SCF lower bounds (median)	Max borrowing rate	U.S. households	$\leq$0.91	$\geq$9%
Quasi-experimental (sufficient statistics applied to published estimates)				
Ganong and Noel (2020)	RD (HAMP mortgages)	Distressed mortgagors	0.76	32%
Indarte (2023)	RKD + ARM IV (bankruptcy)	Homeowners at risk	0.78	29%
Ganong and Noel (2019)	Structural estimation	General consumers	0.52–0.94	6–92%
Laboratory experiments				
Harrison et al. (2002)	Field experiment	Danish adults	0.72	28%
Coller and Williams (1999)	Lab experiment	U.S. students	0.83	20%
Laibson et al. (2024)	Meta-analysis	Various	0.50–0.85	18–100%
Cohen et al. (2020)	Survey of 112 studies	Various	0.82 (median)	22% (median)
Asset market calibrations				
Standard macro calibration	Euler equation	Rep. agent	0.96	4%
Di Tella et al. (2024)	Zero-beta portfolio	Marginal investor	0.92	8%

Notes: β denotes the annual discount factor. Annual rate is $(1/\beta - 1) \times 100$. Estimates from Ganong and Noel (2020), Indarte (2023), and Dobbie and Song (2020) are computed by applying the sufficient-statistics formula of this paper to their reported comparative statics. The Indarte (2023) estimate uses the 83/17 moral-hazard/liquidity decomposition of bankruptcy filing (Tables 1 and 3), with NPV-adjusted coefficients ensuring comparable per-dollar units; $T = 7$ years (median remaining mortgage duration). Estimates from Ganong and Noel (2019) reflect the range between their estimated "low-β" and "high-β" types. Laboratory estimates from Laibson et al. (2024) reflect the range of hyperbolic (β, δ) estimates across multiple studies. The Cohen et al. (2020) row reports the median and mean across 112 point estimates extracted from 222 studies surveyed in that paper; 131 of those studies are lab-based and only 24 use publicly available data.

The estimates from this paper fall in the range implied by quasi-experimental studies and field experiments, and are far from the $\beta \approx 0.96$ of standard macro calibrations. The close agreement between the auto loan estimate ($\beta = 0.79$), the Ganong–Noel HAMP estimate ($\beta = 0.76$), and the Indarte bankruptcy estimate ($\beta = 0.78$) is particularly striking: three different populations (auto borrowers, distressed mortgagors, homeowners at risk of bankruptcy), three different debt types (auto loans, modified mortgages, dischargeable debt), and three fundamentally different identification strategies (logit on administrative data, regression discontinuity, and regression kink design combined with an ARM instrumental variable) all recover annual discount factors in the range 0.76–0.79, implying annual discount rates of 27–32%.

Figure 8 places this paper's estimate in the context of the 112 point estimates cataloged by Cohen et al. (2020). The median prior estimate is 22% per year (implied $\beta = 0.82$), close to this paper's estimate of 27% ($\beta = 0.79$). The consistency across methods—sufficient statistics applied to administrative data, regression discontinuities, randomized experiments, and laboratory elicitations—suggests that high discount rates among borrowing populations are a regularity that no single identification strategy produces on its own. The contribution of this paper relative to most prior estimates is that it uses large-stakes revealed preference from administrative records at population scale, rather than small-sample laboratory or survey elicitations.

Figure 8: Distribution of Discount Rate Estimates Across 112 Prior Studies

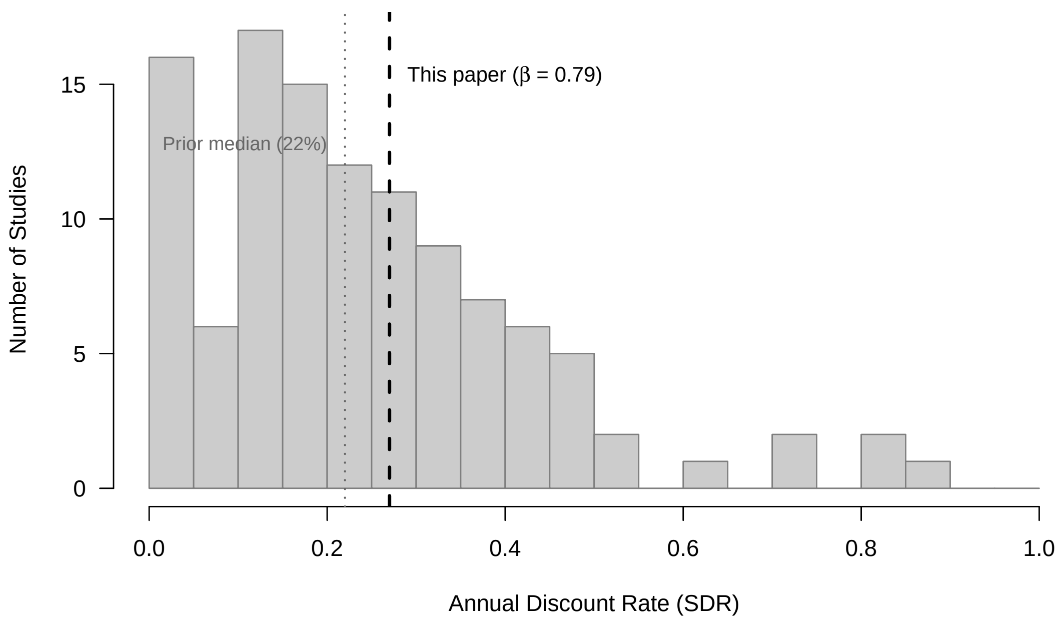

Notes: The histogram shows the distribution of annual subjective discount rate point estimates extracted from the 222-paper survey in Cohen et al. (2020). The dashed line marks this paper's estimate ($\beta = 0.79$, SDR = 27%). The dotted line marks the median of prior estimates (22%).

8 Lower Bounds from Observed Borrowing Rates

One might object that the discount rates estimated in the preceding sections are too high to be consistent with any kind of intertemporal optimization. This section provides a complementary and independent line of evidence showing that they are not. The bounds developed here also do double duty: because their bias runs opposite to the marginal-utility correction of Appendix D, the two methodologies bracket annual discount rates from both sides, confining them to roughly the 20–27% range under any admissible calibration. Using the interest rates at which households actually borrow as revealed-preference lower bounds on their discount rates, and adjusting for consumption-smoothing demand with a calibrated life-cycle income model, I show that the implied lower bounds on impatience are broadly consistent with the sufficient-statistics estimates. This approach requires only that each household holds debt—a far weaker data requirement—and is subject to orthogonal identification concerns.

8.1 Borrowing Rates as Lower Bounds

The key assumption is that households face an upward-sloping credit supply curve and accept offers to borrow at progressively higher rates until the offered rate exceeds their true discount rate. Because credit products are discrete, the highest observed borrowing rate lies weakly below the true discount rate. Figure 9 illustrates the logic schematically.

Figure 9: Illustration of Lower Bound Approach

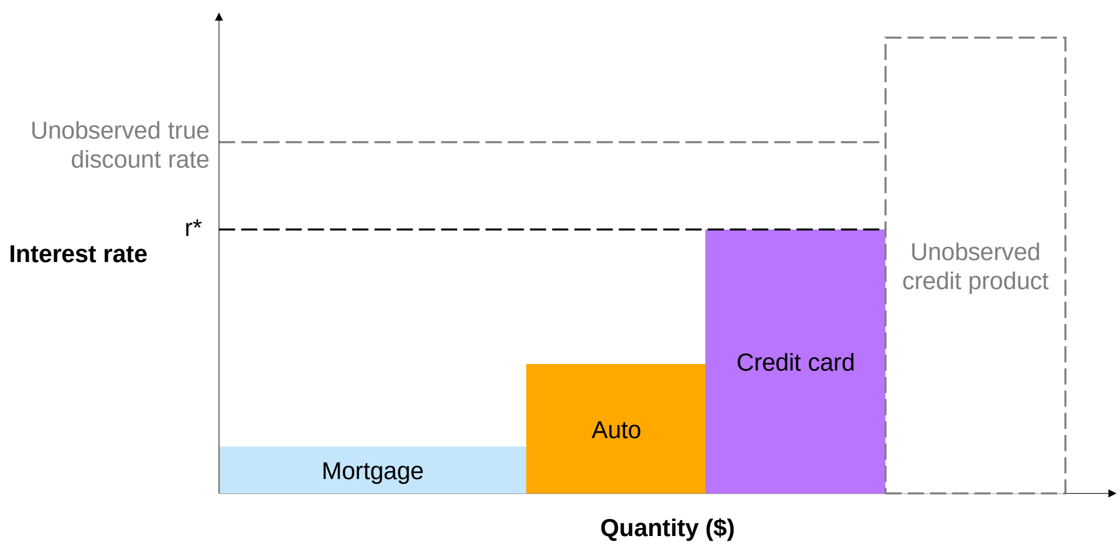

Notes: This schematic illustrates the framework for interpreting observed borrowing rates as lower bounds on discount rates. Because credit supply is upward sloping and products are discrete, the true discount rate lies at or above the highest observed borrowing rate.

For each household in the SCF, I calculate the maximum interest rate on debt that is both currently being repaid and actively accruing interest—excluding, for example, the 67% of student loans currently in forbearance or deferment and any debt with a promotional 0% rate. I include credit card debt, as it is the marginal credit product for most households at short horizons. Figure 10 shows these distributions.

Figure 10: Distribution of Maximum Interest Rate per Household

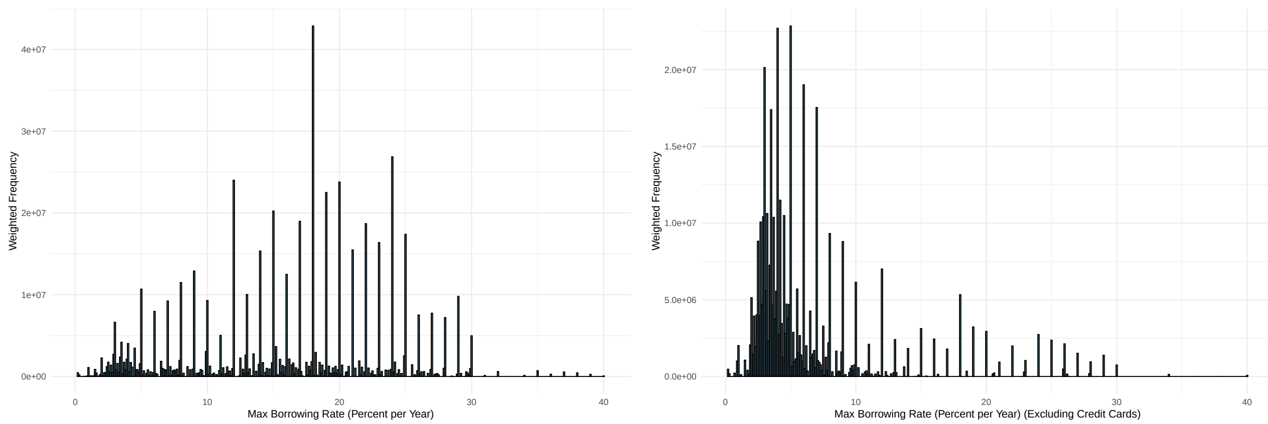

Notes: Panel A includes credit card debt; Panel B excludes it. The highest interest rate on debt currently accruing interest is shown. 86% of Americans pay interest on debt.

The modal maximum borrowing rate is 18%—the dominant default credit card rate, a legacy of state usury laws (Hall (2024)). Even excluding credit cards, over 25% of households borrow above the 5% treasury rate.

8.2 Motivating Evidence: Why Asset Prices Are Uninformative About Most Households

Asset prices correctly identify the discount factor of the marginal investor. But the marginal investor is far from typical. Figure 11 shows that only 12% of households hold any risk-free financial assets, with the top 1% holding 75%. Figure 12 shows the full distribution of financial assets: the bottom third holds nothing, and the top 1% holds 47% of all financial assets. If impatience is negatively correlated with wealth, as is natural to assume, then the marginal investor will be far more patient than the median household. Asset prices reveal her preferences, not theirs.

Figure 11: Risk-Free Asset Holdings Across the Distribution

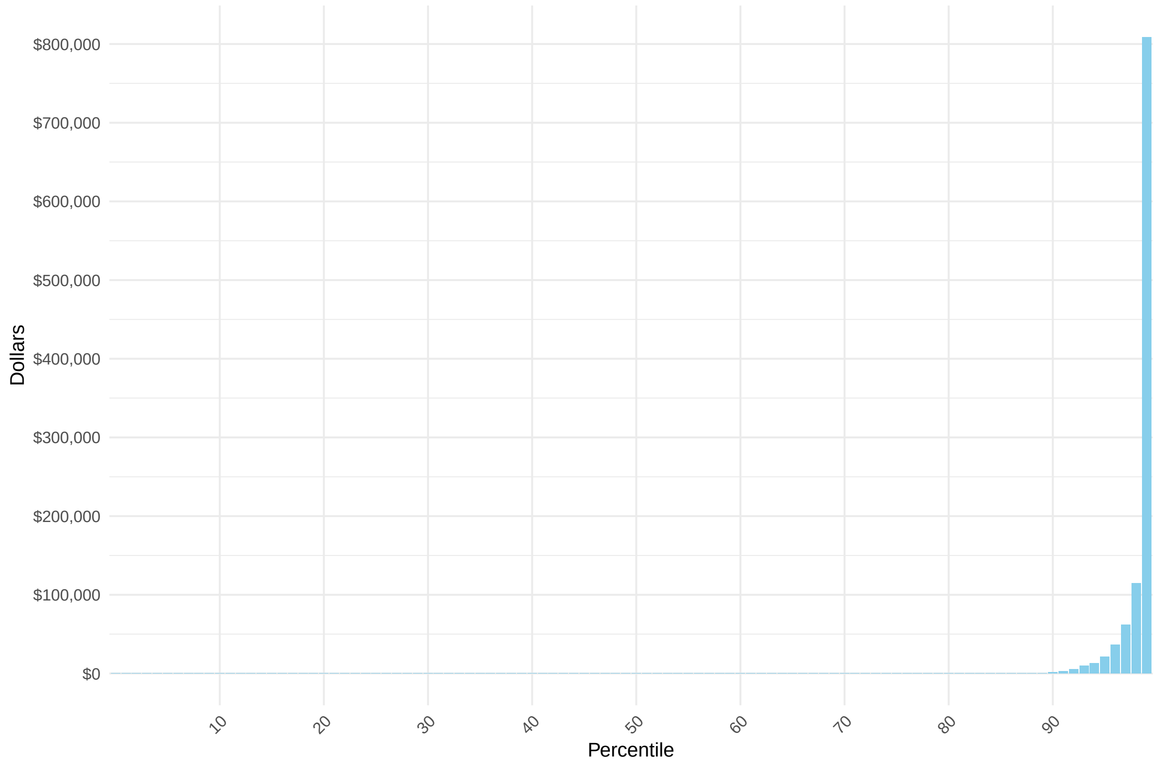

Notes: SCF households sorted by percentile of total financial assets (scaled by survey weights). The y-axis reports average holdings within percentile bins.

Figure 12: Distribution of All Financial Assets by Type

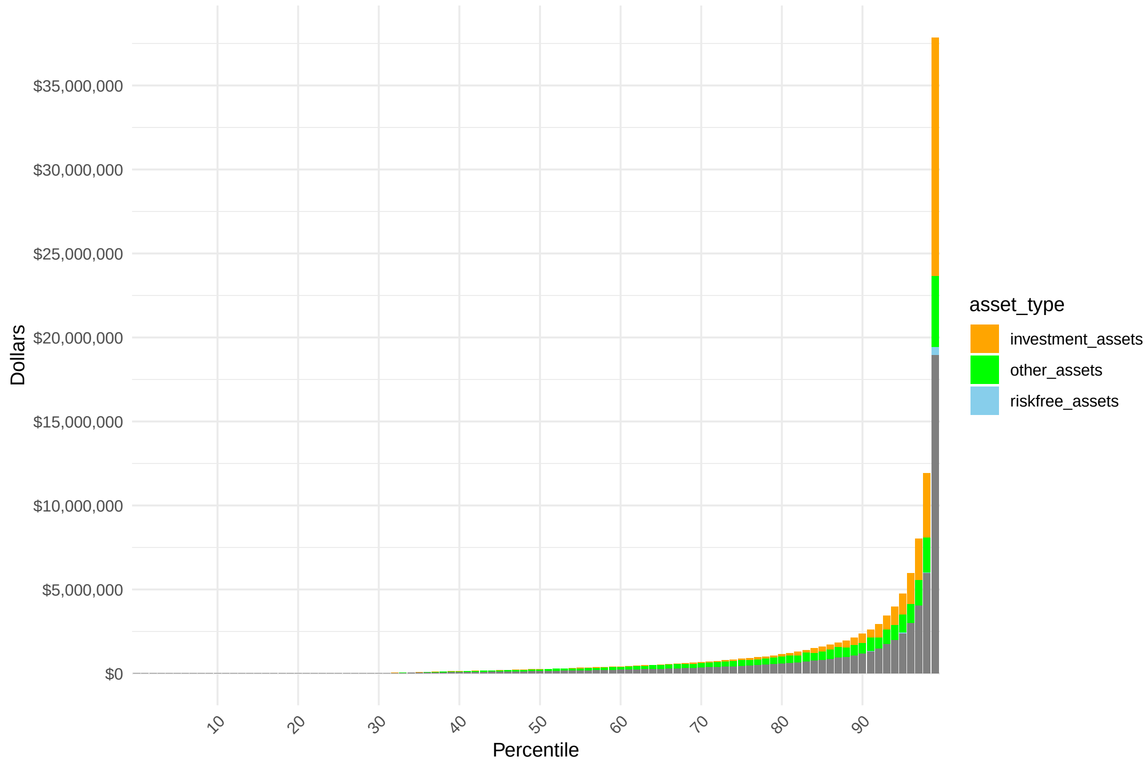

Notes: Assets are classified as risk-free, investment, or other. The bottom third holds no financial assets. The top 1% holds 47% of all financial assets.

Figure 13: Liabilities Across the Asset Distribution

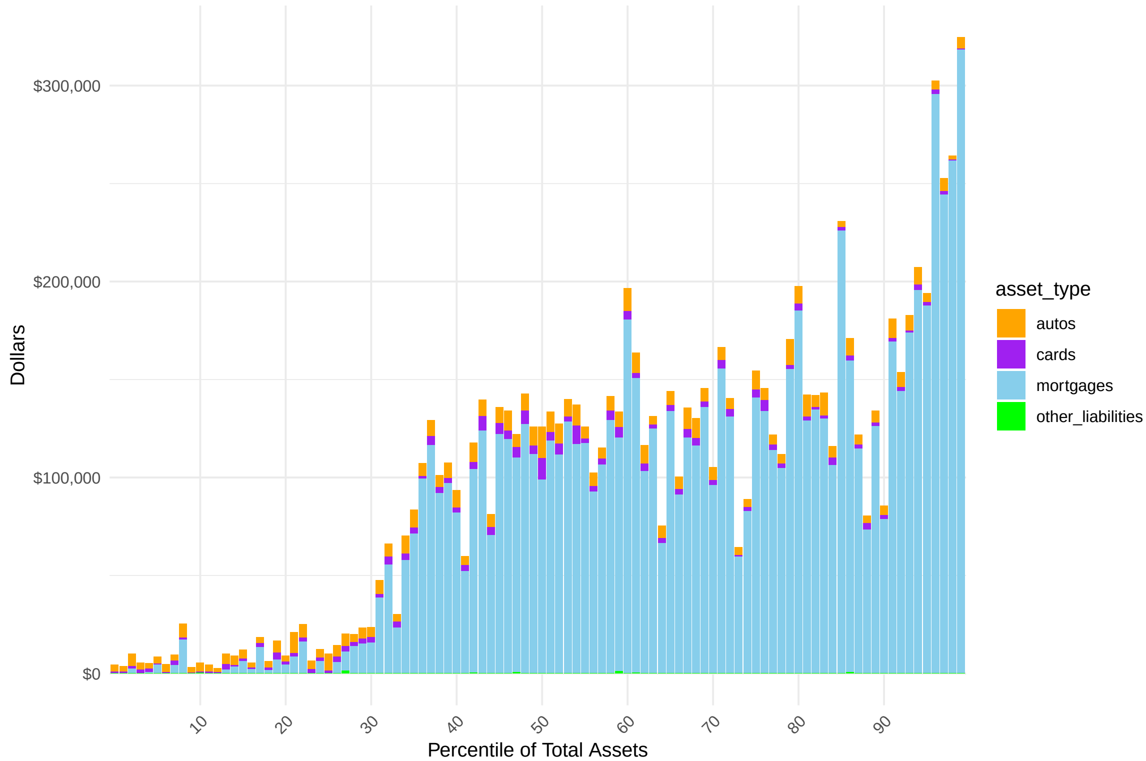

Notes: Same percentile ranks as above, but the y-axis shows average liabilities. Mortgage debt is widespread through the middle; unsecured debts cluster at the bottom.

The contrast between Figures 12 and 13 motivates the paper's core approach: the relevant revealed-preference margin for most households lies on the liability side of their balance sheets. Credit, not investment, is the domain in which they express intertemporal tradeoffs.

8.3 From Interest Rates to Discount Rates: Adjusting for Consumption Smoothing

Observed borrowing rates are lower bounds on r, the household's intertemporal demand—not on β, the pure impatience parameter. The Euler equation decomposes these:

$$\underbrace{r}_{\text{intertemporal demand}} = \underbrace{-\log(\beta)}_{\text{impatience}} + \underbrace{\log E_t\left[U'\left(C_{t+1}\right)\right] - \log U'\left(C_t\right)}_{\text{consumption-smoothing demand}} \tag{11}$$

A household may borrow at high rates not because it is impatient, but because it expects income to grow.[5] To assess this, I calibrate the life-cycle income model of Guvenen et al. (2021) to each SCF household, using their age, current income, and employment status to project expected income growth at 1, 5, and 10-year horizons.

The model produces reasonable estimates: implied income growth expectations correlate with households' subjective reports of optimism (Table 11).

Table 11: Model-Implied Income Growth by Subjective Optimism

Optimism	$E[\Delta L_{t+1}]$	N
>Inflation	0.032	2,900
=Inflation	0.028	5,561
<Inflation	0.027	4,609

Notes: Optimism is from SCF question x7364. $\Delta L_{t+1} \equiv \log L_{t+1} - \log L_t$.

However, income growth expectations explain almost none of the variation in borrowing rates. Table 12 reports regressions of maximum borrowing rates on model-implied expected income growth and subjective optimism. The R^2 is below 2% in all specifications.

Table 12: Income Growth Expectations and Borrowing Rates

Variables	max(r)	max(r ex cards)	max(r)	max(r ex cards)
$E[\Delta L_{t+1}]$	43.44***	−27.06***	17.75***	−7.612***
	(8.940)	(4.896)	(3.415)	(2.517)
Optimism	−0.104	0.035	−0.074	0.089
	(0.259)	(0.229)	(0.106)	(0.078)
Constant			14.48***	4.585***
			(0.264)	(0.195)
Fixed-effects				
Age	Yes	Yes	No	No
Fit statistics				
Observations	13,069	13,069	13,069	13,069
R^2	0.019	0.022	0.002	0.001

Signif. Codes: ***: 0.01, **: 0.05, *: 0.1

This is a key result. Under CRRA utility, the Euler equation implies $r = -\log\beta + \gamma \cdot$

[5] Steinbeck's diagnosis of why Depression-era America produced so few self-described proletarians makes the same point about expectations: "Everyone was a temporarily embarrassed capitalist" (Steinbeck, 1960). The more commonly quoted "temporarily embarrassed millionaires" is a later paraphrase (Wright, 2004).

$E_t[\Delta \log C_{t+1}]$, so the coefficient on expected income growth in a regression of borrowing rates should equal γ.[6] Instead, the R^2 is negligible, and the estimated coefficients are either wrong-signed (excluding credit cards) or implausibly large (including credit cards). Heterogeneity in impatience, not in income expectations, is the dominant source of variation in household borrowing rates. Subjective optimism, controlling for model-implied objective expectations, adds no explanatory power.

Using the life-cycle model to adjust for consumption smoothing under hand-to-mouth consumption with $\gamma = 1.9$ (Laibson et al. (2024)), the median implied upper bound on β is 0.91 (Figure 14), indicating at least 9% annual discounting. This estimate is conservative: even with $\gamma = 5$, which is at the upper end of estimates considered reasonable in the literature, the implied upper bound on β remains below 0.95 for most households. Approximately 2% of households have implied β below 0.75, indicating at least 25% annual discounting.

[6] This specification assumes log utility ($\gamma = 1$) when entering expected income growth linearly. More generally, with $\gamma \neq 1$, the coefficient should equal γ, and the R^2 should be correspondingly higher for any reasonable γ. For γ up to 5 or even 10, the consumption-smoothing term remains far too small to rationalize the observed borrowing rates.

Figure 14: Distribution of Implied Upper Bounds on β

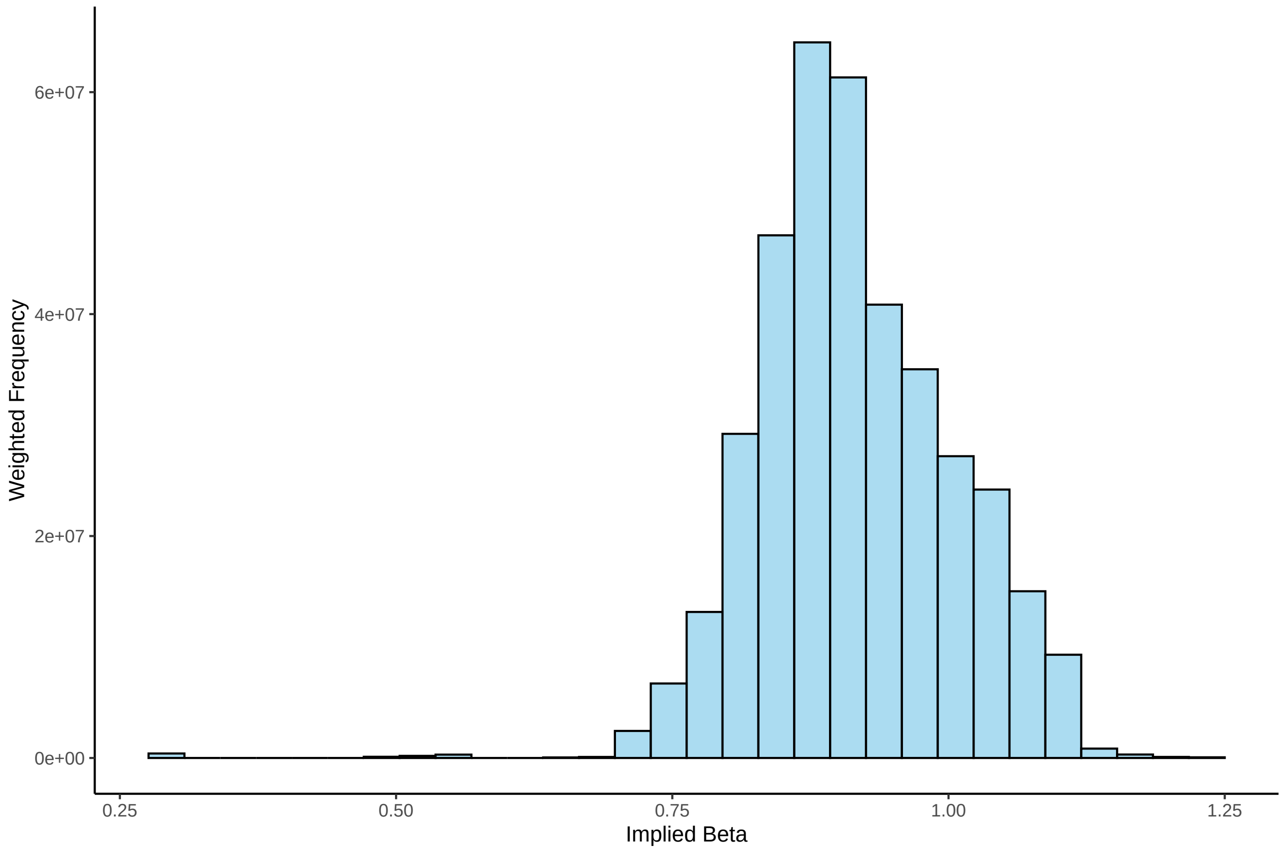

Notes: Upper-bound βs calculated using Equation (11) with $\gamma = 1.9$.

9 Discussion

9.1 What Does β Capture?

The sufficient-statistics estimate of $\beta = 0.79$ is a revealed-preference object, not a primitive preference parameter. Several forces are subsumed within it:

1. Pure time preference: the intrinsic rate at which the household discounts future utility.

2. Liquidity concerns: households facing binding liquidity constraints will behave as if more impatient, even if their underlying time preference is moderate.

3. Probability distortions: if borrowers overweight the likelihood of favorable future income

realizations, they will default less strategically than their true discount rate would predict.

4. Beliefs about future income: pessimism about future earnings would inflate the effective discount rate by reducing the perceived benefit of continued repayment.
5. Dynamic exclusion costs: if the perceived cost of default (σ) is lower than assumed—for example, because borrowers underestimate credit-market exclusion—the estimated β will be biased downward.

This conflation may be a feature rather than a defect. The paper's goal is revealed-preference measurement: what effective rate do households discount the future at, as revealed by their behavior? For many applications in public finance and credit policy, this reduced-form object is precisely the relevant one. A government evaluating whether short-term payment relief will prevent default does not need to decompose β into its constituent channels—it needs to know how borrowers actually trade off present and future.

That said, the decomposition matters for normative analysis. If a low β reflects liquidity constraints rather than preferences, then policies that relax those constraints (emergency savings programs, income insurance) could raise effective patience without changing underlying preferences. Separating these channels requires additional variation—for example, exogenous shifts in liquidity holding income constant, or experimental manipulation of beliefs about future income.

The evidence in this paper provides some guidance. The finding that β varies sharply with income but not with credit score is more consistent with liquidity-driven impatience than with a purely taste-based explanation: income determines the tightness of constraints, while credit score (conditional on income) captures historical repayment behavior and risk characteristics.

9.2 Rate of Time Preference vs. Discount Rate

It is important to distinguish between two related but distinct objects. The rate of time preference is a primitive preference parameter governing how individuals intrinsically value present versus future utility. The discount rate is the rate at which individuals actually discount the future, reflecting both time preference and the economic environment (liquidity constraints, income expectations, uncertainty). The estimates in this paper identify the latter. The former requires additional structural assumptions or variation to recover.

9.3 Extensions: Non-Time-Separable Utility

The baseline model assumes time-separable utility. In the Epstein-Zin-Kreps-Porteus (EZKP) framework, utility is:

$$U_t = \{(1-\beta)C_t^{1-\frac{1}{\psi}} + \beta \left(E_t U_{t+1}^{1-\gamma}\right)^{\frac{1-\frac{1}{\psi}}{1-\gamma}}\}^{\frac{1}{1-\frac{1}{\psi}}} \tag{12}$$

Under lognormal assumptions, the risk-free rate becomes:

$$r_f = -\log\beta + \frac{1}{\psi}E_t\left[\Delta \log c_{t+1}\right] - \frac{\theta}{2\psi^2}\sigma_c^2 + \frac{\theta-1}{2}\sigma_w^2$$

The main effect is that the expected-income-growth term is now scaled by $1/\psi$ (the inverse IES) rather than γ (risk aversion). To assess robustness, consider the range of values for $1/\psi$ that the literature considers reasonable. Even among advocates of EZKP preferences, estimates of ψ cluster between $2/3$ and 2 (Epstein and Zin (1991); Bansal and Yaron (2004)), implying $1/\psi \in [0.5, 1.5]$. For a household borrowing at 20% per year, even with $1/\psi = 1.5$ and expected consumption growth of 3% per year, the implied β is at most $\exp(-(0.20 - 1.5 \times 0.03)) \approx 0.86$—still far below standard calibrations. The consumption growth required to rationalize a 20% borrowing rate with $\beta = 0.96$ would be $(0.20 + \log(0.96))/1.5 \approx 16\%$ per year, which is implausible for any sustained period. Since the R^2 of income growth in explaining borrowing rates is negligible (Table 12), this adjustment has minimal quantitative impact on the estimates.

10 Conclusion

This paper develops and applies a new method for measuring household discount rates. The key intellectual move is to replace "marginal willingness to save"—the traditional identifying variation from asset markets—with "marginal willingness to avoid default." In an endogenous default model, the relative sensitivity of default to short-run versus long-run payment obligations is a sufficient statistic for the discount factor. The identifying identity requires no functional form for utility and no consumption data; the one approximation involved—a ratio of marginal utilities—is disciplined by measuring its key input, the persistence of the marginal borrower's distress state, directly in the panel (annualized 0.992 across 21.4 million delinquency transitions); at the measured persistence the approximation is nearly exact, and under any admissible literature calibration it moves the headline estimate by less than two percentage points (Appendix D). It identifies the discount factor for precisely the households whose preferences asset-pricing data cannot reveal.

The empirical findings are consistent across settings. Across auto loans, mortgage modifications, and bankruptcy filings, borrowing households exhibit annual discount factors near 0.76–0.79, implying discount rates in the 27–32% range at short horizons—substantially above standard calibrations and bracketed from below by the borrowing-rate evidence. Revealed rates decline modestly and monotonically with income and with credit score, and the score gradient is far too small to account for the score–default relationship, which is therefore mostly attributable to income risk and costs of default. A complementary lower-bound analysis using SCF borrowing rates, adjusted for consumption-smoothing demand, corroborates these magnitudes and shows that income growth expectations explain almost none of the cross-sectional variation in borrowing rates.

These findings have practical implications. For credit policy, they imply that short-term payment relief will be far more effective at preventing default than equivalent-cost principal reductions. For public finance, they suggest that the gap between official social discount rates and the revealed preferences of many citizens is large, with consequences for the evaluation of policies whose costs and benefits are separated in time. For structural modeling, they provide new moments—distributional estimates of time preferences linked to observables—that can discipline heterogeneous-agent models in which the discount factor distribution has previously been calibrated or indirectly inferred.

References

Arrow, Kenneth J., Maureen L. Cropper, Christian Gollier, Ben Groom, Geoffrey M. Heal, Richard G. Newell, William D. Nordhaus, Robert S. Pindyck, William A. Pizer, Paul R. Portney, Thomas Sterner, Richard S. J. Tol, and Martin L. Weitzman, "Should Governments Use a Declining Discount Rate in Project Analysis?," Review of Environmental Economics and Policy, July 2014, 8 (2), 145–163.

Bansal, Ravi and Amir Yaron, "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," The Journal of Finance, 2004, 59 (4), 1481–1509. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-6261.2004.00670.x.

Blundell, Richard, Luigi Pistaferri, and Ian Preston, "Consumption Inequality and Partial Insurance," American Economic Review, 2008, 98 (5), 1887–1921.

Board of Governors of the Federal Reserve System, "Survey of Consumer Finances," 2022. Public microdata.

Chetty, Raj, "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods," Annual Review of Economics, September 2009, 1 (Volume 1, 2009), 451–488.

Cohen, Jonathan, Keith Marzilli Ericson, David Laibson, and John Myles White, "Measuring Time Preferences," Journal of Economic Literature, June 2020, 58 (2), 299–347.

Coller, Maribeth and Melonie B. Williams, "Eliciting Individual Discount Rates," Experimental Economics, December 1999, 2 (2), 107–127.

DeMarzo, Peter, Zhiguo He, and Fabrice Tourre, "Sovereign Debt Ratchets and Welfare Destruction," Journal of Political Economy, October 2023, 131 (10), 2825–2892.

Dobbie, Will and Jae Song, "Targeted Debt Relief and the Origins of Financial Distress: Experimental Evidence from Distressed Credit Card Borrowers," American Economic Review, April 2020, 110 (4), 984–1018.

Epstein, Larry G. and Stanley E. Zin, "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis," Journal of Political Economy, April 1991, 99 (2), 263–286.

Flodén, Martin and Jesper Lindé, "Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance?," Review of Economic Dynamics, 2001, 4 (2), 406–437.

Ganong, Peter and Pascal Noel, "Consumer Spending during Unemployment: Positive and Normative Implications," American Economic Review, July 2019, 109 (7), 2383–2424.

_ and _ , "Liquidity versus Wealth in Household Debt Obligations: Evidence from Housing Policy in the Great Recession," American Economic Review, October 2020, 110 (10), 3100–3138.

Guvenen, Fatih, "Reconciling conflicting evidence on the elasticity of intertemporal substitution: A macroeconomic perspective," Journal of Monetary Economics, October 2006, 53 (7), 1451–1472.

_ , "An Empirical Investigation of Labor Income Processes," Review of Economic Dynamics, 2009, 12 (1), 58–79.

_ , Fatih Karahan, Serdar Ozkan, and Jae Song, "What Do Data on Millions of U.S. Workers Reveal About Lifecycle Earnings Dynamics?," Econometrica, 2021, 89 (5), 2303–2339.

Hall, Joseph, "Credit Cards and Retail Firms: Historical Evidence from the U.S.," Working Paper, 2024.

Harrison, Glenn W, Morten I Lau, and Melonie B Williams, "Estimating Individual Discount Rates in Denmark: A Field Experiment," American Economic Review, November 2002, 92 (5), 1606–1617.

Hendren, Nathaniel and Ben Sprung-Keyser, "A Unified Welfare Analysis of Government Policies," The Quarterly Journal of Economics, March 2020, 135 (3), 1209–1318.

House, The White, "Biden-Harris Administration Releases Final Guidance to Improve Regulatory Analysis | OMB," November 2023.

Indarte, Sasha, "Moral Hazard versus Liquidity in Household Bankruptcy," The Journal of Finance, October 2023, 78 (5), 2421–2464.

Laibson, David, Sean Chanwook Lee, Peter Maxted, Andrea Repetto, and Jeremy Tobacman, "Estimating Discount Functions with Consumption Choices over the Lifecycle," The Review of Financial Studies, July 2024, p. hhae035.

Lustig, Hanno and Yi-Li Chien, "The Market Price of Aggregate Risk and the Wealth Distribution," February 2005.

Rabin, Matthew, "Risk Aversion and Expected-Utility Theory: A Calibration Theorem," Econometrica, 2000, 68 (5), 1281–1292.

Steinbeck, John, "A Primer on the Thirties," Esquire, June 1960.

Storesletten, Kjetil, Christopher I. Telmer, and Amir Yaron, "Cyclical Dynamics in Idiosyncratic Labor Market Risk," Journal of Political Economy, 2004, 112 (3), 695–717.

Tella, Sebastian Di, Benjamin Hébert, Pablo Kurlat, and Qitong Wang, "The Zero-Beta Interest Rate," 2024.

Wright, Ronald, A Short History of Progress, Toronto: House of Anansi Press, 2004.

Appendices

A Proofs

This appendix provides a complete derivation of the comparative statics in equation (5).

A.1 Setup and Default Threshold

The borrower defaults in period 1 by comparing:

$$V_1^P(y_1) = u(y_1 - m) + \beta \left[\int_{y_2^*}^{\infty} u(y_2 - (L - m)) dF(y_2 \mid y_1) + \int_0^{y_2^*} (u(y_2) - \sigma) dF(y_2 \mid y_1) \right]$$

and

$$V_1^N(y_1) = u(y_1) + \beta(u(y_2) - \sigma).$$

The threshold y_1^* satisfies $V_1^P(y_1^*) = V_1^N(y_1^*)$, and $p_1 = F(y_1^*)$, so:

$$\frac{\partial p_1}{\partial x} = f(y_1^*) \frac{\partial y_1^*}{\partial x}, \qquad x \in \{L, m\}. \tag{13}$$

A.2 Differentiating the Threshold Condition

Defining $\Delta(y_1; L, m) \equiv V_1^P(y_1) - V_1^N(y_1) = 0$ at $y_1 = y_1^*$, implicit differentiation gives:

$$\frac{\partial y_1^*}{\partial x} = -\frac{\partial \Delta / \partial x}{\partial \Delta / \partial y_1}.$$

A.3 Derivative with Respect to y_1

$$\frac{\partial}{\partial y_1} \Delta(y_1; L, m) = u'(y_1 - m) - u'(y_1) \equiv u'(C_1^P) - u'(C_1^N) < 0.$$

A.4 Derivative with Respect to L

Only the continuation value depends on L:

$$\frac{\partial \Delta}{\partial L} = -\beta(1 - p_2) E\left[u'(C_2^P(y_2)) \,\middle|\, y_2 > y_2^* \right].$$

Combining:

$$\boxed{\frac{\partial p_1}{\partial L} = f(y_1^*) \beta (1 - p_2) \frac{E[u'(C_2^P(y_2)) \mid y_2 > y_2^*]}{u'(C_1^{P*}) - u'(C_1^{N*})}.}$$

A.5 Derivative with Respect to m

The payment m affects both current consumption and the future balance:

$$\frac{\partial \Delta}{\partial m} = -u'(C_1^P) - \frac{\partial \Delta}{\partial L}.$$

Therefore:

$$\boxed{\frac{\partial p_1}{\partial m} = f(y_1^*) \frac{u'(C_1^P)}{u'(C_1^{P*}) - u'(C_1^{N*})} - \frac{\partial p_1}{\partial L}.}$$

A.6 Discussion

Increasing L worsens future consumption while holding current consumption fixed; increasing m worsens current consumption while reducing future obligations. The ratio of these channels maps cleanly to β.

B The T-Period Model: Interior Default and the Discount Curve

This appendix extends the two-period model to T discrete periods, establishes that default probabilities remain strictly interior at every horizon, and shows that a menu of payment-schedule perturbations identifies the entire discount function—of which the two-period sufficient statistic of the main text is a single point.

B.1 Setup and recursive formulation

A contract specifies payments $m_t > 0$ at dates $t = 1, \ldots, T$, with balance $L = \sum_t m_t$. Income y_t is drawn i.i.d. from F with continuous density f and full support on $(0, \infty)$ (all results below go through with Markov income; see Section B.4). Consumption equals income net of any payment, utility u is strictly increasing, strictly concave, and satisfies $u(0^+) = -\infty$ (e.g. CRRA with $\gamma \geq 1$). Default is absorbing: a borrower who defaults at t consumes income thereafter and incurs a punishment flow $\sigma > 0$ in each period $s \in \{t, \ldots, T\}$. All agents are identical after T, so post-T values are normalized to zero.

Let $W_t = \sum_{k=0}^{T-t} \beta^k (\mathbb{E}u(y) - \sigma)$ denote the defaulted value and

$$V_t(y) = \max \left\{ u(y - m_t) + \beta \mathbb{E} V_{t+1},\ \ u(y) - \sigma + \beta W_{t+1} \right\}, \qquad V_{T+1} \equiv 0, \tag{14}$$

the value of good standing. Write the surplus from staying as $\Delta_t(y) = K_t - g_t(y)$, where $g_t(y) \equiv u(y) - u(y - m_t) > 0$ is strictly decreasing on (m_t, ∞) and $K_t \equiv \sigma + \beta(\mathbb{E}V_{t+1} - W_{t+1})$. The borrower defaults iff $y < y_t^* = g_t^{-1}(K_t)$, and $p_t = F(y_t^*)$. Two observations deliver the key recursion. First, the expected value of the default branch of (14) equals W_t itself. Second, $V_t(y) = [u(y) - \sigma + \beta W_{t+1}] + \Delta_t(y)^+$, so $\mathbb{E}V_t - W_t = \mathbb{E}[\Delta_t(y)^+] \geq 0$. Therefore

$$K_t = \sigma + \beta\, \mathbb{E}\left[(K_{t+1} - g_{t+1}(y))^+ \right], \qquad K_T = \sigma. \tag{15}$$

Proposition 1 (Interior default probabilities at every horizon) Under the maintained assumptions, $K_t \geq \sigma > 0$ for every t, the threshold $y_t^* \in (m_t, \infty)$ exists and is unique, and $p_t \in (F(m_t), 1) \subset (0, 1)$—for any horizon T and any payment schedule.

Proof. $K_t \geq \sigma$ is immediate from (15). Since $u(0^+) = -\infty$, g_t maps (m_t, ∞) onto $(0, \infty)$ bijectively, so y_t^* exists and is unique; $p_t < 1$ because y_t^* is finite and F has unbounded support, and $p_t > F(m_t) > 0$ because default is forced for $y \leq m_t$. □

The recursion (15) is the economic heart of the result: a borrower who stays today is not committing to repay the remaining $T - t$ installments—she buys one period of good standing plus the option to reconsider, and the max operator truncates the continuation comparison at zero. The remaining payment burden therefore cannot accumulate into the current default incentive. Dropping the option is not an innocuous simplification:

Lemma 2 (Committed-repayment comparisons degenerate) Suppose the date-t decision instead compares default against the value of committing to all remaining payments, so that $\tilde{K}_t = \sigma + \sum_{k=1}^{T-t} \beta^k (\sigma - \mathbb{E} g_{t+k}(y))$. If $\mathbb{E} g_s(y) > \sigma(1+\epsilon)$ for all s and some $\epsilon > 0$, there exists $\bar{k}(\epsilon, \beta)$ such that $\tilde{K}_t \leq 0$—and hence $\tilde{p}_t = 1$—whenever $T - t \geq \bar{k}$, while $\tilde{K}_T = \sigma$ keeps late-period probabilities interior.

Proof. Each term in the sum is bounded above by $-\epsilon \sigma \beta^k$; the partial sums cross $-\sigma$ once $\epsilon \sum_{k \leq \bar{k}} \beta^k > 1$, which occurs at finite $\bar{k}$ for any $\beta \in (0, 1]$. □

The committed-repayment model thus produces exactly the pattern of probabilities collapsing to 0 or 1 as the horizon grows. A second, independent source of degeneracy is a punishment that vanishes at maturity: if $\sigma = 0$, nothing enforces the final payment, $p_T = 1$, and by (15) the contract unravels backward period by period, as in a finitely repeated game. Any specification whose punishment is proportional to the remaining balance inherits this endgame problem; a flow or fixed punishment component disciplines it.

B.2 Payment-schedule sensitivities

Fix $(m_1, \ldots, m_T)$ and let $\Delta u_1' \equiv u'(y_1^* - m_1) - u'(y_1^*) > 0$.

Proposition 3 (Sensitivities to the payment schedule) Under exponential discounting,

$$\frac{\partial p_1}{\partial m_1} = f(y_1^*) \, \frac{u'(y_1^* - m_1)}{\Delta u_1'}, \tag{16}$$

$$\frac{\partial p_1}{\partial m_s} = f(y_1^*) \, \beta^{s-1} Q_s \, \frac{\mathbb{E}[u'(y - m_s) \mid y > y_s^*]}{\Delta u_1'}, \qquad s \geq 2, \tag{17}$$

where $Q_s = \prod_{k=2}^{s} (1 - p_k)$ is the probability of good standing through date s.

Proof. $\partial p_1 / \partial m_s = -f(y_1^*)(\partial \Delta_1 / \partial m_s) / \Delta u_1'$ by the implicit function theorem. For $s = 1$ only current paying utility depends on m_1. For $s \geq 2$, Δ_1 depends on m_s only through $\beta \mathbb{E} V_2$, and induction using value matching at each future threshold (the boundary terms vanish because $S_t = N_t$ at y_t^*) gives $\partial \mathbb{E} V_t / \partial m_s = -\beta^{s-t} \big(\prod_{k=t}^{s} (1 - p_k) \big) \mathbb{E}[u'(y - m_s) \mid y > y_s^*]$ for $t \leq s$. □

With $T = 2$, $m_1 = m$, and $m_2 = L - m$, equations (16)–(17) reduce exactly to equation (5) of the main text.

B.3 A menu of experiments identifies the discount curve

Let the borrower discount date-$(1+k)$ utility by a general discount function $D(k)$, $D(0) = 1$, and suppose either discounting is exponential or the borrower is naive. Replaying the proof of Proposition 3 yields

$$\frac{\partial p_1/\partial m_s}{\partial p_1/\partial m_1} = D(s-1)\, Q_s\, \frac{\mathbb{E}[u'(y-m_s) \mid y > y_s^*]}{u'(y_1^* - m_1)}, \qquad s = 1, \ldots, T. \tag{18}$$

Proposition 4 (Identification of the discount curve) If an experimenter can perturb each payment m_s separately and observes the initial default response to each perturbation together with the survival probabilities Q_s (directly observable from repayment data), then under the marginal-utility approximation of the main text the discount function is identified pointwise: $D(s-1) \approx (\partial p_1/\partial m_s)/(Q_s\, \partial p_1/\partial m_1)$. With $T \geq 3$ the exponential null is testable: $D(k) = \beta^k$ holds iff $\log[(\partial p_1/\partial m_s)/(Q_s\, \partial p_1/\partial m_1)]$ is linear in s.

Corollary 5 (Quasi-hyperbolic discounting) Under $D(k) = \beta\delta^k$ for $k \geq 1$, perturbations at two distinct future horizons $s, s' \geq 2$ point-identify both parameters: $\delta = (D(s-1)/D(s'-1))^{1/(s-s')}$ and $\beta = D(s-1)/\delta^{s-1}$. A two-period contract offers a single future horizon, hence one moment, and cannot separate β from δ: the two-period statistic is one point on the discount curve, not the curve.

Remark 6 (What the experiments must be) Identification requires ceteris paribus perturbations of individual payments within a fixed population. Comparing borrowers who selected into different contract lengths does not qualify: such populations may differ in their entire discount functions, confounding the curve with selection. The menu describes a feasible-in-principle design—randomized payment holidays or step-ups at prespecified horizons within a homogeneous pool—that, to our knowledge, has not been run. The theory is ready for the experiment; the experiment has not been done.

Remark 7 (Sophistication) The envelope step uses indifference at future thresholds under the deciding self's preferences. Under exponential discounting or naiveté this coincides (up to positive scaling) with the date-1 self's criterion and the boundary terms vanish. Under sophisticated time-inconsistent discounting, date-t's threshold is not optimal for date-1's preferences, and (17) acquires wedge terms of the form $f(y_t^*)\, \omega_t\, (\partial y_t^*/\partial m_s)$, where ω_t is the date-1 valuation of the branch gap at date-t's margin; $\omega_t = 0$ for all t iff $D(t-1+k) = D(t-1)D_t(k)$, i.e. exponential discounting. Estimating the curve on sophisticated agents therefore requires bounding these wedges or modeling the intrapersonal game.

B.4 Toward continuous time

Index periods by length h: payment flow m, stigma flow σ, discounting $e^{-\rho h}$. With income i.i.d. across periods, inspection of (15) shows $K_t = O(h)$ while $g_t(y) = u'(y)mh + O(h^2)$, so the threshold converges to a nondegenerate limit and the per-period default probability stays bounded away from zero—while the number of periods per unit of calendar time grows like $1/h$. The cumulative default probability over any fixed window therefore converges to one: an i.i.d.-income model cannot refine the period while holding the calendar fixed. Interior default at fine time resolution requires persistent income, so that threshold crossings arrive at a finite hazard. This is a distinct degeneracy mechanism from Lemma 2, and it disappears once income follows a discretized diffusion or an AR(1) with autocorrelation $e^{-\kappa h}$.

With persistent (Markov) income, Proposition 3 generalizes directly—the proof uses only the law of iterated expectations—with survival and marginal utility no longer factoring:

$$\frac{\partial p_1}{\partial m_s} = \frac{f(y_1^* \mid \mathcal{F}_1)}{\Delta u_1'} e^{-\rho(s-1)h} \, \mathbb{E}\left[\mathbf{1}\{\text{good standing through } s\} \, u'(y_s - m_s h) \,\middle|\, y_1 = y_1^*\right].$$

Fixing calendar horizons $\tau = sh$ and letting $h \to 0$, ratios of these sensitivities identify $e^{-\rho(\tau-\tau')}$: the continuous-time analogue of the sufficient statistic, in which "changing the monthly payment" becomes changing the payment flow rate $m(\tau)$.

In the full diffusion formulation ($dy = \mu(y)d\tau + \eta(y)dB_\tau$), the repaying value solves $\rho V = u(y - m(\tau)) + V_\tau + \mu V_y + \frac{1}{2}\eta^2 V_{yy}$ in the continuation region, the defaulted value W solves the same equation with flow $u(y) - \sigma$, and the free boundary $b(\tau)$ satisfies value matching and smooth pasting. Because both the remaining payment burden and the remaining stigma shrink linearly as $\tau \to T$, the boundary converges to the interior flow comparison $u'(b)m = \sigma$ at maturity. One subtlety distinguishes this case from discrete time: smooth pasting sets $\Delta_y(b(\tau), \tau) = 0$, so the marginal-defaulter denominator $\Delta u_1'$ vanishes and sensitivities must instead be computed by perturbing the value-matching/smooth-pasting pair jointly, which yields a first-order boundary shift $\delta b = -\varepsilon\, \psi_y(b)/\Delta_{yy}(b)$ for a payoff perturbation $\varepsilon\psi$. The h-period persistent-income model above sidesteps this machinery and is the recommended vehicle for quantitative work; the formal diffusion treatment is left for future research.

C Aggregation Under Heterogeneity

The model of the main text assumes a homogeneous population. Suppose instead the population is a mixture of types $i = 1, \ldots, n$ with weights ω_i, each with its own discount factor β_i, utility u_i, and income distribution F_i, hence its own sensitivities $A_i = \partial p_1^i / \partial L$ and $B_i = \partial p_1^i / \partial m$. Pooled estimation recovers the population-average derivatives $\bar{A} = \sum_i \omega_i A_i$ and $\bar{B} = \sum_i \omega_i B_i$ (exactly, for a linear probability specification; up to the usual first-order approximation, for average marginal effects from a pooled logit), and the estimator is $\hat{\beta} = \bar{A}/(\bar{A} + \bar{B})$.

Proposition 8 (Exact aggregation) Define each type's statistic $\beta_i^{stat} = A_i/(A_i + B_i)$ and its total default sensitivity

$$s_i \;=\; A_i + B_i \;=\; f_i(y_i^*) \, \frac{u_i'(C_1^{P*})}{u_i'(C_1^{P*}) - u_i'(C_1^{N*})} \;>\; 0.$$

Then, exactly,

$$\hat{\beta} = \frac{\sum_i \omega_i \, s_i \, \beta_i^{stat}}{\sum_i \omega_i \, s_i} \;\in\; \left[\min_i \beta_i^{stat},\; \max_i \beta_i^{stat}\right].$$

Proof. $\bar{A} = \sum_i \omega_i s_i \beta_i^{stat}$ and $\bar{A} + \bar{B} = \sum_i \omega_i s_i$; divide. The expression for s_i is the sum of the two sensitivity formulas in equation (5), in which the $\partial p/\partial L$ terms cancel. □

The pooled estimator is therefore a sensitivity-weighted average of type-level statistics: each type is weighted by its population share times (i) the density of its members at its own default margin, $f_i(y_i^*)$, and (ii) the responsiveness of its marginal defaulter, the marginal-utility ratio in s_i. Types far from the default margin contribute no identifying variation and receive no weight—the analogue of complier weighting in local average treatment effects. $\hat{\beta}$ is thus best interpreted as the discount factor of the population at the default margin, which is the relevant population for default-related policy. The direction of the wedge relative to the population-weighted mean depends on the covariance of s_i with β_i and is a quantitative question.

The companion script `Code/aggregation_simulation.R` solves the exact two-period model for CRRA types ($\gamma = 2$) with lognormal income under the paper's punishment convention and verifies these results numerically (output in `Output/aggregation_simulation.csv`). In 18 two-type configurations spanning $\beta_{lo} \in \{0.3, 0.5, 0.6\}$, $\beta_{hi} \in \{0.9, 0.95\}$, and mixture shares $\{0.25, 0.5, 0.75\}$: the identity of Proposition 8 holds to machine precision (maximum error 2.8×10^{-17}); the wedge between $\hat{\beta}$ and the population-weighted mean ranges from $+0.001$ to $+0.011$ in statistic units (at this calibration the patient type carries the larger s_i); and a pooled logit on 200,000 simulated households adds a functional-form difference of about 0.013 relative to exact pooled derivatives. The simulation also quantifies the marginal-utility correction in the exact ratio identity, to which we return in ongoing work.

D The Marginal-Utility Correction: Calibration and Bounds

The exact identity behind equation (7) is $A/(A + B) = \beta\lambda$ with λ defined in equation (6). This appendix calibrates λ, states the resulting corrected point estimates, and shows that large corrections are refuted both by internal consistency and by the borrowing-rate bounds of Section 8. Simulation code: `Code/mu_correction_simulation.R`.

D.1 Why λ could differ from one, and what disciplines it

The denominator of λ evaluates marginal utility at the consumption of the marginal period-1 defaulter; the numerator averages marginal utility over future repayers conditional on being at that margin today. Two forces govern the gap. Income persistence: if income is persistent, a borrower at today's margin is still relatively poor at maturity, and the two marginal utilities largely cancel; if income were i.i.d. across the two dates, future repayers would be drawn from the unconditional (richer) distribution and λ would fall far below one – in a hand-to-mouth i.i.d. benchmark it reaches 0.19. Consumption insurance: households smooth, so consumption dispersion – and hence marginal-utility dispersion – is smaller than income dispersion.

Both forces are measured by large literatures. Annual persistence of the idiosyncratic component of labor income is estimated at $\rho \approx 0.99$ under the restricted-income-profiles specification and $\rho \approx 0.82$ under heterogeneous profiles (Guvenen, 2009), $\rho \approx 0.91$ in PSID data (Flodén and Lindé, 2001), and $\rho \approx 0.95$ in Storesletten et al. (2004); administrative earnings data confirm the high persistence of the persistent component (Guvenen et al., 2021). For insurance, Blundell et al. (2008) estimate that consumption loads with coefficient ≈ 0.64 on permanent income shocks (and only ≈ 0.05 on transitory shocks), with substantially lower pass-through for wealthier households.

D.2 Measured persistence of the default-margin state

The grid below spans the range of persistence estimates from the earnings literature, but the object λ actually requires is more specific than the persistence of labor income in the population: it is the persistence of the economic state of a borrower at the default margin. The credit-bureau panel measures this state's persistence directly. I classify every auto tradeline with a positive past-due balance and no chargeoff at one archive snapshot, and follow it to the next snapshot (mean gap 4.0 months): across 21,350,801 such delinquency-spell transitions, 90.0 percent remain delinquent, 9.8 percent charge off, and only 0.26 percent cure. The implied annualized persistence of the distress state is $\rho = 0.992$ (replication: `Code/pace/48_cure_rates.R`, `results/cure_rates.csv`). Because snapshots are quarterly, spells that begin and cure entirely within a gap are missed, so the measured cure rate is a lower bound; even an order-of-magnitude undercount of cures would leave $\rho \geq 0.97$, within the top rows of the grid below. The measurement selects the high-persistence region of the calibration grid — the region in which the correction vanishes — and it does so for exactly the population and margin the estimator uses.

D.3 Calibration

I embed both forces in the two-period model: $\log y_2 = \tilde{\rho} \log y_1 + \varepsilon$ with $\tilde{\rho} = \rho^T$ for $T = 3.2$ years, stationary innovation variance, and consumption $c(y) = y^{\phi}$ with required payments reducing consumption one-for-one (debt service is not insurable, consistent with the liquidity interpretation of default). Thresholds and λ are computed exactly by the same routines as the baseline model. Table A1 reports λ, the annual correction factor $\lambda^{-1/T}$, and the corrected β for $\gamma = 2$; the replication file contains $\gamma \in \{1, 3\}$.

Table A1: The marginal-utility ratio λ and corrected estimates ($\gamma = 2$)

	λ			corrected β (measured: 0		
Annual persistence ρ	$\phi = 1.00$	$\phi = 0.64$	$\phi = 0.40$	$\phi = 1.00$	$\phi = 0.64$	$\phi =$
0.82 (Guvenen, 2009)	0.72	0.64	0.51	0.88	0.91	0.
0.91 (Flodén and Lindé, 2001)	0.92	0.84	0.72	0.81	0.83	0.
0.95 (Storesletten et al., 2004)	1.04	0.95	0.84	0.78	0.80	0.
0.99 (measured, Appx. D.2; Guvenen, 2009)	1.09	1.00	0.97	0.77	0.79	0.

Notes: λ computed exactly in the two-period model with persistent income ($\tilde{\rho} = \rho^{3.2}$) and consumption function $c(y) = y^{\phi}$; $\phi = 1$ is hand-to-mouth, $\phi = 0.64$ the permanent-shock pass-through of Blundell et al. (2008), $\phi = 0.40$ a more fully insured benchmark. Corrected $\beta = (\hat{R}/\lambda)^{1/T}$ with $\hat{R} = 0.79^{3.2}$. Bold: preferred cell, at the state persistence measured in the credit-bureau panel (Appendix D.2).

Three results. First, at the persistence measured in the panel (Appendix D.2; the $\rho = 0.99$ row) the correction vanishes: $\lambda = 1.00$ at the Blundell et al. (2008) pass-through, and the corrected estimate equals the baseline $\beta = 0.79$. Under the lower persistence preferred in parts of the earnings literature ($\rho = 0.95$, $\phi = 0.64$, $\gamma = 2$) the correction remains small: $\lambda = 0.95$, corrected $\beta = 0.80$, a 25% annual discount rate against the uncorrected 27%. Second, the sign of the correction is ambiguous at high persistence – at $\rho = 0.99$ hand-to-mouth, $\lambda > 1$ and the correction raises measured impatience – so the simplified estimator is not systematically biased toward finding impatience in the empirically relevant region. Third, across every admissible cell of the grid the corrected annual discount rate remains between 19% and 31%: no calibration of the correction brings the estimates near standard macroeconomic values.

D.4 Bounds

Two independent restrictions cage λ without taking a stand on the preferred cell. Internal consistency: since $\beta \leq 1$, identification requires $\lambda \geq \hat{R} = \hat{\beta}^T = 0.47$; calibrations violating this (in the grid, precisely the low-persistence/high-insurance corners that would produce large corrections) are refuted by the data themselves. The borrowing-rate bracket: Section 8 bounds discount rates from below using observed borrowing behavior, a methodology whose bias runs in the opposite direction from the marginal-utility correction. The corrected sufficient-statistic estimates (≈25%, upper range 31%) and the borrowing-rate lower bounds (≈16–20%) bracket the truth from two sides: annual discount rates for these households lie in roughly the 20–27% range under any admissible calibration, and the paper's central claim – impatience far exceeding standard calibrations – holds throughout the bracket.

Table A2 assembles the full correction ladder, applying in sequence the timing-dispersion (Jensen) adjustment, the persistence adjustment at the measured value, and the default-option adjustment; the corrected annual rate under each cumulative combination remains far from standard calibrations.

Table A2: The correction ladder, applied cumulatively

ϕ	face value	+ timing dispersion (Jensen, $\sigma_\Delta = 1.61$)	+ persistence (measured $\rho = 0.992$)	+ default option ($h = 5.7\%$/yr)
0.5	27.1%	30.1%	29.6%	22.2%
0.7	27.1%	30.1%	29.8%	22.4%
0.9	27.1%	30.1%	30.0%	22.6%

Notes: Corrections applied left to right, cumulatively, at three values of the consumption-insurance pass-through ϕ. Replication: `Code/pace/50_ladder_apply.R`.

Because the preferred correction is a single factor applied to all subgroups, the heterogeneity estimates of Section 5 shift level by 1.6% but preserve all cross-group comparisons; subgroup-specific corrections (allowing λ to vary with income via subgroup payment-to-income ratios) are a straightforward extension of the replication file and left for the estimation on the full panel.

E Instrument Diagnostics

Table A3 reports the state of the lender-menu instrumental-variables program described in the identification discussion. Panel A shows the first stages. Panel B applies the sufficient statistic on the identical sample under the fixed-effects baseline, under lender fixed effects, and under control-function specifications using the first-stage residuals; the control-function rows fail sign and support diagnostics in every variant, which is the basis for the paper's statement that the within-lender design remains ongoing work.

Table A3: Lender-menu instruments: first stages and specification grid

Panel A: first stages (leave-one-out lender menus; clustered by lender)					
outcome	instrument	coefficient	(s.e.)	F	
$\log m$	z_{dur}	-0.234	(0.052)	20	
$\log m$	z_{size}	0.587	(0.039)	221	
$\log L$	z_{dur}	-0.024	(0.008)	9	
$\log L$	z_{size}	0.987	(0.005)	47,288	
Panel B: the statistic under each specification (identical sample)					
specification	γ_m	(s.e.)	γ_L	(s.e.)	$\widehat{\beta}$
FE baseline	0.297	(0.067)	0.417	(0.093)	0.843
CF duration	-2.791	(0.334)	-0.103	(0.185)	0.347
CF both	-2.634	(0.326)	1.083	(0.499)	–
FE + lender FE	0.483	(0.045)	0.586	(0.054)	0.826
CF duration + lender FE	-2.668	(1.189)	1.129	(0.196)	–
CF both + lender FE	-2.625	(1.150)	3.522	(0.537)	–

Notes: Instruments are leave-one-out lender-year means: the share of originations with terms of 72 months or longer (z_dur) and mean log loan size (z_size), computed among lenders with at least 500 originations in the year. Fixed effects: origination-year × state, with lender fixed effects added in the lenderFE rows. Standard errors clustered by lender. A dash denotes a specification whose sensitivity ratio falls outside the unit interval. Replication: `Code/pace/65_menu_leniency_iv.R`.

F The Estimator Under Liquidity-Driven Default: Monte Carlo Evidence

This appendix reports the simulation exercise referenced in Section 4.4. Its purpose is to characterize what the sufficient-statistics estimator returns when default is not a forward-looking choice, and thereby to establish which questions the estimator can and cannot answer on its own.

F.1 Design

A population of 400,000 loans is simulated with contract terms drawn from $\{36, 48, 60, 72\}$ months, principals lognormal, and note rates decreasing in a credit-score index; monthly payments follow the annuity formula. Borrowers inhabit one of two worlds. In the preference world, a borrower defaults in month k when the δ-discounted present value of the remaining payment stream exceeds a default cost drawn i.i.d. each month around a persistent individual level (δ is the monthly equivalent of the annual factor β_{true}; the cost level is drawn independently of the loan's own terms). In the liquidity world, default arrives as a hazard proportional to the current payment burden raised to the power 1.5; the future stream never enters the decision. A mixture economy assigns share s of borrowers to the liquidity world. The estimator applied to each simulated cross-section is the paper's two-logit statistic verbatim, with score-ventile and cell fixed effects mirroring the empirical specification. In the calibrated surface of Table A5, the default-cost level in every (β_{true}, s) cell is tuned so the simulated lifetime default rate matches the empirical 5.1 percent.

F.2 Results

Table A4: Mixture economies: the estimator under increasing liquidity shares ($\beta_{true} = 0.95$)

Liquidity share s	$\hat{\beta}$	implied annual rate	default rate
0.00	0.887	12.8%	0.013
0.25	0.838	19.3%	0.025
0.50	0.811	23.3%	0.037
0.75	0.793	26.1%	0.048
1.00	0.783	27.6%	0.059

Notes: Each row simulates a mixture economy with true annual discount factor 0.95 for preference-world borrowers and liquidity share s. Replication: `Code/mc_liquidity.R`.

Table A5: The default-rate-matched surface: $\hat{\beta}$ across true patience and contamination

	liquidity share s				
β_{true}	0	0.10	0.20	0.35	0.50
0.55	0.853	–	–	0.830	–
0.62	0.860	–	–	0.837	–
0.70	0.866	–	–	0.839	–
0.76	0.874	–	–	0.848	–
0.80	0.876	0.871	0.860	0.847	0.828
0.84	0.879	0.872	0.864	0.850	0.829
0.88	0.884	0.875	0.867	0.849	0.831
0.92	0.888	0.879	0.869	0.856	0.836
0.96	0.892	0.884	0.873	0.855	0.836
0.99	0.896	0.885	0.876	0.859	0.841

Notes: Every cell's default-cost level is calibrated so the simulated default rate matches the empirical 5.1%. Dashes: cells outside the extended grid. Replication: `Code/mc_liquidity_bound.R`; surfaces in `results/mc_bound_surface.csv` and `_ext.csv`.

Three findings. First, the estimator cannot acquit itself: in the pure-liquidity economy ($s = 1$, Table A4) the estimated discount factor does not collapse toward zero but settles at 0.78 — a mechanical variance share arising because total obligations $L = mT$ load on the payment m that liquidity-driven default does respond to. An estimate in this range is therefore consistent, on its own, with either interpretation, which is why Section 4.4 rests on direct evidence. Second, contamination biases the estimate toward impatience monotonically, but modestly: at the ceiling admissible share $s = 0.35$, the depression in $\hat{\beta}$ is at most 0.05 across every row of Table A5. Third, and most consequentially, the simulated estimator attenuates toward patience in every world it is given: across true discount factors from 0.55 to 0.99 and contamination shares up to 0.50, no cell of the matched surface produces an estimate as low as the empirical 0.787. Within this model class, patient preferences plus admissible contamination cannot manufacture the measured impatience; the measured estimate requires the preference component itself to discount steeply. The compression of the surface also means the exercise cannot be inverted for a sharp point estimate of the contamination bias; a structurally estimated data-generating process would be required, and the categorical conclusion above is what the design supports.

G Data Construction and Estimation Procedures

This appendix specifies every construction step at reconstruction grade. A reader with the raw inputs and this appendix can rebuild the pipeline without the replication code; the code (file names cited in each exhibit's notes) implements exactly what follows.

G.1 Auto-loan estimation sample

The unit of observation is a loan in the merged auto-loan file, one row per loan carrying origination fields and terminal-snapshot fields. Filters, applied in order: credit score strictly between 0 and 851; income strictly positive. A loan is a default if its chargeoff amount is nonmissing and positive. The monthly payment m is the scheduled payment amount; total obligations are $L = m \times$ terms; the baseline fixed-effects cell is the five-digit zip code concatenated with the archive month. Dates parse as month–day–year; approximately 1.6 percent of origination dates fail to parse and drop where a date is required. Months paid equals the number of whole months between the origination date and the last-payment date. Loans recorded as defaults with months paid outside $[0, \text{terms})$ are excluded from all timing calculations. Rows whose payment-by-terms product overflows to missing drop from the L regression only, matching the originally estimated specification. The resulting sample contains 6,315,095 loans with a 5.1 percent default rate.

G.2 The two-logit estimator

Two logistic regressions with high-dimensional fixed effects: default on $\log m$, and default on $\log L$, each absorbing score-ventile and zip-by-month fixed effects, with classical standard errors. With $\widehat{\gamma}_m$ and $\widehat{\gamma}_L$ the two slope coefficients, $\widehat{R} = \widehat{\gamma}_L / (\widehat{\gamma}_L + \widehat{\gamma}_m)$ and $\widehat{\beta} = \widehat{R}^{1/\widehat{\Delta T}}$, where $\widehat{\Delta T}$ is the defaulter mean of $(\text{terms} - \text{months paid})/12$. Standard errors follow the delta method (Appendix K). Subgroup estimates repeat the identical procedure on the subgroup, re-cutting score ventiles within it.

G.3 Fixed-term contract families

Families are extracted from the monthly trade-level archives by flag conditions evaluated at each snapshot: first mortgages require the first-mortgage flag with monthly terms between 120 and 480; home equity requires the home-equity-loan flag with terms between 12 and 360; personal installment requires the installment flag with the auto, first-mortgage, home-equity, student-loan, medical, and revolving flags all zero and terms between 6 and 120; autos require the auto flag with terms between 12 and 96. Each loan is keyed by hashed consumer identity, hashed institution identity, and origination date; origination fields (payment, high credit, terms, first archive) freeze at first observation, and terminal fields (balance, chargeoff, last payment, narrative codes) overwrite at each subsequent observation. Date fields in the archives are packed variable-width integers in month–day–year order and parse after zero-padding to eight digits; consumer keys are handled as character strings throughout because their values span the full 64-bit integer range. The family-level statistic uses A equal to the defaulter sum of remaining scheduled payments, with payment values restricted to the interval $(0, 10^5)$, and the uniform default cost B equal to the defaulter sum of terminal balances; $\beta = (A/(A+B))^{1/\Delta T}$ with ΔT the defaulter mean remaining horizon in years. Within-person comparisons restrict both families to consumers whose hashed identity appears in each.

G.4 Delinquency-spell persistence

A spell is a tradeline with positive past-due amount, no chargeoff, and positive balance at an archive snapshot. Its outcome at the next available snapshot (mean gap 4.0 months) is cured when the past-due amount equals zero with no chargeoff, chargeoff when the chargeoff amount is positive, still delinquent when the past-due amount remains positive, and unobserved otherwise; unobserved transitions are excluded. The per-gap persistence is the still-delinquent and chargeoff share among observed outcomes, annualized as $\rho_{\text{annual}} = \rho_{\text{gap}}^{12/\text{gap}}$.

G.5 Cross-loan event study

A default event is the first archive month in which any auto tradeline of a consumer shows a positive chargeoff amount. For each event the pipeline locates the archive snapshots nearest one month before and six months after the event, requiring each within two months of its target. Other tradelines open and clean at the pre-snapshot are classified at the post-snapshot; a tradeline is paying when its last-payment date falls within three months of the post-snapshot month. Origination harvesting collects, at first observation, the origination fields of every tradeline of an event consumer opened within 48 months of the event.

G.6 Varying-coefficient specification and imputed rates

The varying-coefficient logits interact $\log m$ and $\log L$ with standardized score, standardized log income, and census region, absorbing score-ventile, income-ventile, and zip-by-month fixed effects; the defaulter-side horizon regression uses the same covariates. Surfaces evaluate on a score-decile by income-quartile by region grid, with uncertainty from 500 parametric-bootstrap draws of the estimated coefficient vectors. Imputed note rates invert the annuity identity $m = Lr/(1-(1+r)^{-T})$ for the monthly rate by bisection on $(10^{-7}, 1)$, annualize by multiplying by twelve, and discard values above 0.6; post-default comparisons take within-consumer means of imputed rates on originations before and after the consumer's event month.

H Additional Figures

Figure A1: Distribution of Liabilities (Sorted by Total Liabilities)

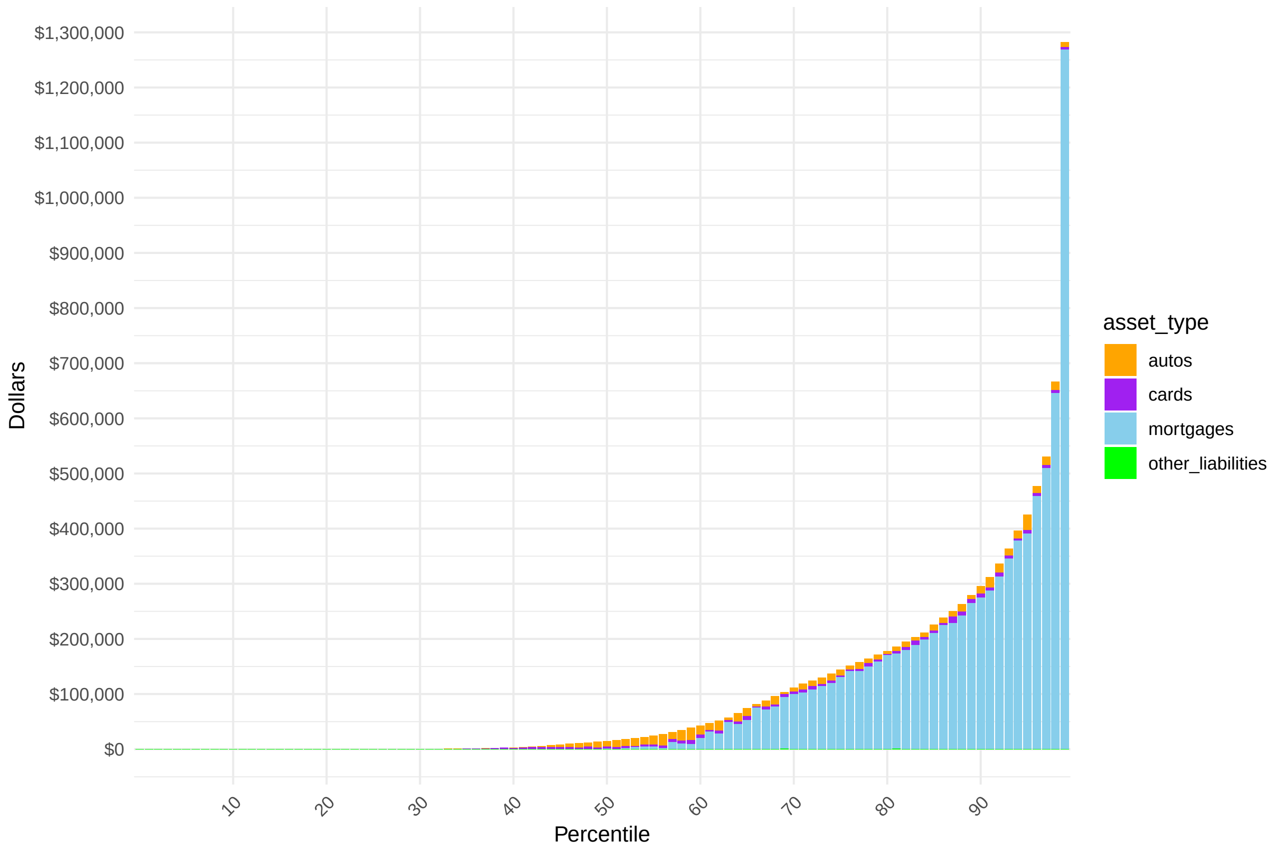

Notes: The x-axis is scaled by SCF survey weights.

I Lifecycle Model Calibration

Following Guvenen et al. (2021), annual earnings for workers aged 25–65 are:

$$L_{it} = (1 - \nu_t^i)e^{(g(t)+\alpha^i+\beta^i t+z_t^i+\varepsilon_t^i)} \tag{19}$$

where $g(t)$ is a quadratic age profile, α^i and β^i are idiosyncratic level and slope parameters, z_t^i is a persistent component, ε_t^i is a transitory shock, and ν_t^i captures nonemployment duration. The estimated common component is:

g_0	10.7
g_1	0.0537
g_2	−0.000718

All remaining parameters are taken from the calibration in Guvenen et al. (2021). For SCF application, I estimate each earner's age-adjusted idiosyncratic earnings position and assign α^i and β^i consistent with their position in the marginal distributions. The state variable is estimated as $E[Z_t^i] = \log L_{it} - g(t) - \alpha^i - \beta^i t$, and future earnings at horizon k are projected as $E_t[L_{t+k}] = \exp(\rho^k Z_{it} + \alpha^i + \beta^i(t+k) + g(t+k))$.

J Conditional Default Rate Calculations

The empirical conditional default probability is:

$$\text{CDR}_{d,m} = \frac{W_{d,m}}{O_d - \sum_{k=1}^{m-1} W_{d,k}}$$

where $W_{d,m}$ is defaults in decile d at month m, and O_d is total originations. I condition on both contract term and credit score deciles to address the concern that heterogeneity in ex-ante risk could spuriously generate a declining conditional default curve. Even after this conditioning, late default is extremely rare:

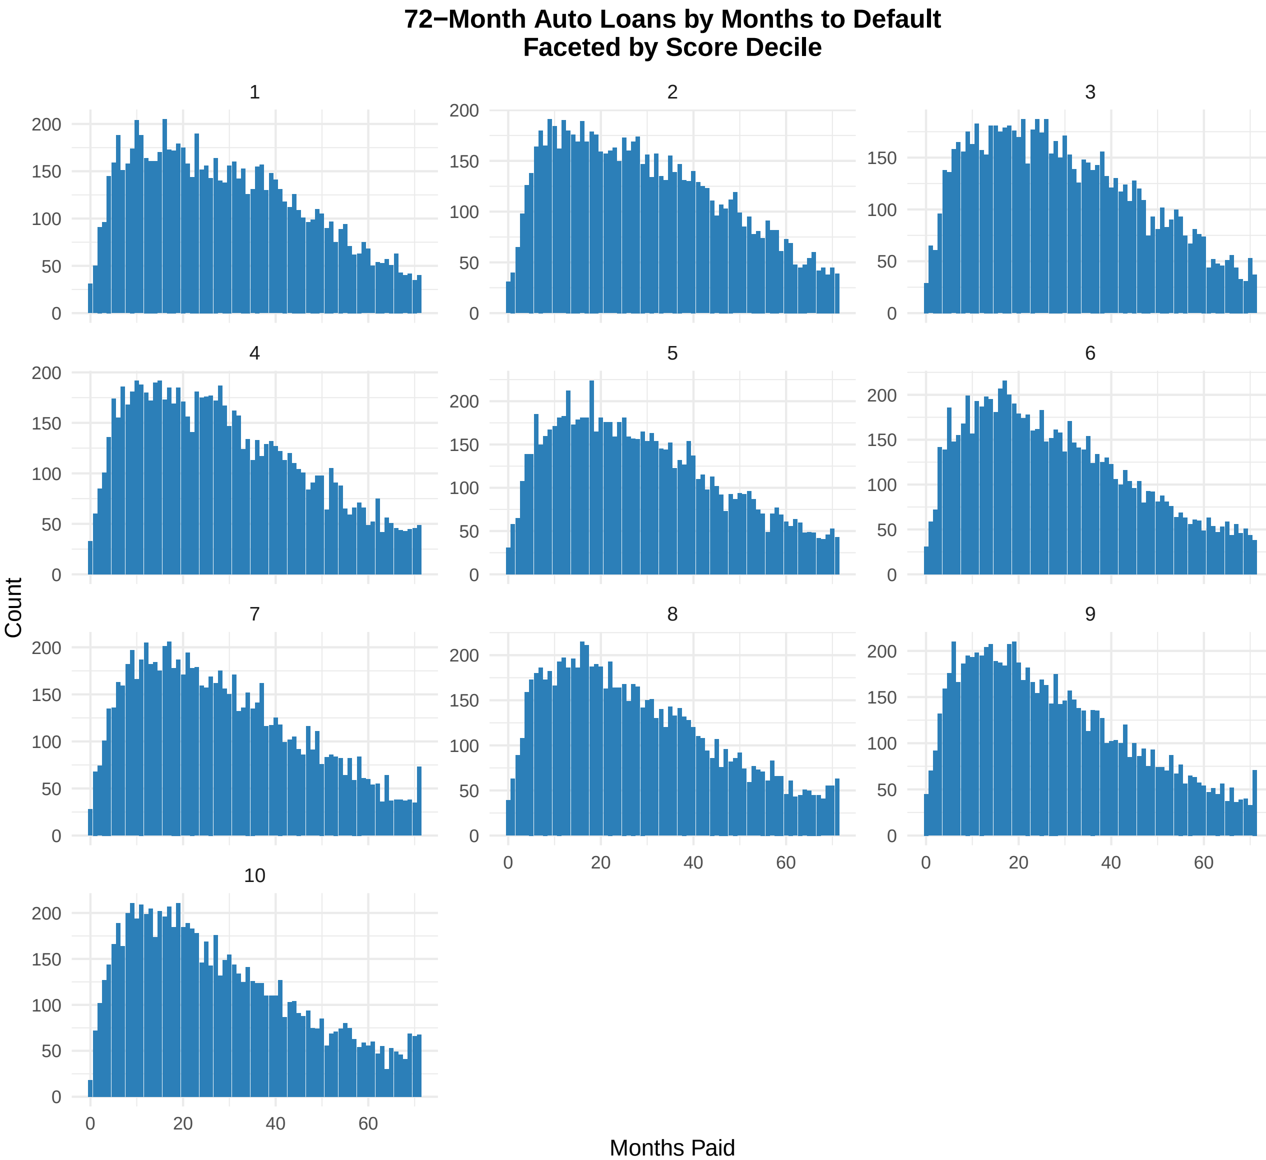
72−Month Auto Loans by Months to Default
Faceted by Score Decile
1
2
3
4
5
6
7
8
9
10
Count
Months Paid
0
20
40
60
0
50
100
150
200

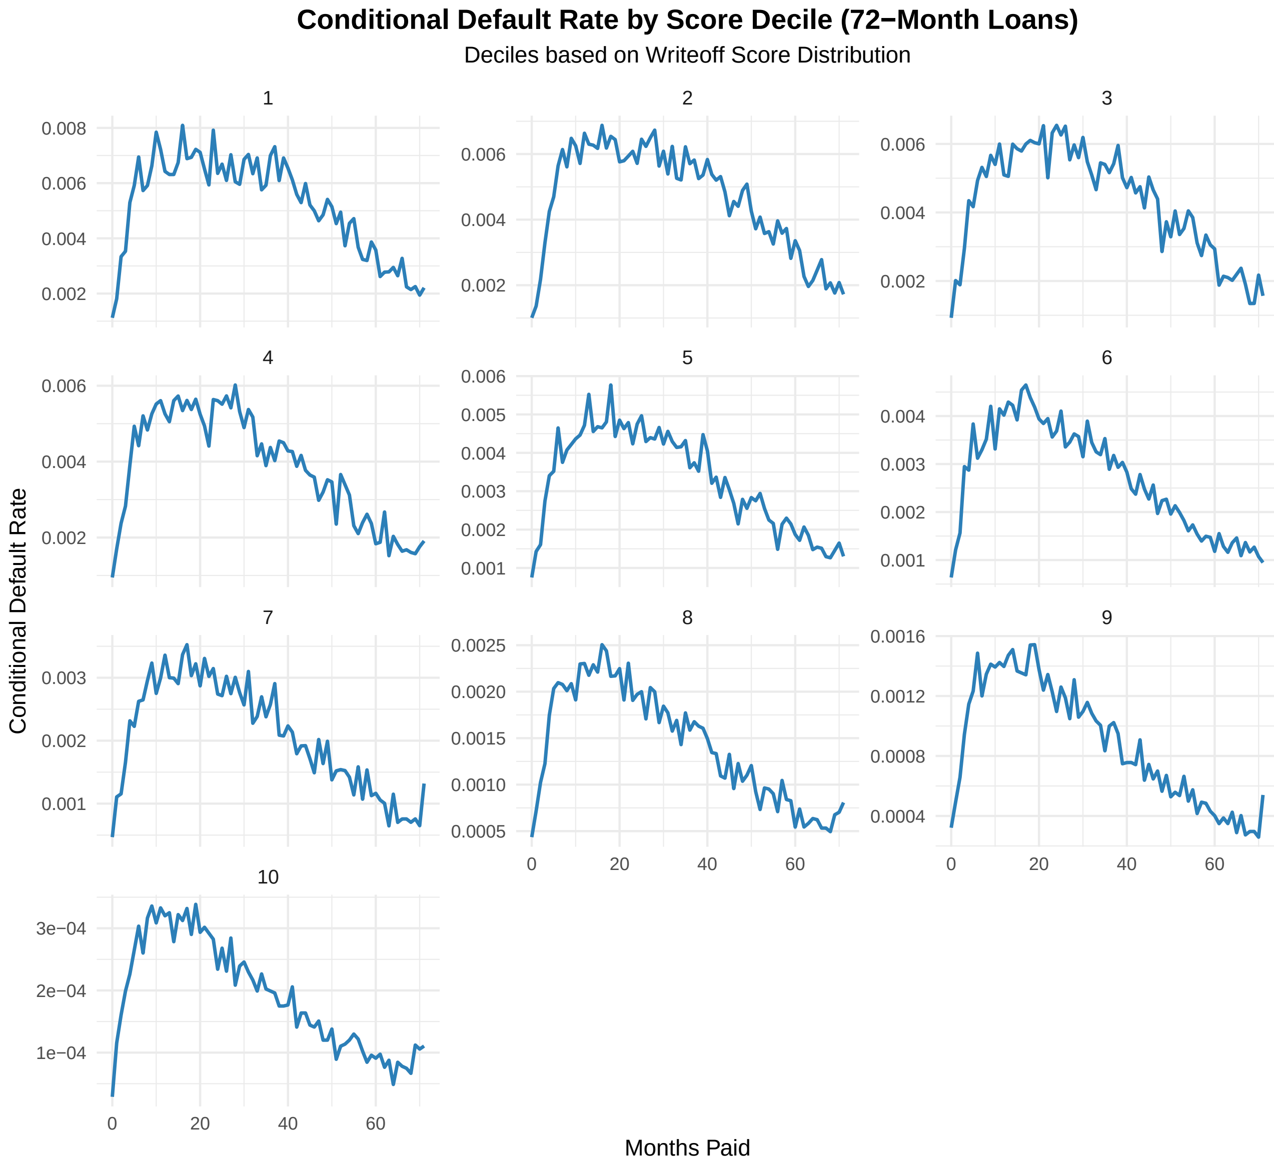

K Delta Method Calculations

Let $A \equiv \partial p_1/\partial L$, $B \equiv \partial p_1/\partial m$, $R \equiv A/(A+B)$, so $\beta = R^{1/T}$. The gradient:

$$\frac{\partial \beta}{\partial A} = \beta \cdot \frac{1}{T} \cdot \frac{B}{A(A+B)}, \tag{20}$$

$$\frac{\partial \beta}{\partial B} = \beta \cdot \frac{1}{T} \cdot \left(-\frac{1}{A+B}\right), \tag{21}$$

$$\frac{\partial \beta}{\partial T} = \beta \cdot \left(-\frac{\ln R}{T^2}\right). \tag{22}$$

Since (A, B, T) are estimated independently, the standard error is:

$$\mathrm{se}(\hat{\beta}) = \beta\sqrt{\frac{1}{T^2}\left(\frac{B}{A(A+B)}\right)^2\sigma_A^2 + \frac{1}{T^2}\left(\frac{1}{A+B}\right)^2\sigma_B^2 + \frac{(\ln R)^2}{T^4}\sigma_T^2}.$$