Revealed Preference Discount Rates*

Joseph P. Hall†

Abstract

Time preferences are central to the study of finance, but our empirical knowledge of discount factors remains limited and heavily shaped by asset-pricing evidence from the small subset of households who hold significant financial wealth, or from lab experiments of unrepresentative populations in unusual conditions. This paper argues that meaningful evidence about intertemporal preferences can be extracted from the credit market and, perhaps surprisingly, from default decisions. I develop a unified framework that uses (i) observed interest rates as revealed-preference lower bounds on constrained households' intertemporal marginal rates of substitution, and (ii) comparative statics of repayment and default from experimental and quasi-experimental settings as sufficient statistics for underlying impatience. I apply these methods to administrative credit-report data covering millions of auto loans, as well as a meta-analysis to existing experimental and quasi-experimental studies. Time preference heterogeneity emerges as large, systematic, and economically significant. The bottom half of the wealth distribution appears dramatically more impatient than standard calibrations assume, with implied annual discount factors often far above 20% annually in low-income populations. I discuss theoretical interpretations, implications for structural modeling, and consequences for policies targeted at liquidity- or credit-constrained households.

JEL: D14, D15, E21, E43, G10, G40, G51, H23, H55, H63

*This version: November 21, 2025. I thank Amit Seru, Greg Buchak, Benjamin Hébert, Jeffrey Zweibel, Hanno Lustig, Sudheer Chava, Xindi He, Indrajit Mitra, and the entire Georgia Tech Financial Services Innovation Lab. This paper subsumes and expands upon the content of my prior paper titled "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

I ask in this paper what we can learn about households' subjective time value of money from the liability side of household balance sheets when they do not have the opportunity to equalize their expected utility across time.

The subjective time value of money, also known as the discount factor of utility, is a primitive object in the study of finance. Any consumption-based asset pricing model with separable utility[1] assumes that individuals maximize a function of the form

$$\max_{C} \mathrm{E}_0 \left[\sum_{t=0}^{\infty} \beta^t \cdot U\left(C_t\right) \right] \tag{1}$$

Where the choice variable C is the consumption strategy which produces the consumption series C_t, and $U()$ is the single-period utility function. The subjective discount factor β plays a crucial role in all such models, and along with the consumption and/or wealth growth processes, defines the market price of risk-free debt. Indeed, in representative-agent models where there is only a single discount factor to measure, the price of risk-free debt r is given by the Euler equation

$$\frac{E_t\left[U'\left(C_{t+1}\right)\right]}{U'\left(C_t\right)} \beta\left(1 + r_{t+1}\right) = 1 \tag{2}$$

The price of risk-free government bonds is therefore naturally used as a chief data moment to match when inferring the discount rate β, as in Auclert et al. (2019), Lustig and Chien (2005), Bansal and Yaron (2004), Alvarez and Jermann (2001); far too many to be exhaustive. Other approaches in the asset pricing literature such as matching the capital-output ratio (Guvenen (2006)) or the zero-beta interest rate (Di Tella et al. (2024)) likewise rely on aggregate asset market data where the representative investor is rich. Indeed, in the United States, only 6.4% of households hold government savings bonds, the asset used to infer the social discount rate recommended by the White House (Board of Governors of the Federal Reserve (2022), The

[1] I discuss non-separable utility functions in Section 6.

White House (2023)). As I show in Section 2, only 12% of households participate in any kind of risk-free financial market investment (excluding transaction accounts), and 75% of risk-free investments are owned by the top 1%.

Macroeconomic research has largely considered the actual time preferences of infra-marginal investors, who are largely hand-to-mouth, as irrelevant – as Auclert et al. (2019) note, public asset market prices are almost totally invariant to the exact discount rates for the less patient and less wealthy segment of households (borrowers), so long as they are below those of the patient and wealthy who are the marginal investors in such markets. Essentially, the time preferences of constrained agents are not identified in common heterogeneous agent models, as constrained agents have no opportunity to express their impatience in the presence of credit constraints.

Why is this the case? Crucially, laws permitting the default of debt, which may be called "bankruptcy" or "limited liability" laws, make it impossible for households to issue risk-free debt. Households who are net borrowers instead borrow with risky, defaultable debt. Thus, it is infeasible for the vast majority of households to sell risk-free debt, and most choose not to purchase it either – retirement savings is mostly done in equity, not debt (Board of Governors of the Federal Reserve (2022)). Therefore the price of risk-free debt is uninformative about the time preferences of all but the richest households, other than to establish a lower bound on their impatience which may be arbitrarily far away from their actual time preferences.

How then should we estimate the time preferences of households who live close to hand-to-mouth and have negative net asset positions for most of their lives? Other approaches include laboratory experiments designed to elicit discount rates of participants (Coller and Williams (1999), Harrison et al. (2002) and many, many others). An obvious criticism of such studies is that convenience samples in university laboratory studies are not representative of overall populations. Furthermore, the behavior of individuals in laboratories does not consistently predict their behavior in everyday life. For these reasons governments by and large do not use the discount rates which result from such laboratory procedures in policy analyses.

Papers which do rely on real-stakes revealed preference do consistently find extremely high rates of time preference over short time periods, often of over 50% in hyperbolic (β, δ) models as reviewed by Laibson et al. (2024). While Ganong and Noel (2019) do use a heterogeneity exercise to estimate that about 1/3 of consumers in their data are extremely impatient "low-β" types, they stop short of producing full distributional estimates of time preferences. My estimates in Sections 3 are therefore useful for mapping out the full distribution.

Why is this exercise important? Accurately characterizing the discount rates of representative populations, not just representative populations, is broadly and deeply important for public finance, climate finance, household finance, and ethics. The United Kingdom and France both publish discount rate schedules at time horizons from 1 to 350 years to be used in cost-benefit analyses of proposed policies (Arrow et al. (2014)); the United States uses a constant discount rate tied to the market prices of government bonds, and consequently recently *lowered* the advised government discount rate from 3% to 2% per year (The White House (2023)). Hendren and Sprung-Keyser (2020) use 3% as their base discount rate for evaluating the Marginal Value of Public Funds (MVPF) for various policies, and in robustness exercises calculate MVPFs for discount rates up to 15% per year. They find that this difference is material for the cases of early childhood vs. adult education: early childhood education creates far superior value at a low discount rate, but adult education triumphs at a high discount rate.

Understanding the degree of impatience, *per se*, and its distribution in the population is important as well for analyzing public policy problems such as bargaining between two different parties who may have different levels of impatience – Rubinstein (1982) shows that in an alternating offer negotiation, the more patient party wins more surplus. This may be applied, for example, to understanding collective bargaining between an impatient labor union and a patient employer. It is also critical for welfare analysis of any intertemporal government policy – for example, DeMarzo et al. (2023) show that patient citizens are harmed by a government who is more impatient than they are being able to borrow. Welfare analysis of government debt policy (Indarte and Kanz (2023)) and climate policy (Farber and

Hemmersbaugh (1993)) likewise crucially depend on citizens' impatience. Disentangling the different roles of (1) changes in marginal utility from (2) impatience is also important for understanding the demand for financial services in the economy. If household demand for credit is driven by expectations of high consumption growth, it is more variable and dependent on economic conditions. If it is driven by impatience, then it is more primitive and we may expect it to be a more permanent fixture in any counterfactual economy.

I develop two complementary empirical strategies. First, because credit is supplied along an upward-sloping schedule, the highest interest rate at which a household borrows provides a lower bound on its true discount rate. Second, in settings where borrowers experience quasi-experimental variation in short- or long-term payments, the resulting changes in default probability reveal forward discount rates directly. I show that these comparative statics are sufficient to recover impatience without observing consumption, income paths, or full loan histories.

Integrating these approaches yields new empirical estimates of patience across the household distribution. Using data from the Survey of Consumer Finances (SCF), administrative credit-report data, and estimates from quasi-experimental studies, I find striking heterogeneity. Among credit-constrained households—those most likely to be directly affected by public policy interventions involving liquidity, repayment terms, or debt relief—discounting is extremely steep. These results have direct implications for consumer-credit regulation, optimal dynamic taxation, and the design of social insurance programs. They also provide new structural moments for calibrating or estimating life-cycle models that incorporate borrowing, liquidity constraints, and default.

The remainder of the paper proceeds as follows. Section 2 documents the limited participation of U.S. households in financial markets and illustrates why asset prices provide little information about the majority of households' discount factors. Section 3 develops the revealed-preference lower-bound approach, drawing on observed borrowing rates. Section 4 introduces a sufficient-statistics method for inferring impatience from repayment and default behavior. Subsequent sections apply these methods to auto-loan data, explore heterogeneity, and discuss extensions. The paper concludes by considering the broader implications of steep impatience for models of household behavior and for policy.

2 Financial Market Participation in the United States: The Case Against Interpreting Financial Market Prices as Discount Rates

A natural starting point is to ask: for which households can asset prices reveal time preferences? Standard macro-finance arguments rely on interpreting observed risk-free or risky returns as the implicit discount rates of the marginal investor. But this logic depends on broad participation, or at least on marginal pricing being representative of the population whose preferences we seek to identify. For households that do not save significant amounts in asset markets, interest rates on Treasury bills or returns on equity convey little about their intertemporal tradeoffs. The price of risk-free bonds (or any other publicly traded asset) is not informative about the time preferences of households near the median of the distribution because the typical household does not participate in the risk-free bond market. I demonstrate just how few households can be said to have any participation in risk-free asset markets using the 2022 Survey of Consumer Finances, the most comprehensive representative sample that exists of households' total financial positions in the United States (Board of Governors of the Federal Reserve (2022)). The Survey of Consumer Finances is ideal for the purposes of this exercise due to its comprehensive information about household balance sheets, including 179 variables that I use in this paper. It is also ideal for estimating the entire distribution of households in the United States due to its rigorous design, implementation, and inclusion of precise survey weights to match the Current Population Survey on several dimensions.

One major drawback of the SCF is that it does not include detailed repayment and default information of loans. It only asks survey participants whether they are behind schedule on each loan, and if they have declared bankruptcy, neither of which are adequate measures of default on particular loans. I address this in Section 4 with a sufficient statistics approach that can impute discount rates using already-reported estimates from other papers that study the impact of contract modifications on default rates. These other papers (e.g. Ganong and Noel (2020) and Dobbie and Song (2020)) are interested in estimating these causal effects in the context of direct policy outcome evaluation; my model and framework allows them to be understood in the broader context of time discounting.

Another limitation of relying on the SCF is that geographic identifiers are redacted for privacy. This leaves me unable to correlate my measures of impatience and its distribution with other geographic variables such as demographics, political attitudes, and other economic outcomes not measured in the SCF. This is a potentially fruitful avenue to pursue in future work. Older versions of the SCF, e.g. from 1983, do include these geographic identifiers, although the Federal Reserve Board of Governors does warn that the survey design of the SCF does not necessarily produce representative samples state-by-state, only within each of the nine major survey districts which comprise the country.

I calculate total risk-free asset holdings as the sum of certificates of deposit and all federal bonds. I use the SCF survey weights to sort households into population percentiles of total holdings of risk-free assets.

Figure 1: Financial Asset Holdings Across the Distribution

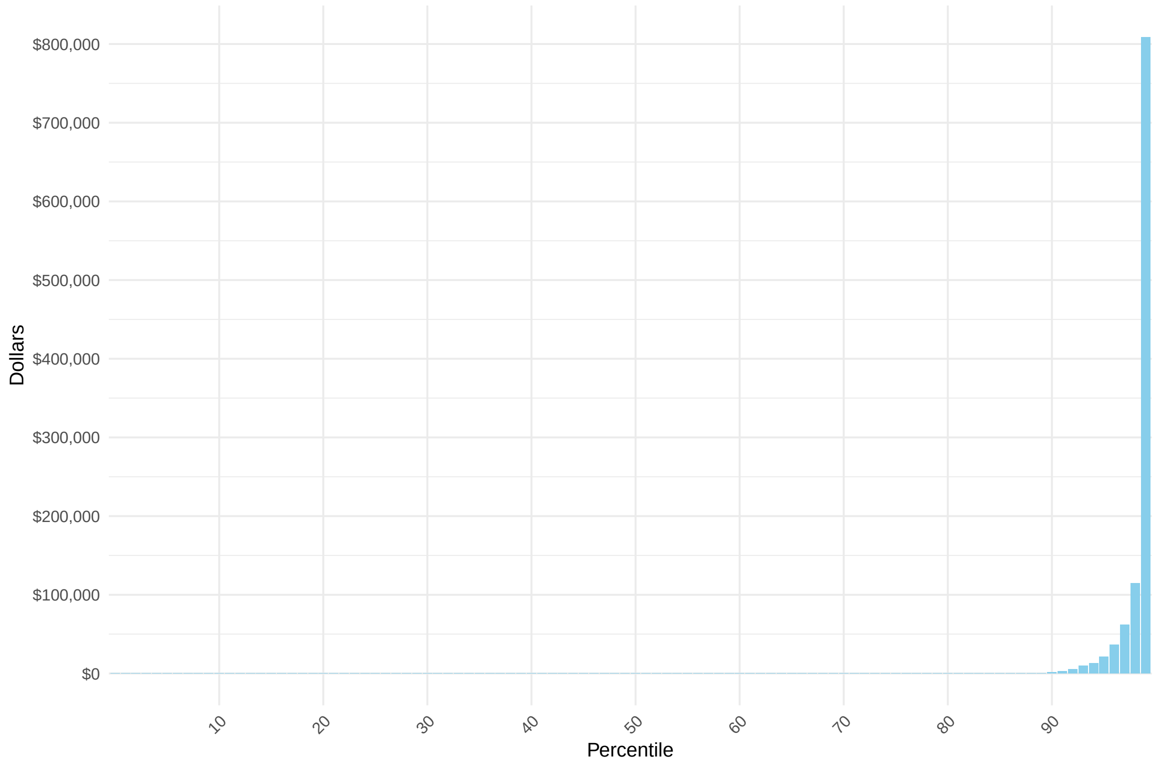

Notes: This figure plots SCF households sorted by the percentile of total financial assets (scaled by survey weights). The y-axis reports average holdings within percentile bins, disaggregated by asset category. Investment assets are defined as liquid, market-traded assets with long-term cash flows. Risk-free assets include checking and savings accounts. Other assets include home equity and annuities.

Figure 1 shows the distribution of risk-free asset holdings. Only 12% of households participate in this market whatsoever, with 75% of risk-free assets being held by the top 1%.

If we are to infer discount rates from asset market data at all, the problem can scarcely be fixed by trying to infer time preferences from public financial asset market data at all, even including risky assets. If impatience is negatively correlated with wealth (a natural assumption), then investors will be few and patient while borrowers are many and impatient. The marginal investor, the one who is indifferent between borrowing and saving, will therefore be much wealthier and much more patient than the average person. Di Tella et al. (2024) calculate the “Zero-Beta Rate” as the return on the lowest-variance portfolio of stocks which has a zero beta to the market and argue that this rate better represents the true market price of intertemporal substitution, that is, the rate at which the investor with the dollar-weighted median demand for borrowing is indifferent between borrowing and lending. They estimate that this rate has averaged 8.3% per year since 1973 and is highly volatile. There are at least two reasons, however, why their approach is unfit for the purposes of this present exercise. For one, the market price of intertemporal substitution is a combination of impatience and consumption smoothing demand as Equation 2 shows. Secondly, this procedure is silent on the distribution of time preferences across the entire population and can only identify the demand for intertemporal substitution of the marginal investor in asset markets. This may explain why the zero-beta rate that they estimate is so extremely volatile in the time series: changes in net worth of different segments of the population cause investors to switch from borrowers to savers or vice versa, shifting the identity of the marginal investor to one who is much more or much less patient. Finally, the estimates from this exercise are likely uninformative about primitive characteristics such as impatience, the fundamental object of interest of this paper, because they are so volatile. I take the view that these highly volatile discount rate estimates are most likely the result of the identity of the marginal investor shifting with shifts in economic conditions and the wealth distribution.

Figure 2: Distribution of Investment, Risk-Free, and Other Assets

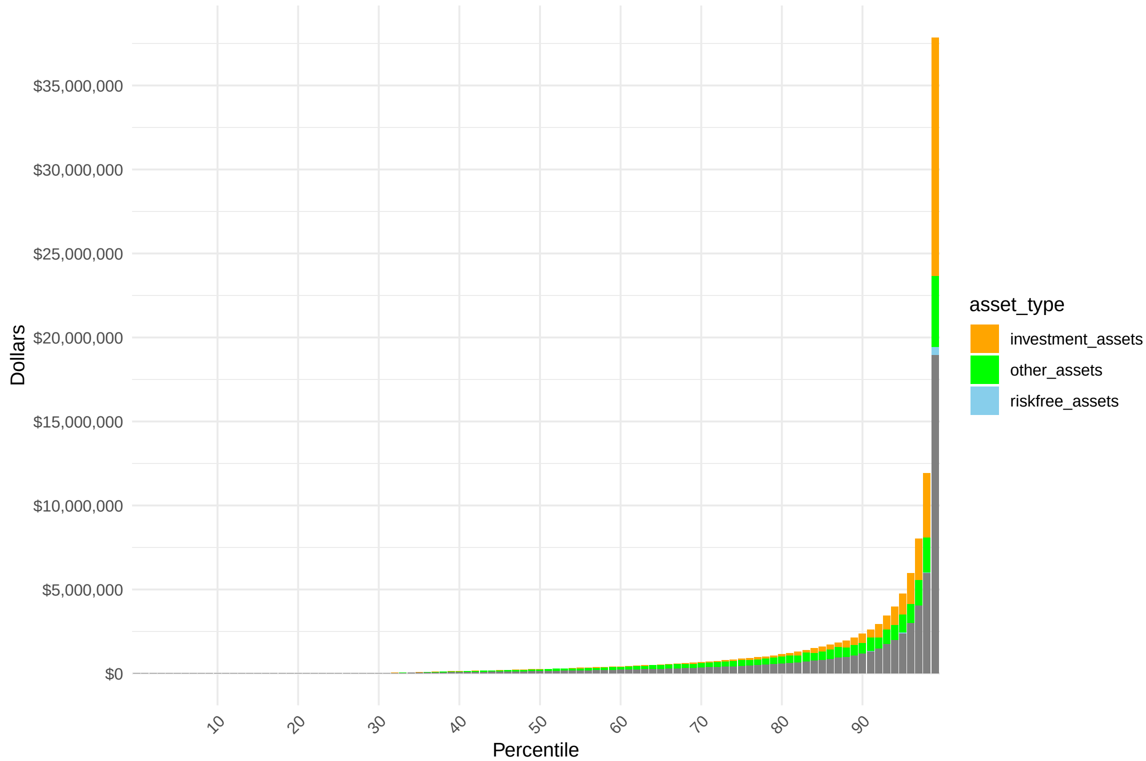

Notes: This figure displays the same percentiles as Figure 1, but separates investment assets from risk-free and other assets. Again, survey weights scale both the percentile ranks and the average values within each bin. Investment assets are extremely concentrated among the wealthiest households.

How un-representative is the marginal investor in asset markets? Figure 2 shows the distribution of all financial assets held by households. I exclude pensions and social security as they are not tradable. I classify financial assets into "risk free", "investment", or "other" categories. Investment assets are distinguished from other assets by having long-term cash flows and being traded in liquid markets. These mostly include stocks and bonds, including assets held in vehicles such as Individual Retirement Accounts. Other assets are mostly comprised of home equity, informal loans, and annuities. The bottom third of households hold no financial assets whatsoever, and the middle third only barely hold small amounts of non-investment assets. The top 1% holds 47% of all financial assets.

We should conclude from this exercise that the time preference rates implied by the returns of public assets are completely uninformative about the preferences of roughly 30-70% of the

U.S. population.

Figure 3: Liabilities Across the Asset Distribution

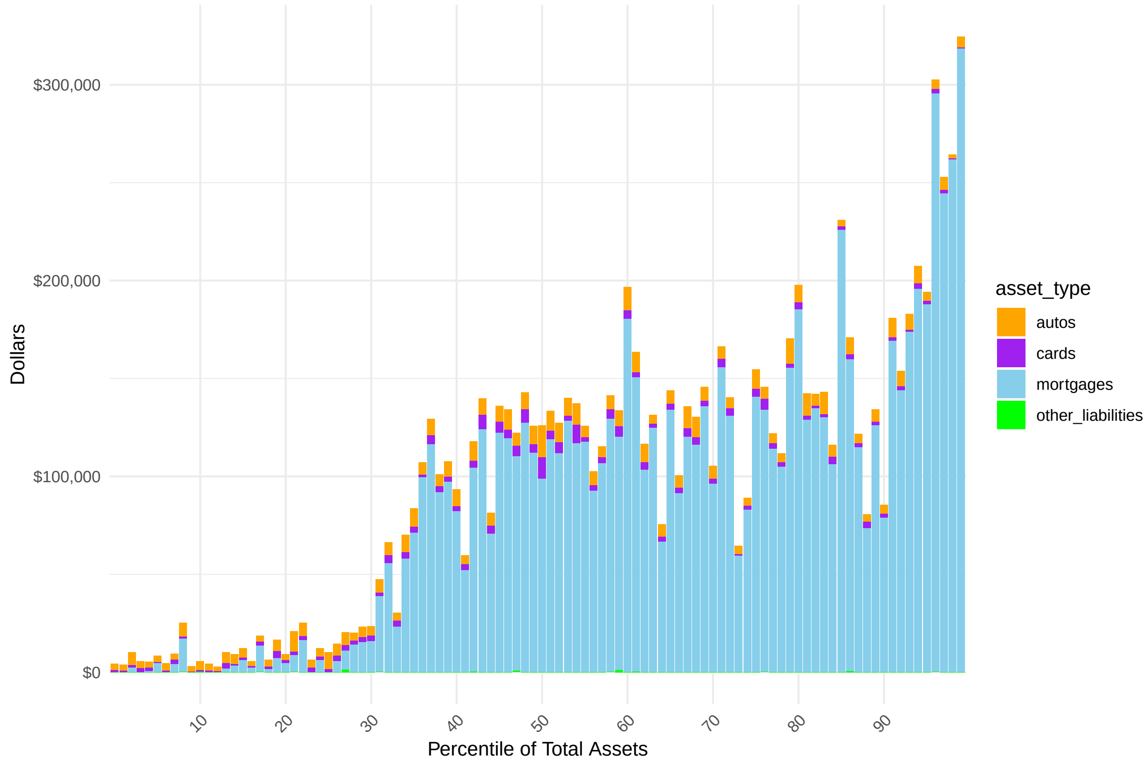

Notes: Households are sorted along the x-axis by total financial assets, as in Figures 1 and 2, but the y-axis displays average liabilities. Mortgage debt is widespread through the middle of the distribution, while unsecured debts cluster among the bottom percentiles.

Instead, in this paper I study the information contained in loan transactions and loan defaults as an alternative source of information about the discount rates of these households. Figure 3 shows, for households sorted into the same percentile ranks as in Figure 3, their liabilities instead of their assets.[2] For credit card balances, I only include The center and left half of this graph is far more populated than Figures 1 and 2, chiefly with mortgage debt, motivating my decision to utilize information contained in the debt market, rather than the asset market, to estimate the time preferences of the full distribution of household time preferences.

[2] I confine the analogous graph sorting households by total liabilities to Appendix B. It is uninformative for this paper's main exercise because households may have few liabilities either because they are poor and not creditworthy or because they are rich and not liquidity constrained.

3 Converting Interest Rates to Discount Rates: The Role of Income Expectations and Simple Lower Bounds For the Latter

The interest rates at which borrowers transact does contain direct revealed-preference information about intertemporal substitution despite not being sufficient to point-identify it. Due to classic finance frictions including asymmetric information, lenders require higher interest rates to extend credit to borrowers perceived as riskier or more likely to default. A household accepts a loan at the highest rate that does not exceed its true discount rate. Because credit products are discrete and vary across lenders, we do not observe the precise point at which a household becomes indifferent. Nevertheless, the maximum observed interest rate among a household's active debts provides a lower bound on its underlying impatience.

Staying true to the objective of this paper of relying on revealed preference, I utilize the interest rates that households borrow at as lower bounds on their prices of intertemporal substitution. The intuition for this approach is illustrated in Figure 4 schematically. Each household faces a menu of borrowing opportunities at different maturities and rates. As the interest rate rises, fewer households are willing to borrow at the margin. The highest rate at which a household holds debt that is actually accruing interest is, therefore, the highest observed point where we can rule out the possibility that the borrower is less patient than that.

Essentially, I assume that individuals face an upward sloping credit supply curve, and accept offers to borrow at progressively higher interest rates until they are offered a loan at higher than their true discount rate. Since credit products are discrete, the true discount rate is therefore unobserved.

For each household in the SCF, I calculate their maximum borrowing rate as the highest interest rate they are currently paying on debt which is accumulating interest. I include only loans on which payments are active, not those currently in forbearance – this excludes 67%

Figure 4: Illustration of Lower Bound Approach

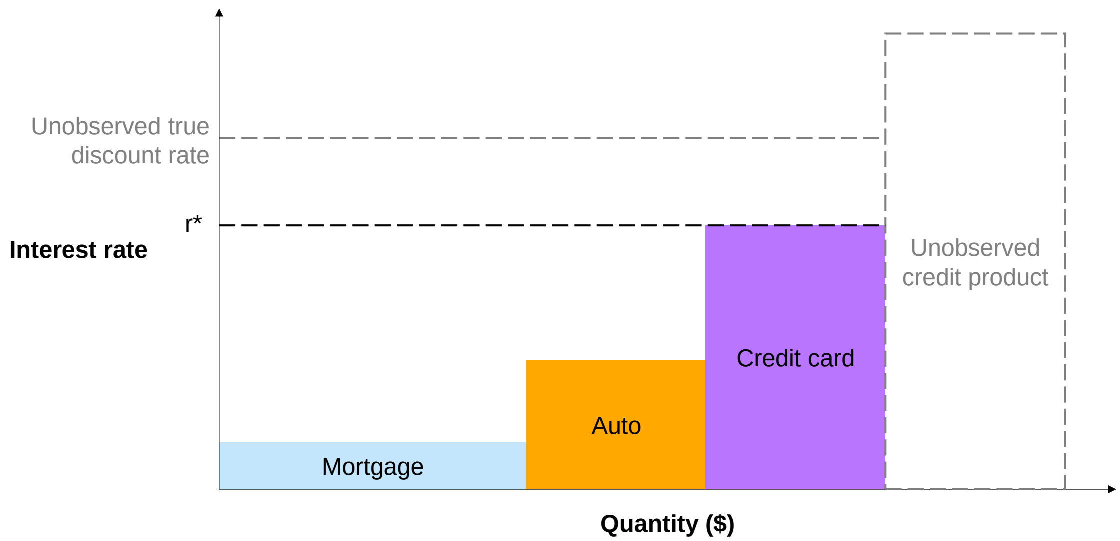

Notes: This schematic figure illustrates the conceptual framework for interpreting observed borrowing rates as lower bounds on households' discount rates. Because the supply of credit is upward sloping and loan products are discrete, the true discount rate lies at or above the highest observed borrowing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

of student loans. I exclude outlier observations of loans with interest rates above 50% per year. Unlike Catherine et al. (Forthcoming), who are mainly interested in long-term discount rates for the purposes of valuing Social Security benefits, I do include credit card debt in this analysis as it is the marginal credit product used by most households in the sample, and therefore a valid lower bound for the short-term true discount rate in the model of Figure 4.

Figure 5 shows the distribution of maximum borrowing rates per household. Since credit card debt in particular continues to represent a puzzle in the literature, in particular the fact that households with substantial liquid savings often do simultaneously hold credit card debt (Gross and Souleles (2002)), I calculate these distributions both including and excluding credit cards. Holding firm to the revealed-preference objective of this paper, I do consider credit card debt going forward as the marginal borrowing rate for most consumers at short time horizons and a valid lower bound for their short-run discount rates.

Figure 5: Distribution of Maximum Interest Rate per Household

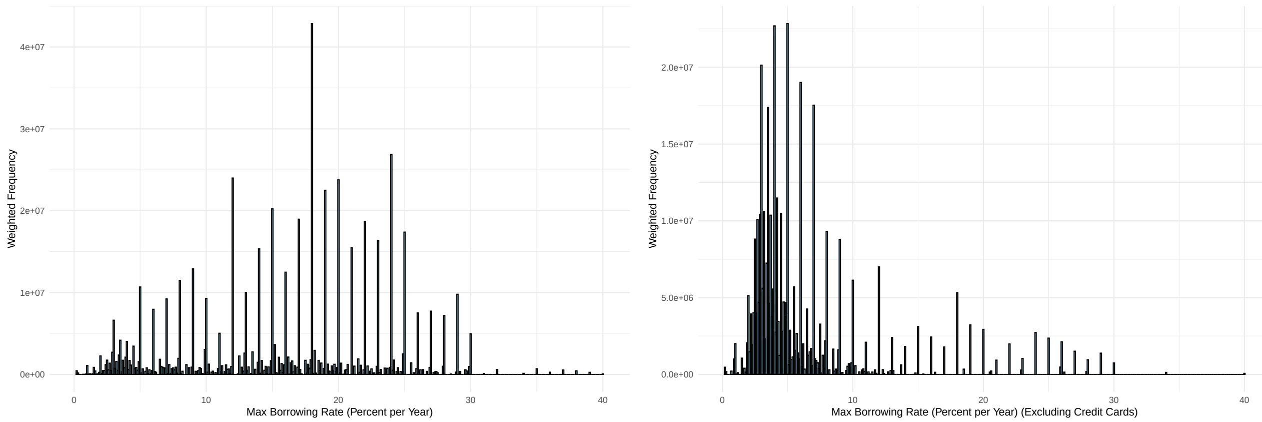

Notes: Panel A shows the distribution of maximum interest rates including credit card debt; Panel B shows the distribution excluding credit card debt. Debts which are not currently being paid, which include 67% of student loans, are not included in this graph. The highest interest debt that a household is currently paying and accruing interest is shown. Represented in these chart are the 86% of Americans who are paying interest on debt.

The modal maximum borrowing rate is 18%, which continues to be the dominant default credit card interest rate in the marketplace since it was the most common maximum legal interest rate in the United States going back to when credit card interest rates were regulated at the state level (Hall (2024)). Even excluding credit cards completely, over 25% of households borrow at more than the current treasury rate of 5% per year.

Borrowing rates also display a term structure – they are not constant over time. Indeed, the UK and French governments explicitly use downward-sloping time discount rates in cost-benefit analysis, although the United States does not (Arrow et al. (2014), The White House (2023)). The "yield curve" (the term structure of risk-free interest rates) is a key object of study in macro-finance and generally slopes upward, unlike the UK and French governments' discount rate schedules which are downward-sloping. Taken at face value, this result would appear to imply that the yield curve's downward slope is due to a term premium demanded by lenders; not to demand by borrowers.

To understand the term structure of household discount rates, I construct a product-level dataset of all household debt products in the Survey of Consumer Finance. Table 1 summarizes this product-level dataset in terms of the number of debt products per household.

Figure 6: Distribution of Maximum Interest Rate per Household

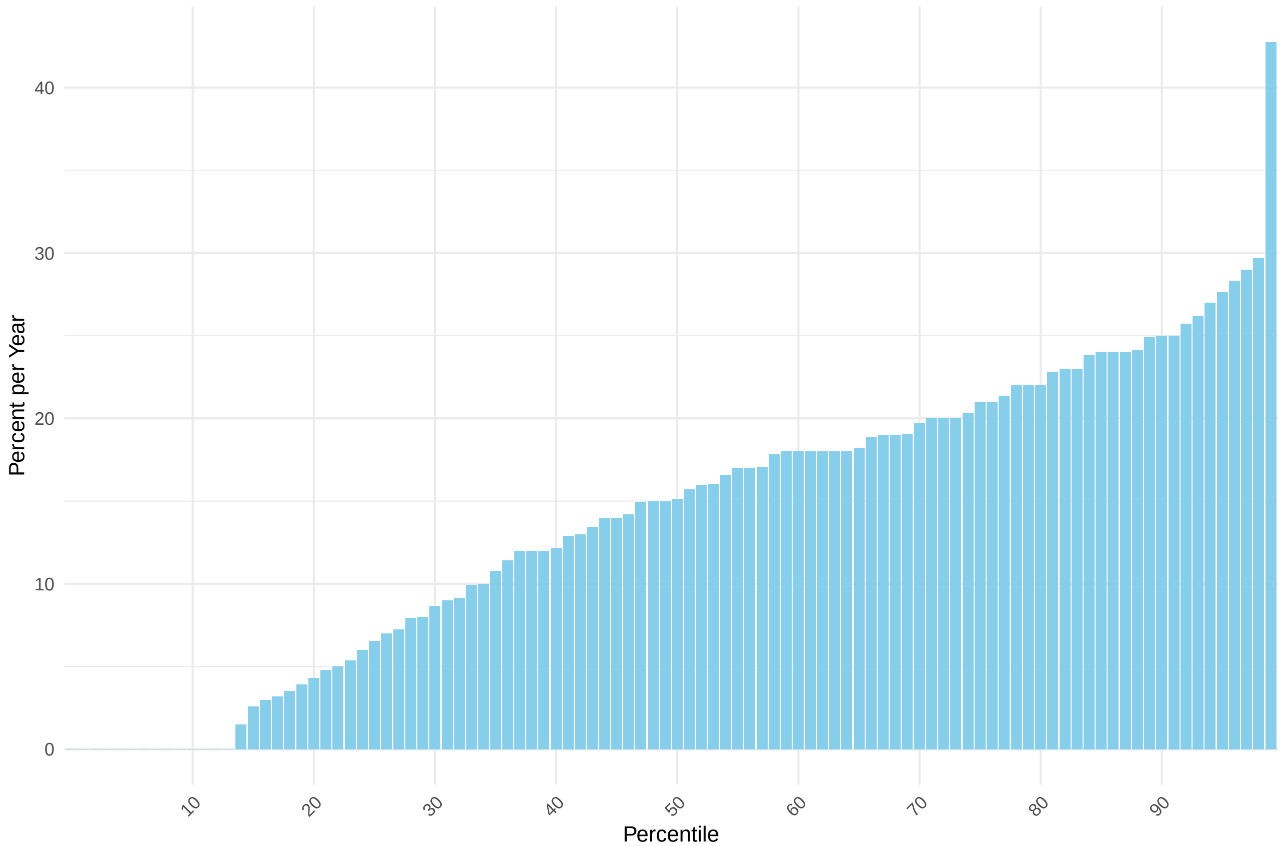

Notes: Debts which are not currently being paid, which include 67% of student loans, are not included in this graph. The highest interest debt that a household is currently paying and accruing interest is shown.

Of households with debt, the majority have only one or two debt products, but a small tail of households have as many as 11 different debt products.

It is critical to remember that these are lower bounds for r and not β, as used in Equation 2. Section 3 clearly explains and estimates the relationship between r and β, the difference between them being consumption-smoothing demand (the term $\frac{E_t[U'(C_{t+1})]}{U'(C_t)}$ in Equation 2. I estimate the missing term using the life-cycle model of Guvenen et al. (2021) in the upcoming subsection.

Table 1: Households by Number of Debt Products

Products	Households	Proportion
1	5533	0.35
2	4615	0.29
3	2851	0.18
4	1574	0.10
5	765	0.05
6	336	0.02
7	146	0.01
8	78	0.00
9	30	0.00
10	5	0.00
11	10	0.00

Notes: All households in the SCF with debt who are making payments and accruing interests are included in this table.

Table 2: Debt Products by Type

Product	Avg. Balance	Avg. Rate	Term (Yrs)	Proportion
Credit Cards	$4694	16%	<1	0.31
Mortgages	$212911	3.6%	26	0.23
Auto Loans	$16173	5.7%	5.2	0.23
Education Loans	$30674	5.5%	16	0.12
Consumer Loans	$5919	7.6%	3.9	0.04
Timeshare Loans	$320192	4.6%	23	0.03
Lines of Credit	$70384	6.5%	<1	0.03
Home Improvement Loans	$24651	4.6%	8.6	0.01

Notes: All household debt liabilities in the SCF on which payments are being made and interest is accruing are included in this table.

3.1 Relating Interest Rates to Discount Rates: The Role of Expected Consumption Growth

Recall that the Euler Equation 2 relates expected consumption growth to the pure impatience parameter β and the representative-agent interest rate r. Rearranging this equation slightly and taking logs, I can see how the lower bound spot rates, which are analogous to r, are related to impatience (β):

Figure 7: Household Debt by Duration and Interest Rate

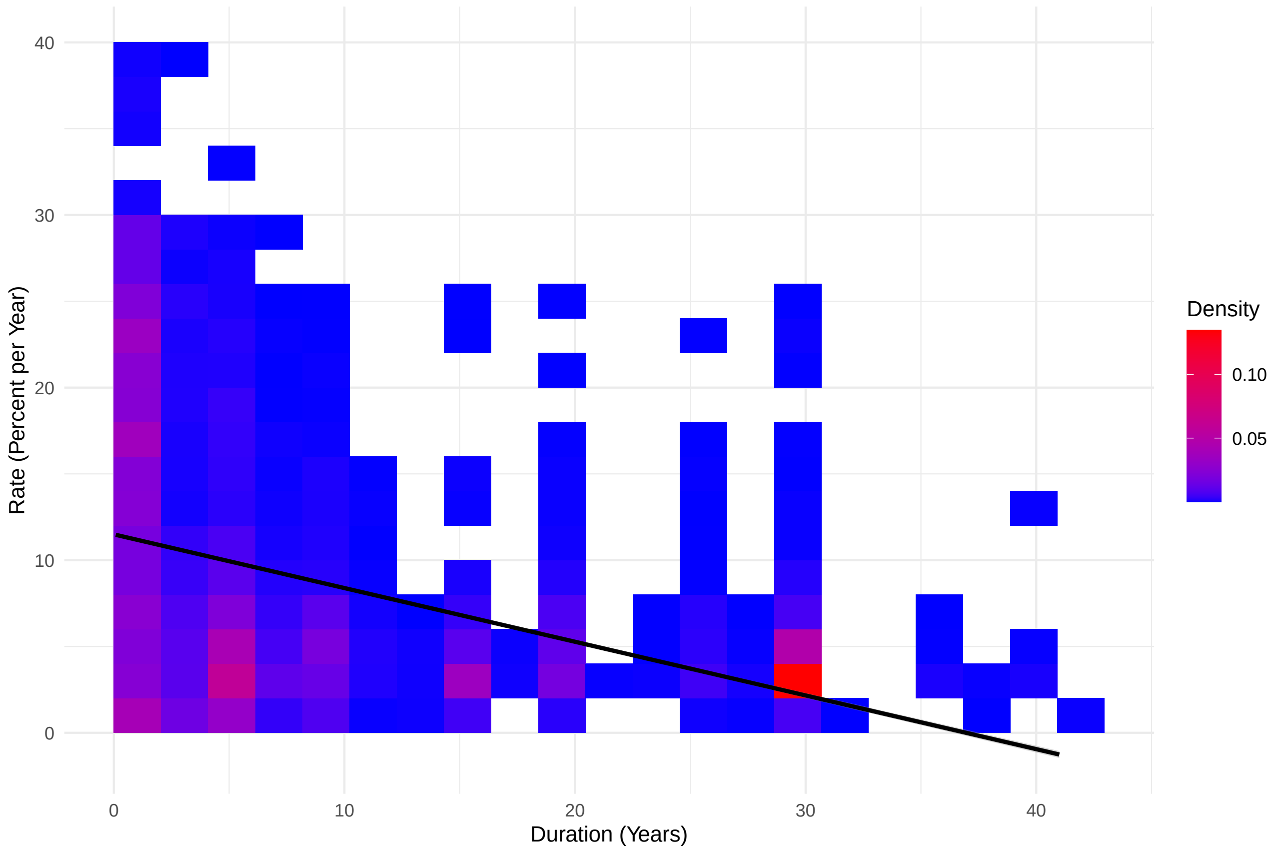

Notes: This heat map shows the observed interest rates and maximum durations of every household debt product in the 2022 Survey of Consumer Finances as well as an OLS trendline. Densities are scaled by survey weights.

$$\underbrace{r}_{\text{intertemporal demand}} = \underbrace{-\log(\beta)}_{\text{impatience}} + \underbrace{\log E_t\left[U'\left(C_{t+1}\right)\right] - \log U'\left(C_t\right)}_{\text{smoothing demand}} \tag{3}$$

Equation 5 shows the relationship between impatience *per se*, intertemporal demand; which intuitively the market interest rate above which an individual wishes to be a saver and below which an individual wishes to be a borrower), and consumption smoothing demand. I adjust these estimates for consumption smoothing demand using the detailed and realistic life-cycle model of Guvenen et al. (2021), which is also used by Catherine et al. (Forthcoming) to estimate the distributional effects of Social Security on wealth inequality. Intuitively, individuals may demand credit because they are simply impatient, or because they expect to be wealthy tomorrow. Indeed, John Steinbeck famously coined the saying that the poor in America "see themselves as temporarily embarrassed millionaires". The life-cycle model is designed to adjust the estimates of r to give estimates of β which are robust to Steinbeck's critique.

Specifically, for workers between ages 25 and 65, annual earnings L_{it} are given by:

$$L_{it} = (1 - \nu_t^i) e^{(g(t) + \alpha^i + \beta^i t + z_t^i + \varepsilon_t^i)} \tag{4}$$

Where the variables therein are given by the following processes:

$$\text{Persistent component} \qquad z_t^i = \rho z_{t-1}^i + \eta_t^i \tag{4.1}$$

$$\text{Innovations} \qquad \eta_t^i \sim \begin{cases} \mathcal{N}(\mu_{\eta,1}, \sigma^2_{\eta,1}) & \text{with prob } p_z \\ \mathcal{N}(\mu_{\eta,2}, \sigma^2_{\eta,2}) & \text{with prob } 1 - p_z \end{cases} \tag{4.1}$$

$$\text{Initial Condition} \qquad z_0^i \sim \mathcal{N}(0, \sigma^2_{z,0}) \tag{4.3}$$

$$\text{Transitory shock} \qquad \varepsilon_t^i \sim \begin{cases} \mathcal{N}(\mu_{\varepsilon,1}, \sigma^2_{\varepsilon,1}) & \text{with prob } p_\varepsilon \\ \mathcal{N}(\mu_{\varepsilon,2}, \sigma^2_{\varepsilon,2}) & \text{with prob } 1 - p_\varepsilon \end{cases} \tag{4.4}$$

$$\text{Nonemployment duration} \qquad \nu_t^i \sim \begin{cases} 0 & \text{with prob } 1 - p_\nu(t, z_t^i) \\ \min\{1, \text{Exp}\{\lambda\}\} & \text{with prob } p_\nu(t, z_t^i) \end{cases} \tag{4.5}$$

$$\text{Prob. of Nonemp. shock} \qquad p_\nu^i(t, z_t) = \frac{e^{a + bt + cz_t^i + dz_t^i t}}{1 + e^{a + bt + cz_t^i + dz_t^i t}} \tag{4.6}$$

In the SCF, I observe an individual at a point in time at a certain age, labor income, and employment status, allowing me to estimate the distribution of individuals' expected earnings at the 1 year, 5 year, and 10 year horizons. This process is explained in detail, and the calibration is provided, in Appendix C. Essentially, I first reverse-engineer the persistent state variable z_t^i from the observed values of income. When an individual is unemployed I impute it using the expected income variable in the SCF, which essentially is expected income in the individual's counterfactual state of employment. I estimate the age-specific average function $g(t)$ as a quadratic equation of log unemployment-adjusted income.

In principle, estimating the idiosyncratic level and slope parameters α^i and β^i separately from the state variable z^i_t requires observing the full history of workers' earnings paths. Lacking that information in the SCF, I take a conservative approach which both (1) minimizes the variance of workers' expected earnings path and (2) maximizes heterogeneity between workers, in order to give this model the best possible chance to match the observed high and variable interest rates at which households are observed to borrow. Details are given in Appendix C. Intuitively, I matching each households' deviation from the age-specific unemployment-adjusted earnings function and assign them idiosyncratic level and slope parameters α^i and β^i consistent with their relative position in the distribution and the unconditional marginal distributions of those variables. The remaining parameters are taken directly from the calibration in Guvenen et al. (2021). The result of this is an extremely sophisticated projection of the expected future income of every household in the SCF conditional on their primary earner's age, the household's present income, and employment status.

I perform several exercises to validate that this model produces reasonable implied estimates of income growth. Even though the model only takes as an input households' income and primary-earner's age, the model-implied estimates correlated with households' subjective levels of optimism about income growth according to the SCF. Table 3 shows the model-implied income growth estimates disaggregated by households' reported levels of optimism, which are reported as either above, equal to, or less than inflation to the SCF surveyors.

Table 3: Model-Implied Income Growth Expectations by Subjective Optimism Levels

Optimism	$E[\Delta L_{t+1}]$	N
>Inflation	0.0315	2900
=Inflation	0.0279	5561
<Inflation	0.0274	4609

Notes: The categorical "Optimism" variable is taken from SCF question x7364 and asks whether the respondent expects their family income to grow more, the same, or less than inflation over the next year. ΔL_{t+1} is an abbreviated notation meaning $\log L_{t+1} - \log L_t$.

As can be seen from the table, there is a 40 basis-point gap, or 15% of the base value, in the model-implied income growth expectations of those who report having high income growth expectations over those who do not.

Figure 8: Model-Implied $E[\Delta L_{t+k}]$ by Age and Income

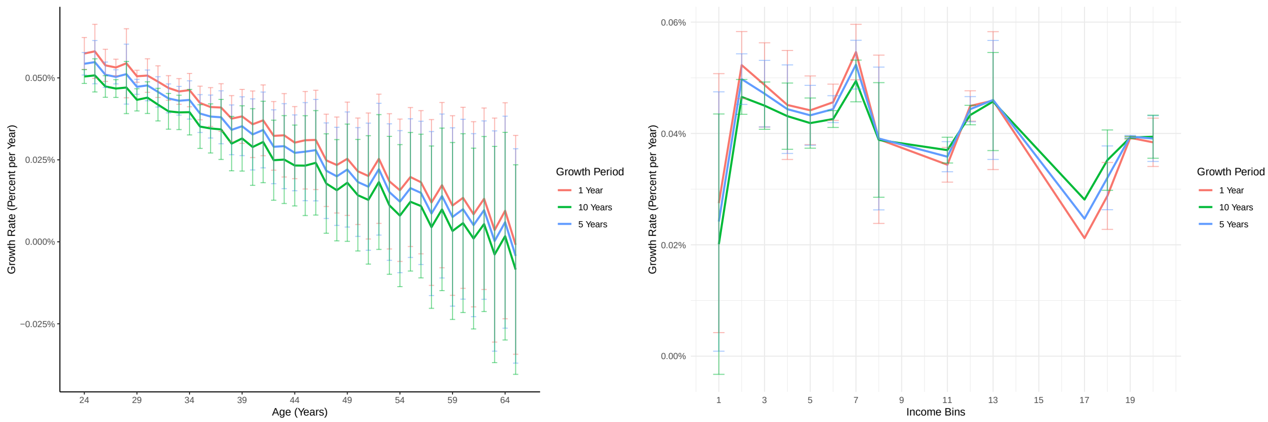

Notes: Lines show the implied expected growth rate of log income at the k-year horizon, denoted $E[\Delta L_{t+k}]$, aggregated unconditionally by age and by income ventiles. Standard deviations of each variable within each age group and income group are depicted as error bars. 1y, 5y, and 10y implied expectations are shown.

I also plot the level and spread of these estimates by each of age and income unconditionally. Figure 8 shows in two panels the distribution of these estimates. There is a clear and consistent downward slope with regards to age, as is to be expected. Within each age group, there is a dispersion of expectations, with standard deviations depicted in error bars. The pattern with regards to income ventiles is less consistent. In all cases, the expectations at each horizon are tightly correlated. Overall, these validation exercises confirm that the model produces estimates of expected income growth as a function of age and income that are consistent with theory and with subjective expectations of income growth.

I now turn to the question of whether these income growth expectations predict borrowing behavior. Equation 5 predicts that intertemporal demand r is is the sum of the pure impatience parameter β and the demand for intertemporal smoothing, which is governed by expected income growth. If individuals have CRRA preferences with risk-aversion parameter γ, then equation 2 can be rewritten as

$$-\log \beta = r - \gamma(E_t[\Delta L_{t+1}]) \quad (5)$$

Thus, if I empirically estimate the equation

$$r_{i,t+1} = \alpha_0 + \alpha_1 E_t[\Delta L_{t+1,i}] + \varepsilon_{it}$$

The empirical estimate of α_1 should exactly line up with the risk-aversion parameter γ, and all remaining heterogeneity in ε_{it} can be attributable to heterogeneity in the impatience parameter β. Table 4 reports the results of this exercise. I use the 1-year model-implied growth rate in all specifications because, as Figure 8 shows, the model-implied growth rates at 1-10-year horizons are all tightly correlated.

The first thing to observe is that the R^2 of these estimates is extremely low, no higher than 2.2% R^2 even when including age fixed effects, with within-R^2 being less than 1% both when predicting maximum interest rate on any debt and interest rate on non-credit-card debt at the individual level in the SCF. The second thing to observe is that it does not matter qualitatively whether or not age fixed effects are included in the regression in terms of the sign and significance of the estimates. The most interesting finding is that the sign of the estimate flips depending on whether credit card debt is included. When credit card debt is included, households who choose to borrow at the highest interest rates are indeed those with the highest expected labor income growth, as would be predicted by consumption smoothing. However, the magnitude of these estimates, which range from 43 to 18 depending on whether age fixed-effects are included,[3] are far out of scale with reasonable expectations of the risk-aversion parameter γ.[4] When credit cards are excluded, the relationship goes the opposite direction from theory: those borrowing at higher rates have *lower* expected labor

[3] Theory implies that age should not be controlled for; and due to the construction of the model estimates, including age strongly reduces the coefficients because it is tightly correlated with the constructed measure of income growth expectations. It is included for completeness and to potentially control for age-specific heterogeneity in beliefs not captured by the model. In any case it is qualitatively irrelevant and the interpretation of the results is not affected by its inclusion or exclusion.

[4] in Epstein-Zin utility, discussed in Section 6.1, replace "risk-aversion parameter γ" with "inverse-IES parameter $1/\psi$" in this sentence.

Table 4: Model-Implied Expectations, Not Stated Expecatations, Predict 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	-0.1044	0.0346	-0.0739	0.0885
	(0.2591)	(0.2294)	(0.1063)	(0.0784)
Constant			14.48***	4.585***
			(0.2643)	(0.1948)
Fixed-effects				
Age	Yes	Yes	No	No
Fit statistics				
Observations	13,069	13,069	13,069	13,069
R^2	0.01876	0.02189	0.00214	0.00083
Within R^2	0.00781	0.00555		

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

income growth. In any case, the extremely low R^2 of these regressions indicate an extremely poor fit of the consumption-smoothing model for explaining household borrowing rates, leaving heterogeneity in impatience as the remaining source of variation.

One criticism of this approach is that individuals' expectations may not be either irrational, or that they may have private information unobserved to the econometrician. However, the inclusion of the optimism variable from the SCF shows that households' subjective beliefs about income growths, controlling for the model's objective estimates, appear to be entirely irrelevant.

With estimates of the distribution of L_{t+1} in hand I can return to equation 2 and empirically calculate the β implied by r. I calculate β using a no-smoothing benchmark, that is, calculating the demand for consumption smoothing from the initial condition of $C_t = L_t$, that is, hand-to-mouth consumption. Using the life-cycle model's numerical estimates of $E[L_{t+1}]$ and σ_L,[5] I can calculate the upper bound of β_i (equivalently, the lower bound of impatience) $\beta_{UB,i}$ for each observation in the SCF. I use a reasonable relative risk aversion parameter of 1.9 from

[5] To be conservative, i.e. avoid estimating a β further from 1 than reality, and to avoid making a strong assumption about risk preferences, I use $\sigma_L = 0$ as the base case in this estimation.

Laibson et al. (2024).

Figure 9: Implied Distribution of Upper-Bounds on Impatience βs

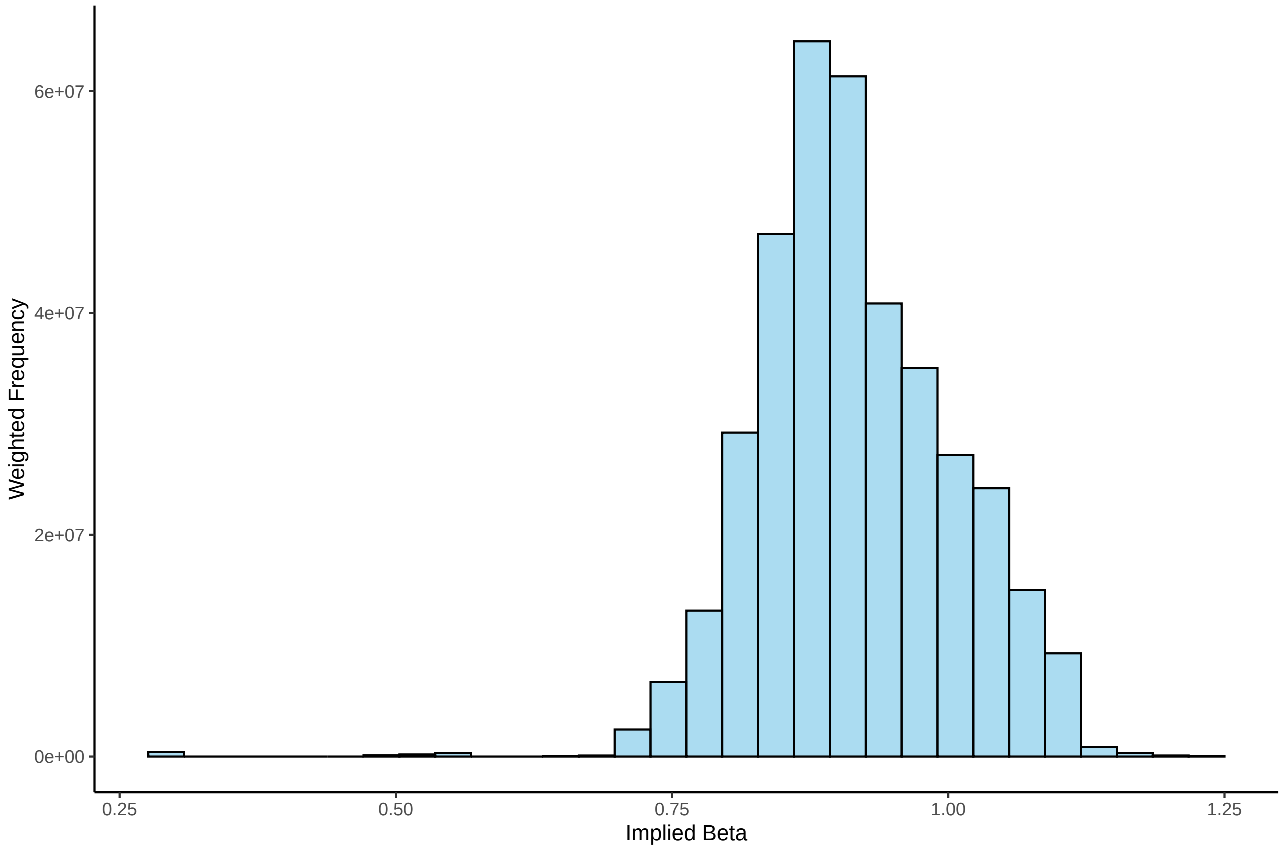

Notes: Implied upper-bound βs are calculated using Equation 5 using $\gamma = 1.9$

Figure 9 shows the distribution of implied upper bound βs, equivalently described as lower bounds of impatience. Interestingly, 20% of households borrow at *lower* interest rates than their implied consumption-smoothing demand would suggest, suggesting according to this procedure that they have *negative* rates of time preference. While this may seem counterintuitive, consider that Epstein and Zin's first attempt to fit their model to asset market data (Epstein and Zin (1991)) resulted in them estimating negative rates of time preference, and that the marginal investor in asset markets is likely drawn from the most patient segment of the distribution. The median implied upper bound β is 91%, indicating at least 9% annual discounting of utility. 2% of people have implied βs below 75%, indicating that they discount utility by at least 25% per year. In the next Section, I describe and give the results of a procedure with more extensive data and identification requirements that gives point estimates instead of upper bounds.

4 Point Estimates: A Sufficient Statistics Approach Using Default Decisions

In this section, I ask what we can learn about households' impatience from debt repayment decisions, rather than debt transaction decisions which, as discussed in Section 3, can only hope to provide a lower bound to impatience even in the best possible world. I show how in a simple endogenous default model with separable utility, the causal effect on default decisions of identified shocks to repayment obligations at two different time horizons is sufficient to point-identify the subjective rate of time preference (β) between those two different points in time. The economic intuition of this approach is that if borrowers are very impatient, they will be very sensitive to short-run changes in their payment schedule; conversely, a patient borrower will strategically default only when there are long-term benefits from doing so. While interest rates provide valuable lower bounds, they do not point identify forward-looking discount factors. For this, we must turn to settings in which borrowers' repayment obligations are experimentally or quasi-experimentally manipulated. In such environments, repayment or default responses to changes in short- and long-term payments reveal the relative weight households place on the near versus distant future.

This section develops a unified sufficient-statistics framework that interprets such repayment sensitivities as direct evidence of impatience. The intuition is simple. Suppose a borrower faces a repayment obligation m in the first period and $(L - m)$ in the second period of a two-period debt contract. If we perturb m or L slightly—for instance, through policy interventions that shift borrowers' repayment schedules—the borrower's default decision may change. The magnitudes of these changes encode the tradeoff the borrower makes between immediate consumption relief and future repayment relief governed by the discount factor.

This model has several powerful advantages. For one, it delivers point estimates instead of lower bound. Unlike laboratory experiments, it utilizes revealed preference with real stakes and realistically common time horizons. Unlike approaches based on asset market data, it can be calculated for the majority of U.S. households who do not meaningfully participate in financial markets. It explicitly accounts for the endogeneity of household default and its timing, and indeed uses that information to obtain the estimate. It requires no functional form assumption for the utility function, only requiring that it be increasing, concave, and time-separable.[6] This model is particularly powerful in its ability to relate the discount rate to readily-measurable moments which are indeed measured in existing papers in the literature. So far I have identified two suitable papers which measure and report the required estimates with valid causal identification strategies: Dobbie and Song (2020) and Ganong and Noel (2020).

Dobbie and Song (2020) use an explicit randomized controlled trial which varies short-term and long-term payments due on individuals' modified credit card repayment plans. Ganong and Noel (2020) utilize a sharp regression discontinuity design in the HAMP mortgage modification program which affected mortgage repayments due at different time horizons. I interpret their results through the lens of this unifying model. The model compares the effect of shocks to repayment obligations at two different time periods to impute an implied forward rate between those two time periods. I show that the empirical size of these two effects are sufficient statistics for the implied forward rate between those two time periods. Standard errors are also easily calculated with the delta method.

This method can be used to calculate implied time preferences for any paper which reports the effect on repayment of two or more shocks to payments owed at two or more time horizons and does not require the full data. It also accommodates any increasing concave utility function, not requiring any functional form such as CRRA or EZKP utility as I assume

[6] I discuss extensions to non-time-separable utility functions in Section 6.1. The assumption of time-separability is not restrictive, but additional assumptions are needed when the intertemporal elasticity of substitution is divorced from risk-aversion.

in Sections 3 and 6.1.

4.1 Two-Period Endogenous Default Model

The basis of the model is a 2-period endogenous default model which features the same basic trade-off as Indarte and Kanz (2023). The household weighs penalty of default (legal, social, or dynamic) against the value of extra consumption from foregoing debt payments. The model is set in 2 discrete time periods $t \in \{1, 2\}$, which represent the two dates at which payments are due that can be potentially modified. The starting balance of the loan at $t = 1$ is defined as L, and the individual payment size, assumed to be constant, is denoted m (for "monthly payment"). The second period payment is therefore $L - m$ in the two-period case. In this base case, $L = 2m$ (constant monthly payments).

I assume only one dimension of uncertainty to make the model tractable: period 1 income is uncertain and is distributed $y_1 > 0 \sim F(y_1)$. Note that no specific assumption on the income process is necessary, other than that it is sufficiently uncertain to generate default: if income is too low, the marginal utility is high, and eventually exceeds the punishment of default. The punishment for default is equal to σ, as in Indarte and Kanz (2023). This is a comprehensive sum of all punishments that befall a borrower who defaults: dynamic exclusion from the credit market, social stigma, and attempts by the lender to garnish wages.

The borrower has a strictly increasing and strictly concave utility function over consumption, additive across period with a time discount factor of β

$$
\begin{aligned}
V_2^P &= u(C_2^P) \\
V_2^N &= u(C_2^N) - \sigma \\
V_1^P &= u(C_1^P) + \beta \left[\int_0^{y_2^*} V_2^N dF(y_2|y_1) + \int_{y_2^*}^{\infty} V_2^P dF(y_2|y_1) \right] \\
V_1^N &= u(C_1^N) + \beta V_2^N
\end{aligned}
\tag{6}
$$

Where y_2^* is defined implicitly by the borrower's endogenous default threshold in period

2. The budget constraints are:

$$
\begin{aligned}
C_1^P &= y_1 - m \\
C_1^N &= y_1 \\
C_2^P &= y_2 - (L - m) \\
C_2^N &= y_2
\end{aligned} \tag{7}
$$

Here P and N denote payment and non-payment, respectively. The borrower defaults if the utility gain from avoiding payment exceeds the punishment σ, which summarizes legal, social, and dynamic consequences of default.[7]

Let y_2^* be the income threshold at which the borrower is indifferent in period 2:

$$
u(C_2^P(y_2^*)) = u(C_2^N(y_2^*)) - \sigma. \tag{8}
$$

Similarly, the borrower defaults in period 1 if the value of non-payment today plus the continuation value of facing the threshold y_2^* tomorrow exceeds the value of payment today plus discounted continuation value. The probability of default in period 1 is therefore

$$
p_1 = F_1(y_1^*),
$$

where y_1^* is determined implicitly by

$$
u(C_1^{P*}) + \beta \mathbb{E}[V_2^P] \; = \; u(C_1^{N*}) + \beta \mathbb{E}[V_2^N]. \tag{9}
$$

Equations (8) and (9) jointly determine default behavior. I do not need closed-form solutions for y_1^* or y_2^* to study comparative statics. Instead, I focus on how small changes in m and L perturb p_1 and thereby reveal the discount factor β.

[7] This formulation follows a large literature modeling default as the result of a discrete comparison between current utility gains and continuation-value losses from dynamic exclusion.

4.2 Default Sensitivities as Sufficient Statistics

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

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

Under the envelope condition for the default threshold in period 1, the comparative statics of y_1^* follow from differentiating (9). Straightforward algebra (with full derivation in Appendix A) yields the two key sensitivities:

$$\begin{aligned} \frac{\partial p_1}{\partial L} &= f(y_1^*)\,\beta(1-p_2)\,\frac{\mathbb{E}\left[u'(C_2^P(y_2)) \mid y_2 > y_2^*\right]}{u'(C_1^{P*}) - u'(C_1^{N*})}, \\ \frac{\partial p_1}{\partial m} &= f(y_1^*)\frac{u'(C_1^P)}{u'(C_1^{P*}) - u'(C_1^{N*})} - \frac{\partial p_1}{\partial L}. \end{aligned} \tag{10}$$

These expressions capture the essence of the identification strategy. A marginal increase in m reduces C_1^P, thereby increasing current marginal utility. A marginal increase in L increases the future payment burden and thus reduces the continuation value for borrowers who expect to repay in period 2. The ratio of these two effects reveals the intertemporal preference parameter that equates marginal utilities of consumption across periods.

These comparative statics can be measured empirically and mapped to the primitives of the model. In particular, I can isolate the discount rate β by rearranging the previous equations as such:

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

This simplifies extremely elegantly for the simplifying case of $y_2 = y_1$ (or, more generally, $E[u'(y_2 - m)|y_2 > y_2^*] = u'(y_1 - m)$), that is, the case where expected marginal utility tomorrow is exactly the same as marginal utility today. This implies no late default ($p_2 = 0$), which is not only an extremely convenient assumption but very nearly empirically true, as shown in

Appendix E. It also implies that short-term expectations of income growth are close to zero, which is also empirically consistent with the data in Section 3. This simplification implies that:

$$E\left[u'\left(c_2^P\right)|y_2 > y_2^*\right] = u'\left(c_1^P\right)$$

This can easily be adjusted by putting bounds on $\frac{E\left[u'\left(c_2^P\right)|y_2>y_2^*\right]}{u'\left(c_1^P\right)}$, for example using the life-cycle model in Section 3. Quantitatively, expectations of income growth are not high over the length of auto loans (approximately 5 years) either with rational expectations or in perceptions. The result of the most simple assumption (that the period one income shock is the only source of uncertainty), is that I obtain the surprisingly simple expression:

$$\frac{\frac{\partial p_1}{\partial L}}{\frac{\partial p_1}{\partial m} + \frac{\partial p_1}{\partial L}} = \beta \tag{11}$$

This maps neatly to intuition: if default decisions are entirely forward-looking and strategic, with borrowers who are not liquidity constrained and do not value present and future consumption differently, then the observed sensitivity of default will be purely a function of the total loan balance and the sufficient statistic formula will evaluate to $\beta = 1$: perfect patience. On the other extreme, if default decisions are entirely myopic and borrowers do not consider the future whatsoever, the default decision is purely a function of the monthly payment and the sufficient statistic formula will valuate to $\beta = 0$: perfect myopia. The formula smoothly nests all intermediate cases, where borrowers put some weight on current conditions and some weight on future conditions – the intermediate point is their discount rate.

The assumption of no default in period 2 is palatable in large part for empirical, not only theoretical reasons: late default, especially on auto loans, is extremely rare. Section E shows these calculations. Not only are defaults extremely rare in the late months of auto loan contracts, but they are rare conditionally on narrow credit score categories and attrition. One might worry that the default rate as a function of time outstanding on the loan might slope downwards for spurious reasons, including the elimination of high default risks from the population early. The calculations in the appendix show that this is not the case: even conditioning on narrow credit score groups and adjusting for attrition to calculate true conditional default rates, very few borrowers default late in the life of the contract.

4.3 Recovering the Annual Discount Factor

The model is set in two periods, and the formula produces the revealed-preference discount rate between those two periods. If the term structure slopes down (for example, if it is hyperbolic as in Laibson et al. (2024)), then this multi-period discount rate, a single number, is not sufficient to identify the entire term structure. To proceed, I proceed under the assumption of constant and time-consistent discount rates, and transform this from a cumulative rate over a variable time period into an annual discount rate, I adjust it as:

$$\beta_{annual} = \beta_{cumulative}^{1/T} \tag{12}$$

Where T is calculated as the time between average default and average loan maturity.[8] This definition of T is chosen to best represent the economic tradeoff that the borrower makes, weighing in their decision both current consumption at the time at which the default decision is made and marginal future consumption at the loan end date affected by variation in the total loan balance.

I calculate standard errors using the delta method, incorporating sampling-based uncertainty in my estimates of both comparative statics and T. The estimand is $\beta_{cumulative}$, which I hereafter call simply β:

$$\beta = \left(\frac{\frac{\partial p_1}{\partial L}}{\frac{\partial p_1}{\partial m} + \frac{\partial p_1}{\partial L}} \right)^{1/T} \tag{13}$$

In Appendix F, I provide the full delta method calculations.

[8] In the auto loan sample (introduced later in Section 4.5) this is 3.2 years unconditionally.

4.4 Application in Experimental and Quasi-Experimental Settings

Armed with the sufficient-statistics formula (13), we can reinterpret results from quasi-experimental studies through the lens of time preferences. Two prominent settings allow direct computation of A and B: **?**, who study credit card borrowers, and **?**, who study mortgage borrowers.

At first glance, these studies report seemingly contradictory findings. Dobbie and Song estimate that reducing short-term payments has almost no effect on default probability, while long-term payment reductions matter substantially. Ganong and Noel report almost the opposite: short-term payment relief affects repayment behavior far more than long-term mortgage forgiveness. Using their point estimates, however, and controlling for the vastly different time horizons between the respective shocks, both studies imply surprisingly similar discount factors.

Let T_{DS} and T_{GN} denote the effective horizons in each study. Using the reported comparative statics:

$$\beta_{\text{GN}} = 0.76 \quad \text{per year,}$$
$$\beta_{\text{DS}} = 0.58 \quad \text{per year.}$$

Despite surface-level differences, both populations appear highly impatient. Credit card borrowers are even more impatient, consistent with intuition: such borrowers tend to be liquidity constrained, have limited access to low-cost credit, and may face greater income volatility.

These findings demonstrate the power of the sufficient-statistics approach. Without needing full structural models or microdata, and without imposing parametric assumptions on income dynamics or utility curvature, we can infer long-run forward discount factors for heterogeneous populations.

In Section 4.5, I apply this approach to auto loan data, where differences in maturity and payment amounts across borrowers provide the variation required to estimate A and B

directly.

4.5 Application to Auto Loan Data

Auto loans offer an attractive empirical environment for applying the sufficient-statistics method. Maturities vary substantially across borrowers, even conditional on credit score, and contractual payments are highly salient. Moreover, auto loans constitute one of the most common forms of installment credit in the United States, with wide participation across the income distribution. For many households—especially those with limited liquid wealth—auto loans represent the primary medium through which repayment and default decisions reveal intertemporal tradeoffs. I use a 1% random sample of U.S. credit reports obtained through a partnership with one of the nation's leading credit bureaus. The sample includes over six million auto loans originated across the past decade. Table 5 summarizes the key variables: loan amounts, balances, monthly payments, maturities, credit scores, ages, and imputed incomes.

Table 5: 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 repayment horizon is typically five years (60 months), and the three most popular maturities of 36, 60, and 72 months account for over half of observations. Figure 10 shows origination volumes by year, indicating steady and large coverage of this data since 2010.

Figure 10: Auto Loan Originations Over Time

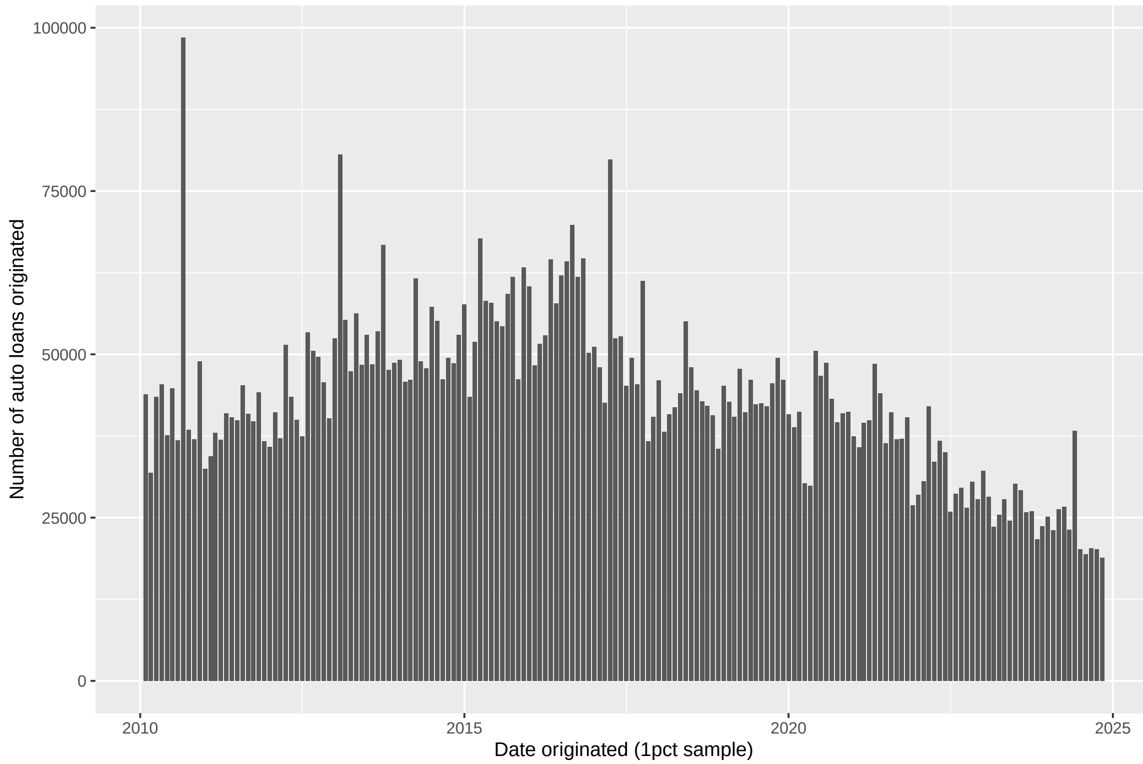

Notes: Number of auto loans originated in the 1% credit bureau sample, plotted by origination year.

To identify the comparative statics $\partial p_1/\partial m$ and $\partial p_1/\partial L$, I leverage heterogeneity in contractual maturity across borrowers with the same credit score. A key identifying assumption, necessary to proceed, is the following exclusion restriction:

> *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 given lenders' reliance on standardized underwriting algorithms that map credit scores to allowable loan structures. While more creditworthy borrowers may indeed receive different contract terms, these differences are mediated almost entirely by credit score, which is observed. Conditional on score, residual sorting across maturities must be ruled out.

Under this assumption, variation in m and L across similar-score borrowers yields quasi-experimental shifts in repayment burdens. Because households with different maturities but identical credit scores face similar income risk, any systematic difference in default hazard can be attributed to the intertemporal allocation of payments rather than underlying heterogeneity.

The assumption is essentially an exclusion restriction that pre-existing creditworthiness only affects maturity and loan size through observable credit score. That is, maturity and loan size are as-if randomly assigned conditional on credit score: while better borrowers may systematically acquire different loan terms (which is certainly the case), the credit score is an accurate summary statistic for the borrower's creditworthiness as it is used to define the loan terms. This is a good assumption because lenders typically use standardized underwriting technologies to decide loan terms based on verifiable information (See, e.g., Yannelis and Zhang (2023)), although this cannot itself be observed in the data per se. The result of this assumption is that conditioning on credit score, the only remaining unknown factor which decides whether the individual will default on their loan is the income shock that they experience after taking out the loan, which lenders and borrowers are no more able to predict than the bank. I document the extent to which both monthly payment size and total loan size are correlated with default rate conditional on narrow credit score categories and zip-by-month fixed effects. In principle, this correlation may be driven by adverse selection or moral hazard: the estimation in this section is exactly correct in the limiting case where the correlation is entirely due to the moral hazard (that is, larger loans cause borrowers to be more willing to decide to default) that I model, and not at all due to adverse selection (worse borrowers take larger loans). The model explicitly calculates the discount rate under this assumption of moral hazard, precluding the possibility of adverse selection. This is arguably a strong assumption, as it precludes the possibility of privately-known information by the borrower or lender that would result in larger or longer loans being systematically assigned to borrowers whose true creditworthiness is better than their credit score suggests. I argue that despite this potential defect, this assumption is the most palatable one consistent with being able to estimate the model consistently and robustly across time, geography, and credit score categories. To bolster the plausibility of this assumption I show that the unconditional default rate as a function of maturity of auto loan contracts varies substantially but not in any systematic way. Lenders and borrowers must come to a joint agreement on the size of the loan and the maturity of the loan as a function of the credit score – this endogenous process removes the average unconditional correlation between default rate and default probability while still leaving room for idiosyncratic factors to shift the loan size and maturity (in a mean-zero way), creating an ex-post correlation between payment size/loan size and default rate. The assumption is also bolstered by a simple, directional incentives argument: if lenders could privately observe worse (better) credit risks conditional on credit score, they would want to give those borrowers *smaller* (larger) loans, which is the opposite of what I observe in the data.

The difficulty of separately identifying adverse selection from moral hazard has been recognized widely in the literature at least since Ciappori and Salanie (2000), and the gold standard approach (Karlan and Zinman (2009) requires experimental variation which is infeasible to collect at scale in the U.S. population. Even if the Karlan and Zinman test could be performed on a representative sample of US borrowers, it would still be insufficient to identify adverse selection at the market outcome level – this test still only identifies the narrow adverse selection present in a single instance of the takeup decision of a single loan. In the absence of such experimental data and faced with this identification problem, I find the assumption of no adverse selection conditional on credit score to be the most palatable assumption possible in this setting to answer the question at hand, which is the change in default rate caused by higher loan balances and higher monthly payment. I am encouraged in believing that the remaining variation in monthly payment and total balance, after controlling for credit score and zipcode-by-month fixed effects, is likely to be purely idiosyncratic by Figure 11, which shows a highly idiosyncratic relationship between unconditional default rate and loan terms. I therefore proceed under the necessary assumption that lenders set loan terms in such a way as to remove the correlation between loan term and default rate by, for example, higher downpayments for less creditworthy customers, and that differences in contract monthly payment and origination balance which are not explained by narrow credit score fixed effects are responsible for changes in default rate that are not explained by

Figure 11: Idiosyncratic Unconditional Relationship Between Maturity and Default

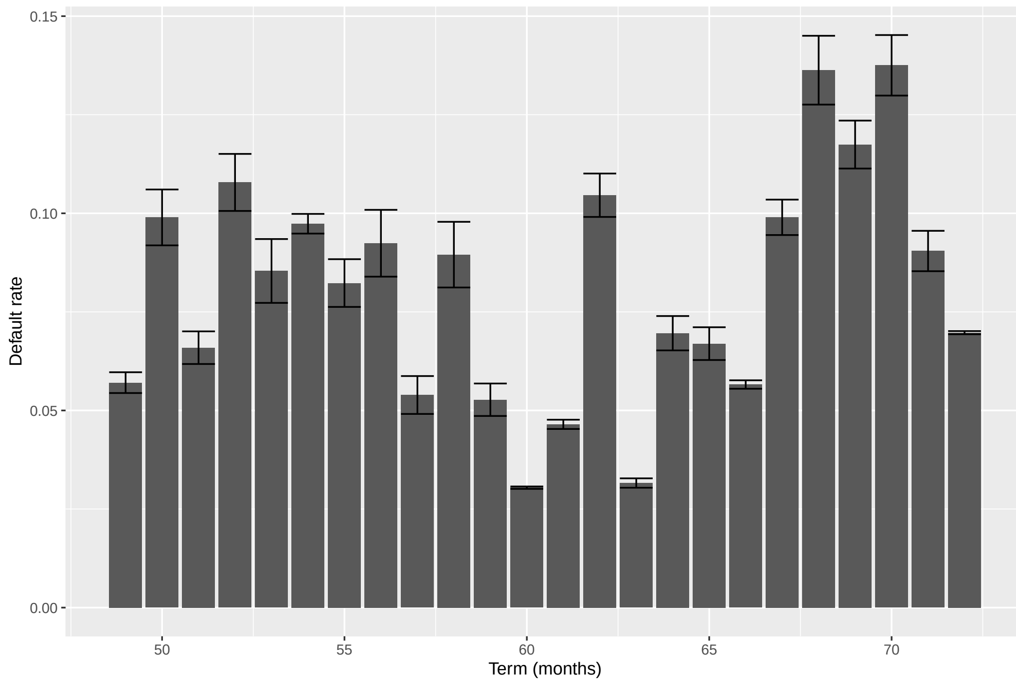

a pre-existing and unobservable correlation between creditworthiness and contract terms. Furthermore, even if there does exist adverse selection in the direction consistent with the results in Table 6, if the degree of adverse selection is equal at all maturities it will cancel out due to the division operation in Equation 11. Therefore what is needed to violate this exclusion restriction and bias the estimates is maturity-varying adverse selection in the opposite direction predicted by lender incentives.

Table 6 reports estimates from two separate logit probability models in which the dependent variable is an indicator for default. Each specification includes an identical set of fixed effects—score-ventile fixed effects and zipcode–by–month fixed effects—which flexibly absorb borrower risk classification and local economic conditions at the time of loan origination.

The first regression estimates the sensitivity of default to the short-term required payment burden. Specifically, I estimate

$$\text{Default}_i = \beta_S \log(\text{MonthlyPayment}_i) + \gamma_{v(i)} + \delta_{z(i),t(i)} + \varepsilon_i, \tag{14}$$

where $\gamma_{v(i)}$ denotes score-ventile fixed effects and $\delta_{z(i),t(i)}$ denotes zipcode–month fixed effects. The coefficient β_S measures the percentage-change sensitivity of default with respect to the current monthly payment obligation. The estimate of $\beta_S = 0.50$ (s.e. 0.0083) indicates a strong and precisely estimated association: borrowers facing higher required monthly payments are substantially more likely to default, holding constant long-term obligations and all fixed effects.

The second regression instead focuses on the long-term debt burden by replacing monthly payments with the logarithm of total payments owed over the life of the loan:

$$\text{Default}_i = \beta_L \, \log(\text{TotalPayments}_i) + \gamma_{v(i)} + \delta_{z(i),t(i)} + \varepsilon_i. \tag{15}$$

Here β_L captures the sensitivity of default to the overall, long-horizon repayment obligation independent of the short-term payment schedule. The estimate $\beta_L = 0.4283$ (s.e. 0.0058) likewise shows a strong positive association: larger total repayment obligations are linked to a higher probability of default.

Across both specifications, the inclusion of rich fixed effects ensures that identification comes from within–zipcode–month and within–score-ventile variation in payment obligations. The results highlight that default risk is highly responsive to both short-term liquidity pressure (monthly payment) and long-term financial burden (total payments). The somewhat larger short-run coefficient suggests that the immediate payment obligation exerts a slightly stronger effect on default behavior than the overall debt load, though both channels are quantitatively important.

Table 6: Estimates of Default Sensitivity to Short and Long-Term 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 x 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*

Here, I report the results of the sufficient statistics exercise deployed over the entire auto loan dataset, whereby I use the multiple regression analysis described above to estimate Equation 11 with standard errors as described in Appendix F:

Table 7: Estimates and Standard Errors of β

Variable	Estimate	Std. Error	Variable Description
$\frac{\partial p_1}{\partial m}$	0.500	0.008	Default sensitivity to monthly payment
$\frac{\partial p_1}{\partial L}$	0.428	0.006	Default sensitivity to total payments
T	3.223	0.004	Average term (in years) between default and loan end date
β	0.787	0.003	Derived annualized discount rate parameter: $\beta = \left(\frac{\frac{\partial p_1}{\partial L}}{\frac{\partial p_1}{\partial L} + \frac{\partial p_1}{\partial m}}\right)^{1/T}$

4.6 Heterogeneity

5 Heterogeneity in Estimated Discount Factors

The preceding section establishes that, on average, auto loan borrowers exhibit steep impatience. But the distribution of β is highly heterogeneous. Understanding this heterogeneity is crucial: it informs which households are most sensitive to liquidity shocks, which are most in need of consumption smoothing, and how credit policies redistribute welfare across different types.In this section, I use variation across income, credit score, and geography to construct conditional distributions of discount factors. This approach yields a granular mapping from borrower characteristics to impatience that is both intuitive and empirically rich.

This section presents heterogeneity analyses of the estimated discount factor β across (i) credit score ventiles, (ii) income ventiles, and (iii) U.S. states. For each subgroup, the following procedure is applied independently:

- Define a partition of the data (score ventile, income ventile, or state).
- Within each partition, recompute score ventiles to preserve fixed-effects structure.
- Estimate two logistic fixed-effects models of default: one on $\log(\text{scheduled payment})$ and one on $\log(\text{total payments})$.
- Compute the derivative sensitivities of default with respect to payment size and loan length.
- Estimate the average time to default and the implied intertemporal horizon ΔT.
- Obtain β and its delta-method standard error.
- Plot the results with point estimates and ± 2 standard-error bands.

Figure 12 shows heterogeneity by credit score ventile, Figure 13 shows heterogeneity by income ventile, and Figure 14 shows heterogeneity across states. The full set of state-level estimates is reported in Table 8.

Figure 12: Estimated Discount Factor β by Credit Score Ventile. Error bars show ± 2 standard errors.

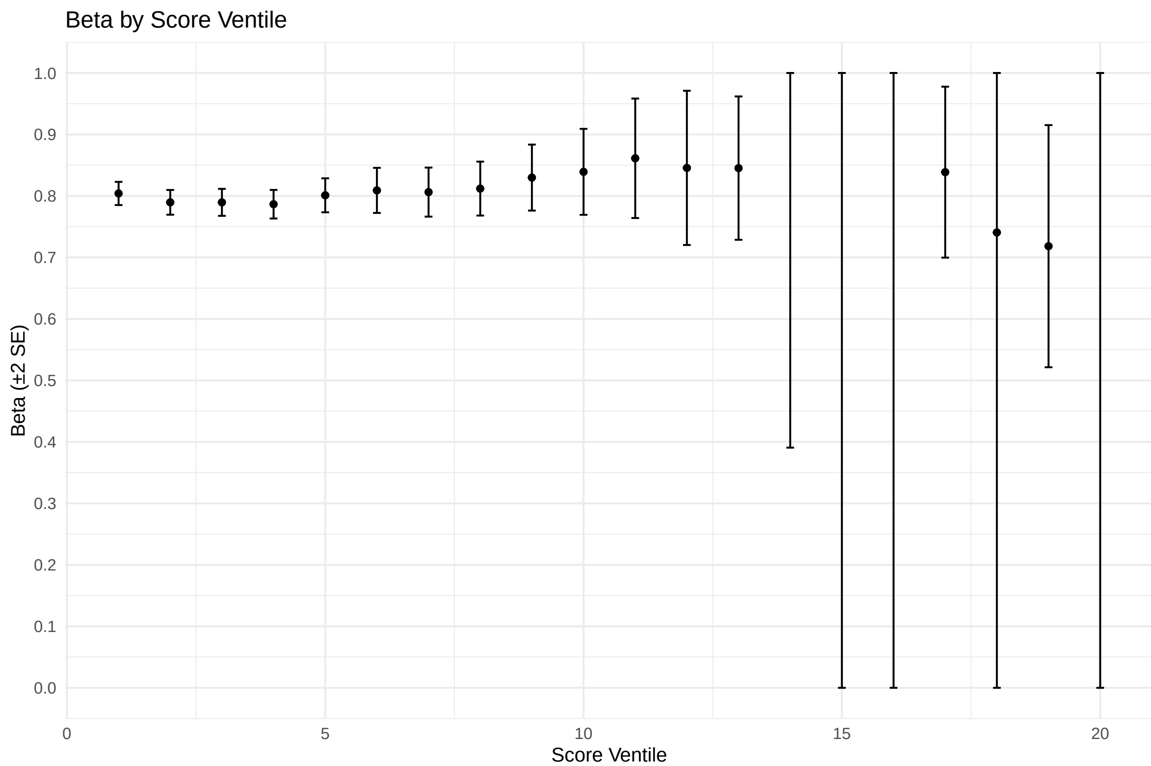

Figure 13: Estimated Discount Factor β by Income Ventile. Error bars show ± 2 standard errors.

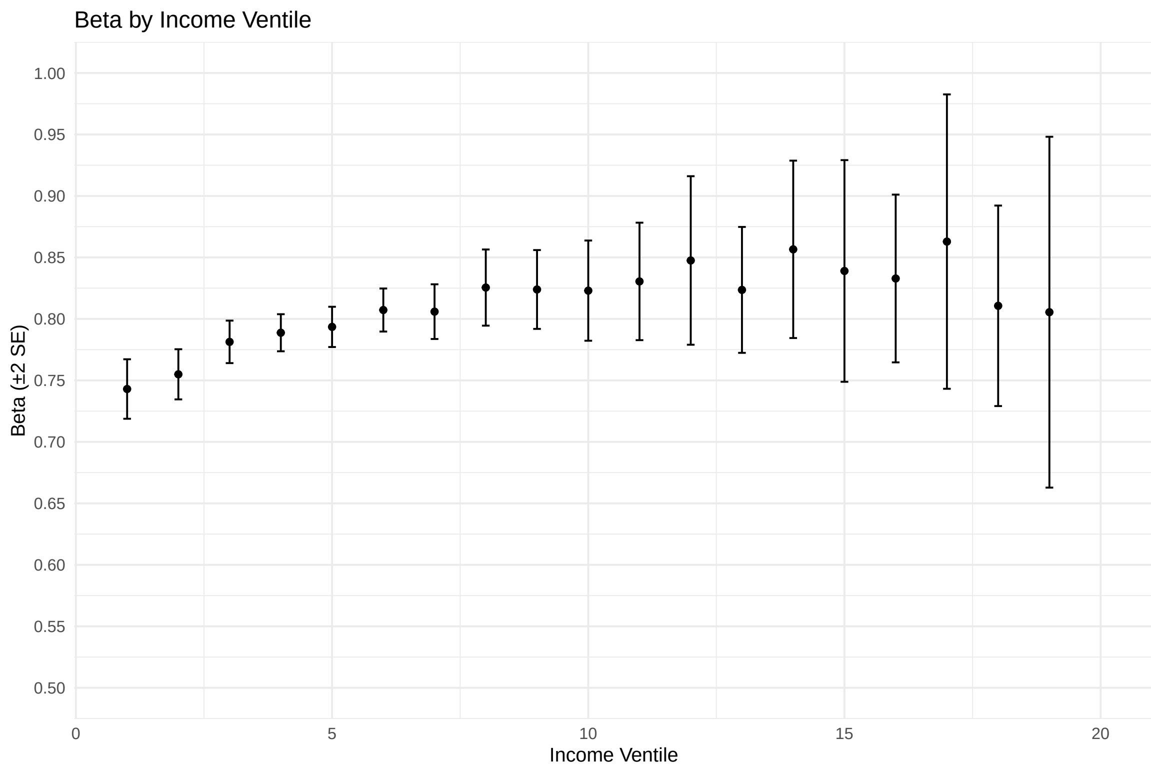

Figure 14: Estimated Discount Factor β by State (Sorted by Point Estimate)

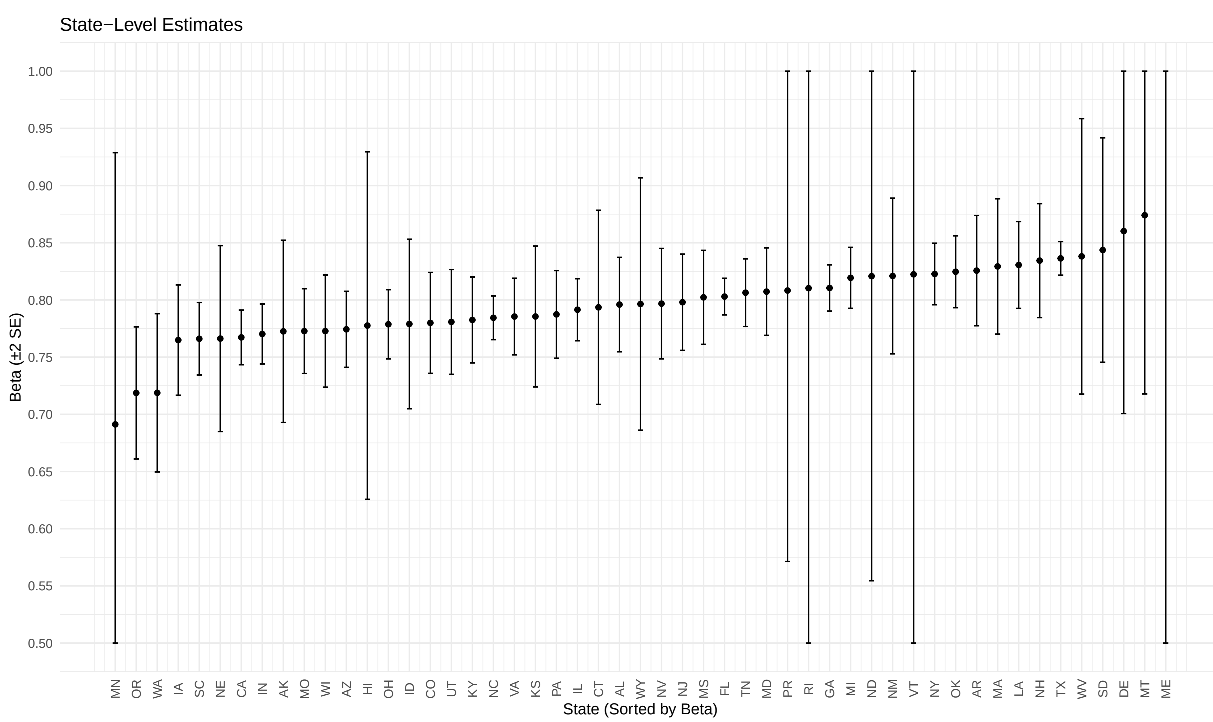

Table 8: State-Level Discount Factor Estimates

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

In the upper score ventiles, the β estimates exhibit extremely wide error bands. This is mechanically driven by the fact that high-score borrowers rarely default in the sample. With so few observed defaults, the estimated sensitivity of default risk to payment size becomes highly uncertain, which translates directly into imprecise β estimates. This is not so bad of a defect, however, considering that for very safe credit risks, this sufficient statistics method is not actually necessary: without default, credit is not rationed, and the marginal interest rate that such individuals borrow or save at is able to reveal their discount rates.

The heterogeneity patterns reveal several noteworthy findings. Firstly, the slope of β across credit scores is relatively flat. Despite the large variation in default rates across scores, the estimated discount factor shows almost no systematic slope with respect to credit score. This suggests that the well-established empirical fact that low-score individuals default more frequently *cannot* be attributed to greater impatience. Instead, other forces—such as higher volatility in income, lower financial buffers, or lower perceived penalties of default—must account for the strong relationship between score and default risk. Impatience does not appear to be the primary mechanism.

Secondly, the slope of β across income ventiles *is* very sharp. Lower-income borrowers are estimated to have meaningfully lower β, consistent with substantially higher impatience (or higher effective discount rates). This upward gradient in β is statistically significant – I can, for example, reject the null that the first and third ventiles have the same discount rate.
Overall, the results imply that differences in default rates across *credit scores* do not reflect systematic differences in time preference, but that differences across *income levels do* reflect economically meaningful variation in discount factors. Low-income individuals discount the future more heavily, consistent with tighter liquidity, higher risk exposure, or higher marginal utility of present consumption.

These findings have implications for credit modeling, welfare analysis, and the design of repayment contracts. In particular, they suggest that variation in impatience is correlated more strongly with income—a measure of economic constraint—than with credit score, which more directly reflects historical credit behavior but is not found using this methodology to be correlated with impatience.

6 Extensions

6.1 Non-Time-Separable Utility

The baseline model analyzed in Section 4 assumed that utility is time-separable with identical utility functions in each period. However, numerous failures of time-separable utility functions in explaining asset pricing phenomena have led to the use of more sophisticated models where utility is not time separable, the most popular of which is the Epstein-Zin-Kreps-Porteus ("EZKP" hereafter) utility function (Kreps and Porteus (1978), Epstein and Zin (1991)).

The EZKP utility function is given by

$$U_t = \{(1-\beta)C_t^{1-\frac{1}{\psi}} + \beta \left(E_t U_{t+1}^{1-\gamma}\right)^{\frac{1-\frac{1}{\psi}}{1-\gamma}}\}^{\frac{1}{1-\frac{1}{\psi}}} \tag{16}$$

Notationally, θ is also used to denote $\frac{1-\gamma}{1-\frac{1}{\psi}}$ (Campbell (2018), Epstein and Zin (1991)).[9]

In a representative-agent economy characterized by such preferences, as well as homoskedastic and jointly lognormal asset returns and consumption growth (assumptions I make hereafter in this section), the risk-free short rate is

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

Where σ_c^2 is the variance of the innovation in log consumption and σ_w^2 is the innovation in the log return on wealth (Campbell (2018)). Adapting Equation 5 to separate IES and risk aversion requires only a few edits to this equation. For one, the "Steinbeck Critique" term $(E_t[L_{t+1}] - L_t)$ is now scaled by $1/\psi$ instead of γ, and secondarily, an additional term appears to account for the additional impact of future uncertainty:

$$\beta = \exp -(r - \frac{1}{\psi}(E_t[L_{t+1}] - L_t) + \frac{\theta}{2\psi^2}\sigma_L) \tag{17}$$

This equation can be taken to the data in analogous fashion to the exercise in Section 3.

7 Conclusion

For such a fundamental object in the study of finance as the time value of money, existing approaches to measuring it and its distribution in large populations overwhelmingly rely

[9] Interestingly, when Epstein and Zin first calibrated this model using GMM, they consistently obtained *negative* estimates of time preference ($\beta > 1$) and expressed their puzzlement at this finding (Epstein and Zin (1991)).

on market prices of instruments traded almost entirely by the richest people in the world. This paper develops new empirical methods for measuring household time preferences in environments where traditional asset-market approaches fail. Because most households do not meaningfully participate in the market for risk-free or even risky financial assets, the prices of publicly traded instruments reveal almost nothing about the intertemporal preferences of the median or modal American household. Instead, the relevant revealed-preference margin for the vast majority of citizens lies on the liability side of their balance sheets. Credit, not investment, is the primary domain in which households express their intertemporal tradeoffs.

The central contribution of this paper is to articulate and implement two complementary identification strategies that leverage this fact. The first strategy uses the interest rates households are *willing to pay* as lower bounds on their subjective discount rates. Because credit products are discrete and rationed, observed borrowing rates lie weakly below individuals' true marginal rates of intertemporal substitution. The second strategy uses the *sensitivity of default to short- vs. long-run changes in repayment obligations* to point-identify the discount factor. In an endogenous-default model with concave and time-separable utility, the ratio of these repayment sensitivities is a sufficient statistic for the discount rate. The sign, magnitude, and interpretation of these sensitivities are transparent, robust, and directly connected to empirical moments measured in existing quasi-experimental and experimental studies.

Across both approaches, I find that most households are substantially more impatient than is typically assumed in macroeconomic, household finance, and public finance applications. The median household discounts future utility at rates far exceeding those implied by Treasury yields or representative-agent calibrations. Even after adjusting for consumption-smoothing motives using a state-of-the-art income process, impatience remains large. These findings support the view that a significant share of intertemporal credit demand reflects true impatience rather than expectations of rising future income or volatility-driven smoothing motives.

A second major finding concerns the *distribution* of time preferences. The heterogeneity analysis reveals two sharp patterns. First, impatience rises steeply as income falls: low-income households value current consumption far more heavily relative to the future. Second, impatience is nearly *orthogonal* to credit score, once one conditions on income. This distinction is conceptually important. Credit scores predict *default risk*, but default risk is not the same as impatience – indeed, the results here suggests that they are not even correlated in the population. This iimplies that the link between credit score and default emerges largely through other factors which may include income volatility, lack of access to savings technologies, lack of financial education, or differing nonpecuniary costs of default – not through differences in discount factors. These results suggest that many models and policy narratives which attribute high default rates among low-score borrowers to high impatience may be misspecified. The data support a more nuanced decomposition: low income predicts impatience whereas low score predicts risk.

The implications of these findings are wide-ranging. First, the results have strong implications for public finance. The United States federal government, following OMB guidelines, applies a social discount rate near 2–3 percent. But if a large share of the population has discount factors near 0.8 or below at annual frequency, then policies involving intertemporal redistribution—from the present to the future, or vice versa—have sharply different welfare implications depending on whose preferences are considered normative. For example, as Hendren and Sprung-Keyser (2020) emphasize, the relative value of early-childhood versus adult human-capital policies depends critically on the chosen discount rate. If actual household discounting is far higher than official rates, then policies delivering longer-run payoffs may appear more attractive to planners than they do to the citizens they aim to serve. This raises normative questions about paternalism, welfare weights, and the appropriate aggregation of heterogeneous time preferences.

Second, the results speak directly to the design of consumer credit markets. If a substantial share of households discounts the future extremely heavily, then long-dated repayment relief (e.g., mortgage principal reduction, long-horizon loan extensions) may have limited effects on default behavior, whereas short-term payment relief may have large effects. The contrast between the credit-card and mortgage modification results—once adjusted for horizon lengths—supports this view. This has operational consequences for delinquency management, loss mitigation, and the structuring of hardship programs. Understanding impatience is essential for designing contracts that support sustainable repayment among financially constrained borrowers.

Third, the findings speak to broader questions in macroeconomics and political economy. If the representative citizen is much more impatient than the representative investor, then political pressure to favor present consumption over future investment—including impatience for taxation, infrastructure maintenance, climate mitigation, or public debt reduction—may not be a puzzle at all. Rather, such preferences may reflect the intertemporal tastes of the median voter. The divergence between asset-market discount rates (revealed by patient wealthy investors) and population discount rates (revealed in repayment behavior) has deep implications for how economists interpret intertemporal policy choices in democratic societies.

Fourth, the analysis informs ethical debates about intergenerational equity and the appraisal of long-run public projects. Climate policy, infrastructure renewal, and investments in basic research involve tradeoffs spanning decades. If policymakers adopt social discount rates far below the revealed-preference rates of the median household, they face a philosophical tension: which preferences—those of patient asset holders or those of impatient borrowers—should guide social decision-making? This issue is inherently normative, but empirical measurement of heterogeneity in impatience, as provided here, is essential for grounding such normative debates in the lived behavior of actual households.

Finally, the analysis suggests several fruitful directions for future research. One avenue is to combine the sufficient-statistics approach with structural heterogeneity models to jointly estimate discount factors, risk preferences, and expectations. Another is to investigate the intergenerational transmission of impatience, including whether discount rates correlate with family background, financial education, or exposure to financial shocks. A third is to embed the repayment-based discount-rate estimates into heterogeneous-agent macroeconomic models with borrowing constraints, to study how impatience shapes macroeconomic volatility, savings behavior, and the propagation of business cycles. A final and particularly intriguing direction is to explore normative frameworks that incorporate heterogeneous time preferences more explicitly, acknowledging that societies may face genuine heterogeneity in how citizens value the future.

In summary, I find that the majority of households are significantly more impatient than would be suggested by treasury rates. Adjusting for consumption-smoothing demand does very little to change this, as most households in surveys are actually quite pessimistic about short- to medium-term income growth. This implies that we should take seriously the ramifications, for political and economic behavior, of the fact that the vast majority of U.S. citizens are far more impatient than the marginal asset market investor. It also implies that when designing asset pricing models to make inferences about the macroeconomy based on asset market data, we should be careful to remember that the marginal investor in these markets is not nearly representative of the U.S. population on an equal-weighted basis. In sum, the evidence in this paper suggests that the time preferences of the median American household differ sharply from those inferred from asset markets, and that impatience plays a first-order role in household financial decisions. Recognizing and appropriately modeling this impatience is crucial for positive and normative research across macroeconomics, household finance, public finance, and political economy. By developing tools to measure discount factors outside asset markets, and by applying them at scale, this paper aims to provide a foundation for a more empirically grounded understanding of intertemporal preferences in the broader population.

References

Alvarez, Fernando and Urban J. Jermann, "Quantitative Asset Pricing Implications of Endogenous Solvency Constraints," *The Review of Financial Studies*, October 2001, *14* (4), 1117–1151.

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. Publisher: The University of Chicago Press.

Auclert, Adrien, Will S. Dobbie, and Paul Goldsmith-Pinkham, "Macroeconomic Effects of Debt Relief: Consumer Bankruptcy Protections in the Great Recession," March 2019.

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.

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

Campbell, John Y., *Financial Decisions and Markets*, Princeton, NJ: Princeton University Press, 2018.

Catherine, Sylvain, Max Miller, and Natasha Sarin, "Social Security and Trends in Wealth Inequality," *The Journal of Finance*, Forthcoming.

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. Publisher: The University of Chicago Press.

Dobbie, Will S. 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. Publisher: The University of Chicago Press.

Farber, Daniel A. and Paul A. Hemmersbaugh, "Shadow of the Future: Discount Rates, Later Generations, and the Environment, The," *Vanderbilt Law Review*, 1993, *46* (2), 267–304.

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.

Gross, David B. and Nicholas S. Souleles, "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data*," *The Quarterly Journal of Economics*, February 2002, *117* (1), 149–185.

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

_ , 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.

Indarte, Sasha and Martin Kanz, "Debt Relief for Households in Developing Economies," 2023.

Kreps, David M. and Evan L. Porteus, "Temporal Resolution of Uncertainty and Dynamic Choice Theory," *Econometrica*, 1978, *46* (1), 185–200. Publisher: [Wiley, Econometric Society].

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.

Rubinstein, Ariel, "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 1982, *50* (1), 97–109. Publisher: [Wiley, Econometric Society].

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

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

Appendices

A Proofs

This appendix provides a complete derivation of the comparative statics that characterize how the first-period default probability responds to changes in (i) the total loan balance L and (ii) the periodic payment m. These results are stated in equation (10) in the main text. Here, I derive them in full, beginning from the definition of the optimal default threshold.

A.1 Setup and Default Threshold

The borrower chooses whether to default in period $t = 1$ 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 period-1 default threshold y_1^* satisfies:

$$V_1^P(y_1^*) = V_1^N(y_1^*). \tag{18}$$

Differentiating this identity with respect to a contract parameter (either L or m) yields the sensitivity of y_1^*, and therefore of the default probability

$$p_1 = F(y_1^*).$$

Because $p_1 = F(y_1^*)$, I have:

$$\frac{\partial p_1}{\partial x} = f(y_1^*) \frac{\partial y_1^*}{\partial x}, \qquad x \in \{L, m\}. \tag{19}$$

Thus, the core task is to compute $\partial y_1^* / \partial x$.

A.2 Differentiating the Threshold Condition

Rewrite (18) as:

$$\Delta(y_1; L, m) \equiv V_1^P(y_1) - V_1^N(y_1) = 0$$

at $y_1 = y_1^*$.

Implicit differentiation gives:

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

I compute each term in turn.

A.3 Derivative with Respect to y_1

Because $C_1^P = y_1 - m$, $C_1^N = y_1$, and future values do not depend on y_1, I have:

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

Evaluated at the threshold:

$$\partial\Delta/\partial y_1 = u'(C_1^{P*}) - u'(C_1^{N*}).$$

Because u is concave, the denominator is negative, as expected.

A.4 Derivative with Respect to L

Only the continuation value depends on L. Differentiating the present-value expression:

$$\frac{\partial V_1^P}{\partial L} = \beta \int_{y_2^*}^{\infty} \left(-u'(y_2 - (L - m))\right)\, dF(y_2 \mid y_1).$$

There is no dependence on L in V_1^N, so:

$$\frac{\partial \Delta}{\partial L} = \beta \int_{y_2^*}^{\infty} \left(-u'(C_2^P(y_2))\right) dF(y_2 \mid y_1).$$

Pulling out the conditional expectation:

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

where $p_2 = F(y_2^*)$ is the second-period default probability.

Plugging into (A.2) and then into (19) yields:

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

This is the first line of equation (10).

A.5 Derivative with Respect to m

The payment m affects both current consumption and the future balance term $L - m$.

Compute each contribution:

(i) Direct effect on C_1^P:

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

(ii) Effect on continuation value through $(L - m)$**:**

$$\frac{\partial}{\partial m} u(y_2 - (L - m)) = +u'(C_2^P(y_2)).$$

Thus:

$$\frac{\partial \Delta}{\partial m} = -u'(C_1^P) + \beta(1 - p_2)E[u'(C_2^P(y_2)) \mid y_2 > y_2^*].$$

Observe that the second term is exactly the negative of the expression for $\partial\Delta/\partial L$:

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

Plugging into (A.2) and (19) yields:

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

This is the second line of equation (10).

A.6 Discussion

The structure of these expressions highlights the intuition: increasing L worsens future consumption while holding current consumption fixed; increasing m worsens current consumption while reducing future obligations. The relative strength of these channels produces a clean revealed-preference mapping into the discount factor.

B Additional Figures

Figure A1: Distribution of Liabilities

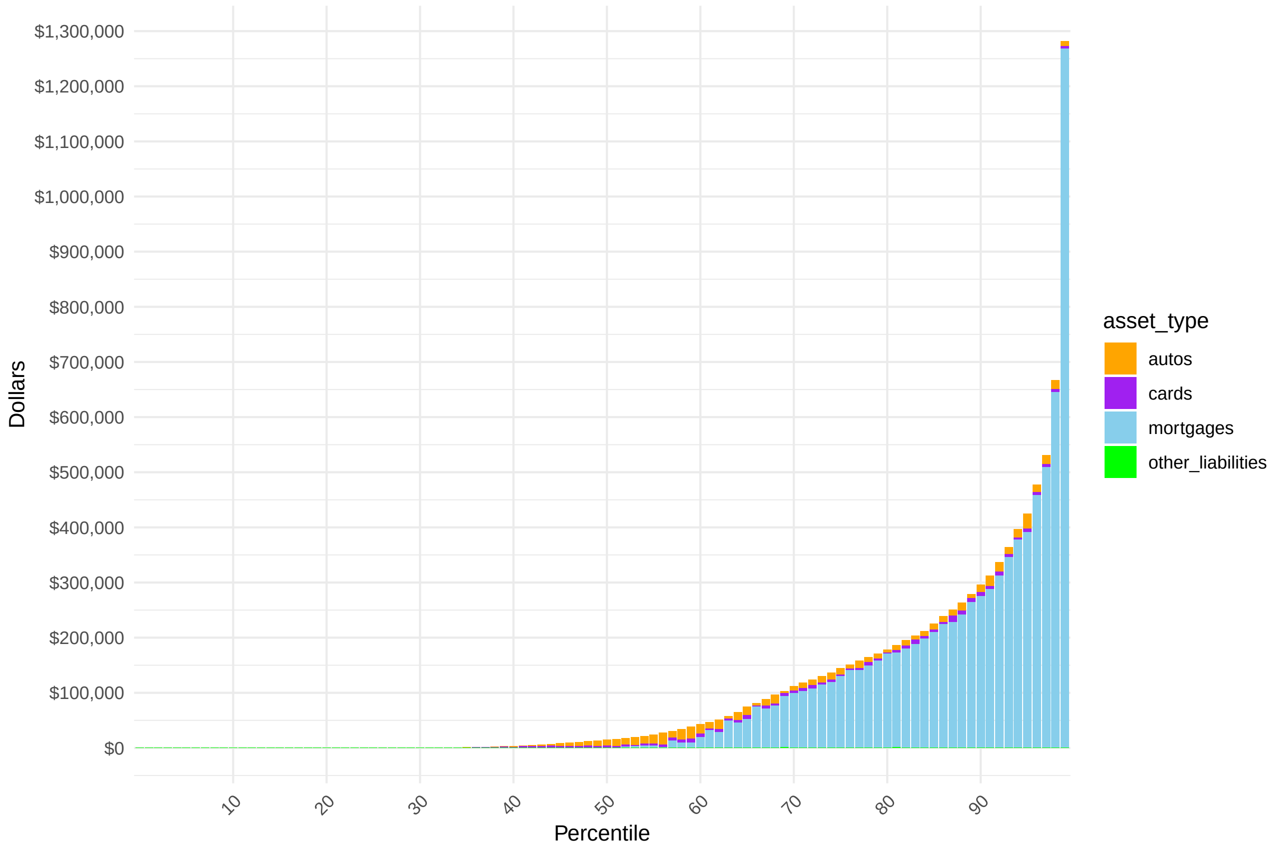

Notes: The x-axis is scaled by SCF survey weights, as are the average values within each percentile bin.

C Lifecycle Model Calibration

First, I calculate the common component $g(t)$ of the log-earnings profile of age for workers between 24 and 66 years of age. $t = \text{age} - 24$. Following Guvenen et al. (2021), I estimate this as a quadratic function of age using Ordinary Least Squares for workers between ages 24 and 66. For unemployed workers, I use expected income (x7362) when available, and exclude workers with zero expected income (i.e. who are not searching for work). The estimating equation is:

$$\log(\text{Income}_{it}) = g_0 + g_1 t + g_2 t^2 + \varepsilon_{it}$$

The estimated parameters are

g_0	10.7
g_1	0.0537
g_2	-0.000718

The parameters for the full life-cycle model estimated in Guvenen et al. (2021), also used in Catherine et al. (Forthcoming), are as follows:

Parameter	**Value**
ρ	0.959
p_z	40.7%
$\mu_{\eta,1}$	-0.085
$\sigma_{\eta,1}$	0.364
$\sigma_{\eta,2}$	0.069
$\sigma_{z,1,0}$	0.714
λ	0.0001
p_ϵ	13.0%
$\mu_{\epsilon,1}$	0.271
$\sigma_{\epsilon,1}$	0.285
$\sigma_{\epsilon,2}$	0.037
σ_α	0.300
$\sigma_{\beta\cdot 10}$	0.196
$\text{corr}_{\alpha\beta}$	0.768
$a_{\nu\cdot 1}$	-3.353
$b_{\nu\cdot t}$	-0.859
$c_{\nu\cdot z_t}$	-5.034
$d_{\nu\cdot t\cdot z_t}$	-2.895
$a_{z_1\cdot 1}$	0.407

Table B1: Guvenen Parameters

For the purposes of applying this process to SCF data at the individual level, it is not possible to exactly estimate the state variable Z_t^i nor the individual idiosyncratic level and slope variables α^i and β^i as doing so would require observing the full history of earnings.

Moreover, for the purposes of this exercise, I am mainly interested in the *expectation* of earnings growth,[10] and in ensuring that (1) I do not underestimate this expectation, nor (2) underestimate its heterogeneity in the population. To ensure this, I calcualate each earners' age-adjusted idiosynratic earnings $\varepsilon_L = \log L_{it} - g(t)$ and then calculate their position in the percentile distribution of ε_L. I then apply these percentiles to the distributions of α^i and β^i to estimate these variables. Finally, I estimate the state variable as

$$E[Z_t^i] = \log L_{it} - g(t) - \alpha^i - \beta^i t$$

And estimated future earnings at the k-year horizon as

$$E_t[L_{t+k}] = \exp(\rho^k Z_{it} + \alpha^i + \beta^i(t+k) + g(t+k))$$

This approach maximally preserves heterogeneity between individuals of the same age while avoiding potential underestimation of future earnings expectation.

[10] Recall the equation $\beta = \exp -(r - \gamma(E_t[L_{t+1}] - L_t - \sigma_L))$ and observe that ignoring σ_L will attenuate my estimates of β, i.e. move them closer to 1.

D Reference Materials

Figure C1: Social Discount Rates

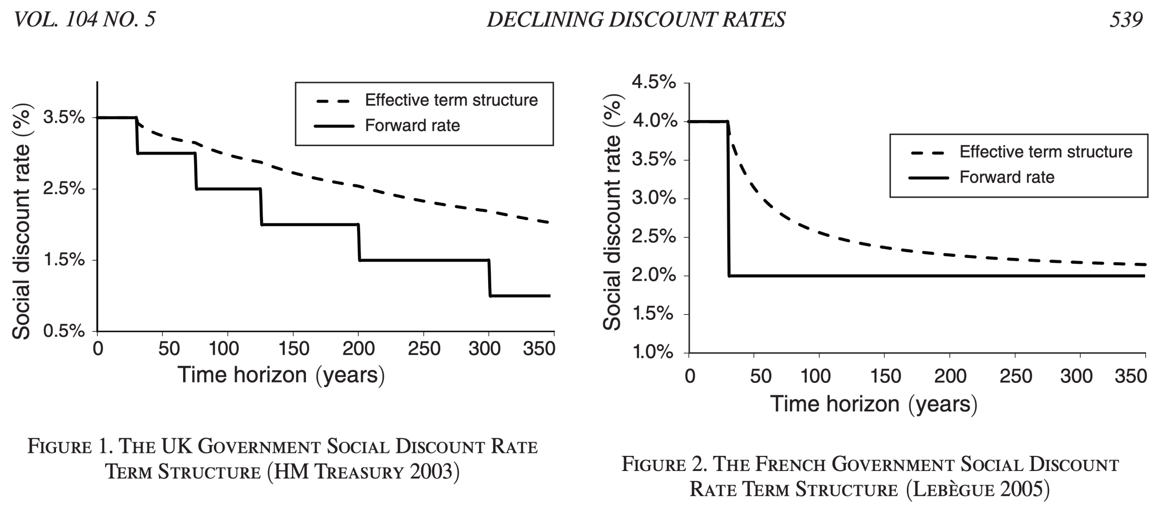

Notes: These rate schedules are used by the UK and French governments, respectively, in cost-benefit analysis. See Arrow et al. (2014).

Figure C2: Humorous Comic from Michael Thrower Chowdhury on X.com

E Conditional Default Rate Calculations

I calculate the empirical conditional default probabilities using the following formula:

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

Table C1: Definition of Terms for the Conditional Default Rate (CDR)

Symbol	Definition
$\text{CDR}_{d,m}$	Conditional Default Rate for score decile d in month m
$W_{d,m}$	Number of loans in decile d that default in month m
O_d	Total number of loans originated in decile d
$\sum_{k=1}^{m-1} W_{d,k}$	Cumulative number of defaults in decile d before month m
$O_d - \sum_{k=1}^{m-1} W_{d,k}$	Number of loans still active at the start of month m (the survivors)

I condition for both contract repayment term (in months) and credit score deciles. I stress the importance of controlling for ex-ante credit risk in this procedure to the the risk of heterogeneity in ex-ante credit risk biasing the estimate of β. Heterogeneity in ex-ante credit risk causes the conditional default curve to slope down for spurious reasons: bad credit risks default quickly, and good credit risks remain in the population, causing the conditional default curve to slope down even if default decisions are completely due to innate characteristics of heterogeneous individuals and do not involve an intertemporal decision whatsoever. To address this concern I estimate the model separately on ex-ante credit risk strata and term lengths. 56% of auto loans in the sample have repayment terms equal to either 36, 60, or 72 months, with 72 months being the most popular repayment period with over 25% of auto loans being exactly 72 months in contract repayment term:

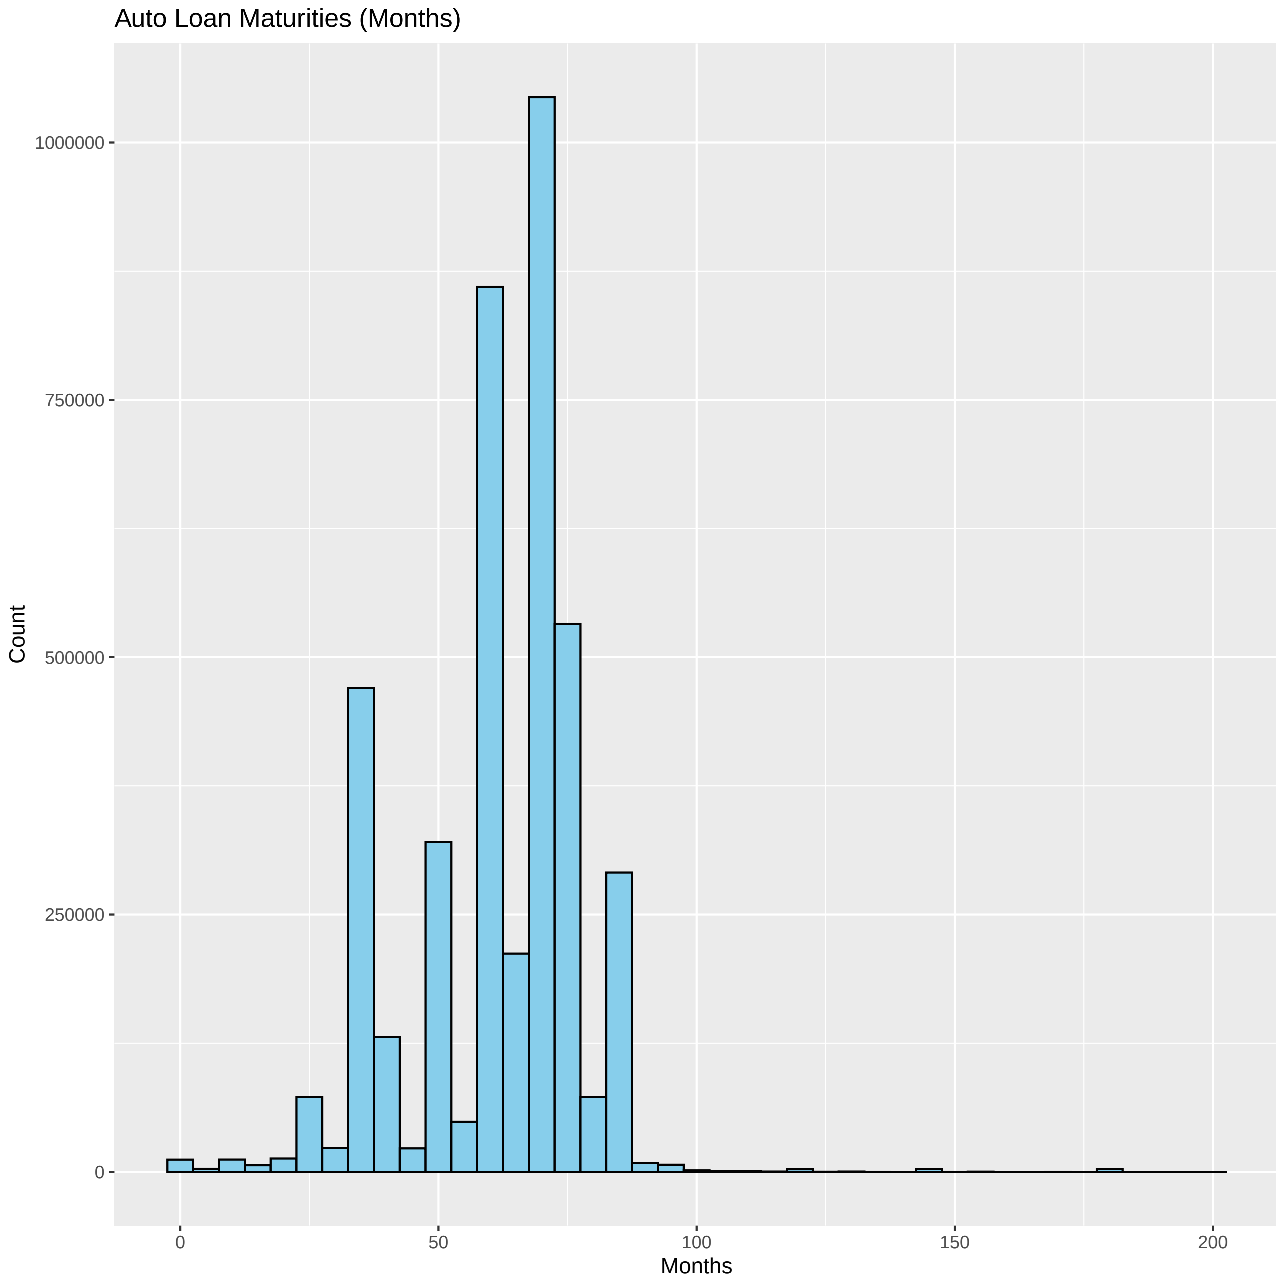

The raw histogram summary of the timing-of-default data for 72-month auto loans without conditioning on credit score:

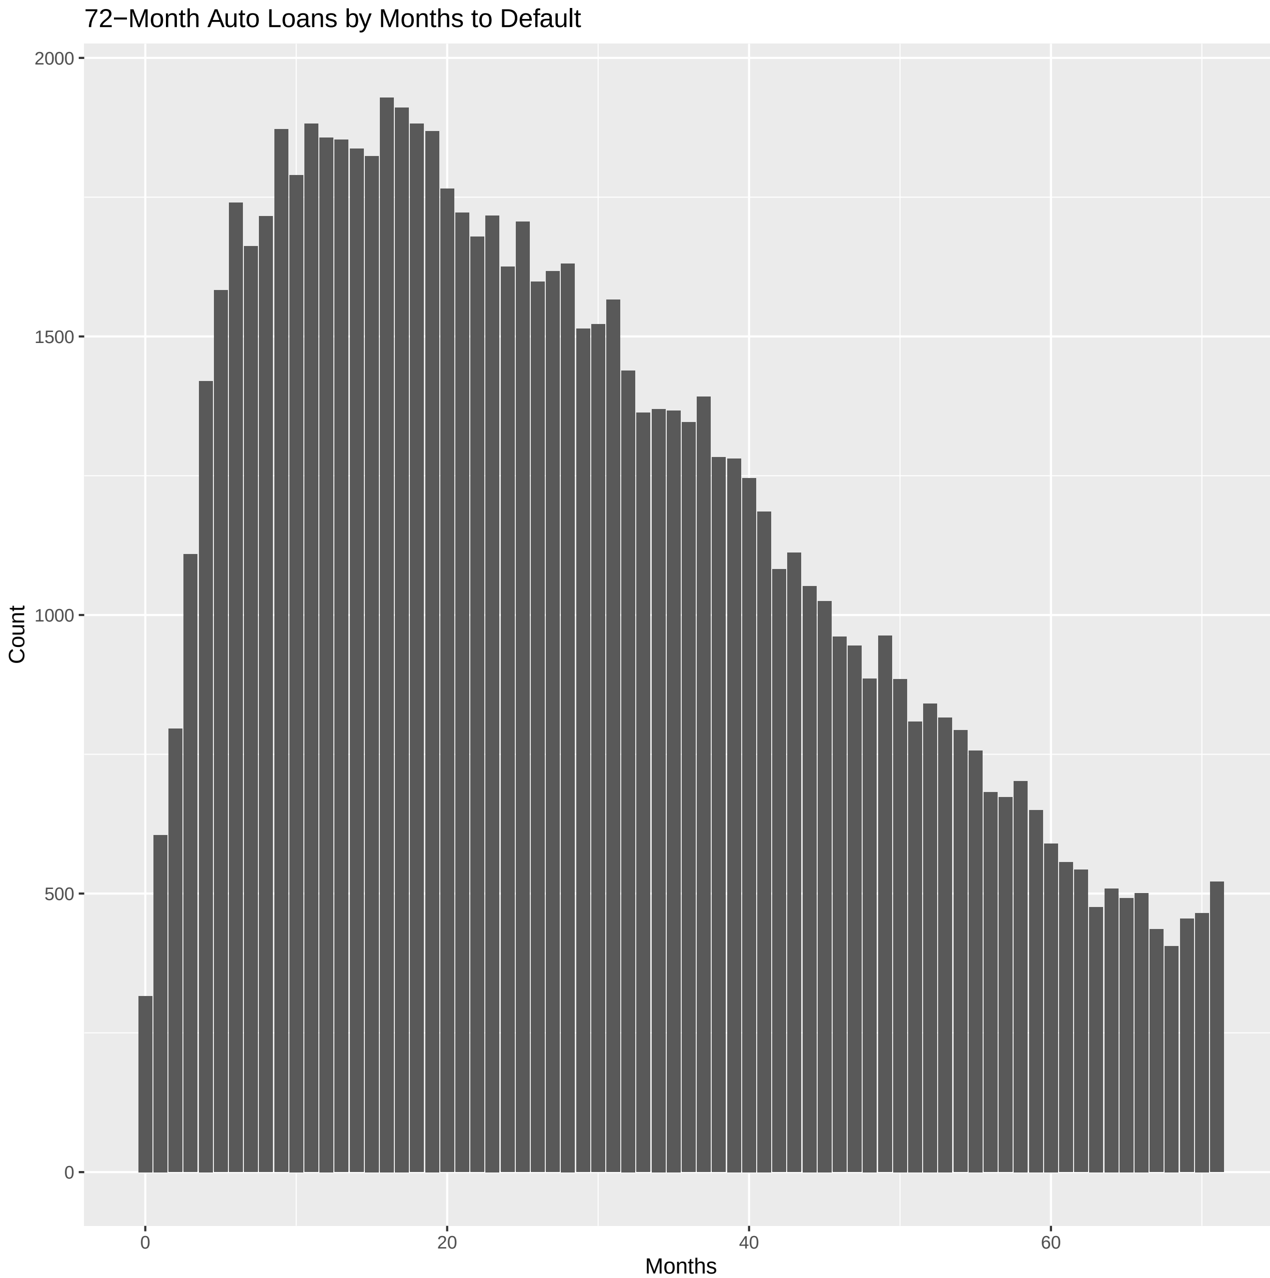

To condition on credit score, I split the data into deciles. To calculate the deciles, I focus on the distribution of ex-post defaulters. After conditioning on credit risk:

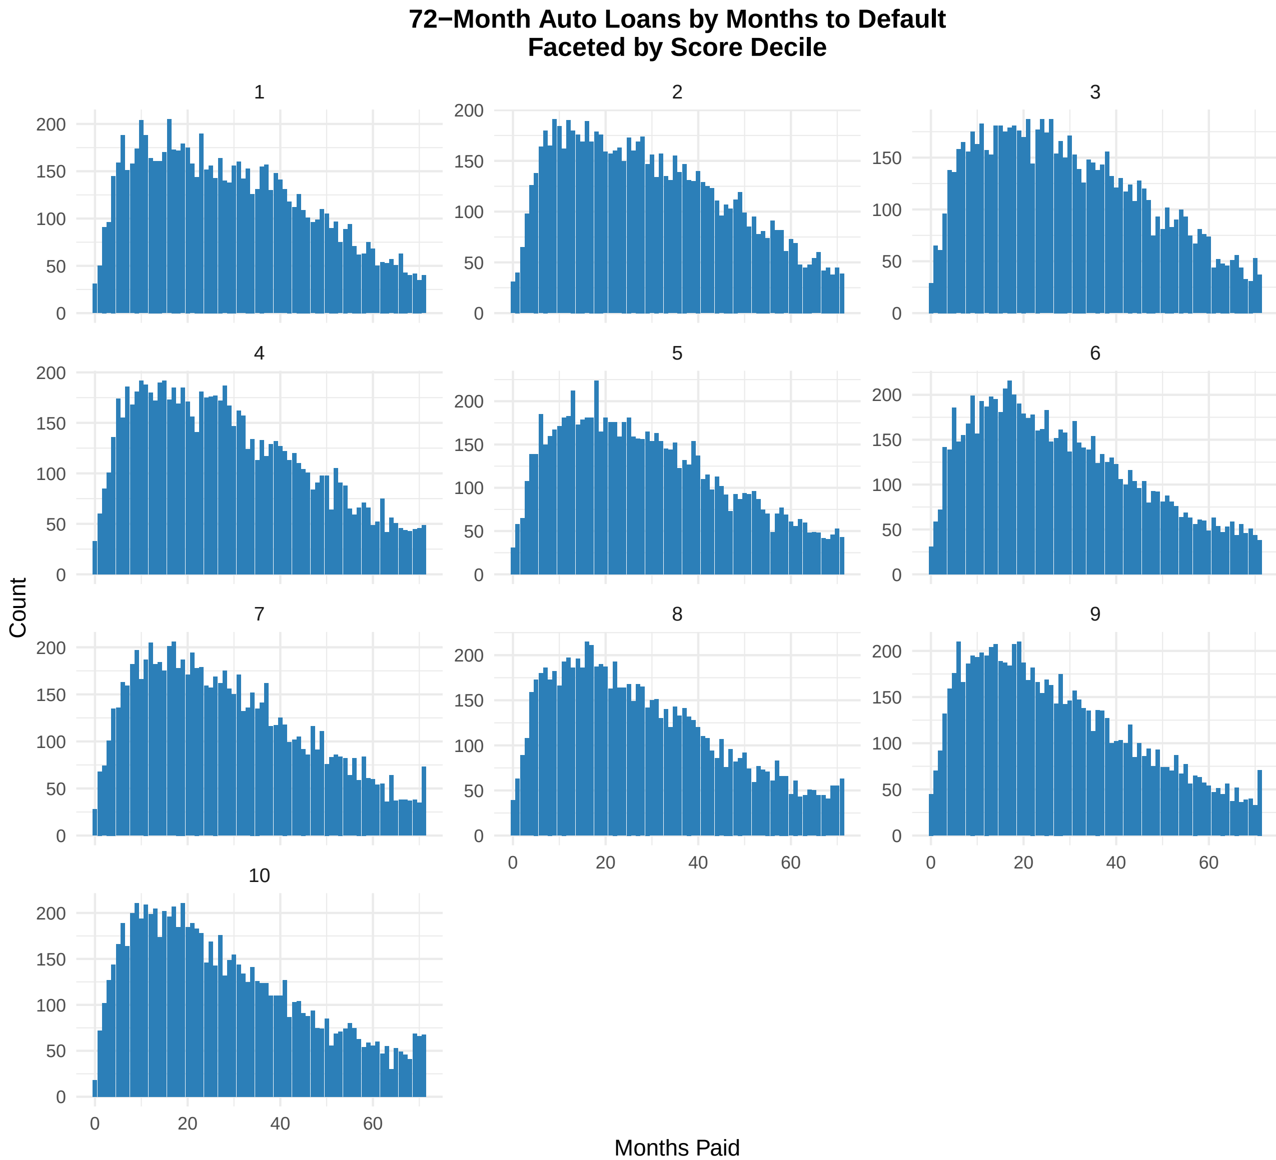

The ventiles are:

Table C2: Score Decile Breakpoints

decile	lower	upper
1	0	466
2	466	505
3	505	529
4	529	549
5	549	566
6	566	583
7	583	602
8	602	624
9	624	657
10	657	850

After transforming them into conditional default rates,

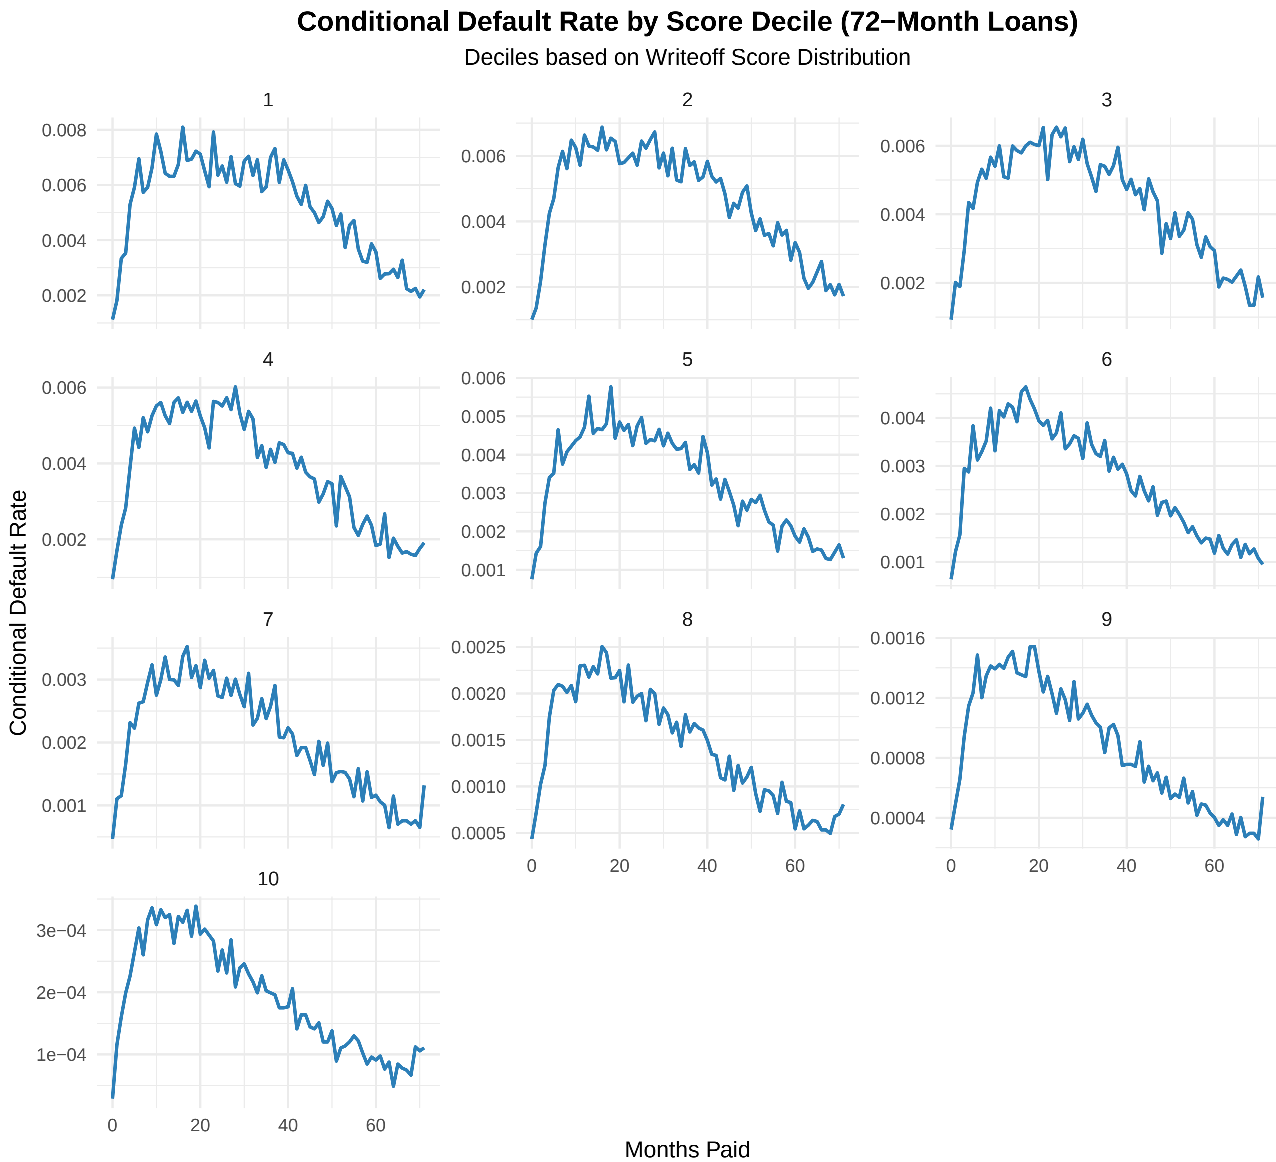

F Delta Method Calculations

I calculate standard errors using the delta method, incorporating sampling-based uncertainty in my estimates of both comparative statics and T. The estimand is $\beta_{cumulative}$, which I hereafter call simply β:

$$\beta = \left(\frac{\frac{\partial p_1}{\partial L}}{\frac{\partial p_1}{\partial m} + \frac{\partial p_1}{\partial L}} \right)^{1/T}. \tag{20}$$

For notational convenience, let

$$A \equiv \frac{\partial p_1}{\partial L}, \qquad B \equiv \frac{\partial p_1}{\partial m}, \qquad R \equiv \frac{A}{A+B}$$

so that $\beta = R^{1/T}$.

Using the delta method, the asymptotic variance of $\hat{\beta}$ is approximated by

$$\widehat{\text{Var}}(\hat{\beta}) \approx \nabla_\theta \beta(\hat{\theta})' \, \widehat{\text{Var}}(\hat{\theta}) \, \nabla_\theta \beta(\hat{\theta}), \qquad \text{where } \theta = (A, B, T)'. \tag{21}$$

The corresponding standard error is

$$\text{se}(\hat{\beta}) = \sqrt{\widehat{\text{Var}}(\hat{\beta})}.$$

The gradient vector $\nabla_\theta \beta$ contains the partial derivatives of β with respect to (A, B, T), which are obtained analytically as:

$$\frac{\partial \beta}{\partial A} = \beta \cdot \frac{1}{T} \cdot \frac{B}{A(A+B)}, \tag{22}$$

$$\frac{\partial \beta}{\partial B} = \beta \cdot \frac{1}{T} \cdot \left(-\frac{1}{A+B} \right), \tag{23}$$

$$\frac{\partial \beta}{\partial T} = \beta \cdot \left(-\frac{\ln R}{T^2} \right). \tag{24}$$

Let the estimated variance–covariance matrix of $(\hat{A}, \hat{B}, \hat{T})$ be

$$\widehat{\text{Var}}(\hat{\theta}) = \begin{pmatrix} \sigma_A^2 & \sigma_{AB} & \sigma_{AT} \\ \sigma_{AB} & \sigma_B^2 & \sigma_{BT} \\ \sigma_{AT} & \sigma_{BT} & \sigma_T^2 \end{pmatrix}.$$

Then, the delta-method variance of $\hat{\beta}$ is

$$\begin{aligned}\widehat{\mathrm{Var}}(\hat{\beta}) \approx\ & \left(\beta \frac{1}{T}\frac{B}{A(A+B)}\right)^2 \sigma_A^2 + \left(\beta \frac{1}{T}\frac{1}{A+B}\right)^2 \sigma_B^2 + \left(\beta \frac{\ln R}{T^2}\right)^2 \sigma_T^2 \\ & + 2\left(\beta \frac{1}{T}\frac{B}{A(A+B)}\right)\left(\beta \frac{1}{T}\left(-\frac{1}{A+B}\right)\right)\sigma_{AB} \\ & + 2\left(\beta \frac{1}{T}\frac{B}{A(A+B)}\right)\left(\beta\left(-\frac{\ln R}{T^2}\right)\right)\sigma_{AT} \\ & + 2\left(\beta \frac{1}{T}\left(-\frac{1}{A+B}\right)\right)\left(\beta\left(-\frac{\ln R}{T^2}\right)\right)\sigma_{BT}. \end{aligned} \tag{25}$$

The three inputs (A, B, T) are estimated independently in separate models. The covariance terms are assumed to be zero, so this simplifies to:

$$\widehat{\mathrm{Var}}(\hat{\beta}) \approx \beta^2 \left[\frac{1}{T^2}\left(\frac{B}{A(A+B)}\right)^2 \sigma_A^2 + \frac{1}{T^2}\left(\frac{1}{A+B}\right)^2 \sigma_B^2 + \frac{(\ln R)^2}{T^4}\sigma_T^2\right]. \tag{26}$$

The corresponding standard error is therefore

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