FEDERAL RESERVE BANK
PHILADELPHIA

Working Papers
RESEARCH DEPARTMENT

Explaining Contract Heterogeneity in the Credit Card Market

Satyajit Chatterjee
Federal Reserve Bank of Philadelphia

Burcu Eyigungor
Federal Reserve Bank of Philadelphia

WP 26-03

PUBLISHED
February 2026

FRB CFI PHILADELPHIA

ISSN: 1962-5361

DOI: https://doi.org/10.21799/frbp.wp.2026.03

Explaining Contract Heterogeneity in the Credit Card Market

Satyajit Chatterjee and Burcu Eyigungor

Federal Reserve Bank of Philadelphia

September 2025

Abstract

Administrative data are used to establish patterns in contract terms, usage, and default rates of anonymized individual credit card accounts. The canonical heterogeneous-agent macro model is extended with a competitive credit card industry and ex-ante heterogeneity to explain these facts, including that the spread on card interest rates is several multiples of default rates. Some model implications of general interest are: (i) a 10 percent cap on credit card interest rates, as proposed in recent legislation, reduces credit limits for risky borrowers and is welfare reducing for them, and (ii) although most people are not liquidity constrained, the average model MPC is in the empirically relevant range because consistency with credit card facts implies people are impatient.

Keywords: Credit limits, interest rate spreads, defaults, heterogeneity, MPC

JEL Codes: D15, E44, G51

Satyajit.Chatterjee@phil.frb.org and Burcu.Eyigungor@phil.frb.org. The authors thank Zach Bethune, Larry Santucci, seminar and conference participants at New York University, Johns Hopkins University, University of Rochester, University of Pennsylvania, FRB Richmond, 2023 CAFRAL Conference, 2024 Texas Monetary Conference, 2024 SED Annual Meeting, and 2024 SITE Meeting for helpful comments. The views expressed in this paper are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

1 Introduction

A striking feature of the US consumer credit card market is that cardholders are quite heterogeneous in their default risk. By one measure, the most risky cardholder is about 20 times more likely to go seriously delinquent than the least risky cardholder. As differences in delinquency risk are also largely predictable, there is a correspondingly wide variation in credit access between the least and the most risky cardholders. In this paper, we aim to understand the nature of differences among individuals that can explain the observed heterogeneity in credit access, credit use, and default risk.

This article makes the following contributions. First, we use high-quality administrative data to document facts about the terms of credit (interest rates and credit limits) on newly originated credit accounts, the subsequent use of these accounts for revolving debts, and the default rates experienced on these accounts. We show that credit terms, credit use, and default rates vary systematically with the cardholder's income and credit score at origination.

Second, we augment the canonical heterogeneous agent model (Imrohoroglu (1989), Hugget (1993), Aiyagari (1994)) with a credit card industry to explain these facts. A credit card is a commitment to provide funds on demand up to a specified limit at a specified interest rate, for as long as the individual does not default. As in competitive search (Moen (1997), Wright, Kircher, Julien, and Guerreri (2021)), card companies post credit card contracts, and individuals search for the best card. There is free entry into the card business.

Third, we investigate whether the canonical setup—where heterogeneity results only from permanent earnings differences and idiosyncratic transitory and persistent shocks to earnings—can explain the patterns we document, and find that it cannot. We argue that this failure suggests that people are different in ways not captured by income differences alone. We demonstrate that allowing people to differ in their discount factors and default costs enables the model to closely explain the observed patterns. The model implies substantial correlated heterogeneity in discount factors and default costs: The more patient individuals have higher default costs, and people are generally estimated to be quite impatient.

Fourth, we demonstrate that our model can largely explain the puzzlingly big gap between the spread on credit card interest rates (i.e., interest rates net of the cost of funds) and credit card default rates. Depending on borrower characteristics such as income and credit score, the spread can be between 3 and 40 times the corresponding default frequency. This divergence was first noted in Dempsey and Ionescu (2025) using the same administrative data we use. They also pointed out that in the standard macro model of consumer debt and default (Chatterjee, Corbae, Nakajima, and Rios-Rull (2007), Livshits, MacGee, and Tertilt (2007)), this ratio is predicted to be 1.

Our resolution of this puzzle rests on two channels. One is that the administrative data allows us to distinguish between balances rolled over with promotional rates (so-called teaser rates) and those rolled over at regular non-promotional interest rates. Once promotions are considered, the effective interest rate paid on credit card debts falls significantly below regular non-promotional rates, especially for individuals with high credit scores. The second channel is the nature of the default behavior on a credit card. On a credit card, default happens only when the debt is close to the card's credit limit. Therefore, the debts that a cardholder pays back are, on average, less than the debt on which a cardholder defaults. Since the cardholder is charged the same interest rate on all debts, the spread on the card *must* exceed the default probability on the card for the card company to break even. The degree to which break-even spreads exceed default probabilities depends on how much revolving debt is serviced on the card, i.e., on the card's anticipated utilization rate. This connection can be approximately summarized by the formula "break-even spreads $\approx$ default rates $\div$ utilization rate." Thus, if the utilization rate on a credit card is 25 percent (a typical figure), the break-even spread on the card will be *four times* the default frequency. This is an important difference in credit card contracts relative to term loans for which the utilization rate is 1 and the break-even spread equals default risk. In summary, the first channel lowers the *observed* spreads, while the second channel raises *model* spreads. Together, the two channels contribute to explaining the high ratio of spreads to default frequency on credit card accounts.

Fifth, we compute the markups implied by competitive search. Due to the costliness of search, once a consumer and a card company come into contact, there is a surplus to be shared between them. Competitive search leads to a division of this surplus. A positive surplus for companies implies that the posted interest rates will always exceed the post-contact zero-profit interest rates. The gap between the two is the markup implied by competitive search. We show that these markups vary across contracts in expected ways, being higher for consumers who really need a card. However, except for the most risky borrowers, the absolute sizes of these markups are small and do not contribute significantly to spreads for most borrowers.

Since our model aligns well with the documented facts, we utilize it to shed light on two key issues. First, we analyze the implications of capping credit card interest rates at 10 percent, as proposed in recent legislation.[1] The motivation for such caps seems to derive from the perception that credit card interest rates are too far above the cost of funds to be competitive. However, as noted, spreads exceeding default probabilities by a large margin can easily arise in a competitive credit card market. Consequently, rate caps could be inefficient. We use our model to demonstrate that a 10 percent rate cap can significantly reduce credit limits for high-risk borrowers. While the cap lowers indebtedness and default frequencies among this group, it takes them longer to obtain a credit card. Overall, the cap is welfare-reducing for this group.

Finally, we explore the model's implications for the magnitude and dispersion of the marginal propensity to consume (MPC). Administrative data and the estimated model provide an alternative method for measuring MPC from liquid assets. We find MPCs to be relatively high, heterogeneous, and in line with other empirical evidence (discussed later). In our model, MPCs are high due to the impatience required to match the observed facts, particularly the high default rate among high-risk borrowers. Furthermore, impatience is also consistent with generally high credit card utilization rates observed in the face of high interest rates. All these facts call for impatience, which leads to the typically high marginal propensity to consume (MPC). Thus, default and utilization facts have a bearing on estimated MPCs, a point that is overlooked in calibrations and estimations that ignore credit market facts.

The paper is structured as follows. Section 2 connects our paper to the literature. Section 3 develops the facts that will be our focus. Section 4 lays out the model environment. Section 5 discusses the calibration of the income process, the setting of parameters set independently, and the estimation strategy to be followed for the remaining parameters. Section 6 presents our estimation results, beginning with the model that has only income heterogeneity. Section 7 proceeds to the main model, which includes preference and default cost heterogeneity. Section 8 is devoted to credit card interest rate spreads. Section 9 concludes by presenting the model's implications for a 10 percent rate cap and the distribution of MPCs.

2 Literature

There is now a sizable and active macro literature on consumer debt and default that began with Athreya (2002), Chatterjee, Corbae, Nakajima, and Rios-Rull (2007) and Livshits, MacGee, and Tertilt (2007) (see Exler and Tertilt (2020) for a survey). This literature has

[1] Bills H.R. 1944 and S.381 introduced in the 119th Congress (2025-2026) titled "10 Percent Credit Card Interest Rate Cap Act."

made progress on endogenizing the default premium on consumer debt, but has typically assumed that borrowing is *via* one-period debt contracts (recent exceptions are noted below). On the other hand, most heterogeneous-agent macro models pay scant attention to consumer credit or consumer default risk. It is common to assume that agents cannot borrow at all, or can borrow at the risk-free rate, or can borrow at an exogenous premium up to an exogenous limit (Deaton (1991a), Deaton (1991b), Aiyagari (1994), Gourinchas and Parker (2002), Kaplan and Violante (2022), among others), all assumptions that are far from reality. Our paper advances both literatures by explicitly modeling the credit instrument that is predominantly used to borrow unsecured in the U.S., namely, a credit card, and showing that the model can successfully fit facts on their terms, usage, and performance gathered from high-quality administrative data.

Our finding that discount factor and default cost heterogeneity are necessary to explain patterns in credit card terms, usage, and performance reinforces earlier findings on the importance of such variation in explaining credit market outcomes for broad groups.[2] Chatterjee, Corbae, Dempsey, and Ríos-Rull (2023) find that heterogeneity in privately observed discount factors is needed to account for the evolution of credit scores over the life cycle. Athreya, Sanchez, and Mustre-del-Rio (2019) find that differences in impatience across people can explain the observed persistence of financial distress. Athreya, Tam, and Young (2009) leverage differences in survival probabilities by educational attainment to generate differences in effective discount factors across members of these groups, and in Athreya, Tam, and Young (2012) complement these differences with differences in default costs across education groups. They also find that a positive association between default costs and effective discount factors is needed to explain their credit market facts.

Regarding spreads on credit cards, Dempsey and Ionescu (2025) have drawn attention to the large gaps between credit card spreads and default frequencies. They approximate a credit card by a long-duration debt contract (Leland (1998), Hatchondo and Martinez (2009)) to study spreads. They find that intermediation costs and the expected duration of debt are key factors in determining spreads. In contrast, we model credit cards as an option to borrow up to a specified limit at a fixed interest rate. We show that wide gaps between spreads and default probabilities result from default behavior on such a contract

[2] There is also a large literature on the implications of heterogeneity for macroeconomics (see, for instance, Krueger, Mitman, and Perri (2016)) which gives great importance to the reasons why consumption differs across people and what it means for macroeconomic policies. This literature has found that the workhorse model with incomplete markets and plausible ex-post heterogeneity in incomes cannot easily account for the observed wealth inequality. But augmented with ex-ante heterogeneity in discount factors (and stochastically finite lives), it can do so — as shown in Krusell and Smith (1998) and Carroll, Slacalek, Tokuoka, and White (2017).

and that the utilization rate on the card (a concept that has no analog in Dempsey and Ionescu (2025)) is key for understanding credit card spreads.

Several recent studies, Bethune, Saldain, and Young (2024), Galenianos, Law, and Nosal (2023), and Raveendranathan, Stefanidis, and Sublet (2025), describe environments in which rate caps are beneficial. The first two studies seek to account for the very high interest rates paid by the highest-risk borrowers and do so by assuming that some of these borrowers face monopoly lenders.[3] Raveendranathan, Stefanidis, and Sublet (2025) have a broader focus but differ from us in that contract terms are determined in a Nash bargain post-match and the lender's bargaining power is calibrated to match observed markups. In contrast, we assume competitive search, so there is no ex-post bargaining. We find that search-induced markups are non-negligible for risky borrowers, but even so, a 10 percent rate cap is welfare reducing for this group.

There are several recent studies that draw out the MPC implications of consumer credit as we do.[4] Using observed interest rates and life cycle variation in credit limits, Fulford and Schuh (2024) estimate a structural model of savings, borrowing, and default to fit life cycle variation in consumer spending, credit card utilization rates, and bankruptcy. They report that the best fit is found for a model in which types differ by discount factors and implies average MPCs in the empirically estimated range. Lee and Maxted (2023) study a model where quasi-hyperbolic consumers (Laibson (1997)) can borrow via contracts that resemble credit cards. Present bias often leads households to borrow small amounts at high interest rates, resulting in high MPC.[5] We do not take interest rates and credit limits as given (they are equilibrium objects) but arrive at empirically relevant MPCs for essentially the same reasons: consistency with facts on credit card debt. Since our approach is more micro-founded, our findings further support the importance of accounting for the facts on consumer credit in the quest to understand high MPCs in the U.S.

Koşar, Melcangi, Pilossoph, and Wiczer (2024) and Boutros and Mijakovic (2024) draw out a different MPC implication of high-cost consumer borrowing: Indebted households'

[3]In a different line of attack, Saldain (2023) focuses on competitive provision of credit to high-risk borrowers who suffer from temptation preferences. Rate caps are potentially beneficial in this situation.

[4]The context is Kaplan and Violante (2014)'s influential study pointing out that average MPC out of tax cuts were much higher, about 25 percent, than what is implied by the classical permanent income hypothesis. But they did not focus on the high-interest borrowing of US households and, instead, offered an explanation for high MPCs of wealthy households based on their desire to hold almost all their financial wealth in high-yielding, illiquid assets. Aguiar, Bils, and Boar (2025) and Gelman (2021) present evidence that the behavior of wealthy households with low liquid assets reflects differences in discount factors or interest rates.

[5]Drozd and Kowalik (2023) postulate quasi-hyperbolic preferences in order to explain industry practice of giving promotional interest rates on credit cards early in the contract, followed by high interest rates later in the contract.

MPC is lowered because they use a portion of any windfall increase in wealth to pay off consumer debt. In our model, too, credit card debts work to lower MPCs, and for this reason, the credit card industry in our model is a (modest) force in favor of smaller MPCs.

Since our model incorporates both credit lines and search, it aligns with several recent studies that feature one or both.[6] Raveendranathan (2020) presented a credit card model of unsecured consumer borrowing, wherein people search for cards, with contract terms determined via Nash bargaining upon contact. Raveendranathan and Stefanides (2025) use a version of this model to study the fall in consumer credit during the Great Recession.[7] Braxton, Herkenhoff, and Phillips (2024) use a model of credit lines to study the impact of insurance afforded by credit cards on the optimal level of public insurance against unemployment risk. Herkenhoff and Raveendranathan (2025) study the welfare consequences of a shift from a pure monopoly to an oligopolistic market structure for card companies in a setup where there are credit lines but no search.[8]

Heterogeneity in credit card lending terms has been studied from different perspectives in prior work. Narajabad (2012) documented the increase in dispersion of credit card interest rates and limits over time in the U.S. and used this fact to explain the secular rise in credit card debt and bankruptcy filing rates. Livshits, MacGee, and Tertilt (2016), Athreya, Tam, and Young (2012) and Sanchez (2018) pursue this topic in greater sophistication. All these studies focused on improvements in information technology as the underlying factor. Herkenhoff and Raveendranathan (2025) show that rising credit limits and defaults also accompany the removal of entry barriers. Galenianos and Gavazza (2022) focus on dispersion in interest rates among observably similar borrowers and show this can occur in a search model if borrowers face attention costs.

3 Credit Card Terms, Usage, and Default Rates

At a broad level, credit card contracts offered must reflect the card companies' assessment of the expected behavior of their intended customers. The goal of the quantitative analysis is to learn how customers must differ for the equilibrium of the model to match what we see the card companies doing. Given this objective, we focus on *newly* originated cards, as we expect contract terms to most closely reflect customer characteristics when a card is first issued.

[6] Significant earlier contributions on credit cards are Mateos-Planas and Ríos-Rull (2007) and Drozd and Nosal (2008).

[7] This framework is also used in the Raveendranathan, Stefanidis, and Sublet (2025) study mentioned earlier.

[8] Ausubel (1991) drew attention to high and sticky credit card interest rates during the 1980s and argued for a failure of competition in the credit card industry. Although markups do not necessarily imply supra-normal profits (Evans and Schmalensee (2005, p.240)), the general view among scholars is that the U.S. credit card industry had significant barriers to entry in its early days, which have since eroded.

We utilize confidential supervisory data collected by the Federal Reserve System in connection with the annual Comprehensive Capital Analysis and Review (CCAR) of large U.S. bank holding companies. One part of these data, called Y-14M, tracks, at a monthly frequency, the terms, usage, and performance of all credit card accounts issued by all U.S. bank holding companies with assets greater than $50 billion. The data covers about 80 percent of all U.S. credit card accounts in existence at a point in time.

We use a 1-in-200 anonymized random sample of all new, revolving, general-purpose, and unsecured credit card accounts originated between June 2014 and May 2015 with one primary account holder.[9] After trimming the sample for outliers,[10] we truncated the sample in two ways. First, accounts with income-at-origination in the top or bottom 7 percent of the income-at-origination distribution were removed. This was done to better align the income distribution in our sample with the distribution implied by estimated earnings processes commonly used in the calibration of macro models. Second, accounts with credit scores below 619, corresponding to the bottom 11.5 percent of the credit score distribution, were removed. The average score of those excluded is 593, which is well within the score band of $580-619$ commonly designated as "subprime." This is done for computational reasons, as explaining the contract terms and behavior of these most risky borrowers would require more heterogeneity than we currently permit in the model.

The contract terms we focus on are credit limits and interest rates paid on revolving balances. The measurement of interest rates is complicated because most new originations come with a promotional period, typically a year, during which the contractual interest rate on the card is low or zero. Furthermore, it is also the case that new originations often start with promotional balances. These starting balances are likely transferred from another card with a higher, non-promotional interest rate. Our data distinguishes between promotional and non-promotional revolving balances, which allows us to construct a balance-weighted monthly effective interest rate at the account level. Since balance transfers come with a one-time balance transfer fee, the person transferring balances pays interest in advance for the transfer amount. To consider this, we assume a promo-

[9] Most credit cards have adjustable interest rates, meaning that the interest rate on revolving balances is a contractually specified spread ("margin") over (usually) the bank prime lending rate. The prime lending rate was constant at 3.25 percent during 2011-2015. Since our quantitative analysis focuses on steady states, we chose our sample period to coincide as much as possible with this period of essentially contractually fixed interest rates. The Y-14M data set begins in June 2012; however, not all variables are reported or available for the first few years.

[10] We excluded accounts for which reported borrower annual income was less than $\$1,000$ (most likely students) and credit-score-at-origination was less than 100 (most likely a reporting error); we also excluded accounts for which credit-limit-at-origination was less than $100.

tional interest rate of 3 percent, rather than zero. We calculate the effective interest rate on the card as the average interest rate paid over 60 months following origination.[11]

Credit limits also change over time, with most changes being increases in the limit. However, the fraction of accounts seeing limit increases each month is small, so limits generally have a long duration. To avoid any potential limit increases resulting from the end of promotions, we measure credit limits 24 months following origination. Correspondingly, we measure utilization rates at 24 months following origination, and we measure default frequency over the same 24-month period. These timing assumptions are maintained when we measure these rates in the model.

Figure 1 is a binned scatter plot of the contract terms ordered by credit score at origination (left panels) and by income at origination (right panels). A dot corresponds to the mean value for a quintile of our trimmed and truncated sample in each figure. The plotted lines are simply lines of "best fit" through these dots.

There is systematic variation in effective interest rates with respect to score and income at origination. Individuals with low scores pay higher interest rates on average, as do those with low income. Interest rates decline with scores and incomes, with the decline being sharper for scores: The mean interest rate in the top quintile of scores is 8 percentage points lower than the mean interest rate in the bottom quintile, while the corresponding difference for income quintiles is only about $2\frac{1}{2}$ percentage points.

There is significant variation in credit limits scaled by (annual) income at origination, and the variations go in opposite directions for scores and income. Individuals with credit scores in the top quintile have an average ratio of approximately 2.5 times that of individuals with scores in the bottom quintile. With regard to income, individuals in the bottom quintile have an average ratio of about 1.4 times that of individuals in the top quintile. This pattern arises because credit limits increase less than proportionately with income.

Figure 2 displays the patterns concerning usage and default. An account is in default if within 24 months of origination the account is reported to be in a debt waiver/bankruptcy or it is 120 days or more past due and has a charge-off flag within a year of its being reported 120 days past due (this so that we don't count delinquent accounts that cure). The top left panel shows that default frequency declines sharply with credit score: The

[11] In computing the time average, we make two further adjustments. First, we deflate all future balances (promotional or not) by the growth in nominal wages (we use average hourly earnings of production and nonsupervisory employees), with June 2014 as the base. Second, contract interest rates on cards change with changes in the bank prime rate. Since our theory abstracts from changes in the risk-free rate, we add or subtract a time-period-specific constant from monthly averages so that the average of all interest rates in each month is the same as the average interest rate in June 2014.

Figure 1
Variation in Contract Terms by Score & Income

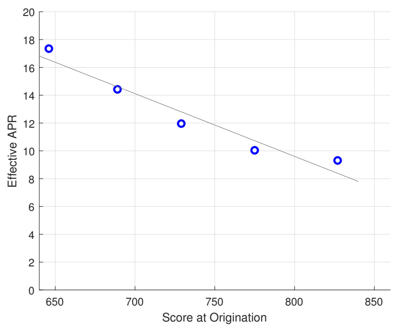

APR by Score

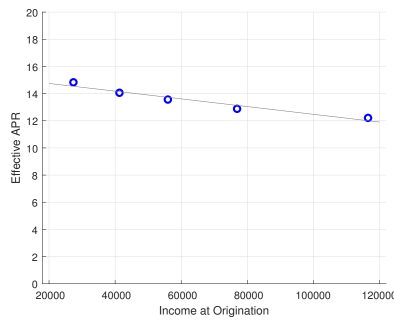

APR by Income

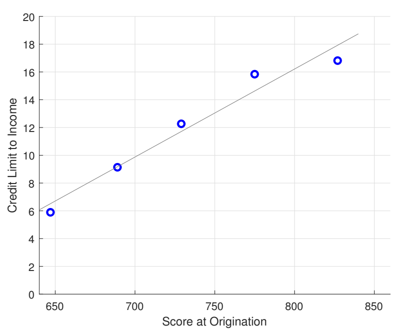

Credit-Limit-to-Annual-Income Ratio by Score

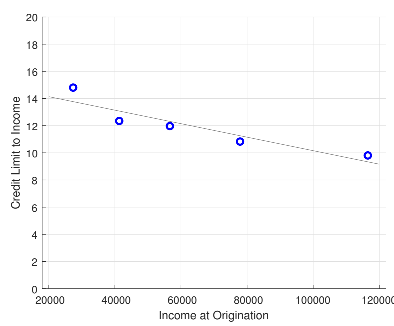

Credit-Limit-to-Annual-Income Ratio by Income

Source: Author calculations based on Federal Reserve Y-14M data

Figure 2
Variation in Default Frequency and Utilization by Score & Income

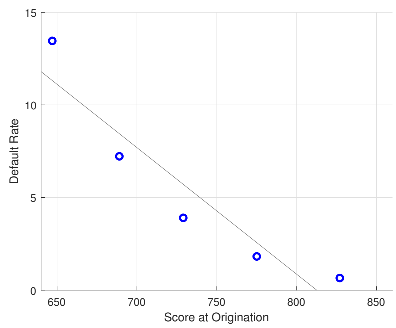

Default Frequency by Score

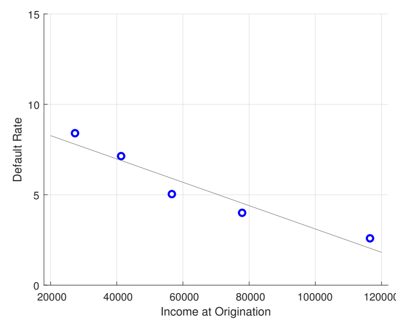

Default Frequency by Income

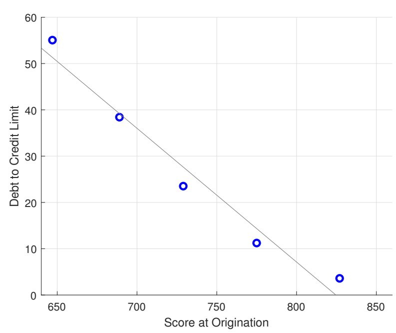

Utilization Rate by Score

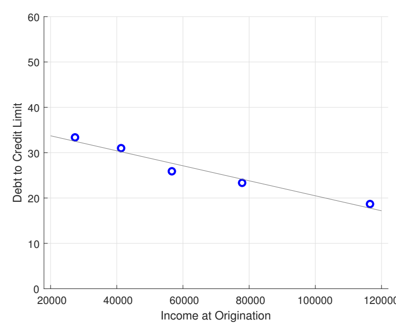

Utilization Rate by Income

Source: Author calculations based on Federal Reserve Y-14M data frequency in the bottom quintile is 13.46 percent, while in the top quintile it is 0.65 percent. The right panel shows that default frequencies also decline with income but less sharply: The default frequency is 8.40 percent in the bottom quintile and 2.59 percent in the top quintile.

The bottom panels display the ratio of revolving balances to credit limits. Individuals in the bottom quintile of scores have a utilization rate of 55.06 percent, while the top quintile has a utilization rate of only 3.60 percent. The utilization rate also varies negatively with income, but less so. Individuals in the bottom quintile of incomes have a mean utilization rate of 33.37 percent, while those at the top income quintile have a mean utilization rate of 18.67 percent.

The patterns displayed in Figures 1 and 2 are the facts we seek to explain in this paper.

4 Model

4.1 Environment

Time is discrete. We use recursive notation, so a variable z_t is denoted by z and z_{t+1} by z'.

An individual's income in any period is given by $P \cdot [y + m]$, where $P > 0$ is a time-invariant permanent component and y and m are random. The component y is persistent and follows a first-order Markov chain $F(y'|y)$. The component m is drawn each period from a uniform distribution with support $[-\lambda y,\ \lambda y]$, $0 < \lambda < 1$; conditional on y, m is transitory.

There are different types of people $i \in 1, 2, \dots, I$. An individual's preferences is a discounted sum of period utilities where the discount factor β^i depends on type. The period utility function is of the CRRA form with curvature parameter $(1 - \gamma)$, $\gamma > 0$. An individual survives from one period to the next with a constant probability $\nu \in (0, 1)$.

We also allow P to vary across people, so an individual is indexed by the pair (i, P). However, since all decision rules can be shown to scale with P, we lightened the notation and developed the theory for $P = 1$. At the end of this section, we introduce variation in P and state the scaling properties.

Individuals can borrow via credit cards. A credit card is a bilateral contract between a card company and an individual that allows the latter to borrow up to $\underline{a} < 0$ at a (gross) interest rate R. Credit cards can differ in these terms. We use $\omega \equiv (\underline{a}, R)$ to denote a card contract and $\Omega \equiv \mathbb{R}_{-} \times R_{+}$ to denote the space of all card contracts. An individual can hold at most one card, and the card's contract is fixed for as long as the individual holds the card.[12] All individuals can save at the common risk-free (gross) interest rate R_f.

A person with a balance on her credit card can default. If a type i person defaults, she (i) loses her card and, in the default period, surrenders a fraction $0 < [1 - \phi^i] \leq 1$ of her transitory income in excess of $-\lambda y$ to the card company, (ii) cannot save in the period of default, and (iii) with probability $(1 - \delta)$ she is shut out of the credit card market in the period following default and, conditional on being shut out, she continues in that state with probability $(1 - \delta)$ in future periods.

Figure 3
Default Cost

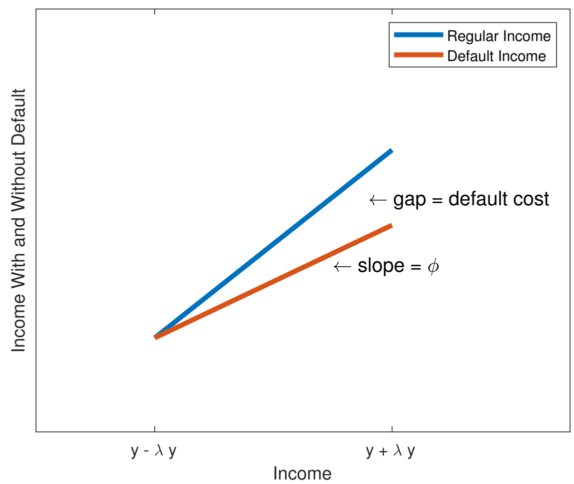

Figure 3 displays the income loss suffered by a defaulter as a function of the transitory component of her income in the default period. Since the loss is $(1 - \phi^i)(m - (-\lambda y)) = (1 - \phi^i)(m + \lambda y)$, it is zero if $m = -\lambda y$ and increases linearly with m to a maximum of $(1 - \phi^i)(2\lambda y)$. The specific form of the default cost ensures that the default decision rule typically has a convenient threshold property: Default is optimal for realizations of m below a cutoff that depends on the level of card balances and value of y.[13]

The income loss from default, which we refer to as the default cost for expository simplicity, deserves comment. The substantive assumption is that the cost increases in

[12] This is a simplification. While the 2009 Credit CARD Act constrains changes to credit card contracts, credit card companies can, and do, change customers' interest rates and credit limits over time.

[13] A convenient threshold property can also be ensured if default costs are modeled as a utility cost that is Gumbel distributed (as in Raveendranathan (2020), Chatterjee, Corbae, Dempsey, and Ríos-Rull (2023) and others). However, this approach introduces another parameter, the variance of the utility shock, which must be estimated. Additionally, Gumbel shocks cause people to borrow small amounts simply for the chance of experiencing positive utility shocks from default. The income-loss approach has the advantage that the variance of the transitory shock is known from estimating the income process and one does not have to confront counterintuitive effects of default-payoff shocks on borrowing decisions.

both income components and may vary across types. Since the default cost is the payment that creditors extract from delinquent borrowers, it is plausible that this payment increases with available income. More broadly, default costs include the monetary cost of the loss in reputation due to failure to repay, and this loss is likely to be higher for individuals with higher incomes.

A person who does not have a card and is not shut out of the card market searches for a card with probability $\mu > 0$. Search happens at the start of a period before the transitory shock m is realized. The searcher searches among contracts offered to people of her type i, her asset position a, and her persistent earnings y. Card companies commit to honoring the terms of posted contracts. These assumptions make our search environment one of competitive/directed search.

To generate active search, we assume that cardholders get separated from their cards with probability ξ.[14] The separation shock occurs before the transitory shock m is realized. Upon separation, if she is not carrying a balance on her card, the individual searches for a new card with probability μ. If she has a balance on her card, her balance is transferred to a prespecified contract $\bar{\omega}^i_y$ if this contract is profitable for the card company. Otherwise, the individual continues on with her existing card.

4.2 Individual's Decision Problem

Let $h^i(a, y, m)$ denote the value of a cardless individual who is not shut out of the card market (she could be in this state because she did not search, or she searched but failed to make contact). Let $S^i(a, y)$ denote the value of a type i person in state (a, y) who is without a card but searching for one. Then $h^i(\cdot)$ solves:

$$h^i(a, y, m) = \max_{a' \geq 0} \frac{c^{1-\gamma}}{1-\gamma} + \nu\beta^i \mathbb{E}_{y'|y}\{\mu S^i(a', y') + (1-\mu)H^i(a', y')\} \tag{1}$$
$$c = y + m + R_f a - a' \text{ and } H^i(a, y) = \mathbb{E}_m h^i(a, y, m).$$

Let $x^i(a, y, m)$ denote the value of an individual who is without a card and who is also excluded from the credit card market (because of a past default). Such a person is also

[14] We make this assumption because separation rates are roughly the same across income quintiles. If we allowed separations to be endogenous, the people who want to drop their existing card and search for a new one are those whose economic circumstances have recently improved, which will lead them to obtain cards with better terms. This would lead to separation rates that rise strongly with income (which is counterfactual).

limited to saving and $x^i(\cdot)$ solves:

$$x^i(a,y,m) = \max_{a' \geq 0} \frac{c^{1-\gamma}}{1-\gamma} + \nu\beta^i \mathbb{E}_{y'|y}\left[\delta\left[\mu S^i(a',y') + (1-\mu)H^i(a',y')\right] + (1-\delta)X^i(a',y')\right] \tag{2}$$

$$c = y + m + R_f a - a' \text{ and } X^i(a,y) = \mathbb{E}_m x^i(a,y,m).$$

Let $v^i(a,y,m;\omega)$ denote the value of an individual in state (a,y,m) who has a credit card with terms ω and who chooses to make payments (if any) on her card. Then, v^i solves:

$$v^i(a,y,m;\omega) = \max_{a' \geq \underline{a}} \left\{\frac{c^{1-\gamma}}{1-\gamma} + \nu\beta^i \mathbb{E}_{y'|y} W^i(a',y';\omega)\right\} \tag{3}$$

$$c = y + m + Qa - a', \text{ where}$$

$$Q = \begin{cases} R \text{ if } a < 0 \\ R_f \text{ if } a \geq 0. \end{cases}$$

Here $W^i(a,y;\omega)$ is a cardholder's *ex-ante* continuation value in the state (a,y) and its recursive form will be given below. Let $a'^i(a,y,m;\omega)$ and $c^i(a,y,m;\omega)$ denote the associated asset and consumption decision rules of a repayer.

Let $v^i_{\text{DEF}}(y,m)$ denote the value of default. Then,

$$v^i_{\text{DEF}}(y,m) = \frac{c^{1-\gamma}}{1-\gamma} + \nu\beta^i \mathbb{E}_{y'|y}\left[\ \delta\left[\mu S^i(0,y') + (1-\mu)H^i(0,y')\right] + (1-\delta)X^i(0,y')\ \right] \tag{4}$$

$$c = y - \lambda y + \phi^i \cdot [m + \lambda y].$$

In the event of default, the individual retains some fraction of her transitory income and pays the rest to her creditors.

Letting $V^i(a,y;\omega)$ denote $\mathbb{E}_m v^i(a,y,m;\omega)$, the recursion for $W^i(a,y;\omega)$ is

$$W^i(a,y;\omega) = \begin{cases} \xi S^i(a,y) + (1-\xi)V^i(a,y;\omega) & \text{if } a \geq 0 \\ \xi^i(a,y)\mathbb{E}_m\left[\max\left\{v^i(a,y,m;\bar{\omega}^i_y), v^i_{\text{DEF}}(y,m)\right\}\right] + & \\ \quad [1-\xi^i(a,y)]\mathbb{E}_m\left[\max\left\{v^i(a,y,m;\omega), v^i_{\text{DEF}}(y,m)\right\}\right] & \text{if } a < 0, \end{cases} \tag{5}$$

where $\xi^i(a, y)$ is the probability that an indebted person will have a balance transfer and will be specified below. For future use, let $d^i(a, y, m; \omega)$ denote the default decision rule: $d^i(\cdot)$ is 1 if $v^i(a, y, m; \omega) < v^i_{DEF}(y, m)$ and 0 otherwise.

Finally, we can specify the recursion for $S^i(a, y)$. Let $f^i(\omega; a, y)$ denote the probability of an individual with characteristics $\{i, a, y\}$ encountering a contract ω. Then,

$$S^i(a, y) = \max\Big\{\max_{\omega\in\Omega}\Big\{f^i(a, y; \omega)\cdot[V^i(a, y; \omega) - H^i(a, y)] + H^i(a, y)\Big\}, H^i(a, y)\Big\}. \quad (6)$$

The inner max chooses over ω for the best contract. The outer max recognizes that the individual need not search and get $H^i(a, y)$.

4.3 Credit Card Companies' Decision Problem

Card companies choose the contract terms to offer to people of different types in different states. Once a contract is accepted, the companies follow through on the terms until the individual separates from the card or defaults on it. We assume that a company's opportunity cost of funds is R_f.

Denote the value to a company of a credit card contract ω held by a person of type i in state (a, y, m) as $\pi^i(a, y, m; \omega)$ and denote $\mathbb{E}_m\pi^i(a, y, m; \omega)$ by $\Pi^i(a, y; \omega)$. Then,

$$\begin{aligned}
&\pi^i(a, y, m; \omega) = \\
&\quad d^i(a, y, m; \omega)\times[1-\phi^i](m+\lambda y) + \big[1 - d^i(a, y, m; \omega)\big]\times \\
&\qquad\Big(-\min\{0, a\}\cdot R(\omega) + \min\{0, a'^i(a, y, m; \omega)\} + \\
&\qquad\quad \frac{\nu}{R_f}\mathbb{E}_{y'|y}\Big[\big[1-\xi^{i\prime}\big]\,\Pi^i(a'^i(a, y, m; \omega), y') + \xi^{i\prime}(-\min\{0, a'^i(a, y, m)\}\cdot R(\omega))\Big]\Big).
\end{aligned} \quad (7)$$

If the individual is carrying a balance on the card and defaults, the card company gets $[1-\phi^i](m+\lambda y)$ and the credit line is closed. If she does not default, the card company receives the (gross) interest payment less the new credit extended, if any. If the individual is not carrying any balances, the company extends new credit, if any. Regardless, the company gets the expected continuation value of the contract discounted by the risk-free rate. The expected continuation value takes into account that the cardholder survives with probability ν and that the company might lose the contract with probability $\xi^i(a', y')$. If there is a balance on the card, then in the event of separation, the company receives the full balance, inclusive of interest; otherwise, it receives nothing.

We can now specify how the balance transfer probability is determined. As noted earlier, for $a \geq 0$, $\xi(a, y) = \xi$ for all y. For $a < 0$, we assume that the (exogenously specified) contract $\bar{\omega}^i_y$ is offered to a type i person in persistent state y regardless of her balances a. Let $\Pi^i(a, y; \bar{\omega}^i_y)$ be the *ex-ante* value to the card company of this card when the person has balance a. Then,

$$\xi^i(a, y) = \begin{cases} 0 & \text{if } \Pi^i(a, y; \bar{\omega}^i_y) < 0 \\ \frac{\Pi^i(a,y;\bar{\omega}^i_y)}{\tau}\xi & \text{if } 0 \leq \Pi^i(a, y; \bar{\omega}^i_y) \leq \tau \\ \xi & \text{if } \tau < \Pi^i(a, y; \bar{\omega}^i_y). \end{cases} \tag{8}$$

Thus, no balance transfer occurs if the contract is expected to incur a loss. To avoid jumps, we assume that the probability of transfer increases continuously as Π rises above 0, reaching ξ as profitability reaches or exceeds the contract posting cost τ.[15]

A card company chooses the profit-maximizing contract ω for people of type i in state (a, y). In making this choice, it takes the contact probability function $q^i(a, y; \omega)$ as given. The profit-maximizing contract solves:

$$\max_{\omega \in \Omega} q^i(a, y; \omega) \cdot \Pi^i(a, y; \omega).$$

Let $\omega^{i*}(a, y; q^i)$ be a contract that attains the maximum, let $\Pi^{i*}(a, y; q^i)$ denote $\Pi^i(a, y; \omega^{i*}(a, y; q^i))$, and let the net expected profit from posting a single $\omega^{i*}(a, y; q^i)$ contract be

$$\eta^{i*}(a, y; q^i) = q^i(a, y; \omega^{i*}(a, y; q^i)) \cdot \Pi^i(a, y; \omega^{i*}(a, y; q^i)) - \tau.$$

Let $e^i(a, y; q^i)$ be the measure of contracts $\omega^{i*}(a, y; q^i)$ posted to match with a type i person in state (a, y). Then,

$$e^i(a, y; q^i) = \begin{cases} 0 & \text{if } \eta^{i*}(a, y; q^i) < 0 \\ \text{indeterminate} & \text{if } \eta^{i*}(a, y; q^i) = 0 \\ +\infty & \text{if } \eta^{i*}(a, y; q^i) > 0. \end{cases} \tag{9}$$

[15] As the card company does not pay a posting cost for balance transfers, the assumption that the probability reaches ξ as profitability reaches τ is for convenience.

4.4 Contact Probabilities and Market Tightness

We follow den Haan, Ramey, and Watson (2000) and assume that the matching function is

$$M\left(B^{i}(y,a;\omega),E^{i}(y,a;\omega)\right)=\frac{B^{i}(y,a;\omega)\cdot E^{i}(y,a;\omega)}{\left[B^{i}(y,a;\omega)^{\zeta}+E^{i}(y,a;\omega)^{\zeta}\right]^{1/\zeta}},\quad \zeta\geq 0,$$

where $B^{i}(a,y;\omega)$ denotes the mass of individuals of type i in state (a,y) who are searching in the submarket ω and $E^{i}(a,y;\omega)$ denotes the mass of contact attempts made by the totality of card companies in the same submarket.

With this matching function, the probability that a credit card company will successfully contact a type i customer in state (a,y) is:

$$q^{i}(a,y;\omega)=\frac{M\left(B^{i}(\cdot),E^{i}(\cdot)\right)}{E^{i}(\cdot)}=\frac{1}{\left(1+\theta^{i}(a,y;\omega)^{\zeta}\right)^{1/\zeta}}=q(\theta^{i}(a,y;\omega)),\tag{10}$$

where

$$\theta^{i}(a,y;\omega)=\frac{E^{i}(a,y;\omega)}{B^{i}(a,y;\omega)}.\tag{11}$$

The ratio $\theta^{i}(a,y;\omega)$ can be interpreted as the "tightness" — from the perspective of card companies — of the submarket. A high value means stiff competition for customers and a low probability of a successful contact. On the other side, the probability that an individual of type i in state (a,y) in the submarket ω will successfully contacts a card company is

$$f^{i}(a,y;\omega)=\frac{M\left(B^{i}(\cdot),E^{i}(\cdot)\right)}{B^{i}(\cdot)}=\frac{\theta^{i}(a,y;\omega)}{\left(1+\theta^{i}(a,y;\omega)^{\zeta}\right)^{1/\zeta}}=f(\theta^{i}(a,y;\omega)).\tag{12}$$

This probability increases with market tightness: In a tight market, one can obtain a credit card quickly.

4.5 Equilibrium

Since there is no interaction between types, the equilibrium can be described in terms of equilibrium for each type $i\in\mathbb{I}$. An equilibrium for type i is a pair of functions $S^{i*}(a,y)$

and $\theta^{i*}(a, y; \omega) \geq 0$ such that

$$S^{i*}(a, y) = \max\left\{ \max_{\omega \in \Omega} \left\{ f(\theta^{i*}(a, y; \omega)) \cdot [V^{i*}(a, y; \omega) - H^{i*}(a, y)] + H^{i*}(a, y) \right\}, H^{i*}(a, y) \right\} \tag{13}$$

$$q^{i*}(a, y; \omega) = \begin{cases} 1 / \left[1 + f^{-1}\left(\frac{S^{i*}(a,y) - H^{i*}(a,y)}{V^{i*}(a,y;\omega) - H^{i*}(a,y;)} \right) \right] & \text{if } V^{i*}(a, y; \omega) > H^{i*}(a, y) \\ 0 & \text{otherwise} \end{cases} \tag{14}$$

$$q^{i*}(a, y; \omega)\Pi^{i*}(a, y; \omega) \leq \tau \text{ for all } \omega \in \Omega \tag{15}$$

$$\left[q^{i*}(a, y; \omega)\Pi^{i*}(a, y; \omega) - \tau \right] \theta^{i*}(a, y; \omega) = 0 \text{ for all } \omega \in \Omega. \tag{16}$$

Equation (13) asserts that equilibrium values of S^i and θ^i must be consistent with individual optimization. The asterisks on V^i, H^i and Π^i indicate that these functions implicitly depend on $S^{i*}(a, y)$ and $\theta^{i*}(a, y; \omega)$.

Equation (14) is the key equilibrium condition for directed search. The top branch asserts that for any contract ω offered to people of type i in state (a, y) that delivers more utility than $H^{i*}(a, y; S^{i*})$, the market tightness must be such that the *ex-ante* value of searching for that contract is the same as the equilibrium *ex-ante* value of active search, $S^{i*}(a, y)$. The bottom branch asserts that any contract that delivers less utility than $H^{i*}(a, y)$ doesn't attract any customers and, so, $q^{i*}(a, y; \omega)$ is zero (in effect, market tightness for such a contract is infinite).

Equation (15) asserts that all feasible contracts earn non-positive expected net profits.

Equation (16) is the complementary slackness condition that ensures consistency with free entry of lenders in any submarket: If the equilibrium market tightness for a contract is strictly positive, the contract must earn zero expected net profits, and if a contract generates negative expected net profits, its equilibrium market tightness must be 0. The conditions (15) and (16) together assert that equilibrium contracts maximize expected net profits.

The equilibrium measure of contactable people, $B^{i*}(a, y; \omega)$, and the aggregate measure of posted contracts, $E^{i*}(a, y; \omega)$, do not appear directly in these equilibrium equations because these quantities effect equilibrium outcomes only through $\theta^{i*}(a, y; \omega)$. They can be backed out from knowledge of equilibrium market tightness function and the equilibrium contract in each market.[16]

4.6 Scaling

We now extend the setup to any $P > 0$. A type-i person's state expands to $\{a, y, m, P\}$, and she searches for a card in the market $\{i, a, y, P\}$. For tractability we aim for the following scaling property: The interest rate offered to $\{i, a, y, P\}$ is the interest rate offered to a $\{i, a/P, y, P = 1\}$ person and the credit limit offered is P times the credit limit offered to the $\{i, a/P, y, P = 1\}$ person.

Proposition 1. *If the default cost of a type i person in income state $\{P, y, m\}$ is $P \cdot [(1 - \phi^i)(m + \lambda y)]$ and the cost of posting a contract in $\{i, a, y, P\}$ market is $\tau \cdot P$, there is an equilibrium with the following properties:*

- $\omega^{i*}(a, y, P) \equiv \omega_P^{i*}(a, y) = \{R^{i*}(a/P, y),\ P \cdot \underline{a}^{i*}(a/P, y)\}$
- $c^{i*}(a, y, m, P; \omega_P^{i*}) = P \cdot c^{i*}(a/P, y, m)$, $a'^{i*}(a, y, m, P) = P \cdot a'^{i*}(a/P, y, m)$, *and* $d^{i*}(a, y, m, P) = d^{i*}(a/P, y, m)$
- $h^{i*}(a, y, m, P) = P^{1-\gamma} \cdot h^{i*}(a/P, y, m)$, $x^{i*}(a, y, m, P) = P^{1-\gamma} \cdot x^{i*}(a/P, y, m)$, $v^{i*}(a, y, m, P) = P^{1-\gamma} \cdot v^{i*}(a/P, y, m)$, *and,* $v_{DEF}^{i*}(a, y, m, P) = P^{1-\gamma} \cdot v_{DEF}^{i*}(a/P, y, m)$
- $\xi^{i*}(a, y, P) = \xi^{i*}(a/P, y)$ *and* $\theta^{i*}(a, y, P; \omega_P^{i*}) = \theta^{i*}(a/P, y; \omega^{i*}(a/P, y))$
- $\pi^{i*}(a, y, m, P; \omega_P^{i*}(a, y)) = P \cdot \pi^{i*}(a/P, y, m; \omega^{i*}(a/P, y))$.

Proof. If default costs and contract posting costs scale with P, the existence of such an equilibrium follows directly from the homogeneity of degree $1 - \gamma$ of the period utility function. $\square$

4.7 Computation

We compute an approximate equilibrium of the $P = 1$ model. First, the space of assets a is approximated by a fine grid. Next, for each $\{i, y\}$ pair, contracts for K asset levels $\{0, a_2^i, a_3^i, \ldots, a_K^i\}$ only are offered. An $\{i, a, y\}$ person is offered the contract for person

[16]For instance, if in steady state individuals of type i with persistent income y and savings a search for contract ω^*, then $B^{i*}(a, y; \omega^*)$ is just μ times the steady state measure of type i households that do not possess a card, are not shut out, and are in state $\{a, y\}$. The aggregate measure of posts in this (sub)market is then $E^{i*}(a, y; \omega^*) = B^{i*}(a, y; \omega^*) \cdot \theta^{i*}(a, y; \omega^*)$.

$\{i, a_k^i, y\}$, where a_k^i is the closest asset level less than or equal to a. For a given $\{i, a_k^i, y\}$, $\underline{a}$ and R are treated as continuous variables, and first-order conditions (FOCs) are used to pin them down. For the equilibrium values of $P \neq 1$ persons, Proposition 1 is used.

We cannot approximate $(\underline{a}, R)$ by grids because the solution is sensitive to the choice of grid points. But treating them as continuous variables leads to a different challenge, since FOCs are a necessary but not sufficient condition for optimality (the card companies' profit function, $q^{i*}(a, y; \omega)\Pi^{i*}(a, y; \omega) - \tau$, is generally nonconcave in the elements of ω). To protect against settling on a local maximum, we search for solutions to the FOCs at many randomly chosen pairs of $(\underline{a}, R)$ and pick the solution with the highest profits.

5 Quantitative Analysis: Preliminaries

5.1 Earnings Process

We calibrate the earnings process based on estimates in Storesletten, Telmer, and Yaron (2004). They estimate a process for annual earnings given by $y(t)z(t)\alpha^i$, where y, the persistent component of earnings, is an AR1 process in logs, z, the transitory component, is i.i.d. with $\ln(z)$ having standard deviation σ_z, and α^i is a permanent individual component with $\ln(\alpha)$ having standard deviation σ_α. Their estimates are reported in Table 1.

Table 1
Earnings Process Parameters

Parameter	Description	Value
ρ	Autocorrelation of log of persistent component	0.938
σ_ε	Std of innovation to log of persistent component	0.138
σ_z	Std of log transitory shock	0.257
σ_α	Std of permanent shock	0.290

Notes: Source: Storesletten, Telmer, and Yaron (2004, p. 705, Table 2, Row E). All parameters are at annual frequency. The value of 0.29 for σ_α is below the 0.301 value reported in their table. A smaller value was chosen to maintain a close match between the distribution of the incomes implied by the process and the distribution of incomes reported in our data.

Our model period is a quarter, and we assume that income is the sum of the persistent and transitory components: $y + m$. We assume that $\ln(y_t) = \rho \ln(y_{t-1}) + \varepsilon_t$ and that m is uniformly distributed with support $[-\lambda y,\ \lambda y]$, given y.[17] We discretize y into 5 levels and assume that it follows a first-order Markov process. We pick λ and the values of the

[17] Note that this assumption is equivalent to assuming that $\ln(z)$ is distributed with support $[-\lambda,\ \lambda]$.

Markov transition matrix so that the model earnings generate a ρ, σ_ε and σ_z at annual frequencies that match the values in Table 1.

5.2 Parameters Set Independently

Table 2 displays the other parameters whose values are set externally. The (real) quarterly risk-free (gross) interest rate R_f was set to 1.005, the probability of survival, ν, was set to 0.9958 per quarter to target average life expectancy of 78.8 years for a 21-year-old. The CRRA parameter γ was set at 2, which is conventional in macroeconomics.

The separation parameter, ξ, was set to 0.022 to match the attrition rate of accounts in good standing in our sample between June 2014 and May 2019.

The probability of search is set at 0.50, which delays getting a new card (following separation or re-entry) by about 6 months. The delay lowers the outside option of searching individuals and increases the profits of card companies. The cost of posting a contract, τ, is set to a low value of 0.001. These settings ensure that in equilibrium, credit card contracts are offered to all $\{i, a, y\}$ triples.

The entry probability following default, δ, is set so that the average period of exclusion from the credit card market following default is 6 years. Since we assume that a person without a card does not look for a card for about half a year, the average effective exclusion from the credit card market following a default is about 6.5 years.

The elasticity of the matching function ζ is set to 1.

The prespecified balance transfer contract $\bar{\omega}^i(y)$ is set to $\omega^{i*}(0, y)$, i.e., to the equilibrium contract offered to $\{i, y\}$ people with zero assets.

Finally, we assume that when a person dies, she is replaced by a newborn who is of the same type, has the same y, and is eligible to search right away.

These parameter choices are held fixed in subsequent analyses.

5.3 Estimation Strategy

The estimation takes as given the number of types I and the population fraction of each type μ^i and locates the $\{\beta^i, \phi^i\}$, $i \in \{1, 2, \ldots, I\}$, pairs that minimize the objective function:

$$w_1 \sum_{j=1}^{5} \left(\text{CL-to-INC}^{\text{model}}_{\text{score}_j} - \text{CL-to-INC}_{\text{score}_j}\right)^2 + w_2 \sum_{j=1}^{5} \left(\text{DEF-FRQ}^{\text{model}}_{\text{score}_j} - \text{DEF-FRQ}_{\text{score}_j}\right)^2.$$

Table 2
Parameters Set Independently

Parameter	Description	Value
λ	Bounds for transitory shock	0.826
R_f	Gross risk-free rate	1.005
ν	Probability of survival	0.9958
γ	Curvature of CRRA utility function	2.0
δ	Entry rate after default	1/24
τ	Cost of posting a contract	0.001
μ	Probability of searching for a card	0.50
ξ	Probability of exogenously separating from card	0.022
ζ	Elasticity of the matching function	1.00
$\bar{\omega}^i(y)$	Balance transfer contract	$\omega^{i*}(0, y)$

Notes: All rates and probabilities are at quarterly frequency.

Here, $\text{CL-to-INC}_{\text{score}_j}$ and $\text{DEF-FRQ}_{\text{score}_j}$ are the (average) credit-limit-to-income ratio and default frequency, respectively, in the jth score quintile. We chose to target a small set of moments likely to be most informative about a person's type. Holding fixed an individual's income process, her default cost is a key determinant of her debt capacity and, therefore, her credit-limit-to-income ratio. The individual's discount factor is a key determinant of her propensity to borrow and, conditional on being in debt, to default. Furthermore, as we will see, the income-based moments are not sensitive to the choice of β and ϕ, so it is not critical to consider them in the estimation. The moments that are not targeted will be used to evaluate the model's performance.

The estimation needs to recognize that our model is about an individual, but our data follows credit card *accounts*, and merging accounts belonging to the same individual is impossible. We scale up the credit-limit-to-income ratios shown in Figure 1 by a factor of 2 to take into account the average number of cards that individuals have.[18] Concerning the interest rate and utilization rates of card accounts, we assume that the average interest rate and utilization rates for the score and income bins are a good proxy for the individual-level averages for people in the corresponding score and income bins.

To implement the estimation, we need a model analog to a credit score. Consistent with real-world credit scores, the model credit score is defined to be the probability of not encountering a default within a two-year horizon. Given decision rules, this probability can be computed recursively as described in Appendix B.

[18]The number of cards the individuals possess increases with credit scores, but we are not taking this into account.

Figure 4

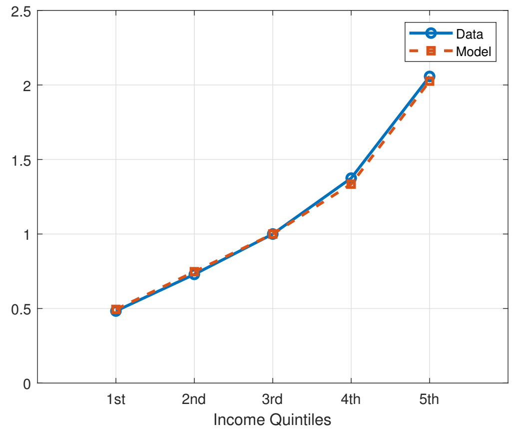

Notes: Data plot is author calculations based on Federal Reserve Y-14M data

Finally, we set $w_1 = 1$ and $w_2 = 1000$, thus giving much more weight to the deviations in default frequencies. This yields a close match between the model and the data in this dimension. This is useful because it allows us to compare how utilization rates, credit-to-income ratios, and interest rates respond to default risk in the model versus the data.

6 Can the Canonical Heterogeneous-Agent Model Explain These Patterns?

The canonical version has a single type, i.e., $I = 1$, with all individuals having a common β and ϕ. To start, Figure 4 plots average incomes in each quintile relative to the average in the middle quintile for both the model and the data. There is a very close fit, which is partly due to the truncation of incomes in our sample. Still, the fit is noteworthy because incomes in our sample are recorded on new origination, while incomes in the model are primarily determined by the income process in Table 1. Furthermore, the set of people looking for a new card in the model is partly endogenous: Some are seeking a new card following a default and re-entry, and among indebted individuals, those who obtain a new card are those for whom the balance transfer is profitable for the card company.

Default Frequency Patterns

The default frequencies for income and score quintiles are compared with the data in Figure 5. The model can explain the fact that default frequencies decline with incomes, and the decline is sharper for scores than for incomes. The fact that default frequencies decline with incomes and scores has a simple explanation. The persistent incomes of people in the top income bin are, on average, higher than those in the lower bins. Individuals with high y expect their y to decline over time, and this induces them to accumulate assets (or deleverage if they have debts). Consequently, they are less likely, as a group, to

Figure 5
Default Frequencies

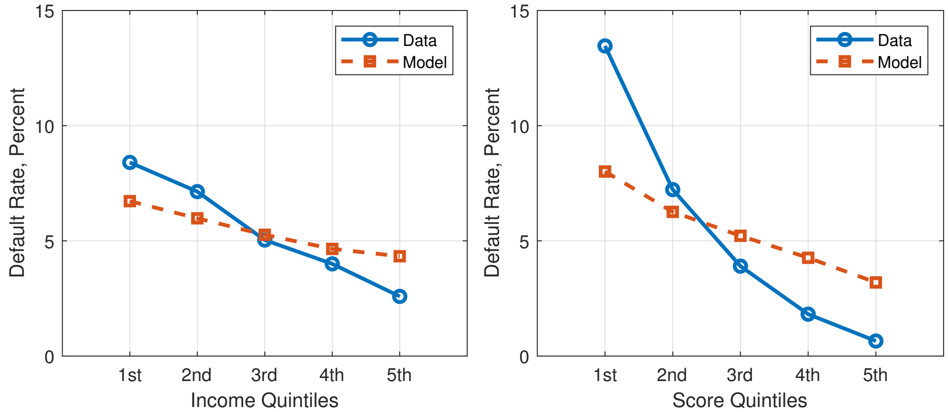

Notes: Data plots are author calculations based on Federal Reserve Y-14M data

get into debt levels that trigger default over a 2-year horizon. Similarly, when people are ranked by their credit scores, the top bin collects individuals with high y *and* high α. As a group, these people are also less likely to end up with defaultable debt levels within a 2-year horizon than people in the lower bins and are less likely to default than those with high incomes alone,[19] which accounts for its steeper slope relative to the income bins line. Another reason for the steeper slope of the score bins line is that the top income bin includes people whose y is low but whose permanent component of income, P, is high. The mixing of high and low y people, even in the top income bin, makes the drop in default frequency shallower for incomes. Such mixing of high and low y individuals is not possible for score bins because the scaling property ensures that assets and debts scale with P and default probability depends on only y and assets or debts *relative* to P. That said, the quantitative fit is imperfect in both panels, even though the estimation seeks to closely target the pattern for score quintiles.

Credit-Limit-to-Income Ratio Patterns

The left panel of Figure 6 shows that the model can capture the decline in this ratio with income quite well, even though this pattern is not a target of the estimation. The model credit-limit-to-income ratio declines with income for two reasons. The most important reason is the nature of the card contract, which commits card companies to a credit limit and an interest rate, regardless of the individual's future y levels. To see why this matters, consider the case where there is no transition between y states, i.e., a person's persistent income y is permanent. The blue line in the right panel of Figure 6 displays

[19] For instance, the high income group would include high y individuals who attained this status recently and, therefore, have not had the time to accumulate a high level of assets.

Figure 6

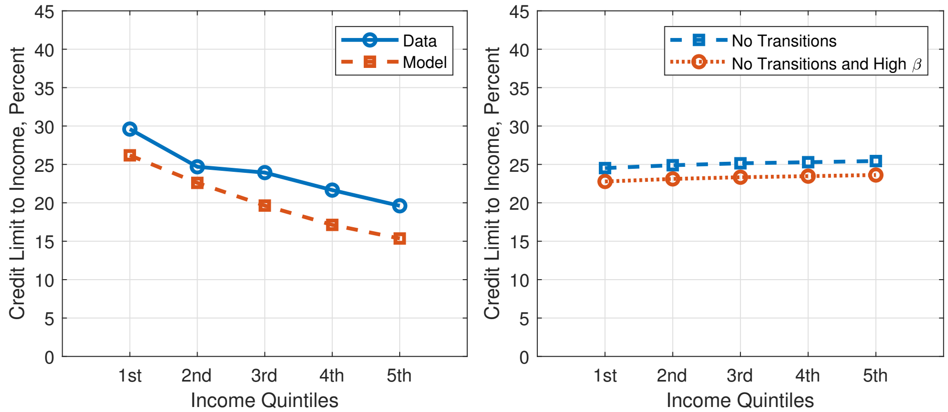

Notes: Data plots are authors calculations based on Federal Reserve Y-14M data

the credit-limit-to-income ratio pattern for the "no transitions" equilibrium (all other parameter values are unchanged). The equilibrium credit-limit-to-income ratio is roughly constant across income bins.[20] But with positive probability of transitions, a (roughly) constant credit limit *ratio* is very sub-optimal for card companies: If an individual's y drops, her expected default cost will drop proportionately and she might be tempted to borrow all the way to her large limit and default. In other words, in the presence of y transitions, the contract needs to ensure that the temptation to default does not increase too much when there is a drop in the persistent component of income. This constrains how much the credit limit can rise with y.

A second reason is that high y individuals tend to save or borrow less, i.e., tend to act patient, making them want to apply for cards with lower credit limits and interest rates. The situation here is similar to insurance: The good risks prefer contracts with high deductibles and low premiums. To confirm this logic, we examine how patience directly affects credit limits. The red line in the right panel of Figure 6 displays the credit-limit-to-income ratios in the no-transition economy for a higher discount factor. Observe that the ratios are systematically lower.

Turning to the relationship between credit-limit-to-income ratios and scores, Figure 7 plots these relationships for the model and the data. There is now a *qualitative* mismatch, not just a quantitative one: In the data, the credit limit ratio increases with credit scores, while in the model it is *decreasing*. This mismatch arises even though the relationship be-

[20]Since the support of the (uniform) distribution of the m-shock is proportional to y, ϕ^i is independent of y, and the utility function is homogeneous (of degree $1-\gamma$), the individual's decision problem is homogeneous in y if the market tightness function $\theta^i(a, y; \omega)$ is homogeneous of degree 0 in y. This latter property does not hold exactly since the contract posting cost τ does not scale with y. But since $\tau \approx 0$, it holds approximately. Consequently, the equilibrium of the no-transitions model features credit limits that approximately scale with y.

Figure 7

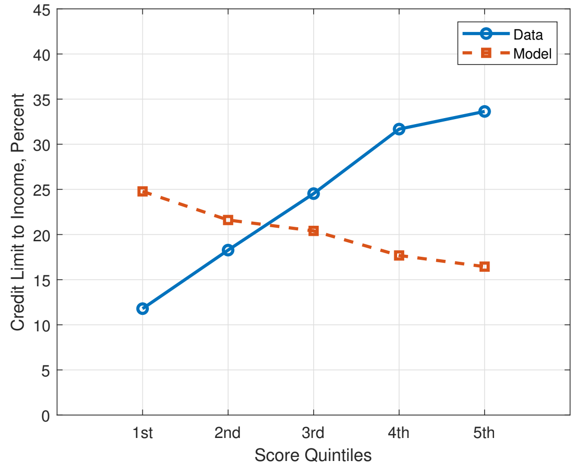

Notes: Data plots are author calculations based on Federal Reserve Y-14M data

tween credit-limit-to-income ratios and scores is a target of the estimation. The estimation cannot find any parameter values that can make the model generate a positive slope.

Why can't the model generate a positive slope? As explained above, the credit-limit-to-income ratio is decreasing in y. However, it is the high-y individuals who have a low probability of default and, therefore, high credit scores. Combining these two facts means that credit scores and credit-limit-to-income ratios *have* to be inversely related. Our inference is that the model is missing non-income factors that affect credit scores and are not strongly correlated with income. This inference is supported by the fact that in the data, the average income in the highest score quintile is 1.26 times the average income in the bottom score quintile, while in the model this ratio is 1.82.[21]

7 The Base Model with Multiple Types

The estimation assumes $I = 3$ with μ^i set to 0.25 for the type with the lowest β and μ^i set to 0.375 for the two remaining types. We departed from a uniform distribution of types to better account for the sharp change in the default behavior between the lowest credit score bin and the next one. By lowering the fraction of the most impatient type, we can make the estimation avail of more heterogeneity at the lower end of the score distribution without increasing the number of types.

[21]Given that high-y people congregate in the top score quintile, this might seem low. However, many high-y individuals have low P, and many low-y individuals have high P. This mixing tends to equalize average incomes across score bins.

7.1 Parameter Estimates and Model Fit

The estimated parameters are reported in Table 3. The estimation calls for a wide variation in discount factors, with all types estimated to be quite impatient. The most impatient type (Type 1) has a 59 percent quarterly discount rate, while the most patient type (Type 3) has a 3 percent quarterly discount rate. Default costs are estimated to be positively related to discount factors. The most impatient type forfeits 23 percent of transitory income to creditors upon default, whereas the most patient type forfeits 72 percent.

Table 3
Estimate of Type Parameters

Type	β	$(1-\phi)$
1	0.63	0.23
2	0.87	0.32
3	0.97	0.72

Figure 8 displays the targeted data moments and their model counterparts. The panels show that the Base Model can explain the score-based patterns in default frequencies (left panel) and credit-limit-to-income ratios (right panel) quite well and does far better than the canonical model (compare with Figure 5 (right panel) and Figure 7).

Figure 8

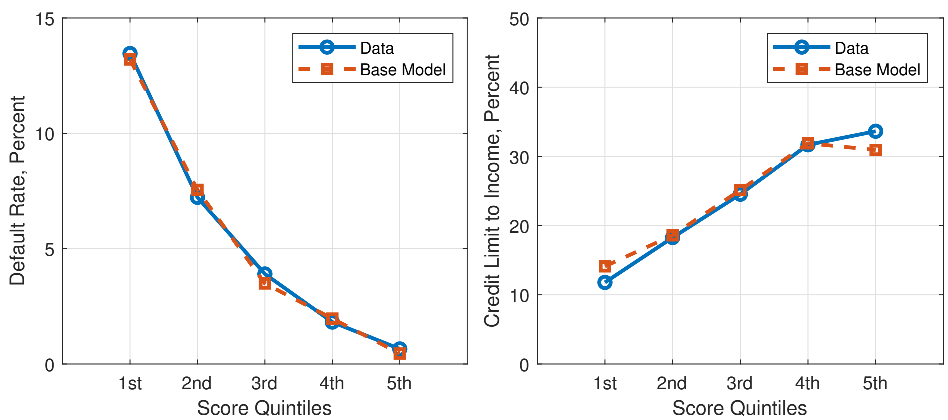

Notes: Data plots are author calculations based on Federal Reserve Y-14M data

To understand how the expanded model can do so much better, consider the implications of discount factor heterogeneity for default frequencies across scores. Since more patient people are less likely to borrow and default, they populate the upper score bins and lower the frequency of default in those bins relative to lower bins. The estimation amplifies this negative gradient by attributing pronounced differences in patience across types, especially between the least and most patient types, thereby getting the *sharp* drop-off in default frequencies across score bins.

Turning to the credit-limit-to-income ratios, it is helpful to recall the effects of income heterogeneity: As in the canonical model, heterogeneity in y makes for a negative slope between scores and credit-limit-to-income ratios. Adding discount factor heterogeneity *strengthens* this tendency because, all else equal, the more patient desire cards with lower interest rates and limits (as we discussed in connection with the red line in the right panel of Figure 6). The estimation overcomes these tendencies by assigning higher default costs to the more patient types. This allows card companies to offer the low interest rates this group wants for much higher credit limits relative to income. Also, higher default costs lower default propensities and is another force depressing the default frequency of high credit score people. Thus, correlated variation in default costs and patience allows the model to fit the pattern in credit-limit-to-income ratios without sacrificing its ability to match the pattern in default frequencies.[22]

Table 4 reports the distribution of types across score quintiles among individuals acquiring new cards in any period. The most impatient type (Type 1) is concentrated in the bottom two quintiles, and the most patient type (Type 3) is mostly concentrated in the top two quintiles. Middle patient types (Type 2) are concentrated in the middle three quintiles.[23]

Table 4
Type Distribution of New Originations

	I	II	III	IV	V	Fraction
Types			Score Quintiles			
1	0.71	0.29	0.00	0.00	0.00	0.28
2	0.00	0.32	0.43	0.23	0.02	0.37
3	0.00	0.01	0.12	0.32	0.55	0.35

Notes: The final column reports the steady state fraction of each type in the pool of households getting new cards.

Figure 9

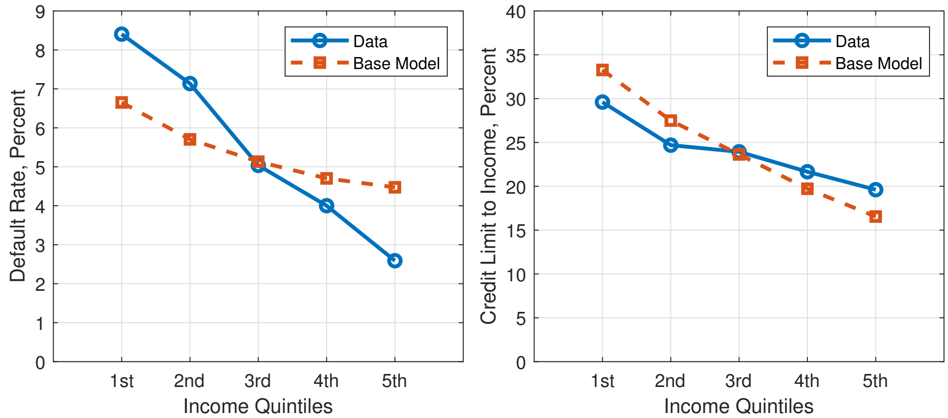

Notes: Data plots are author calculations based on Federal Reserve Y-14M data

7.2 Non-targeted Moments

Figure 9 shows the patterns for income quintiles. The panels show that the Base model can account for the patterns in a qualitative sense, but the quantitative fit is about the same as the canonical model (compare with the left panels of Figures 5 and 6). The drop-off in default frequencies with income is not as sharp as in the data, and the drop-off in credit limit ratios with income is too sharp relative to the data.[24]

Figure 10 shows the patterns in utilization rates for income and score quintiles. The performance of the Base model is impressive: Despite being non-targeted moments, the model can produce a close quantitative match. This is another confirmation that people are, in fact, impatient and borrow at high interest rates.

Figure 11 displays the interest rates by income and score quintiles for the Base model. Remarkably, the model gets the level of interest rates about right. There is a tendency for the model to under-predict interest rates by a modest amount. But this is to be expected, as the model abstracts from the flow costs of running a credit card company.

[22] Another way to view the estimation results is that the pattern in default frequencies (across score bins) identifies the variation in β, and the pattern in the credit-limit-to-income ratio identifies the variation in ϕ.

[23] The table also reports the prevalence of the three types among the people getting new cards. While the most impatient type constitutes 25 percent of the economy, it represents 28 percent of individuals getting new cards. This is due to selection: Among the group searching and obtaining new cards are individuals who defaulted in the past, and the most impatient type is overrepresented in that set.

[24] The fit can be improved if i and P are assumed to be positively correlated. Since high-P individuals are more likely to inhabit the high income bins, the positive correlation will raise the fraction of Type 2 and Type 3 individuals in the upper income bins and lower them in the lower income bins. This will lower default frequencies and raise credit-limit-to-income ratios in the upper income bins, while having the opposite effects in the lower income bins. Furthermore, the improved fit won't interfere with the model's success in terms of score-based patterns or utilization rates, as default frequencies and utilization rates are independent of P due to the scaling properties.

Figure 10
Utilization Rates, Model and Data

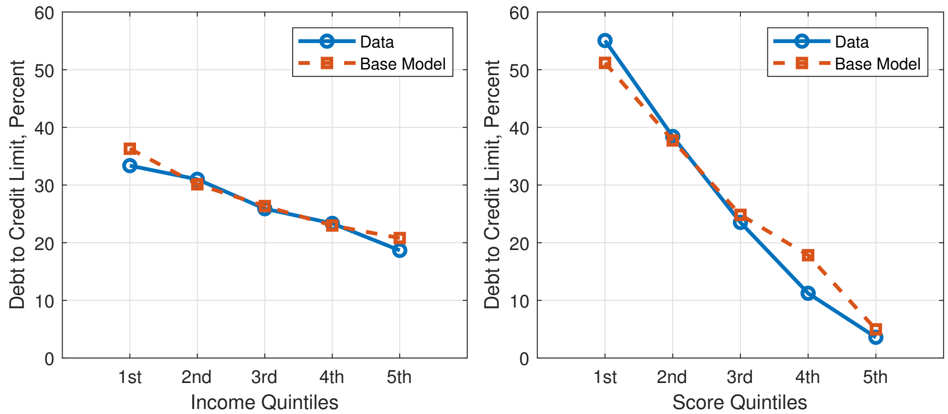

Notes: Data plots are authors' calculations based on Federal Reserve Y-14M data

Figure 11
Utilization Rates, Model and Data

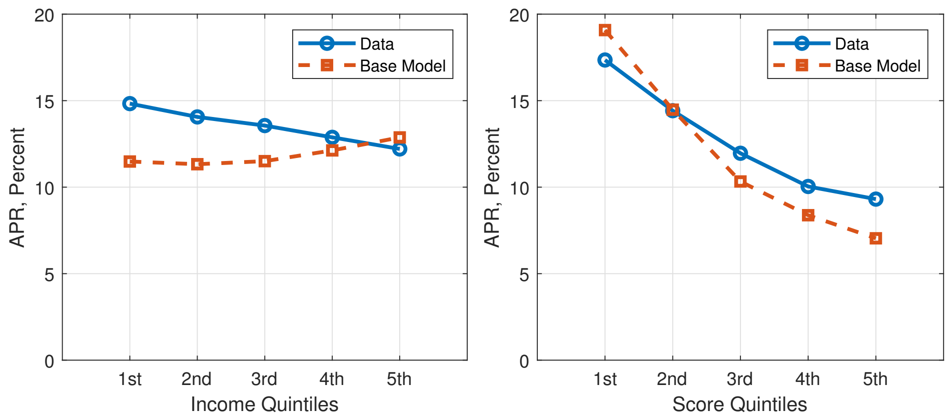

Notes: Data plots are author calculations based on Federal Reserve Y-14M data

The fact that the model is getting the level of interest rates about right means that it is successfully predicting spreads that are well in excess of default probabilities. This is illustrated in Figure 12. The left panel plots the average spread, defined as the difference between contract interest rates and the 2 percent risk-free rate (the opportunity cost of funds in the model), and annualized default frequency by income quintiles. In the standard one-period model of unsecured consumer debt, these two lines should lie on top of each other. Here, there is a gap of about 8 to 9 percentage points.

Figure 12
Spreads and Default Probabilities

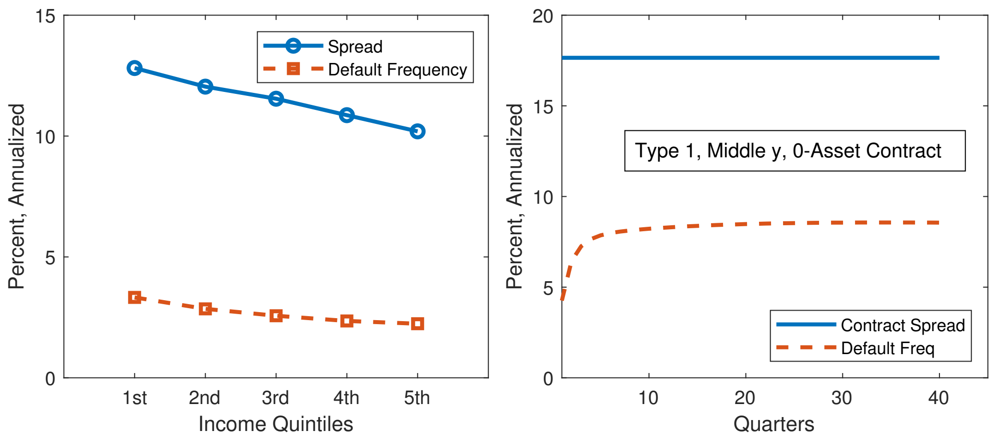

Since income bins contain a mix of contracts offered to people of different types, with varying persistent y's, the right panel focuses on a specific contract for clarity. This contract is provided to a Type 1 with $P = 1$, middle level of y, and no assets. The blue horizontal line, at roughly 17 percent, represents the spread between the contract interest rate and the 2 percent risk-free rate. The red dashed line is the default probability (annualized) on surviving contracts.[25] The spread exceeds default frequency at any time by more than 10 percentage points. Understanding how the model can generate spreads that exceed default probabilities by such a wide margin and still be consistent with competitive provision of cards will be the primary focus of Section 8.

[25]The default frequency is increasing because it takes time for cardholders to accumulate debts. The longer the card lasts (i.e., the longer there is no default or separation), the closer the default rate gets to the long-run default rate on the card.

Figure 13
Contract Terms by Type, Income & Assets

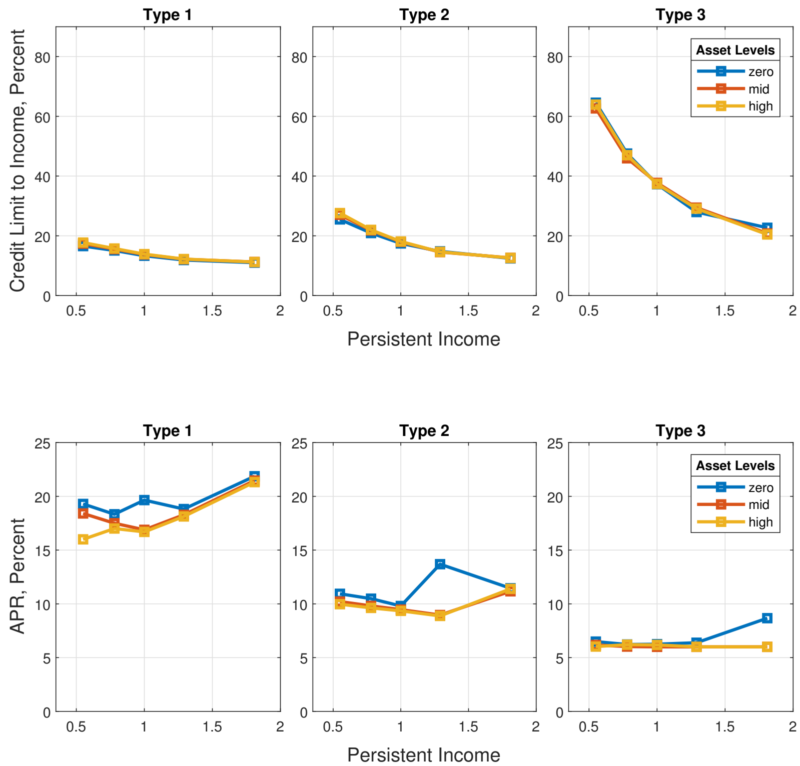

7.3 Equilibrium Contracts

In the Base model, we permit card companies to distinguish between 3 different asset levels for each type. With 5 persistent income states and 3 types, there are 54 distinct contracts offered. Figure 13 displays the contract patterns visually.

Several features stand out. Turning first to the top panel, we see that credit-limit-to-income ratios rise with i, reflecting the rise in default costs with i. Second, for all types, the credit-limit-to-income ratio decreases in y, reflecting that increases in credit limits do not keep pace with increases in persistent earnings levels. Third, the contract credit limits are quite insensitive to asset levels for each type. This is a consequence of the long-duration nature of a credit card contract: In the long run, the distribution of a person's asset position is independent of his or her initial assets. Thus, contract terms are not overly affected by the level of assets at origination. This is also why it is possible to have a sparse asset grid for contracts without affecting the results.

Turning to the bottom panel, we see that contract interest rates decline with i, reflecting the ordering of default frequencies across types (displayed later in the paper). Contract interest rates are mildly sensitive to the assets at origination, with individuals who have zero assets offered somewhat higher rates. Contract interest rates are also not very sensitive to variation in y but tend to increase with y. This tendency is why we see a similar tendency in the left panel of Figure 11. This tendency is puzzling because default rates decline with y. The reason is connected to the gap between spreads and default rates, to which we now turn.

8 The Logic of Credit Card Interest Rates

There are two key aspects to this: the default behavior of individuals who borrow on a credit card and the fact that the credit card interest is not responsive to the amount borrowed on the card.

Default Behavior

Consider an individual who holds a card and has debt. If this person draws a low m, she can buffer her consumption against the low m by shedding her debt at a low cost. On the other hand, if her debt is well below her card limit, she can also buffer her consumption by borrowing more on the card. Default is more likely when the second option is unavailable or its scope is limited, i.e., if she is maxed out on the credit card or carrying balances close to her credit limit.

Figure 14
Default Probability of a Zero Asset, Type 1 Individual

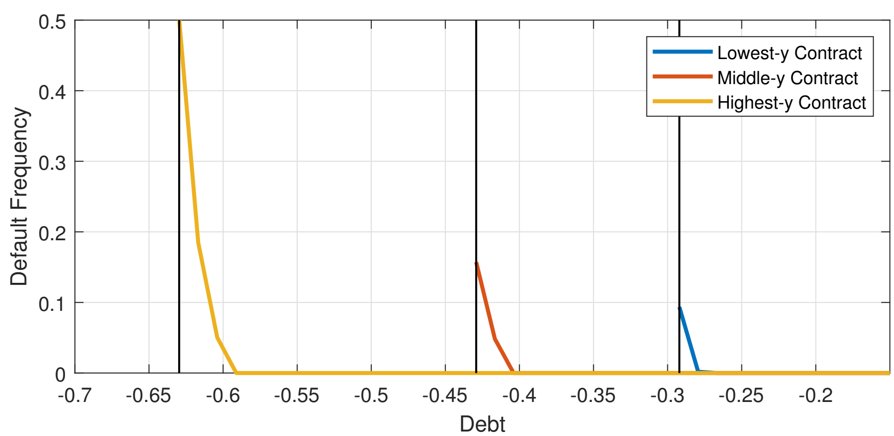

Figure 14 plots the default frequency against the beginning of period debt of a Type 1 individual for three contracts she might hold in equilibrium: zero-asset contracts for the low, middle, or highest y. The solid vertical lines are drawn at the credit limit of each contract. For each contract, the default frequency is zero until the debt level approaches

Table 5
Default Share by Persistent Income and Type

Income	Type 1	Type 2	Type 3	Total
0.55	0.17	0.12	0.04	0.33
0.78	0.15	0.07	0.02	0.24
1.00	0.12	0.05	0.01	0.18
1.29	0.10	0.04	0.01	0.15
1.81	0.08	0.03	0.00	0.10
Total	0.62	0.30	0.07	1.00

the credit limit. Thus, what matters for default is not the level of debt *per se* but how close the cardholder is to her credit limit, i.e., on her utilization rate. In our model, the utilization rate of defaulters is just short of 100 percent: People who default are very close, or at, their credit limit, although not everyone in this situation defaults (the probability of default is generally less than 50 percent even at the limit). These predictions align with the data: In our sample, defaulters have a utilization rate of 89 percent at the time of default, but only a small fraction of borrowers with high utilization rates default.

To round out the discussion of default behavior, Table 5 displays the steady state incidence of default across persistent income levels and types. Default is most prevalent among the most impatient people with the lowest y. For each type, the incidence of default drops quickly with y, and for each y, it drops quickly with type.

Why Spreads Can Far Exceed Default Probability

To understand why credit card spreads can far exceed default probabilities, it is helpful to consider a stylized two-period model. Imagine that a card company has committed to a credit limit L and a gross interest rate R to a large number of cardholders. In both periods, cardholders experience an idiosyncratic shock, which can be either good or bad with probability $(1 - \Delta)$ and Δ, respectively. Among cardholders, a fraction $(1 - \xi)$ never borrows on the card, regardless of their idiosyncratic state. Among the remaining ξ fraction, the pattern of borrowing is as follows. In the first period, a cardholder borrows $z < L$ if she experiences the good shock and borrows the limit L if she experiences the bad shock. In the second period, if her shock is good, she repays her debts. But if she gets the bad shock, she repays if she had borrowed z and defaults if she had borrowed L. Thus, if the borrower is not at the credit limit, she never defaults, and if she is at the credit limit, she defaults with positive probability less than 1. This captures, in a simple way, the default patterns in Figure 14.

Assuming that the card company's opportunity cost of funds is R_f, the probability distribution of its second-period payoff for any given credit card account is:

$$\Pi = \begin{cases} 0 \cdot [R - R_f] & \text{with prob. } (1-\xi) \\ z \cdot [R - R_f] & \text{with prob. } \xi \cdot [1-\Delta] \cdot [1-\Delta] \\ L \cdot [R - R_f] & \text{with prob. } \xi \cdot \Delta \cdot [1-\Delta] \\ z \cdot [R - R_f] & \text{with prob. } \xi \cdot [1-\Delta] \cdot \Delta \\ L \cdot [0 - R_f] & \text{with prob. } \xi \cdot \Delta \cdot \Delta. \end{cases}$$

Given this, the card company's second-period expected profit is

$$\mathbb{E}\Pi = \xi\,[(1-\Delta)\cdot z + \Delta \cdot L \cdot (1-\Delta)]\,R - \xi\,[(1-\Delta)\cdot z + \Delta \cdot L]\,R_f.$$

Rearranging gives

$$[R - R_f] \cdot \frac{\xi[(1-\Delta)\cdot z + \Delta \cdot L]}{RL} = \xi\Delta^2 + \frac{\mathbb{E}\Pi}{RL}. \tag{17}$$

The l.h.s. is the product of the contract spread and the expected revolving debt scaled by the maximum debt that can be owed on the card. The r.h.s. is the sum of the probability of default on the contract, $\xi\Delta^2$, and expected profits on the card, scaled, again, by the maximum that can be owed on the card.

Equation (17) explains why spreads can exceed default probability. Assume that competition has forced $\mathbb{E}\Pi$ to zero. Since the ratio multiplying the spread is less than 1, the zero-profit contract spread will *exceed* $\xi\Delta^2$, the default probability on the contract. This happens because repayments occur on debts that, on average, are smaller than the debt on which default occurs *and* the same interest rate applies to both small and large debts. In this situation, spreads must exceed the default probability for the card company to break even.

Equation (17) helps us understand the spreads we see in our model. Note that the term multiplying the contract spread can be interpreted as the *expected utilization rate* on the card, i.e., the ratio of the average level of revolving balances on which interest is paid and the credit limit (interest inclusive). Denote the expected utilization rate by U to get

$$\frac{[R - R_f]}{\xi \cdot \Delta^2} = \frac{1}{U} + \frac{\mathbb{E}\Pi/RL}{\xi \cdot \Delta^2 \cdot U}. \tag{18}$$

The l.h.s. is the ratio of contract spread to the default probability. On the r.h.s., the first term is the inverse of the expected utilization rate, and the second term is a transform of the expected profit rate on the card. Ignoring the profit term, Equation (18) suggests that spreads scaled by the default probability should be explained by the inverse of the utilization rate on the card.

Figure 15
Spreads, Default Probabilities and Utilization Rates, Model

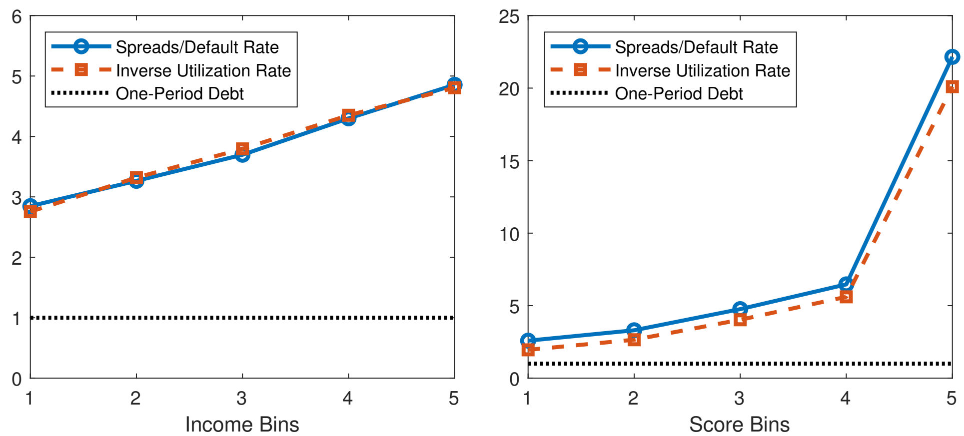

We can test this prediction against our model. The left panel of Figure 15 plots the ratio of contract spreads to default probability for each income bin along with the inverse of the utilization rate for each income bin. The right panel does the analogous plot for score bins. The dotted line at the bottom of both panels represents the prediction of the standard one-period debt model with risk-neutral and competitive lenders, for which this ratio is always 1.

A couple of things stand out. First, spreads at origination are several multiples of the two-year default frequency. In the left panel, the ratio ranges from approximately 3 to 5, and in the right panel, it ranges from approximately 4 to above 23. Second, these multiples line up remarkably well with the inverse of the utilization rate, indicating that the logic for spreads in the simple model is informative about spreads in the Base model.

Figure 16 displays the corresponding plots for our data, assuming a cost of funds of 2 percent. Here, too, the inverse utilization rate has substantial explanatory power, especially for credit score bins. This is evidence that the reason for spreads in the model is an important part of the story for why spreads exceed default probabilities in the real world.

Figure 16
Spreads, Default Probabilities and Utilization Rates, Data

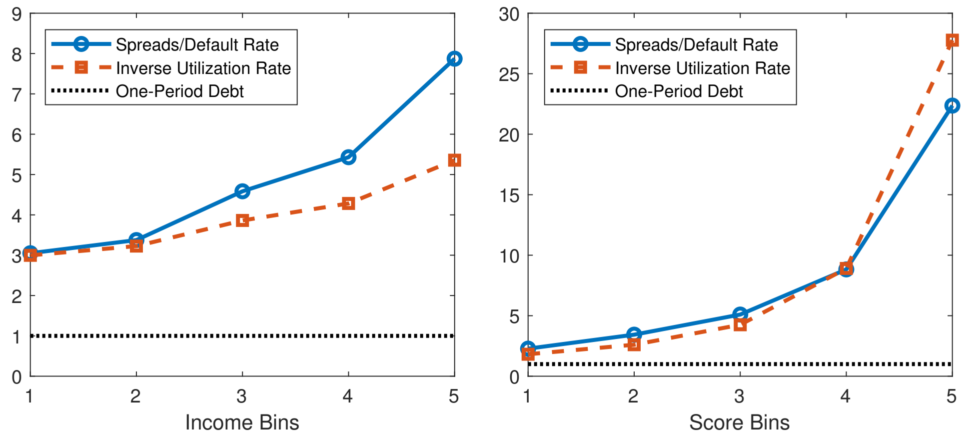

Notes: Data plots are author calculations based on Federal Reserve Y-14M data

How Much Do Search-induced Markups Contribute to Spreads?

The fact that it takes time for an individual to match with a credit card company and it takes both time and resources for card companies to match with customers means that, post-match, there is a surplus to be shared between the customer and the card company. If all the post-match surplus were to go to the customer, the card company would earn zero profits post-match. We define the markup due to search as the difference between the contract rate and the post-match zero-profit interest rate. Table 6 reports these markups for all contracts.

Across contracts, the markup varies from 1.26 to 0.09 percentage points. Looking across cells, we find that asset-poor, low-income, and impatient people suffer higher markups than rich, high-income, and patient individuals. All else equal, markups always decline with higher patience, and within each type, they decline with level of assets. This is because not having the option to borrow is more consequential for impatient and low-asset individuals, and therefore they give up a greater surplus to credit card companies.

Markups vary nonmonotonically with y, reflecting the varying strengths of two countervailing forces: On the one hand, a higher y means more desire to save for precautionary reasons, which makes not having a card less consequential and, therefore, favors lower markups. On the other hand, a higher y means more volatile transitory shocks (the range of variation of the transitory shock increases with y), which makes not having the option to borrow more consequential. Which effect dominates as y varies depends on y and i.

Table 6
Excess of Contract APRs Over Zero-Profit APRs

Type 1				Type 2				Type 3			
$y \downarrow$ $a \rightarrow$	0.00	0.45	1.20	$y \downarrow$ $a \rightarrow$	0.00	0.90	2.40	$y \downarrow$ $a \rightarrow$	0.00	3.00	8.16
0.55	1.26	0.77	0.52	0.55	0.68	0.31	0.21	0.55	0.32	0.11	0.09
0.78	1.08	0.70	0.53	0.78	0.66	0.31	0.21	0.78	0.34	0.12	0.20
1.00	1.13	0.64	0.50	1.00	0.59	0.35	0.22	1.00	0.47	0.16	0.24
1.29	1.02	0.73	0.56	1.29	0.77	0.31	0.21	1.29	0.65	0.20	0.17
1.81	1.05	0.76	0.66	1.81	0.69	0.36	0.31	1.81	0.92	0.25	0.31

Notes: a and y refer to asset and income levels at origination. Each cell reports Gross Contract Interest Rates, Ann. − Gross Zero Profit Interest, Ann.

In magnitude, markups are non-negligible for 0-asset Type 1 individuals, exceeding 1 percentage point. For instance, of the 17 percentage point spread on the Type 1, middle y, and 0-asset contract (Figure 12, left panel), 1.13 percentage points is due to markups. Overall, though, markups are not a large contributor to credit card spreads. This reflects costless and free entry into the card business, as well as low contract posting costs.

9 Implications

9.1 Rate Cap Experiment

This section examines how credit contracts change if contract interest rates are capped at 10 percent. An examination of Figure 13 shows that a 10 percent cap would be strongly binding for Type 1 individuals, marginally binding for Type 2 individuals, and not binding for Type 3 individuals. Furthermore, among Type 1 individuals, contracts are almost independent of the individual's asset level. These observations suggest that the implications of a 10 percent rate cap can be examined by analyzing its impact on 0-asset, Type 1 individuals with varying persistent income levels.

The left panel of Figure 17 plots how credit limits change with the rate cap. The cap leads to a drop in credit limits that ranges between 25 and 31 percent. From the perspective of card companies, reducing the credit limit reduces the loss given default and also reduces the probability of default, as shown in the right panel. The decline in default probabilities is also substantial, ranging between 5 and 7 percentage points.

The left panel of Figure 18 shows the effect of the rate cap on the probability of finding a card conditional on search. The probability declines about 5 to 7 percentage points. The drop in contact probability is a direct consequence of the reduction in the card's profitabil-

Figure 17
Effect of Rate Cap on Credit Limits and Default Probability
Type 1, 0-Asset Contracts

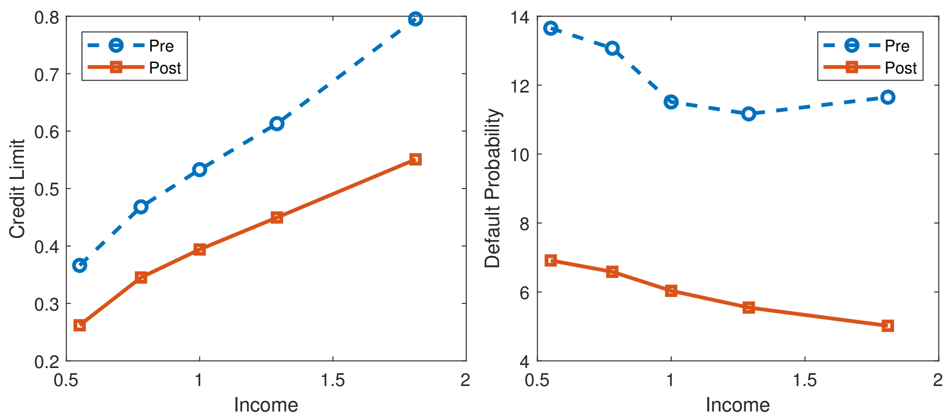

Figure 18
Effect of Rate Cap on Probability of Getting a Card and Welfare of Cardholder

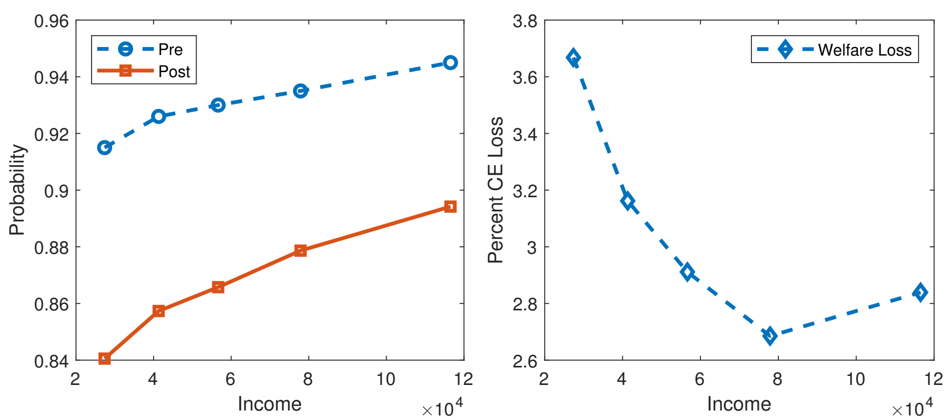

ity. The measure of posted contracts declines to the point where the likelihood of finding a *customer* has risen enough to compensate for the reduction in profitability.

Finally, the right panel plots the consumption equivalent welfare loss across persistent income types. It is the percentage of consumption that a 0-asset, Type 1 individual would be willing to pay not to have the cap imposed on them. The loss ranges between 2.7 and 3.7 percent. The lowest-y individuals are hurt the most.

Since our data sample dropped subprime (the most risky) borrowers, the effects described above apply to the bottom 30 percent of our truncated credit score distribution.[26] Roughly speaking, this is the group of borrowers who fall into the "near-prime," and "prime" categories with scores between 620 and 720. For this group, less credit card indebtedness could lead to better terms on secured loan products, such as mortgages or auto loans. If so, these benefits of the rate cap are not part of our welfare calculation.

9.2 Credit and the Marginal Propensity to Consume

In this section, we briefly report the implications of our model for the distribution of the MPC. MPCs are, of course, an important ingredient in many macroeconomic phenomena (Kaplan and Violante (2022)). Furthermore, it is understood that credit constraints affect MPCs by making individuals engage in buffer-stock/precautionary savings (Deaton (1991a), Deaton (1991b)). Given the endogeneity of credit limits and (borrowing) interest rates in our model, it is interesting to see what our model implies about the distribution of MPCs.

To set the stage, the left panel of Table 7 reports the average assets, in steady state, by types and persistent earnings, where an individual's assets are normalized by her permanent earnings level P. When persistent income y is at its lowest level, all types carry balances on their cards, on average. Note that a more patient type is *more*, not less, indebted. This is a result of endogeneity of credit terms: Holding y constant, more patient types are offered contracts with higher credit limits and lower interest rates (recall Figure 13) as they also have higher default costs.[27] Since a lowest-y person expects her income to rise over time, she has a strong incentive to borrow, which can be met better if she is known to have higher default costs. For other y levels, the buffer-stock savings motive is generally dominant, and average assets tend to increase with type. Holding type fixed,

[26] Since all Type 1 individuals are affected by the rate cap, all individuals in the first quintile of scores and half of all individuals in the second quintile of scores (see Table 4) are affected. Since Type 1 individuals are likely to rank below Type 2 in credit scores, the policy will affect the bottom 30 percent of the truncated score distribution.

[27] Because of history dependence, not everyone with currently low income will have a contract meant for low-income individuals. But persistence of y ensures that most of them will.

average (normalized) assets increase in y, reflecting the precautionary savings motive: Individuals with higher y accumulate less debt or more savings, on average.

Table 7
Mean Assets and Utilization Rates

Income	Mean Assets			Mean Util. Rates		
	Type 1	Type 2	Type 3	Type 1	Type 2	Type 3
0.55	−0.13	−0.16	−0.43	0.62	0.43	0.38
0.78	−0.14	−0.06	0.14	0.56	0.33	0.22
1.00	−0.12	0.07	0.72	0.50	0.26	0.13
1.29	−0.09	0.25	1.48	0.45	0.20	0.08
1.81	−0.05	0.53	2.89	0.41	0.16	0.04

Notes: Mean assets in a cell is the mean of assets normalized by permanent income of all individuals in that cell expressed as a proportion of median quarterly income in the economy. The mean utilization rate is the mean utilization rate of all individuals in a cell, where individuals with nonnegative assets have a utilization rate of zero.

The right panel of Table 7 reports the steady state average credit card utilization rate across y and types. For individuals with balances on their cards, the utilization rate is the ratio of the beginning-of-period balance on the card to the card's credit limit. For those without balances, the utilization rate is zero. For a given persistent income level, the utilization rate is monotonically declining with patience and default costs: More patient/higher default cost individuals hold less debt *relative to their limits*, on average. For a given type, the utilization rate consistently decreases with persistent income levels: As income rises, individuals borrow less frequently, and when they do, they borrow smaller amounts.

Figure 19 shows the distribution of 3-month MPCs implied by this sort of behavior. The mean of the distribution is 0.33 and the median is only slightly lower at 0.30. The distribution is quite dispersed and is consistent with the wide variation in estimated MPCs. Johnson, Parker, Souleles, and McClelland (2013) estimate average 3-month MPC to be in the $0.5-0.9$ range and Orchard, Ramey, and Weiland (n.d.) estimate it to be in the $0.3-0.5$ range, with their preferred point estimate as 0.33 [28] Thus, the model-implied MPCs are in line with empirical evidence.

The top left panel of Table 8 displays how MPCs relate to individual characteristics. For Type 1 individuals, their mean MPCs range between $0.59 - 0.61$, for Type 2, between $0.33 - 0.35$; for Type 3, between $0.09 - 0.13$. Thus, the main variation is in types, with

[28] Orchard, Ramey, and Weiland (n.d.) note that when the change in expenditures is broken down into change in durables and nondurables exependitures, it is mostly durable expenditures that respond.

Figure 19
Distribution of MPCs

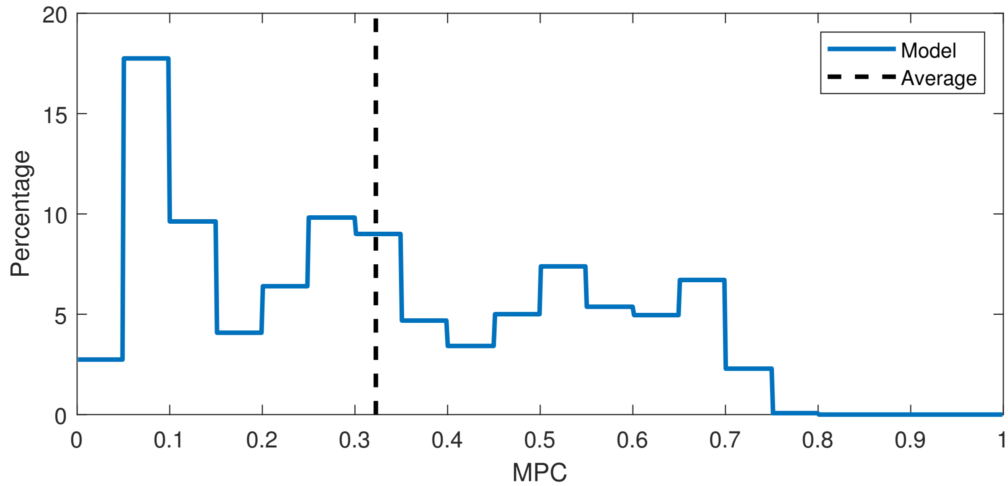

differences in y playing a lesser role. Regarding type, the primary factor is the estimated discount factor for each type. A type's impatience determines the type's average financial wealth. Since the more impatient types have lower wealth, they have higher marginal utility and MPC. Recall that variation in β across types was identified by the pattern in default frequency across credit scores. The implication is that credit market facts matter for understanding MPCs.

The bottom panels of Table 8 separate each type-income group into those with cards and those without. People without cards are mostly newborns or those who have recently lost their card and haven't yet acquired a new one (some are defaulters who are not eligible for a card). Holding type and y constant, the MPC of the latter group is measurably higher than that of the former group, with the sole exception of the most patient and highest y individual. This is because non-cardholders have assets, while many cardholders have debt. All else equal, an indebted person's MPC will be lower because a portion of the windfall gain will be used to pay down debt (Koşar, Melcangi, Pilossoph, and Wiczer (2024), Boutros and Mijakovic (2024)). The difference in MPC is not large because many cardholders do not have any debts and are similar in this respect to non-cardholders.

The top right panel of Table 8 shows how MPCs change if the option to borrow is taken away altogether.[29] Comparing cell by cell, MPCs are now higher, but in most cases, the increase is at most 1 percentage point. This pattern again reflects that MPCs are higher when there is no debt to be paid back. Consistent with this, the most significant increase

[29] The risk-free rate R_f is kept unchanged, so general equilibrium effects are ignored.

Table 8
MPCs by Type and Persistent Earnings Levels

Income	Type 1	Type 2	Type 3	Type 1	Type 2	Type 3
	Card Economy Average: 0.32			No-Cards Economy Average: 0.33		
0.55	0.61	0.35	0.12	0.62	0.37	0.18
0.78	0.61	0.35	0.13	0.62	0.36	0.15
1.98	0.59	0.34	0.11	0.61	0.35	0.13
1.29	0.59	0.33	0.10	0.61	0.33	0.10
1.81	0.60	0.33	0.09	0.61	0.33	0.10
	Card Economy					
	Cardholders, Avg. 0.30			Non-Cardholders, Avg. 0.46		
0.55	0.59	0.34	0.11	0.64	0.41	0.27
0.78	0.59	0.34	0.13	0.64	0.41	0.23
1.98	0.58	0.33	0.11	0.63	0.39	0.17
1.29	0.58	0.32	0.09	0.62	0.37	0.11
1.81	0.59	0.32	0.09	0.62	0.37	0.09

in MPC (4 percentage points) is for lowest-y-Type 3 group, which is also the group that carries the most credit card debt (Table 7).

We round out the discussion of MPCs by comparing the distribution of assets in the model to the distribution of net liquid assets in the data. For this comparison, we use the 2013 Survey of Consumer Finances (SCF). To make the comparison as clean as possible, we eliminated the top and bottom 3 percent of the income distribution, which brought the range of incomes in the survey in line with the range of incomes in our model. We also eliminated people without checking accounts or access to one so that we can consider the pool of people who either have a credit card or could get one. We also adjusted reported debts for underreporting as recommended in Laibson, Lee, Maxted, Repetto, and Tobacman (2024).

Figure 20 displays the distributions in the model and in the data. Several things are to be noted. First, the small blue spike at zero is the mass of people with zero liquid assets in the data. The red dotted spike (slightly offset from zero) is the corresponding mass in the model. There is a close match. Second, the fraction of people within the first bin of liquid assets is much larger in the data than in the model. Most likely, this "spike" is the "wealthy hand-to-mouth" phenomenon highlighted in Kaplan and Violante (2014). Our model cannot speak to this because we haven't modeled high-return but illiquid assets. Finally, the frequency of people with substantial debts is higher in the model than in the data. However, note that our model aligns closely with the utilization rates in our data

Figure 20
Distribution of Liquid Balances, Model and Data

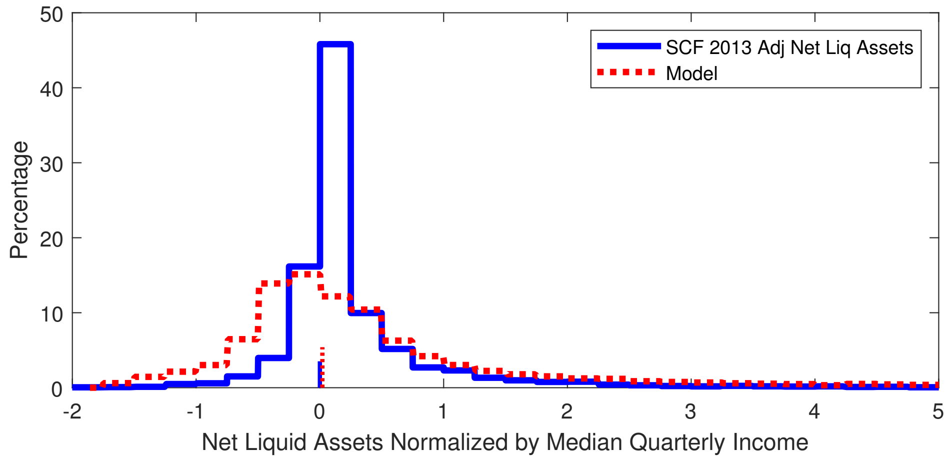

Notes: Data plot is author calculation based on the 2013 Survey of Consumer Finances

(see Figure 10). It is possible that the under-reporting of large credit card debts is even more serious in the SCF.

References

AGUIAR, M., M. BILS, AND C. BOAR (2025): "Who Are the Hand-to-Mouth?," *Review of Economic Studies*, 92(3), 1293–1340.

AIYAGARI, S. (1994): "Uninsured Idiosyncratic Risk and Aggregate Savings," *Quarterly Journal of Economics*, 109, 659–684.

ATHREYA, K. (2002): "Welfare Implications of Bankruptcy Reform Act of 1999," *Journal of Monetary Economics*, 49, 1567–1595.

ATHREYA, K., J. M. SANCHEZ, AND J. MUSTRE-DEL-RIO (2019): "The Persistence of Financial Distress," *Review of Financial Studies*, 32, 3851–83.

ATHREYA, K., X. S. TAM, AND E. R. YOUNG (2009): "Unsecured Credit Markets Are Not Insurance Markets," *Journal of Monetary Economics*.

——— (2012): "A Quantitative Theory of Information and Unsecured Credit," *American Economic Journal: Macroeconomics*, 4(3), 153–183.

AUSUBEL, L. M. (1991): "The Failure of Competition in the Credit Card Market," *American Economic Review*, 81(1), 50–81.

BETHUNE, Z., J. SALDAIN, AND E. R. YOUNG (2024): "Consumer Credit Regulation and Lender Market Power," Staff Working Paper 2024-36, Bank of Canada.

BOUTROS, M., AND A. MIJAKOVIC (2024): "Macroeconomic Implications of Coholding," Staff Working Paper 2024-16, Bank of Canada.

BRAXTON, J. C., K. HERKENHOFF, AND G. PHILLIPS (2024): "Can the Unemployed Borrow? Implications for Public Insurance," *Journal of Political Economy*, forthcoming.

CARROLL, C., J. SLACALEK, K. TOKUOKA, AND M. N. WHITE (2017): "The Distribution of Wealth and the Marginal Propensity to Consume," *Quantitative Economics*, 8, 977–1020.

CHATTERJEE, S., D. CORBAE, K. DEMPSEY, AND J.-V. RÍOS-RULL (2023): "A Quantitative Theory of the Credit Score," *Econometrica*, 91(5), 1803–1804.

CHATTERJEE, S., D. CORBAE, M. NAKAJIMA, AND J.-V. RIOS-RULL (2007): "A Quantitative Theory of Unsecured Consumer Credit with Risk of Default," *Econometrica*, 75(6), 1525–1589.

DEATON, A. (1991a): "Saving and Liquidity Constraints," *Econometrica*, 59(5), 1221–1248.

——— (1991b): *Understanding Consumption*. Oxford University Press.

DEMPSEY, K., AND F. IONESCU (2025): "Default Risk and Pricing in the U.S. Credit Card Market," *Journal of Political Economy: Macro*, accepted.

DEN HAAN, W. J., G. RAMEY, AND J. WATSON (2000): "Job Destruction and Propagation of Shocks," *American Economic Review*, 90(1), 482–498.

DROZD, L., AND M. KOWALIK (2023): "The Conundrum of Zero APR: An Analytical Framework," Federal Reserve Bank of Philadelphia, Working Paper No. 23-06.

DROZD, L., AND J. NOSAL (2008): "Competing for Customers: A Search Model of the Market for Unsecured Credit," Unpublished.

EVANS, D. S., AND R. SCHMALENSEE (2005): *Paying with Plastic: The Digital Revolution in Buying and Borrowing*. MIT Press, 2nd edn.

EXLER, F., AND M. TERTILT (2020): "Consumer Debt and Default: A Macro Perspective," *Oxford Research Encyclopedia of Economics and Finance*.

FULFORD, S. L., AND S. D. SCHUH (2024): "Credit Cards, Credit Utilization and Consumption," *Journal of Monetary Economics*, (in press).

GALENIANOS, M., AND A. GAVAZZA (2022): "Regulatory Interventions in Consumer Financial Markets: The Case of Credit Cards," *Journal of the European Economic Association*, 20(5), 1897–1932.

GALENIANOS, M., T. H. LAW, AND J. NOSAL (2023): "Market Power in Credit Card Markets," *mimeo*.

GELMAN, M. (2021): "What Drives Heterogeneity in the Marginal Propensity to Consume? Temporary Shocks vs Persistent Characteristics," *Journal of Monetary Economics*, 117, 521–542.

GOURINCHAS, P.-O., AND J. A. PARKER (2002): "Consumption over the Life Cycle," *Econometrica*, 70(1), 47–89.

HATCHONDO, J. C., AND L. MARTINEZ (2009): "Long Duration Bonds and Sovereign Defaults," *Journal of International Economics*, 79(1), 117–125.

HERKENHOFF, K. F., AND G. RAVEENDRANATHAN (2025): "Who Bears the Welfare Cost of Monopoly? The Case of the Credit Card Industry," *Review of Economic Studies*, 92(5), 3067–3111.

HUGGET, M. (1993): "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies," *Journal of Economic Dynamics and Control*, 17, 953–969.

IMROHOROGLU, A. (1989): "Cost of Business Cycle with Indivisibilities and Liquidity Constraints," *Journal of Political Economy*, 97(6), 1364–1383.

JOHNSON, D. S., J. A. PARKER, N. S. SOULELES, AND R. MCCLELLAND (2013): "Consumer Spending and the Economic Stimulus Payments of 2008," *American Economic Review*, 103(6), 2530–2553.

KAPLAN, G., AND G. L. VIOLANTE (2014): "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, 82(4), 1199–1239.

——— (2022): "The Marginal Propensity to Consume in Heterogeneous Agent Models," *Annual Review of Economics*, 14, 747–775.

KOŞAR, G., D. MELCANGI, L. PILOSSOPH, AND D. WICZER (2024): "Stimulus Through Insurance: the Marginal Propensity to Repay Debt," Unpublished.

KRUEGER, D., K. MITMAN, AND F. PERRI (2016): "Macroeconomics and Household Heterogeneity," in *Handbook of Macroeconomics*, ed. by H. Uhlig, and J. Taylor, pp. 843–921. Elsevier.

KRUSELL, P., AND A. A. SMITH (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106(5), 867–896.

LAIBSON, D. (1997): "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 112(2), 443–477.

LAIBSON, D., S. C. LEE, P. MAXTED, A. REPETTO, AND J. TOBACMAN (2024): "Estimating Discount Functions with Consumption Choices over the Life Cycle," *Review of Financial Studies*, (forthcoming).

LEE, S. C., AND P. MAXTED (2023): "Credit Card Borrowing in Heterogeneous-Agent Models: Reconciling Theory and Data," *mimeo*.

LELAND, H. E. (1998): "Agency Cost, Risk Management, and Capital Structure," *Journal of Finance*, 53(4), 1213–1243.

LIVSHITS, I., J. C. MACGEE, AND M. TERTILT (2007): "Consumer Bankruptcy: A Fresh Start," *American Economic Review*, 97(1), 402–418.

——— (2016): "The Democratization of Credit and the Rise in Consumer Bankruptcies," *Review of Economic Studies*, 83, 1673–1710.

MATEOS-PLANAS, X., AND J.-V. RÍOS-RULL (2007): "Credit Lines," mimeo.

MOEN, E. R. (1997): "Competitive Search Equilibrium," *Journal of Political Economy*, 105(2), 385–411.

NARAJABAD, B. N. (2012): "Information Technology and the Rise of Household Bankruptcy," *Review of Economic Dynamics*, 15, 526–550.

ORCHARD, J., V. A. RAMEY, AND J. F. WEILAND (n.d.): "Micro MPCs and Macro Counterfactuals: The Case of the 2008 Rebates," *Quarterly Journal of Economics*, forthcoming.

RAVEENDRANATHAN, G. (2020): "Revolving Credit Lines and Targeted Search," *Journal of Economic Dynamics and Control*, 118.

RAVEENDRANATHAN, G., AND G. STEFANIDES (2025): "The Unprecedented Fall in U.S. Revolving Credit," *International Economic Review*, 66(1), 393–451.

RAVEENDRANATHAN, G., G. STEFANIDIS, AND G. SUBLET (2025): "Rate Caps on Revolving Credit Lines," *Journal of Political Economy: Macro*, accepted.

SALDAIN, J. (2023): "High-Cost Consumer Credit: Desperation, Temptation and Default," *mimeo*.

SANCHEZ, J. M. (2018): "The Information Technology Revolution and the Unsecured Credit Market," *Economic Inquiry*, 56(2), 914–930.

STORESLETTEN, K., C. TELMER, AND A. YARON (2004): "Cyclical Dynamics in Idiosyncratic Labor Market Risk," *Journal of Political Economy*, 112(3), 695–717.

WRIGHT, R., P. KIRCHER, B. JULIEN, AND V. GUERRERI (2021): "Directed Search and Competitive Search: A Guided Tour," *Journal of Economic Literature*, 59, 90–148.

10 ONLINE APPENDIX

Appendix A Additional Data Table

Table 9
Terms, Utilization and Performance

Description	I	II	III	IV	V
Mean Bin Scores	647	689	729	775	827
APR (in %)	16.00	13.00	10.00	8.00	7.00
Credit Limit to Income (in %)	5.51	8.73	11.84	15.53	16.67
Default Frequency (in %)	11.56	5.39	2.76	1.17	0.41
Debt to Credit Limit (in %)	53.80	37.20	22.80	10.80	3.42
Relative Bin Incomes	0.86	0.95	1.00	1.08	1.18
Mean Bin Incomes	27378	41299	56632	77880	116518
APR (in %)	12.40	11.60	11.01	10.44	9.91
Credit Limit to Income (in %)	14.28	11.97	11.66	10.57	9.62
Default Frequency (in %)	6.71	5.77	3.95	3.02	1.95
Debt to Credit Limit (in %)	32.13	30.12	25.17	22.72	18.22
Relative Bin Incomes	0.48	0.73	1.00	1.38	2.06

Notes: Estimates are derived from a bin-scatter regression. The APR is at an annual frequency, and credit limits are expressed as a ratio of limits and annual income.

Table 10
Utilization Rates at Default

Description	Data	Model
Util. Rate at Default (in %)	89.00	~ 100.00

Notes: Utilization rate at default is the ratio of debt to credit limit at the time of default.

Table 11
Attrition Rate of Accounts

Description	Overall
Monthly Attrition Rate (in %)	0.70

Notes: The attrition rate is the observed frequency of accounts in good standing that disappear from our sample each month for the period June 2013 through December 2017.

Appendix B Computation of Credit Score

A credit score is the probability, computed at the start of a period, that an individual will not default within the next 8 model periods including the current one.

Let $t = 1, 2, 3, \ldots, 8$ denote the time period that is t periods in the future from the current period, with $t = 1$ being next period and $t = 8$ being the period just beyond the two-year horizon.

Let:

- $k_C^i(a, y; \omega, t)$ be the probability of non-default within the next $8 - t$ periods of a type i individual in state (a, y) and having a credit card contract ω.
- $k_S^i(a, y, t)$ be the probability of non-default within the next $8 - t$ periods of a type i individual in state (a, y) who is without a card but not excluded from searching in the card market.
- $k_X^i(a, y, t)$ be the probability of non-default within the next $8 - t$ periods of a type i individual without a card in state (a, y) and who is excluded from searching in the credit card market.

Set each of these functions identically to 1 for $t + 1 = 8$ and solve for the three functions for $t = 7$ using the following recursions.

If $a \geq 0$:

$$k_C^i(a, y; \omega, t) =$$

$$(1-\nu) + \nu \cdot \left(\begin{array}{c} \xi \cdot \left[\begin{array}{c} \mu \cdot f(\theta^i(a,y))\mathbb{E}_{m,y'} k_C^i \left(a'\left(a, y, m; \omega^i(a,y)\right), y'; \omega^i(a,y), t+1\right) \\ + \left(1 - \mu \cdot f(\theta^i(a,y))\right) \mathbb{E}_{m,y'} k_S^i \left(a'(a,y,m), y'; t+1\right) \end{array} \right] \\ + (1-\xi)\, \mathbb{E}_{m,y'} k_C^i \left(a'(a,y,m;\omega), y'; \omega, t+1\right) \end{array} \right)$$

If $a < 0$:

$$k_C^i(a,y;\omega,t) = \nu \cdot \left(\begin{array}{c} \xi^i(a,y;\omega) \cdot \ \mathbb{E}_{m,y'} k_C^i\left(a'\left(a,y,m;\bar{\omega}_y^i,y'\right);\bar{\omega}^i,t+1\right) \\ +\left(1-\xi^i(a,y;\omega)\right)\mathbb{E}_{m,y'}\left[(1-D(a,y,m;\omega)) \cdot k_C^i\left(a'(a,y,m;\omega),y';\omega,t+1\right)\right] \end{array} \right)$$

$$k_S^i(a,y;t) = (1-\nu) + \nu \cdot \left(\begin{array}{c} f(\theta^i(a,y)) \cdot \mu \cdot \mathbb{E}_{m,y'} k_C^i\left(a'\left(a,y,m;\omega^i(a,y)\right),y';\omega^i(a,y),t+1\right) \\ +\left(1-f(\theta^i(a,y)) \cdot \mu\right)\mathbb{E}_{m,y'} k_S^i\left(a'(a,y,m),y';t+1\right) \end{array} \right)$$

$$k_X^i(a,y;t) = (1-\nu) + \nu \cdot \left(\begin{array}{c} \delta \cdot f(\theta^i(a,y)) \cdot \mu \cdot \mathbb{E}_{m,y'} k_C^i\left(a'\left(a,y,m;\omega^i(a,y)\right),y';\omega^i(a,y),t+1\right) \\ +\delta \cdot \left(1-f(\theta^i(a,y)) \cdot \mu\right) \cdot \mathbb{E}_{m,y'} k_S^i\left(a'(a,y,m),y';t+1\right) \\ +(1-\delta)\,\mathbb{E}_{m,y'} k_X^i\left(a'(a,y,m),y';t+1\right) \end{array} \right)$$

Given the solutions for $t = 7$, solve for $t = 6$ functions using the recursions and so on. The functions for $t = 0$ give the credit score conditional on an individual's type and state.