Explaining Contract Heterogeneity in the Credit Card Market

Satyajit Chatterjee and Burcu Eyigungor
Federal Reserve Bank of Philadelphia

Preliminary
March 2025

Abstract

Administrative data are used to establish patterns in contract terms, usage, and default rates of anonymized individual credit card accounts. The standard heterogeneous-agent macro model augmented with a competitive credit card industry and ex-ante heterogeneity in discount factors and default costs can explain these facts, including the fact of interest rates that are very high relative to the cost of funds. 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 MPC in the model is in the empirically relevant range because people are estimated to be impatient.

Keywords: Credit limits, defaults, financial frictions, heterogeneity, spreads, MPC

JEL Codes: D15, E44, G51

Satyajit.Chatterjee@phil.frb.org and Burcu.Eyigungor@phil.frb.org. The authors thank Zach Bethune, Larry Santucci and seminar and conference participants at New York University, Johns Hopkins University, University of Rochester, University of Pennsylvania, 2023 CAFRAL Conference, 2024 Texas Monetary Conference and 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 30 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 seek to understand the nature of differences across people that can explain the observed heterogeneity in credit access, credit use, and default risk.

We use administrative data on newly originated credit card accounts to document how (i) interest rates and credit limits (contract terms), (ii) revolving balances as a fraction of the credit limit (credit use), and (iii) the frequency of serious delinquency or bankruptcy (default rate) vary with cardholders' incomes and credit scores at the time an account is opened. The following patterns emerge:

- The credit-limit-to-income ratio rises with credit score and declines with income.
- Interest rate declines with credit score and income with the decline being mild for income and sharp for scores.
- The utilization rate declines with credit score and income, with the decline being sharper for scores than for incomes.
- The frequency of default declines with credit score and income and, again, the decline is sharper for scores than for incomes.

The first two facts tell us that individuals with higher credit scores have greater access to credit relative to their earning capacity in comparison to individuals with lower credit scores.[1] They also tell us that borrowing capacity does not keep pace with earnings capacity: higher income individuals are offered higher credit limits at lower interest rates but the ratio of credit limit to income declines. The next fact tell us that lower-income individuals and those with lower credit scores utilize their borrowing capacity more than higher-income individuals and those with higher scores. The final fact tells us that lower-income individuals encounter financial distress more frequently than higher-income individuals. There is a negative relationship between default rates and credit scores as well, which

[1] A credit score is a measure of creditworthiness. It is a negative monotonic transformation of the probability of being 90 days or more past due on some obligation within the next two years. If δ is this probability, the credit score is $a + b \cdot \ln(1 - \delta)$, where a and b are positive numbers. An individual's credit score is estimated based on an their credit history.

shows that credit scores are good at predicting the absence of a serious delinquency on a credit card.

To understand how such facts might emerge, we augment the workhorse heterogeneous-agent macro model with a credit card industry. The industry is composed of competitive card companies that offer credit lines to people. A credit line 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. Card companies post credit card contracts and individuals search for the card that is best for them. There is free entry into the card business. We assume an open endowment economy with an internationally given risk-free rate and abstract from aggregate shocks.

We investigate if heterogeneity in incomes induced by permanent income differences between people and by the realizations of idiosyncratic persistent and transitory income shocks are sufficient to account for the patterns. For this investigation we estimate model parameters to fit the score-based patterns in default frequency and credit-limit-to-income ratios only. We find that the workhorse model has some measure of explanatory power. However, what it fails to explain, even qualitatively, is the strong positive association between credit scores and credit-limit-to-income ratios: The model predicts a *negative* relationship, despite the positive relationship being a target of the estimation. This failure suggests that the workhorse model is missing some source of variation that affects credit scores independently of income.

So, what else could matter for credit scores? Lenders perceive high credit score individuals to be those capable of carrying more debt relative to their incomes without defaulting. Intuitively, this can be the case if some people are known to have higher costs of default than others and/or are known to be more patient than others (since some costs of default are borne in periods subsequent to default, a more patient person would be more deterred by these costs). This intuition motivates our main model wherein individuals differ permanently with respect to default costs and discount factors, in addition to permanent, persistent and transitory differences in earnings.

Our second finding is that when discount factor and default cost parameters are estimated to again fit the score-based patterns in default frequency and credit-limit-to-income ratios, the model gives a remarkably good account of *all* the documented patterns, including the ones that are not directly targeted. The estimation implies heterogeneity in both discount factors and default costs and this heterogeneity is positively related: A more patient person is also a person with higher default costs. We explain why the estimation

attributes heterogeneity in both discount factors and default costs and why the two are positively correlated.

Our third finding concerns the spread of credit card interest rates over the card companies' cost of funds. Absent costs, the standard one-period model of unsecured debt predicts that the spread must equal the probability of default on the card. In reality, spreads far exceed default probabilities, creating the puzzle of "excess spreads". Our model provides a resolution of this puzzle. The excess arises because of the default behavior of individuals borrowing on a credit card: Default only happens when a cardholder is close to the credit limit and gets a bad income shock. This means that debts that are paid back are, on average, less than the credit limit. Since the card company charges the same interest rate on all debts at or below the credit limit, it must charge an interest rate higher than the default rate to break even. Our estimation implies that most of the excess spread is accounted for by this simple logic.

Our fourth finding relates to the closely related matter of whether excess spreads reflect markups. In our model, once a customer matches with a card company, the costliness of search means that there is a surplus to be shared between the two parties. Competitive search leads to a division of this surplus and implies that spreads will always exceed the spread implied by zero expected profits conditional on a match. Thus markups are indeed a second reason for excess spreads. Our model predicts that poor, low-income, and impatient individuals suffer substantial markups relative to rich, high-income and less impatient individuals. But, given competition and the low cost of posting a contract, the absolute sizes of these markups are generally small.

In the rest of the paper, we use the model to shed light on two issues. First, we analyze the implications of capping credit card interest rates at 10 percent, as proposed in recent legislation. The motivation for such caps seem derive from the perception that excess spreads reflect monopoly power. But as already noted, excess spreads can easily arise in a competitive credit card market with a low cost of funds. We show that the cap leads to a significant reduction in credit limits for risky borrowers and is welfare decreasing for them.

Next, we explore the model's implications for the magnitude and dispersion of the marginal propensity to consume (MPC) across people in the economy. Although the vast majority of people in the model are not liquidity constrained, the model generates quarterly MPCs that range from 0.08, on average, for high-income and most patient types to 0.63, on average, for low-income and least patient types. These relatively high MPCs

reflect the impatience required to explain the credit card facts, in particular, the default facts. We also find that availability of credit cards depresses the MPC of the most patient low-income types, as they are the ones that engage most in buffer-stock savings in the absence of a card.

The paper is organized 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 discusses the estimation strategy to be followed to determine the remaining parameters. Section 6 presents our first substantive finding: The workhorse heterogeneous agent macro model with only income heterogeneity cannot explain the patterns we documented in section 3. Section 7 moves on to our main model and presents our second substantive finding: The model with discount factor and default cost heterogeneity can explain the patterns we document. Section 8 presents a simple model and some corroborating evidence to explain why in our model, and in the data, interest rates on credit cards can exceed the default probability on the card by a large amount. Section 9 discusses the implications of our model for the effects of a rate cap and for the distribution of marginal propensity to consume.

2 Contributions and Connections to the Literature

To the best of our knowledge, ours is the first quantitative-theoretic attempt to come to grips with the credit card facts stated in the introduction. But there is a fairly large literature on macro consumer default to which our paper is connected.

Dempsey and Ionescu (n.d.) examine contract interest rates using the same administrative data as us. They document large gaps between contract interest rates and expected default probabilities, which they call "borrowing premia." They sought to account for these premia by adapting long-duration debt models from the sovereign default literature (Hatchondo and Martinez (2009), Chatterjee and Eyigungor (2012), among others). In their accounting, both duration and intermediation costs play key roles. In contrast, we model credit cards as opposed to long-duration debt (and ignore intermediation costs). The distinction matters because the interest rate on a credit card is fixed and does not vary with the amount borrowed, whereas interest rates are responsive to debt in sovereign default style long-duration debt models. As already noted, the constancy of interest rates and the nature of repayment and default behaviors on credit card debt is key for generating excess spreads.

Our finding that discount factor and default cost heterogeneity is needed to explain credit card facts has echoes in the literature. Recently, Chatterjee, Corbae, Dempsey, and Ríos-Rull (2023) estimated the heterogeneity in privately-observed discount factors needed to account for a range of aggregate credit market statistics, including the evolution of credit scores over the life-cycle, and found substantial heterogeneity in discount factors. Our paper differs from theirs in two important ways: we model credit cards as the credit instrument as opposed one-period debt and we abstract from private information. That said, both papers suggest that heterogeneity in discount factors are a partial explanation for the observed heterogeneity in credit scores.[2] Athreya, Sanchez, and Mustre-del-Rio (2019) showed that financial distress is persistent and explain this by differences in impatience across people, with more impatient people encountering financial distress more often. In earlier studies, Athreya, Tam, and Young (2009) used differences in survival probabilities by educational attainment to generate differences in effective discount factors across members of these groups. Athreya, Tam, and Young (2012) allow differences in default costs across education groups and find that a positive association between default costs and educational attainment is needed to explain their credit market facts. Outside of the consumer credit literature, there are studies that have emphasized heterogeneity in discount factors to explain wealth inequality and relatively high marginal propensities to consume (Krusell and Smith (1998) and Carroll, Slacalek, Tokuoka, and White (2017), respectively).

Several recent studies, Bethune, Saldain, and Young (2024), Galenianos, Law, and Nosal (2023), and Raveendranathan, Stefanidis, and Sublet (2023), describe environments in which rate caps are beneficial. The main distinction between these studies and ours is they assume more monopoly power than we do. The first two studies seek to account for the very high interest rates paid by highest-risk borrowers and do so by assuming that some of these borrowers face monopoly lenders.[3] Raveendranathan, Stefanidis, and Sublet (2023) has a broader focus but differs 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) and in our calibration search-induced markups are quite low. As noted already, our model matches the observed interest rates quite well.

[2]It is worth noting that Meier and Sprenger (2012) showed, in an experimental setting, that participants who acted more patiently in lab experiments tended to have higher credit scores.

[3]In a different line of attack, Saldain (2023) focuses on high-risk borrowers but does not invoke monopoly power to explain high interest but assumes that some people suffer from temptation preferences. Rate caps are potentially beneficial in this situation.

There are some recent studies that draw out the MPC implications of consumer credit as we do. Taking as given 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 debt and bankruptcy. They have two main findings: First, the best fit is found for a model in which there are two types who differ in their discount factors and, second, this two-types model generates average MPCs in the empirically estimated range.[4] Lee and Maxted (2023) study a model where consumers can borrow via contracts that resemble credit cards. They abstract from default and preference heterogeneity but focus on the implications of present-bias induced by quasi-hyperbolic or β-δ discounting. Present bias leads many households to borrow small amounts on their cards. Since such borrowing is at high interest rates, implied MPCs are high.[5]

The main difference between these two studies and ours is that we do not take interest rates and credit limits as given and, instead, estimate the underlying heterogeneity to fit the distribution of default frequencies and credit limits, and by extension, interest rates for newly originated credit card accounts. We estimate relatively low discount factors and correlated heterogeneity in default costs. Our estimated discount factors are low because that is what is required to match the distribution of default frequencies. Thus we too arrive at MPCs in the empirically relevant range but for different reasons than in these two studies.[6]

Since our model features both credit lines and search, it relates to a number of recent studies that feature one or both. Raveendranathan (2020) presented a credit card model of unsecured consumer borrowing, wherein people acquire cards via search, with contract terms — the credit limit and the interest rate — determined upon contact.[7] His model is calibrated with life-cycle aggregate credit statistics. A version of this model is used in Raveendranathan and Stephanidis (2022) to study the fall in consumer credit during the

[4]The context is an influential study by Kaplan and Violante (2014) who noted that average MPC out of tax cuts were much higher, about 25 percent, than what would be implied by classical permanent income hypothesis. 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 (2021) and Gelman (2021) present evidence that the behavior of wealthy households with low liquid assets reflects differences in discount factors and/or interest rates.

[5]Drozd and Kowalik (2023) postulate quasi-hyperbolic preferences in order to explain promotional interest rates on credit cards.

[6]Koşar, Melcangi, Pilossoph, and Wiczer (2024) draw out a different MPC implication of high-cost consumer borrowing: Indebted household's MPC is lowered because they use a portion of a marginal increase in wealth to pay off consumer debt. Since interest rates rise with debt, this effect increases with the household's indebtedness. This implies that stimulus payments can attenuate the short-run fiscal multipliers but enhance long-term consumption spending levels.

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

Great Recession. The Raveendranathan, Stefanidis, and Sublet (2023) study also uses this framework. 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.[8] Herkenhoff and Raveendranathan (2021) 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.[9]

Finally, heterogeneity in credit card lending terms has been studied from different perspectives in prior work. Narajabad (2012, Table 1, p. 530) documented the increase in dispersion in credit card interest rates and limits reported by Survey of Consumer Finances respondents from 1989 to 2007 and used this fact to offer an explanation for the rise in credit card debt and bankruptcy filing rates over the same period. This topic is pursued with greater theoretical sophistication in Livshits, MacGee, and Tertilt (2016) and greater quantitative sophistication in Athreya, Tam, and Young (2012) and Sanchez (2018).[10] These studies focused on improvements in information technology as the underlying factor. Herkenhoff and Raveendranathan (2021) show that rising credit limits and defaults also accompany the removal of entry barriers.

3 Credit Card Terms, Usage, and Default Rates

At a broad level, 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 the ways in which customers must differ in order 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 reflect the characteristics of customers most closely when a card is first issued.

We use confidential supervisory data collected by the Federal Reserve System in pursuance of the annual comprehensive capital analysis and review (CCAR) of large U.S. bank holding companies, as required by the 2010 Dodd-Frank Wall Street Reform and Consumer Protection Act. One part of these data, called Y-14M, track, at a monthly fre-

[8] Athreya (2003) is an early quantitative-theoretic study of the interaction between unemployment insurance and the insurance afforded by the bankruptcy option.

[9] Ausubel (1991) drew attention to high and sticky credit card interest rates during the 80s 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 has since eroded.

[10] More recently, a different aspect of the dispersion of credit terms has gathered attention. It has been observed that credit card interest rates offered to observably similar borrowers can be quite dispersed (Stango (2002), Stango and Zinman (2016)). Galenianos and Gavazza (2022) show that this sort of dispersion can occur in a search model if borrowers face attention cost.

quency, 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, which start in June 2012, cover 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 July 2013 and June 2014 with one primary account holder.[11] Following this initial selection, we trimmed the sample for outliers.[12] We also 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 align the distribution of income with the 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.[13]

The contract terms we focus on are credit limits and interest rates actually paid on revolving balances. The measurement of interest rates is complicated by the fact that most new originations come with a promotional period, usually a year, when the contractual interest rate on the card is low, usually zero. Furthermore, it is also the case that new originations often start off with promotional balances. Very likely, these starting balances are balances transferred from another card with a higher, non-promotional, interest rate. Our data distinguishes between promotional and nonpromotional revolving balances which allows us to construct a balance-weighted monthly effective interest rate at the account level. One complication is that balance transfers come with a one-time 3 percent balance transfer fee which means that the account holder pays interest in advance for the transfer amount. To take this into account, we assume that promotional interest rate is 3 percent, rather than zero. We take the average interest rate paid over a 60-month period following origination as the effective contractual interest on the card.

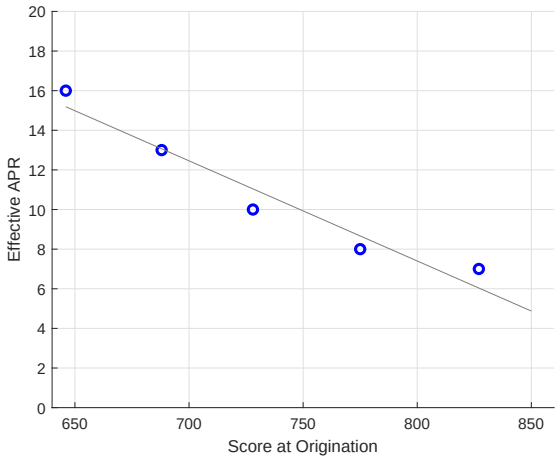
Credit limits also change over time and the vast majority of the changes are limit increases. However, each month only about 2 percent of accounts see a limit increase, so

[11] We chose our sample period with the following consideration in mind. The majority of 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 mid 2012, but not all variables are reported/available for the first few years.

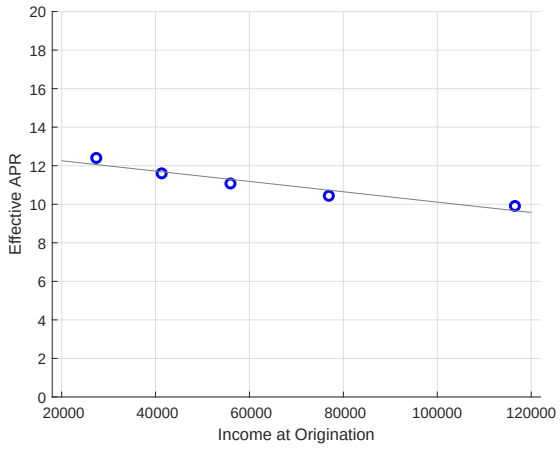
[12] We excluded accounts for which reported borrower annual income was less than $1000 (most likely students) and credit-score-at-origination was less than 100 (most likely a reporting error); we also excluded accounts for which credit-limit-at-origination was less than $100.

[13] The average scores of those excluded is 593, which is well inside the score band of $580 - 619$ commonly designated as the "subprime" category of borrowers. The terms and performance of these most risky accounts seem outside the explanatory scope of our model.

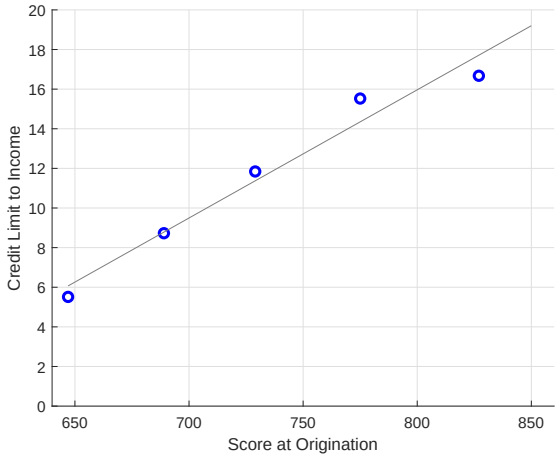
Figure 1
Variation in Contract Terms by Score & Income



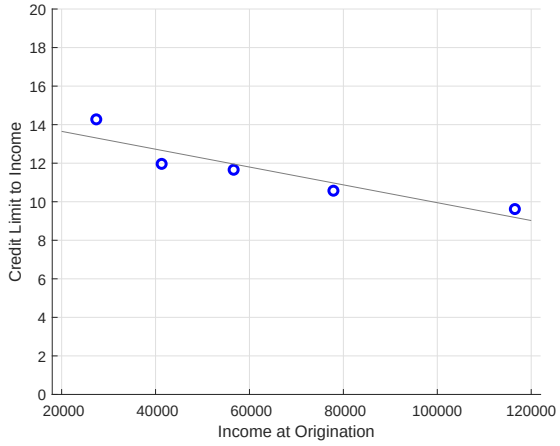
APR by Score



APR by Income



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



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

limits generally have a long duration. To avoid any potential limit increases resulting from 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 24 months following origination. These timing assumptions are also followed 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). In each figure, a dot corresponds to the mean value for a quintile of our trimmed and truncated sample. 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 individuals 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 9 percentage points lower than the mean interest rate in the bottom quintile, while the corresponding difference for income quintiles is only about 2 percentage points.

There is pronounced variation in credit limits scaled by (annual) income at origination. The variations go in opposite directions for scores and income. Individuals with credit scores in the top quintile have a ratio that is, on average, about 3 times the ratio of individuals with scores in the bottom quintile. With regard to income, though, individuals with incomes in the bottom quintile have a ratio that is, on average, about 1.5 times the ratio of individuals with incomes in the top quintile. This pattern arises because credit limits increase less than proportionately with income.

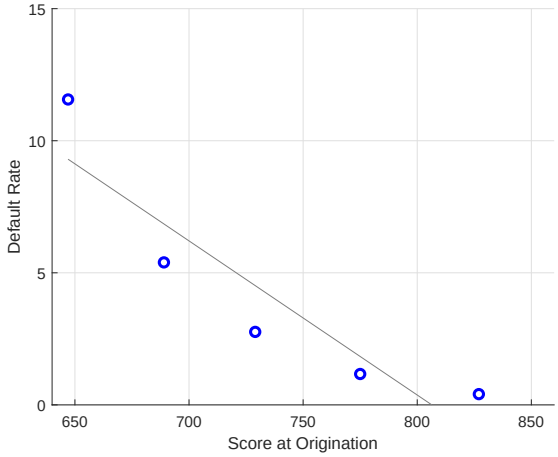
Figure 2 displays the patterns with respect to usage and default. We define an account to be in default if one of the following has occurred within 24 months of origination: (i) The account is reported to be in a debt waiver or in bankruptcy, or (ii) the account is reported 120 days or more past due and the delinquency is not cured in the following 6 months.[14] The top panels display the variation in frequency of default. Default frequency is strongly related to credit score. The default frequency in the bottom quintile of scores is about 11.56 percent while it is 0.41 percent for the top quintile. Default frequencies also decline with income but less so. The default frequency in the bottom quintile of incomes at origination is 6.71 percent, while for the top quintile it is 1.95 percent.

The bottom panels display the ratio of revolving balances to credit limits. Individuals in the bottom quintile of scores have a utilization rate of 53.8 percent while the top quintile has a utilization rate of only 3.4 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 32.1 percent, while those at the top income quintile have a mean utilization rate of 18.2 percent.

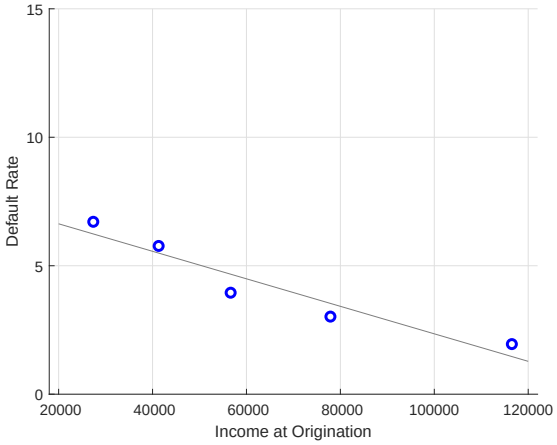
The patterns displayed in Figures 1 and 2 are the facts we focus on in this paper.

[14] A delinquency is cured if the account is not classified as delinquent and does not transition into one of the categories mentioned under (i) in the next 6 months.

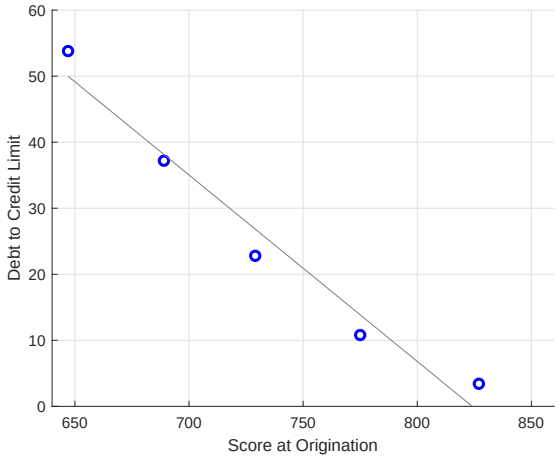
Figure 2
Variation in Default Frequency and Utilization by Score & Income



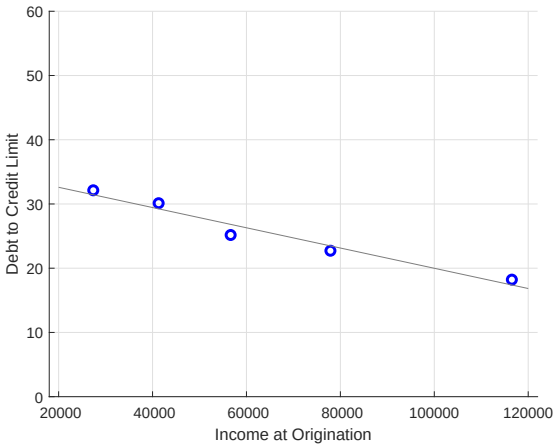
Default Frequency by Score



Default Frequency by Income



Utilization Rate by Score



Utilization Rate by Income

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 $i \in 1, 2, \ldots, I$. 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). But since all decision rules can be shown to scale with P, we lighten the notation and develop 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 with regard to 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.[15] All individuals can save at the common risk-free (gross) interest rate R_f.

A person with a balance on her credit card has the option to default. If a type i person defaults, she (i) loses her card and a fraction $0 < [1 - \phi^i] \leq 1$ of her transitory income component in excess of $-\lambda y$ is taken by the card company in the period of default, (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.

[15] This is a simplification, as credit card companies do change the interest rates and credit limits of customers over time. As of 2009, these contractual changes must comply with restrictions imposed by the CARD Act.

Figure 3
Default Cost

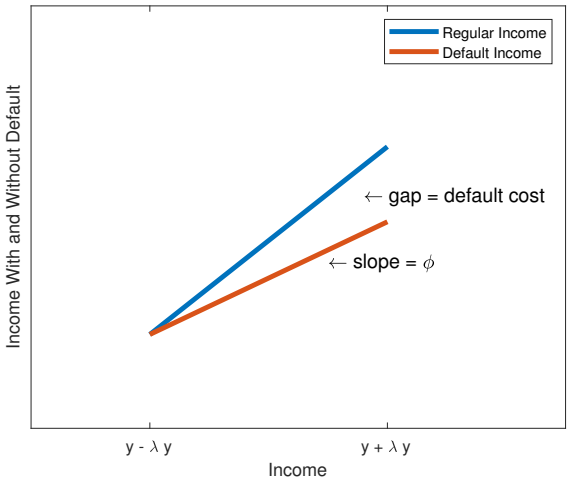


Figure 3 displays the income loss suffered by a defaulter as a function of the transitory component of her income in the period of default. Since the loss is $(1-\phi^i)(m-(-\lambda y))$ $=(1-\phi^i)(m+\lambda y)$, it is zero if $m=-\lambda y$ and increases linearly with m to a maximum of $(1-\phi^i)(2\lambda y)$. 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.[16]

The income loss from default, which for expository simplicity we refer to as the default cost, deserves comment. The substantive assumptions are that the cost is increasing in both components of income and potentially different across types. Since the default cost is the payment creditors extract from delinquents, it is plausible that this payment is rising in 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 higher income people.

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. An individual who can search searches for contracts being offered to people of her type i, her asset position a, and her persistent earnings y. Within each such $\{i, a, y\}$ market there are submarkets offering different contract terms. Card companies

[16] 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, it is known that utility shocks distort saving/borrowing behavior because people are tempted to borrow small amounts simply for the chance of enjoying positive utility shocks from default. The distortion is not too troubling if the variance of the utility shock is low but in the current context the variance will have to be large to explain the high default rates of low-income and low-credit score people. The income loss approach is cleaner in this respect.

commit to honoring the terms of posted contracts, if accepted. These assumptions make our search environment one of directed search.

To generate active search, we assume that card holders get separated from their card with probability ξ.[17] 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 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

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

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 card holder'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 $v^i_{DEF}(y, m)$ denote the value of default. Then,

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

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

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

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,

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

The inner max chooses over ω for the best contract. The outer max recognizes that the individual has the option to 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 individual states. Once a contract is accepted, the companies simply follow through on the terms until the individual separates from the card or defaults on it. We assume that companies' 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,

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

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 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 card holder 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 lose money. 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.[18]

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,

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

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.

[18] 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.

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 contact 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 is increasing in market tightness: In a tight market one can get a credit card quickly.

4.5 Equilibrium

Since there is no interaction between types, 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

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

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

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

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

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 nonpositive expected net profits.

Equation (16) is the complementary slackness condition that ensures consistency with free entry of lenders in any submarket: if 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 contracts actually offered in equilibrium 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.[19]

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

[19]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 who 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 $E^{i*}(a, y; \omega^*) = B^{i*}(a, y; \omega^*) \cdot \theta^{i*}(a, y; \omega^*)$.

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. □

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 $\{i, a_k^i, y\}$, where a_k^i is the closest asset level less 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 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 turns out to be 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 with respect to 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 the pick the solution with the highest profits.

5 Quantitative Analysis: Preliminaries

5.1 Earnings Process

We calibrate the earnings process based on estimates provided 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.[20] We discretize y into 5 levels and assume that it follows a first-order Markov process. We pick λ and the values of the 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 independently. The (real) risk-free (gross) interest rate R_f was set at 1.02 percent annualized and the probability of survival, ν, at 0.82 percent annualized. The CRRA parameter γ was set at 2, which is conventional in macroeconomics.

The probability of search following a separation, μ, is set at 0.50, which means that people without cards cannot look for one for about half a year, on average. The delay in

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

getting a new card lowers the outside option of card holders and increases the profits of card companies. The cost of posting a contract, τ, is set to a low value of 0.001. These choices for μ and τ ensure that a profitable credit card exists for 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, on average, the effective exclusion from the credit card market following a default is about 6.5 years.

The separation parameter, ξ was set to 0.022. We pin this down from the observed frequency, between June 2014 and May 2019, of accounts in good standing that disappear from our sample.

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

Finally, 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.

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	0.998
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.

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:

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

Here, CL-to-INC$_{\text{score}_j}$ and 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 that are 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, of her credit limit-to-income ratio. And, 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 take them into account in the estimation. The moments that are not targeted will be used to evaluate the performance of the model.

The estimation needs to recognize that our model is about an individual but our data set follows credit card *accounts* and it is not possible to merge accounts belonging to the same individual. We scale up the credit-limit-to-income ratios shown in Figure 1 by a factor of 2.2 to take into account the average number of cards that individuals have.[21] With regard to the interest rate and utilization rates of card accounts, we proceed on the assumption 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 of 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.[22]

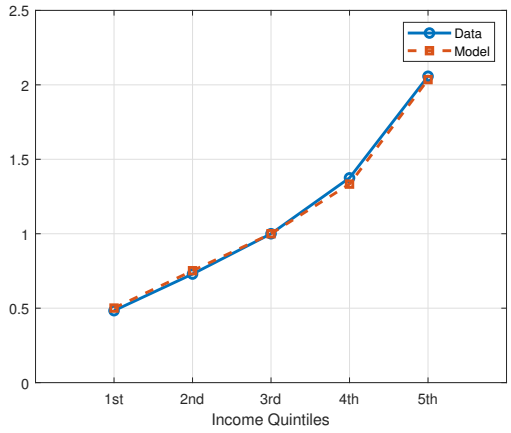
Finally, we set $w_1 = 1$ and $w_2 = 1000$, thus giving much more weight to the deviations in default frequencies. We do this because we want the model to generate the observed default frequencies for the score bins so that we can usefully compare how utilization rates, credit-to-income ratios and interest rates respond to default risk in the model versus the data.

6 Can the Workhorse Model Explain These Patterns?

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

[22] Note that in our setup decision rules do not depend on the credit score, *per se*. This is a consequence of our simplifying assumption that a person's type is observable to creditors. If a person's type is hidden information, then, as shown in Chatterjee, Corbae, Dempsey, and Ríos-Rull (2023), credit scores will matter for decision rules.

Figure 4



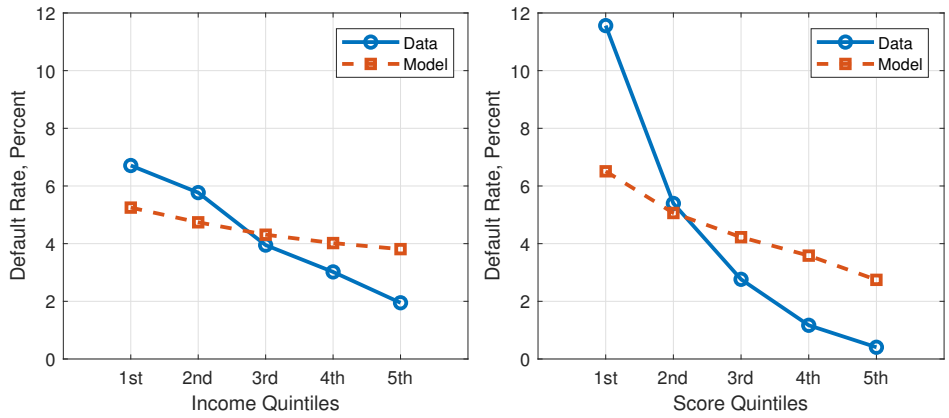
The workhorse version has a single type, i.e., $I = 1$, and all individuals have a common β and ϕ. To start off, Figure 4 plots average incomes in each quintile relative to the average in the middle quintile for both the model and the data. There is a very close fit. A close fit is not guaranteed, as incomes in the data are incomes recorded on new origination, while incomes in the model are largely 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 looking for a new card following a default and re-entry, and, among indebted people, the ones getting a new card are those for whom the balance transfer is profitable for the card company. That said, the concordance is helped by the fact that our data sample considers only individuals in the 7th to 93rd percentiles of incomes-at-origination.

Default Frequency Patterns

The default frequencies for both income and score quintiles are compared to the data in Figure 5. The model can explain the fact that default frequencies decline with incomes and the decline is sharper for scores as compared to incomes. But the quantitative fit is clearly imperfect in both panels, even though the estimation seeks to target the pattern with respect to score quintiles very closely.

The fact that default frequencies decline with incomes has a simple explanation. The persistent incomes of people in the top income bin is, on average, higher than the persistent incomes of people in the lower bins. Individuals with high persistent incomes, i.e., with high y, expect their y to decline over time. This induces them to accumulate assets or deleverage if they have debt. Consequently, they are less likely, as a group, to get into debt levels that trigger default over a 2-year horizon.

Figure 5
Default Frequencies



Why is the decline in default frequencies sharper for scores than for incomes? One reason is that when people are ranked by their credit scores, the top bin collects individuals with high y *and* high a. As a group, these are people who are even less likely to end up with defaultable debt levels within a 2-year horizon than people who just have high y. For instance, the group of high y individuals will include those who attained this status recently and, therefore, have not had the time to accumulate a high level of assets.

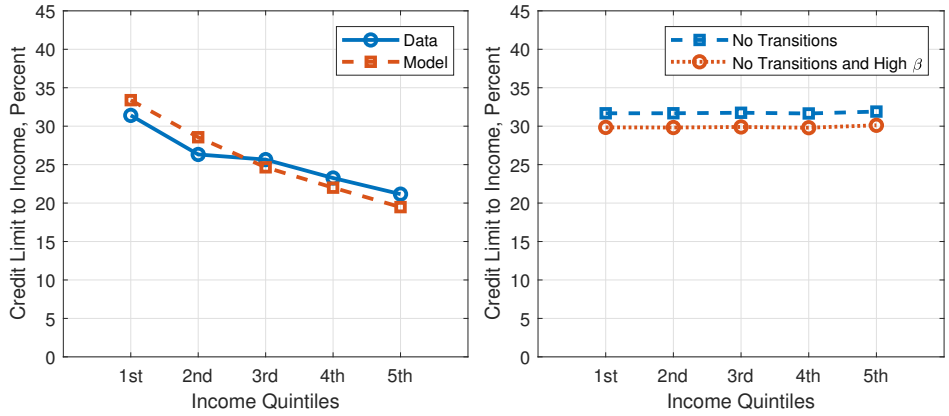
Another reason for the decline in default frequency to be shallower for incomes than for scores is that the top income bin includes people whose y is low but whose permanent component of income, P, is high. The mixing of high and low y people even in the top income bin makes the drop in default frequency shallower for incomes. Note that such mixing of high and low y individuals is not possible for score bins because the scaling property ensures that assets or debts scale with P and default probability depends on only y and debts *relative* to P.

Nevertheless, the model severely overpredicts the default frequency of people in the top score quintiles and severely underpredicts it for the bottom quintile. Given a single dimension of heterogeneity (incomes), the model attempts to get the pattern right by making people very impatient — the estimated value of β is 0.42, annualized. Impatience helps to get the default frequency in the bottom quintile closer to the observed value but at the expense of overpredicting default frequencies for a majority of the sample.

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. The credit limit to income ratio declines with income in the model for two reasons. The most important reason is the nature of the card contract,

Figure 6



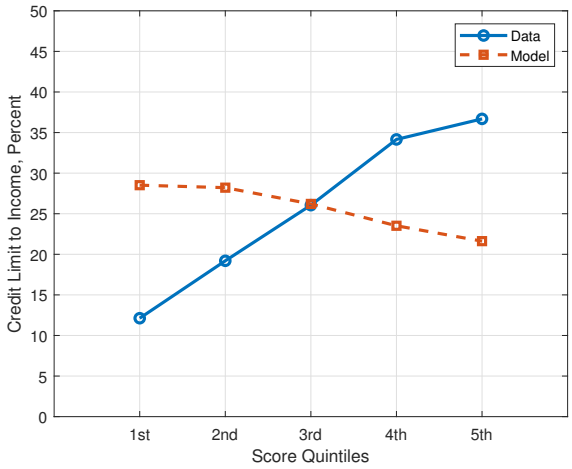
which commits card companies to a credit limit and an interest rate, regardless of the individual's future income states.

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 actually permanent. The blue line in the right panel of Figure 6 displays the credit-limit-to-income ratio pattern for the "no transitions" equilibrium (all other parameter values are held constant). The equilibrium credit-limit-to-income ratio is roughly constant across income bins.[23] But with positive probability of transitions, a (roughly) uniform 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, which makes them want to apply for cards with lower credit limits and lower interest rates. The situation here is similar to that in insurance: the good risks prefer contracts that have 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.

[23] 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 is not exactly true since the contract posting cost τ does not scale with y. But since $\tau \approx 0$, it is almost true. Consequently, the equilibrium of the no-transitions model features credit limits that approximately scale with y.

Figure 7



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 is increasing with credit scores while in the model it is *decreasing*. This mismatch arises even though the relationship between credit-limit-to-income ratios and scores is a target of the estimation. Evidently, the estimation is unable to find any parameter values that can get the model to 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. But it is the high-y individuals who have a low probability of default and, therefore, high credit scores. The combination of these two facts means that credit scores and credit-limit-to-income ratios *have* to be inversely related.

The most plausible explanation for this failure is that the model is missing non-income factors that affect credit scores and that are not strongly correlated with income. This inference is also supported by the fact that in the data, the average income in the highest score quintile is 1.35 times the average income in the bottom quintile while in the model this ratio is 1.54.[24]

To summarize: With income as the only source of heterogeneity, the model can qualitatively match the relationships between income-at-origination and the frequency of default and between income-at-origination and the credit-limit-to-income ratios. The model can also match, qualitatively, the observed relationship between default frequency and credit

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

scores. But it fails to match the positive association between credit-limit-to-income ratios and credit scores.

7 Expanding to Multiple Types

Following some experimentation, 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.

7.1 Parameter Estimates and Model Fit

The estimated parameters are reported in Table 3. The estimation calls for a very 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 least impatient type (type 3) has a 3 percent quarterly discount rate. Default (garnishment) costs are estimated to be positively related to discount factors. The most impatient type forfeits 20 percent of their transitory income to creditors upon default whereas the least impatient type forfeits 75 percent.

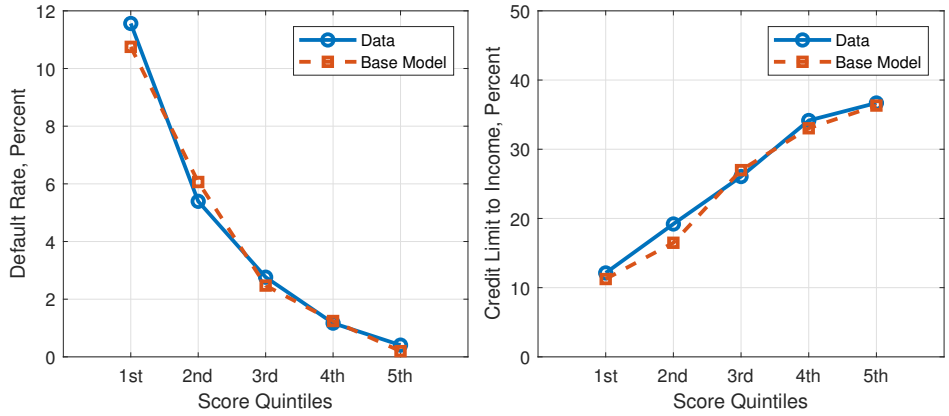
Table 3
Estimate of Type Parameters

Type	β	$(1-\phi)$
1	0.63	0.20
2	0.86	0.29
3	0.97	0.75

Figure 8 displays the targeted data moments and their model counterparts. For comparison purposes, the panels also show the estimated patterns for the workhorse model, labeled IO (for "income only"). 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 workhorse model.

To understand why the Base model can hit the default frequency targets, consider discount factor heterogeneity. This type of heterogeneity induces heterogeneity in default frequency: More patient people are less likely to borrow and default. When the higher discount factor individuals search for and acquire new cards, they show up in the upper

Figure 8



score quintiles and reduce the default frequencies in those bins relative to the lower bins. The estimation amplifies such differences by attributing pronounced differences in patience across types — especially, between the least and most patient types — and thereby gets the *sharp* drop-off in default frequencies across score bins. To verify this logic, 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 least impatient type (Type 3) is mostly concentrated in the top two quintiles. Middle patient types (Type 2) are concentrated in the middle three quintiles.[25]

Table 4
Type Distribution of New Originations

	I	II	III	IV	V	Fraction
Types	Score Quintiles					
1	0.70	0.30	0.00	0.00	0.00	0.29
2	0.00	0.31	0.42	0.26	0.00	0.36
3	0.00	0.00	0.14	0.30	0.56	0.35

To understand why the Base model can hit credit-limit-to-income ratio targets it is helpful, again, to consider the role of discount factor heterogeneity. As previously discussed, more patient people desire (and get) cards with lower limits and lower interest rates. Since the more patient types will populate the upper score bins, discount factor heterogeneity *strengthens* the negative relationship between credit limit ratios and scores induced by heterogeneity in y. Thus, in this instance discount factor heterogeneity makes the fit worse, not better. But the estimation can overcome this problem by assigning

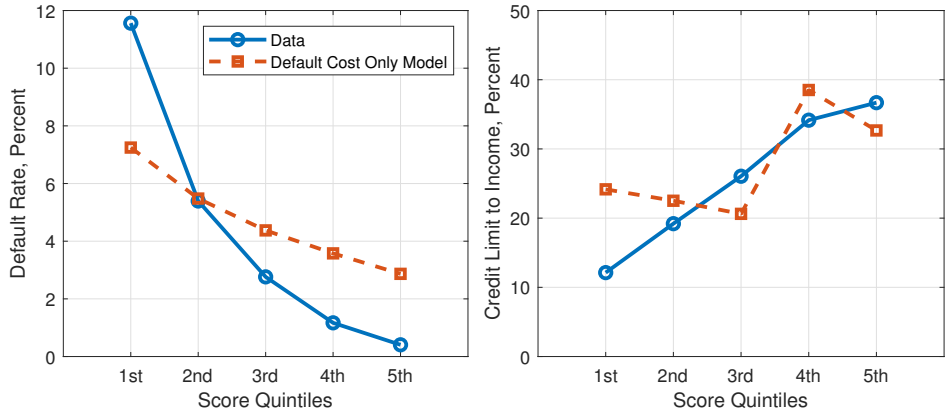
[25]The table also reports the prevalence of the three types among the people who are getting new cards. While the most impatient type constitutes 25 percent of the economy, it represents 29 percent of individuals getting new cards. This is because of selection: Among the group that is searching and obtaining new cards are individuals who defaulted in the past and the most impatient type is over-represented in that set.

higher default costs to the more patient individuals. Doing so raises the credit-limit-to-income ratio among the more patient types and further lowers their default propensities. The positive association between default costs and patience allows for an improved fit with respect to credit-limit-to-income ratios without interfering with the fit with respect to default frequencies.

Another way of viewing the estimation results is that the pattern in default frequencies (across score bins) identifies the variation in β and the pattern in credit-limit-to-income ratio identifies the variation in ϕ. This is completely in accord with the intuition that led us to consider variations in discount factors and default costs.

The fact that variation in default costs can raise credit limit ratios and lower default probabilities across score bins raises the question whether variation in default costs alone can fit the patterns reasonably well. Figure 9 displays the fit for default-cost-only model (labeled DCO). The model cannot get the sharp drop in default frequencies across score bins — which we might have expected — but it fails also to get the monotonic relationship between credit- limit-to-income ratios and scores.

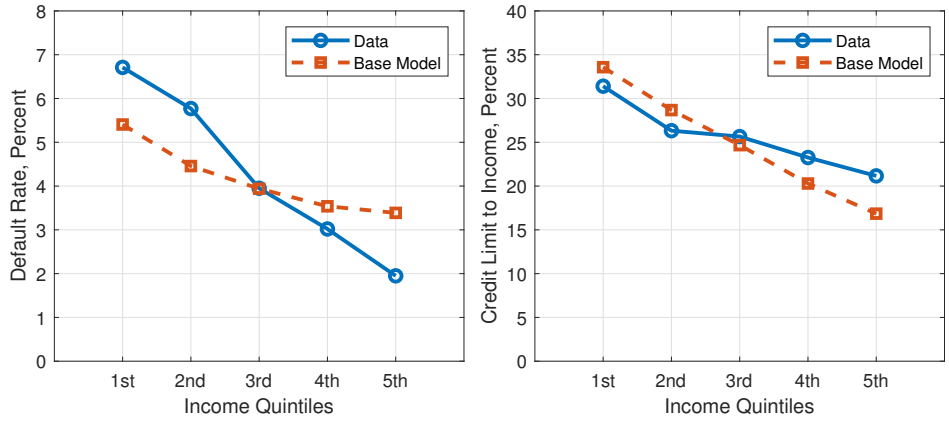
Figure 9



7.2 Non-targeted Moments

Figure 10 shows the patterns with respect to income quintiles. The panels show that the Base model can continue to account for the patterns in a qualitative sense but the quantitative fit is about the same as the workhorse model. The drop-off in default frequencies with respect to income is not as sharp as in the data and the drop-off in credit limit ratios with respect to income is too sharp relative to the data. These facts reiterate our earlier

Figure 10



statement that income-based patterns in default frequency and credit-limit-income ratios are insensitive to variations in discount factor and default costs.[26]

Figure 11 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 (have low β) and borrow at high interest rates. In contrast, the workhorse model substantially over-predicts utilization rates in both panels. As explained previously, the estimated discount factor in the workhorse model is very low and the over-prediction is a direct consequence of that.

Figure 12 displays interest rate on card 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-predicts interest rates by a modest amount, but since the model abstracts from the flow costs of running a credit card company, some under-prediction is to be expected.

The fact that the model is getting the level of interest rates about right means that it is successfully predicting very large "excess spreads". This is illustrated in Figure 13. The left panel plots the average spread, defined as the difference between contract interest rates and the 2 percent risk-free rate (which is the opportunity cost of funds in the model), and annualized default frequency by income quintiles. In the standard one-period model

[26]The quantitative miss of the Base model in Figure 10 could be alleviated if $\imath$ 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 and have opposite effects on the lower income bins. Furthermore, the improvement in fit won't interfere with the model's success with regard to the score-based patterns or the patterns with respect to utilization rates because default frequencies and utilization rates are independent of P, by virtue of the scaling properties.

Figure 11
Utilization Rates, Model and Data

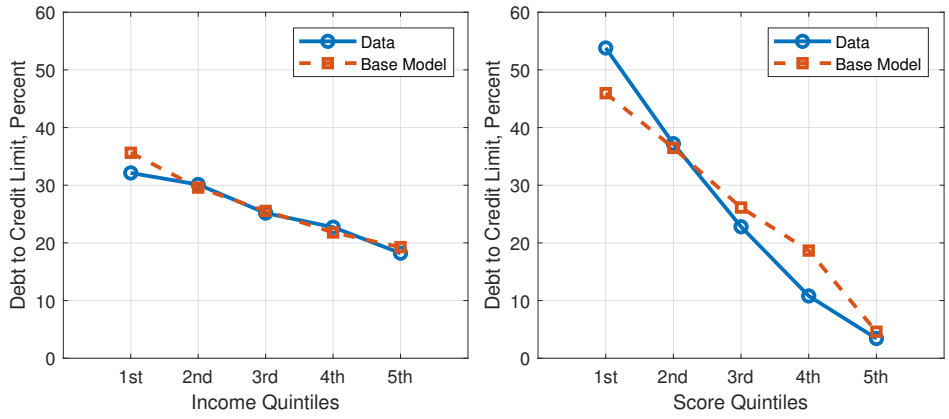
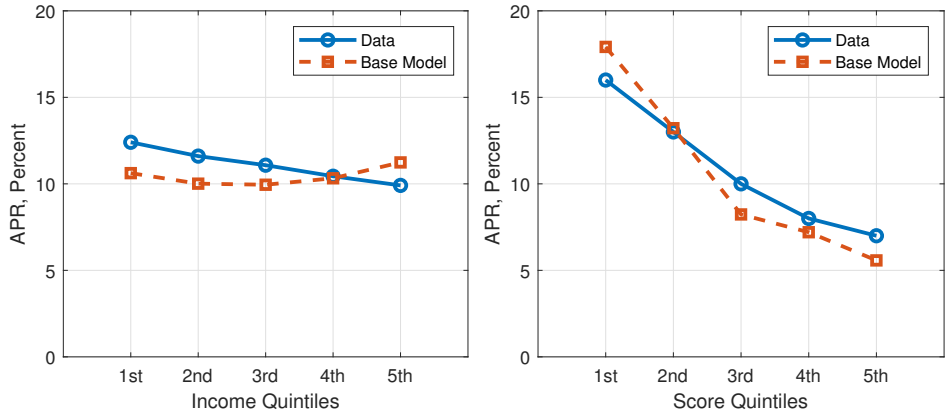
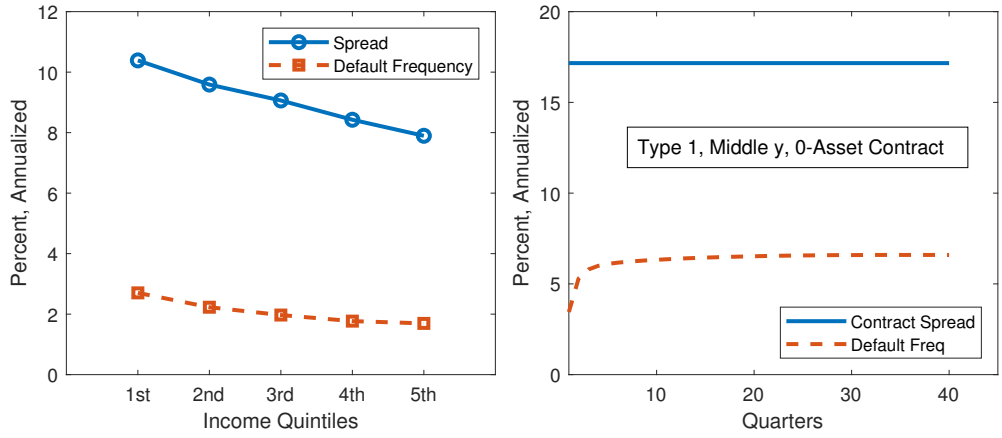


Figure 12
Interest Rates, Model and Data



of unsecured consumer debt, these two lines should lie on top of each other. Here, there is a gap of about 7 to 8 percentage points.

Figure 13
Spreads and Default Probabilities



Since income bins contain a mix of contracts offered to different types and to people with different peristent y levels, for clarity the left panel zeros in on a specific contract. It is the contract offered to a Type 1 individual with the middle level of y and no assets. The blue horizontal line at roughly 17 percent is 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.[27] The spread exceeds default frequency at any time by more than 10 percentage points.

Understanding how the model can generate such large excess spreads and be consistent with free entry in the card business will be the main focus of the following section of the paper. We end this section by describing the range of equilibrium contracts offered in this model.

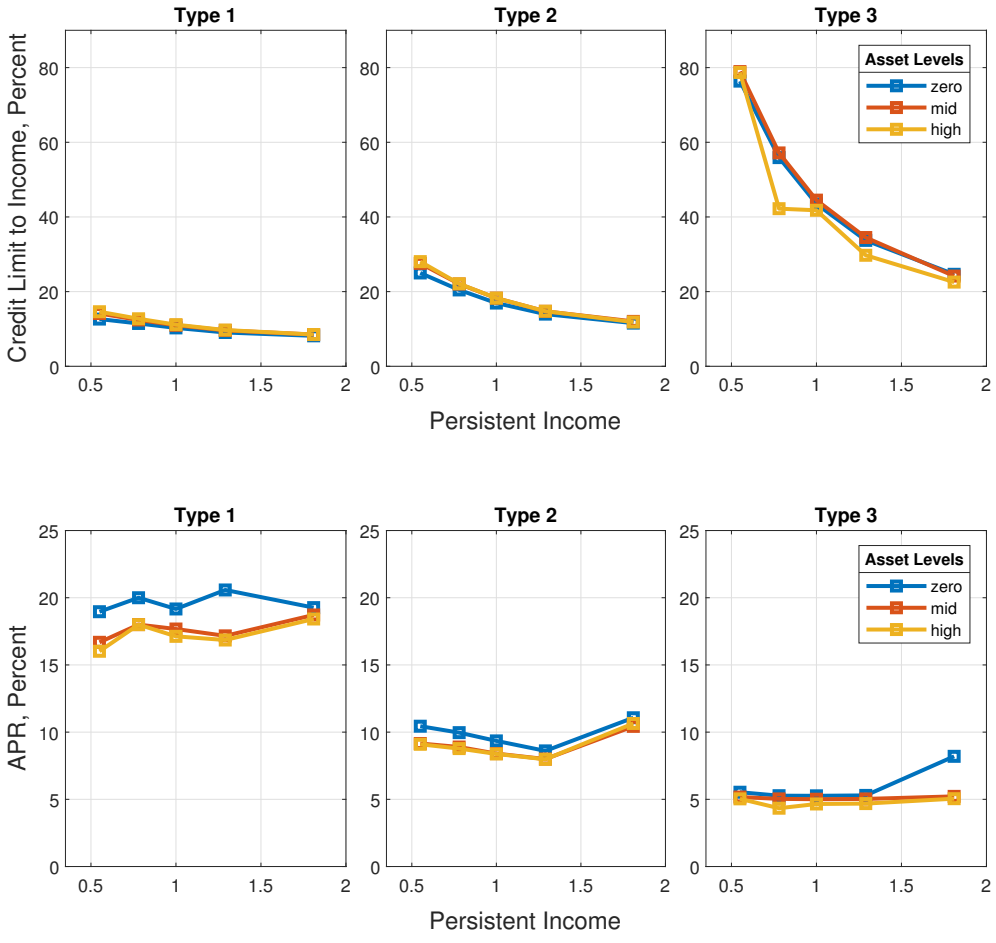
7.3 Equilibrium Contracts

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 14 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. For any given persistent income level y, Type 3 is offered higher

[27] 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 14
Contract Terms by Type, Income & Assets



limits than Type 2 and Type 2 is offered higher limits than Type 1. This ordering reflects the fact that default costs rise with i. Second, for all types, the credit-limit-to-income ratio is decreasing in persistent income y. This reflects the logic discussed previously: credit limits do not keep pace with earnings because card companies protect themselves from the possibility of a decline in y. Third, for each type, the contract credit limits are quite insensitive to asset levels. This is a consequence of the long-duration nature of a credit card contract and the fact that 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 assets at origination. This is also the reason why it is possible to have a sparse asset grid for contracts without affecting the results.

Turning to the bottom panel we see that, again, there is marked variation in contract interest rates across types: Type 3 is offered better rates than Type 2 and Type 2 is offered better rates than Type 1. The ordering reflects the ordering of default frequencies across types (displayed later in the paper). Contract interest rates are mildly sensitive to assets at origination, with zero-asset individuals offered somewhat higher rates. This lack of sensitivity reflects, again, the long-duration nature of card contracts. Contract interest rates show a tendency to increase with y but the overall variation is mild. The tendency for contract interest rates to rise with y is the reason we see a similar pattern in the left panel of Figure 12. This pattern is puzzling because default rates do tend to decline with y. It turns out that the reason for this pattern is intimately connected to the reason for excess spreads in the model. It is to this we now turn.

8 The Logic of Credit Card Interest Rates

We now turn examining how the base model accounts for credit card interest rates. There are two key parts to this. First is the default behavior of individuals who borrow on credit card and second is the crucial fact that credit card interest rates depend on the borrower's income and type but not on how much the individual borrows (subject to the limit) 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, i.e., defaulting (the cost of default is a proportion $(1-\phi^i)$ of the *excess* of m over $-\lambda y$ and gets closer to zero as m gets closer to its lower bound). 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 15
Default Probability of a Zero Asset, Type 1 Individual

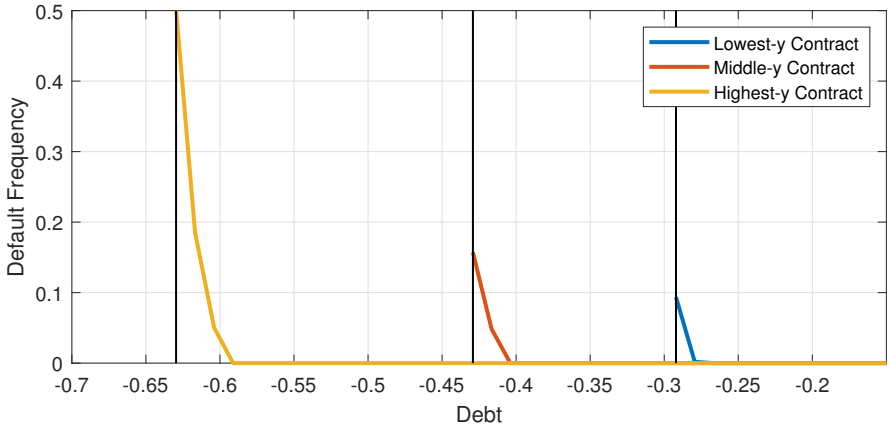


Figure 15 plots the default frequency against the beginning of period debt of a Type-1 individual for three different contracts she might hold in equilibrium: zero-asset contracts for the low, middle or highest y. The solid vertical lines are drawn at the credit limit of each contract. For each contract, default frequency is zero until the debt level gets close to the credit limit and then it becomes positive. 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 are in line 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.

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

Contract Spreads:

To understand the determination of spreads, consider a stylized two-period example. 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 idiosyn-

Table 5
Default Share By Persistent Income and Type

Income	Type 1	Type 2	Type 3	Total
0.55	0.17	0.12	0.05	0.34
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	0.14
1.81	0.08	0.03	0.00	0.10
Total	0.61	0.30	0.08	1.00

cratic shock which can be either good or bad with probability $(1-\Delta)$ and Δ, respectively. Among cardholders, a fraction $(1-\xi)$ never borrow 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 cardholders borrows z if she has the good shock and borrows the limit L if she has 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 but defaults if she had borrowed L.

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:

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

The card company's second-period expected profits is

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

Rearranging gives

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

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, $\xi \cdot \Delta^2$, and expected profits on the card scaled, again, by the maximum that can be owed on the card.

Equation (17) gives us some insight into how excess spreads can arise on credit cards. Assume for the moment that competition forces $\mathbb{E}\Pi$ to be zero. Since the ratio multiplying the contract spread is clearly less than 1, the zero-profit contract spread will *exceed* $\xi \cdot \Delta^2$, the default probability on the contract. In other words, even a competitive credit card contract will feature a *positive* excess spread. 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.

Does equation (17) help us understand the excess spreads we see in our model and in the data? To explore this, note that the term multiplying the contract spread can be interpreted as the *expected utilization rate* on the card, i.e., the average level of revolving balances carried on the card as a ratio of 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 expected profit rate term, the equation (18) suggests that excess spreads scaled by the default probability should be explained by the inverse of the utilization rate on the card. This is a prediction of this model that we can take to our model-generated data.

The left panel of Figure 16 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 toward the bottom of the both panels is the prediction of the one-period debt model for which this ratio is always 1. The model output lines up remarkably well with the predicted relationship, especially for income bins. Thus, the model-generated data accords with the insight afforded by the simple model. We can, of course, ask the same question regarding our sample data. Figure 17 shows the corresponding plots. The inverse utilization rate has substantial explanatory power.

Figure 16
Spreads, Default Probabilities and Utilization Rates, Model

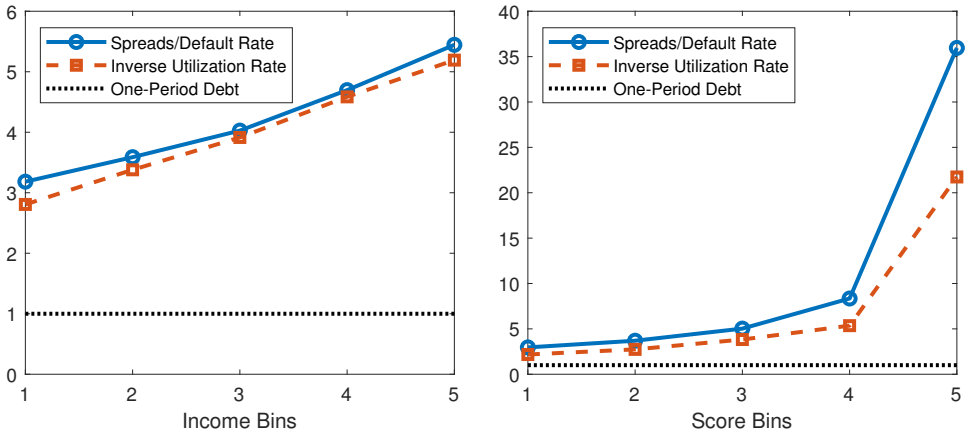
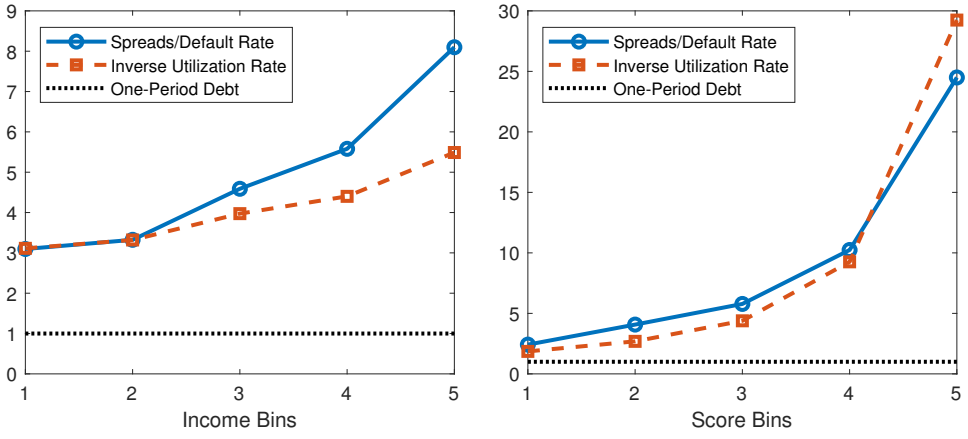


Figure 17
Spreads, Default Probabilities and Utilization Rates, Data



Ex-post Markups:

How much of the contract spread results from the *ex-post* monopoly power of card companies? For each $\{i, a, y\}$ contract, Table 6 reports the excess of the contract interest rate over the interest rate that would give zero expected profits post-match (holding fixed the terms of all other equilibrium contracts). If all of the post-match surplus were to go households, the interest rate they would pay would be this zero-profit rate. Across contracts, this markup varies from 2.53 to 0.10 percentage points. Thus, the size of the markups, alternatively, the surplus given up by the household, is typically small, reflecting both competition and the low cost of posting contracts.

Still, it is worth examining what our theory predicts about the division of the surplus. Cell by cell, markups almost always decline from the most to the least impatient types. The most impatient types are hurt the most from not having the option to borrow and, therefore, give up the most surplus. Within each type, markups are generally declining in assets because the need to borrow is lower and not having a credit card is, correspondingly, less consequential. Variation in markups with respect to y, however, is generally nonmonotonic. We suspect that the nonmonotonicity reflects the varying strengths of two opposing 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, lowers 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 both y and i. For Types 2 and 3, it is the volatility effect that dominates for the highest y individuals: For these two types, markups go up sharply (in relative terms) as y rises from the next-to-highest to the highest level.

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	2.53	1.34	1.04	0.55	1.17	0.47	0.38	0.55	0.44	0.18	0.18
0.78	2.31	1.39	1.07	0.78	1.09	0.44	0.40	0.78	0.46	0.19	0.35
1.00	1.99	1.26	1.00	1.00	1.04	0.40	0.34	1.00	0.55	0.23	0.16
1.29	2.02	1.15	0.93	1.29	0.88	0.44	0.39	1.29	0.63	0.28	0.10
1.81	1.64	1.22	0.99	1.81	1.11	0.57	0.51	1.81	1.45	0.50	0.53

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.

9 Implications

9.1 Rate Cap Experiment

In this section, we examine how the credit contracts change if contract interest rates are capped at 10 percent. An examination of 14 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, the baseline contract is almost independent of the individual's asset level. These observations indicate that the implications of a 10 percent rate can be studied by examining how it affects Type 1 individuals with zero assets.

The left panel of Figure 18 plots how credit limits change with the rate cap. The cap leads to a drop in credit limits of about 15 to 18 percent. From the perspective of card companies, reducing the credit limit reduces the loss given default (because individuals default at or close to the credit limit) and reduces the probability of default on the card as well, as shown in the right panel. The decline in default probabilities is substantial, ranging between 50 and 53 percent (rounded) and results from the cardholder being less indebted.

The left panel of Figure 19 shows the effect of the rate cap on the probability of finding a card conditional on search. In percentage point terms, the probability declines about 4 or 5 percentage points for the top four persistent income groups, but it declines by 8 percentage points for the lowest persistent income group. The drop in contact probability is a direct consequence of reduction in profitability of the card. The measure of posted contracts declines to the point where the probability 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 zero-asset Type 1 individual would be willing to pay to not have the cap be imposed. The loss ranges between 1.4 and 2.4 percent. The lowest y individual is hurt the most.

Figure 18
Effect of Rate Cap on Credit Limits and Default Probability of Type 1-Zero-Asset Contracts

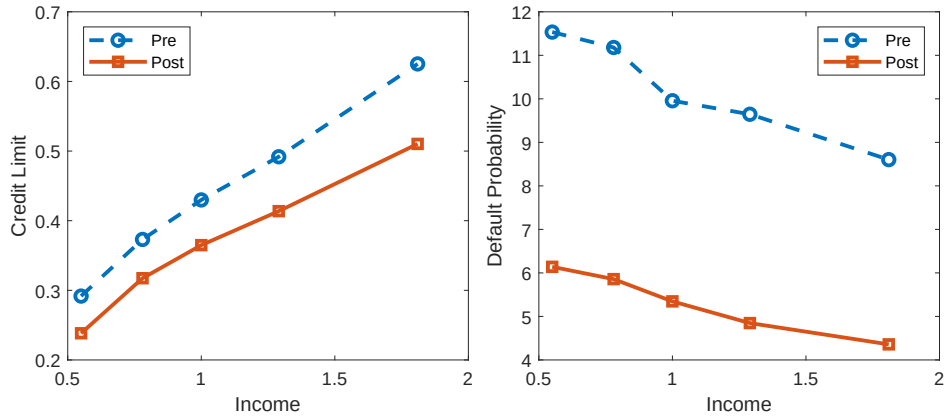
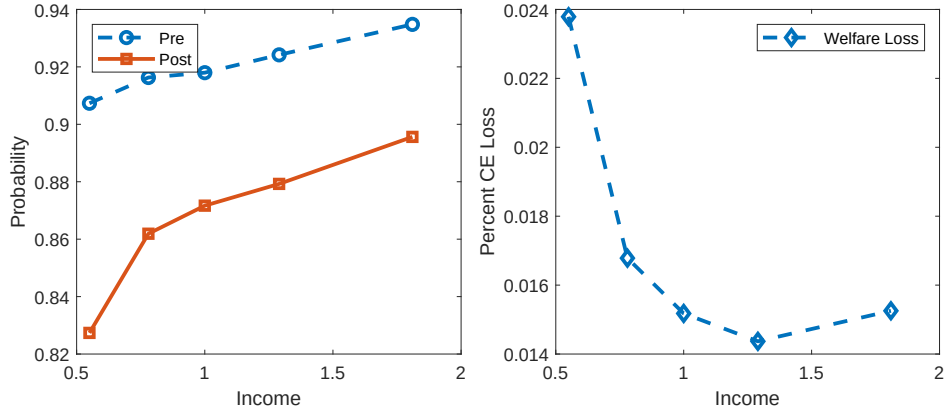


Figure 19
Effect of Rate Cap on Probability of Getting a Card and Welfare of Cardholder



9.2 Credit Cards and the Marginal Propensity to Consume

The market arrangement in our model resembles an Aiyagari-style model with a borrowing constraint and idiosyncratic income shocks. An important difference, though, is that individual's $(\underline{a}, R)$ contract depends strongly on her type and her persistent earnings at the time she accepted the card.

Given each type's discount factor and the borrowing interest rates they face, each type would want to borrow all the way to the limit, if their earnings were constant. But earnings risk and the implied risk of default provides a reason to limit borrowing and, perhaps, accumulate savings.

Table 7 reports the long-run (steady state) average assets — normalized by individual-level permanent earnings level P — by types and persistent earnings. When persistent income y is at its lowest level, all types are carrying balances on their cards, on average. For each type, the mean (normalized) assets in increasing in y. The positive association reflects the precautionary savings motive: Regardless of type, 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	−0.07	−0.16	−0.58	0.54	0.43	0.43
0.78	−0.07	−0.05	0.06	0.48	0.32	0.24
1.00	−0.04	0.08	0.72	0.44	0.26	0.15
1.29	0.00	0.26	1.58	0.39	0.20	0.08
1.81	0.64	0.54	3.23	0.35	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 positive assets have a utilization rate of zero.

Variation with respect to type depends on the person's persistent income level. For the lowest level, a more patient type is *more*, not less, indebtedness, on average. This is because holding income constant, more patient types are offered contracts with higher credit limits and lower interest rates as they have higher default costs.[28] A low y individual expects her income to rise over time and, so, has a desire to borrow. More patient

[28] 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.

types can indulge this desire more effectively and are more indebted, on average. For the three y levels between the lowest and the highest, the precautionary motive is dominant and mean assets increase with type. For the highest y, mean assets decline from Type 1 to Type 2, but increase from Type 2 to Type 3.

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

We now turn to the distribution of MPC implied by this sort of behavior. The top left panel of Table 8 displays the mean MPC for various cells in the steady state of the Base model. The MPCs range between $0.60 - 0.62$ for Type 1, between $0.33 - 0.36$ for Type 2, and between $0.08 - 0.12$ for Type 3. The key determinant of this pattern is the estimated discount factors of each type.

The model implies a significant difference in MPC between individuals who possess a credit card (the bottom left panel) and those who do not (the bottom right panel), with the largest difference for the least impatient type (Type 3). People without cards are mostly people who have defaulted in the past and are in exclusion. Since they cannot borrow, they face a hard liquidity constraint and engage in buffer-stock savings. This depresses their consumption and increases their MPC. As the more patient types engages more in this behavior, the MPC is elevated the most for Type 3 and the least for Type 1.

The top right panel shows how MPCs change if the option to borrow is taken away altogether.[29] Comparing cell by cell, we see that the MPCs are now higher but only very modestly. Evidently, if borrowing is not an option people accumulate more buffer-stock savings balances and their MPC essentially resembles their MPCs in the credit card economy. Still, the credit card industry is a modest force in favor of lower MPCs.

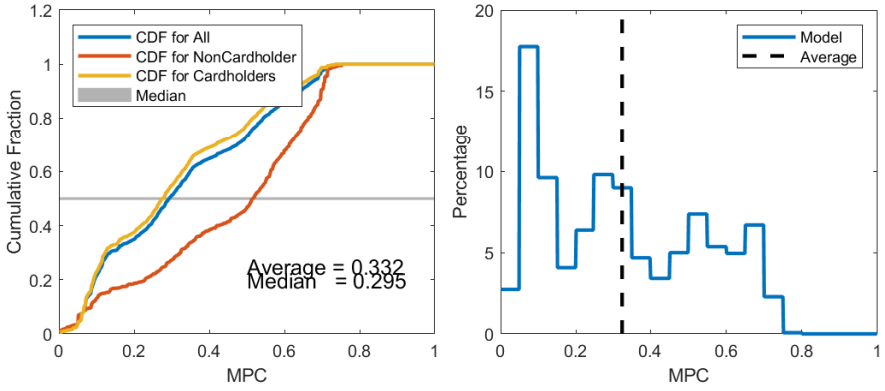
Finally, with a few exceptions, MPCs decline with higher persistent income levels for all types in all four panels. The reason is that higher y is associated with less debt or more savings (as shown in the top panel of Table 7). Notice also that this pattern of declining MPC is most pronounced for the most patient types, who are the ones most able to save.

[29] The risk free rate R_f is kept unchanged, so general equilibrium effects are being 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, Avg. 0.32			No-Cards Economy, Avg. 0.33		
0.55	0.62	0.36	0.12	0.62	0.37	0.16
0.78	0.61	0.35	0.12	0.62	0.36	0.14
1.98	0.61	0.34	0.11	0.61	0.35	0.11
1.29	0.60	0.33	0.08	0.61	0.34	0.09
1.81	0.61	0.34	0.08	0.61	0.34	0.08
	Card Economy					
	Cardholders, Avg. 0.30			Non-Cardholders, Avg. 0.45		
0.55	0.61	0.35	0.11	0.64	0.41	0.26
0.78	0.61	0.35	0.12	0.64	0.40	0.21
1.98	0.60	0.34	0.10	0.63	0.39	0.15
1.29	0.60	0.33	0.08	0.62	0.37	0.09
1.81	0.60	0.33	0.08	0.62	0.37	0.08

Figure 20
Distribution of MPCs

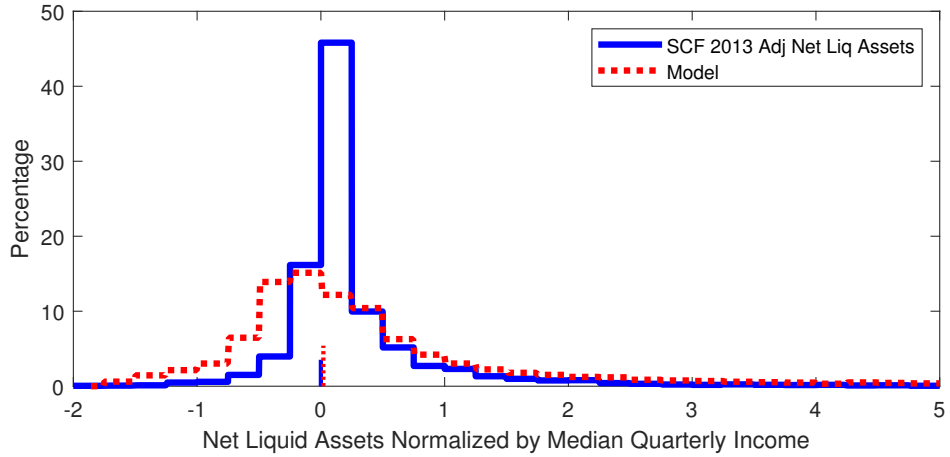


How do these MPCs compare to estimates of MPCs? Johnson, Parker, Souleles, and McClelland (2013) estimate 3-month MPC to be in the $0.5-0.9$ range and Orchard, Ramey, and Weiland (2023) estimate it to be in the $0.3-0.5$ range, with their preferred point estimate as 0.33. Thus the model-implied average MPC of 0.32 is in line with empirical evidence.[30] Figure 20 displays aspects of the distribution of MPCs (left panel) and its overall distribution (right panel).

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. To make the comparison as clean as possible, we

[30]Orchard, Ramey, and Weiland (2023) note that when the change in expenditures is broken down into change in durables and nondurables exependitures, it is mostly durable expenditures that responds.

Figure 21
Distribution of Liquid Balances, Model and Data



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 21 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. But we note that our model matches the utilization rates in our data quite well (see Figure 11). It is possible that the under-reporting of large debts is even more serious in the SCF.

References

AGUIAR, M., M. BILS, AND C. BOAR (2021): "Who Are The Hand-to-Mouth?," Unpublished.

ATHREYA, K. (2003): "Unemployment Insurance and Personal Bankruptcy," *Economic Quarterly, Federal Reserve Bank of Richmond*, 89(2), 33–53.

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.

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., AND B. EYIGUNGOR (2012): "Maturity, Indebtedness and Default Risk," *American Economic Review*, 102(6), 2674–2699.

DEMPSEY, K., AND F. IONESCU (n.d.): "Default Risk and Pricing in the U.S. Credit Card Market," Unpublished.

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.

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.

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 (2021): "Who Bears the Welfare Cost of Monopoly? The Case of the Credit Card Industry," Unpublished.

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.

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

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

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*.

LIVSHITS, I., J. C. MACGEE, AND M. TERTILT (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.

MEIER, S., AND C. D. SPRENGER (2012): "Time Discounting Predicts Creditworthiness," *Psychological Science*.

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 (2023): "Micro MPCs and Macro Counterfactuals: The Case of the 2008 Rebates," Unpublished.

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

RAVEENDRANATHAN, G., G. STEFANIDIS, AND G. SUBLET (2023): "Rate Caps on Revolving Credit Lines: Who Benefits," *mimeo*.

RAVEENDRANATHAN, G., AND G. STEPHANIDIS (2022): "The Unprecedented Fall in U.S. Revolving Credit," Unpublished.

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.

STANGO, V. (2002): "Pricing with Consumer Switching Costs: Evidence from the Credit Card Market," *Journal of Industrial Economics*, 50, 475–492.

STANGO, V., AND J. ZINMAN (2016): "Borrowing High versus Borrowing Higher: Price Dispersion and Shopping Behavior in the US Credit Card Market," *Review of Financial Studies*, 29, 979–1006.

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

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 %)	22.21	20.60	18.68	16.79	16.54
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 %)	19.76	19.45	19.11	18.60	17.93
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 annual frequency and credit limits are expressed as a ratio of annual income.

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.

Appendix C Finding the Default Thresholds

Suppose that, conditional on repayment, a type i individual in state (a, y, m) chooses the asset level a'_k for $m \in [m_k, m_{k+1}] \subset [-\lambda y, \lambda y]$. We need to determine if default is a better choice than repayment.

For $\gamma = 2$ case considered in this paper, the lifetime utility conditional on repayment is:

$$-\frac{1}{y + Ra - a'_k + m} + \nu\beta^i \mathbb{E}_{y'|y} W^i(a'_k, y'),$$

and conditional on default it is:

$$-\frac{1}{y - \lambda y[1 - \phi^i] + \phi^i m} + \nu\beta^i \mathbb{E}_{y'|y} \left[\delta \left[\mu S^i(0, y') + (1 - \mu) H^i(0, y')\right] + (1 - \delta) X^i(0, y')\right].$$

For ease of exposition, denote the repayment continuation value as N (no default), the default continuation value as D, $y + Ra - a'_k$ as A and $y - \lambda y[1 - \phi^i]$ as B.

The following cases can arise.

1. If $N > D$ and $A + m_k > B + \phi m_k$ then for all $m \in [m_k, m_{k+1}]$ the individual repays.
2. If $D > N$ and $B + \phi m_{k+1} > A + m_{k+1}$ then for all $m \in [m_k, m_{k+1}]$ the individual defaults.
3. Otherwise, we need to find m values where there could be a switch from default to repayment or vice versa. To locate these switch points or thresholds, set the utility from repayment equal to the utility from default:

$$-\frac{1}{A + m} + N = -\frac{1}{B + \phi m} + D$$

Subtracting the r.h.s. from the l.h.s. leads to a quadratic equation in $am^2 + bm + c = 0$, where $a = \phi(N - D)$, $b = (1 - \phi) + (N - D)A\phi + (N - D)B$, and $c = (A - B) + AB(N - D)$.

Denote the two roots of by $m_1 < m_2$.

If $N - D > 0$, the coefficient $a > 0$ and the quadratic is U-shaped. It follows that for all $m \in (m_1, m_2)$, the value of the quadratic is negative, i.e., value of default exceeds the value of repayment. Therefore, the individual will default for all $m \in [m_k, m_{k+1}] \cap (m_1, m_2)$. If there is no overlap then there is repayment for all $m \in [m_k, m_{k+1}]$. Otherwise, there can be 1 or 2 thresholds where there is a switch between default and repayment.

If $N - D < 0$, the coefficient $a < 0$ and the quadratic is inverted-U-shaped and, so, for all $m \in (m_1, m_2)$ the value of the quadratic is positive, i.e., the value of repayment exceeds the value of default. Therefore, the individual will *repay* for all $m \in [m_k, m_{k+1}] \cap (m_1, m_2)$. If there is no overlap between the two intervals, the

individual will default for all $m \in [m_k, m_{k+1}]$. Otherwise, there can be at 1 or 2 thresholds where there is a switch between default and repayment.