Heterogeneity in the Credit Card Market

Satyajit Chatterjee and Burcu Eyigungor

Federal Reserve Bank of Philadelphia

Preliminary Draft, April 2023

Abstract

A relatively new administrative data set is used to establish patterns regarding contract terms, usage, and performance of anonymized individual credit card accounts. Accounts are distinguished by account holders' income and credit score at origination. It is found that contract interest rates decline with credit score and income while credit-limit-to-income ratios rise with credit score but decline with income. Utilization and default rates decline sharply with credit score and income. The workhorse heterogeneous agent macro model augmented with a credit card industry is developed to explain these patterns. If individuals differ by discount factors and default costs, the model can account for almost all of these patterns. The model predicts large spreads between interest rates and default frequencies, as observed, despite the competitive provision of credit card contracts. It implies that the credit card industry lowers the average MPC of consumers, especially for low-income individuals.

Keywords: Credit limits, defaults, financial frictions, heterogeneity

JEL Codes: D15, E44, G51

Satyajit.Chatterjee@phil.frb.org and Burcu.Eyigungor@phil.frb.org. The authors thank Larry Santucci and seminar participants at NYU, JHU and Rochester 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.

Contents

1 Introduction

A striking feature of the US consumer credit market is that borrowers are quite heterogeneous in terms of default risk. By one measure, the average default risk of the most risky borrowers is, on average, 30 times that of the least risky borrowers. As these differences in default risk are largely predictable, there is a correspondingly wide variation in credit access between the least and the most risky borrowers.

In this paper, we seek to shed light on the fundamental determinants of the observed differences in default risk and credit access. To this end, we first document patterns in newly originated credit accounts with respect to (i) contract terms, namely, interest rates and credit limits, (ii) the utilization rate of accounts, i.e. revolving balances as a fraction of the credit limit, and (iii) the frequency of default, i.e. of serious delinquency and bankruptcy. We use a relatively new confidential administrative data that follows anonymized individual credit card accounts over time. This is data collected by the Board of Governors of 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.[1]

We distinguish accounts by the two characteristics of account holders that must be reported for every newly issued card. These are the *credit score* and *income* of the account holder at the time the card is issued. An account holder's credit score is, in essence, an estimate of the probability that the account holder will not experience a delinquency on some financial obligation within the next two years.[2] The estimate is based on an individual's credit history maintained by credit bureaus. It is a measure of an individual's creditworthiness that is widely used in the consumer credit industry.

The patterns we focus on are the variation in contract terms, utilization rates, and default rates with respect to the credit score and income of the individual recorded at the time of origination of the account. The following patterns emerge:

- The credit-limit-to-income ratio rises with credit score and declines with income.

[1] One part of this data, called Y-14M, reports on the terms, usage, and performance of individual credit card accounts managed by the reporting banks.

[2] More precisely, a credit score is a negative monotonic transformation of the probability of being 90 days or more late on a financial obligation. If δ is the probability of delinquency within the next two years, the credit score is $a + b \cdot \ln(1 - \delta)$, where a and b are positive numbers.

- The interest rate (on cards) declines with credit score and income, but the decline is relatively mild
- The utilization rate declines with credit score and income
- Default rates — measured as the frequency of a serious delinquency — defined as being 120 days or more late on payments due, or a bankruptcy within a two-year horizon — declines with incomes and credit scores.

The first two facts tell us that individuals with higher credit scores have greater to access to credit relative to their earning capacity in comparison to individuals with lower credit scores, but the interest rate offered to higher credit score individuals is not much lower than that those offered to lower credit score individuals. 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 individuals with low incomes and/or low credit scores utilize their borrowing capacity more than individuals who are richer and/or have higher scores. The final fact tells us that poorer individuals encounter financial distress more frequently than richer individuals. There is a negative relationship between default rates and credit scores as well which shows that credit scoring is effective in predicting serious delinquency or bankruptcy.

To understand how such facts might emerge, we begin with the workhorse macro model with uninsured idiosyncratic income shocks and augment it with a credit card industry. The industry is composed of card companies who 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. Individuals acquire credit cards via directed search: Card companies post credit card contracts and individuals search for the card that is best for them. We assume free entry into the card business so card companies earn zero profits in expectation. We assume an open endowment economy with an internationally given risk-free rate. We abstract from aggregate shocks.

We first investigate if the ex-post heterogeneity in income induced by idiosyncratic income shocks is sufficient to account for all of the observed patterns. We find that the answer is No, although the workhorse model comes quantitatively close to explaining all of the income-based patterns mentioned above and most of the credit score-based patterns. What it fails to explain is the strong positive association between credit scores and credit-limit-to-income ratios: The model predicts a strong *negative* relationship. As will become clear, this prediction is basic to this class of models and cannot be circumvented.

The failure, we argue, is an indication that the observed variation in credit scores cannot result solely from income differences between individuals and that some other source (or sources) of heterogeneity must also be present. Indeed, incomes and credit scores are less strongly related to each other in the data than in the model with ex-post income heterogeneity only. This difference is suggestive of some variation among individuals that make poor individuals have high credit scores.

Lenders perceive high credit score people as capable of carrying more debt relative to their incomes without defaulting. This could be an equilibrium outcome if some people have higher credit scores because they are more patient and/or have higher default costs. This motivates our main model specification wherein individuals differ with respect to both discount factors and default costs.

Our second finding is that the model with discount factor and default cost heterogeneity can explain almost all the patterns we observe. If we restrict people to have a common discount factor, the model comes surprisingly close to the patterns to seek to explain. In particular, it does much better than the workhorse model with only income heterogeneity. In contrast, if we restrict people to have a common default cost but different discount factors, the equilibrium has the same drawback as the workhorse model: The relationship between credit scores and credit-limit-to-income ratio remains negative.

In the balance of the paper, we use the (quantitative) model with both default cost and discount factor heterogeneity to shed light on two issues that has attracted the attention of prior researchers.

The first issue relates to the relationship between interest rates and default probabilities on credit card accounts. In the data, the spreads between credit card interest rates and card companies' cost of funds seem far in excess of default probabilities on card accounts. We show that our model can replicate this fact even though there is no flow cost of servicing a credit card account and each account earn zero (expected) profits in equilibrium. The "excess" spread is explained by the default behavior of individuals. If an individual's debt is not close her credit limit then it is optimal, in the face of an income shortfall, for her to increase the balance on her card with the expectation that she will reduce her balance when her income recovers. In contrast, if her debt is already at, or close to, her credit limit, she might choose to buffer her consumption against an income shortfall by defaulting. The key consequence of this behavior is that default occurs on card balances that are, on average, much *larger* than balances on which individuals make payments. Since the same interest rate applies to debts of all sizes below the credit limit, the spread between a card interest rate and the risk-free rate must exceed the default probability on the card for the company to break even.

The second issue relates to the implications of a functioning credit card industry for the distribution of MPCs across the population. The first implication is that the discount factors required to explain the credit card facts — especially the default facts — imply that all individuals are quite impatient. This impatience manifests itself in an average MPC of about 0.22. A second implication is that if the credit card industry did not exist, individuals with lower incomes and discount factors would have higher MPCs, while the MPCs of people with higher discount factors and incomes would be largely unaffected. Thus a functioning credit card industry pulls down the MPC of poor and impatient people closer to the MPC of richer and patient people. Overall, the existence of credit card industry lowers average MPC in the economy.

2 Contributions and Connections to the Literature (Incomplete)

Our paper belongs to the quantitative-theoretic macroeconomics literature on consumer borrowing and default initiated in Athreya (2002), Chatterjee, Corbae, Nakajima, and Rios-Rull (2007) and Livshits, MacGee, and Tertilt (2007). Up to now, this literature — recently surveyed in Exler and Tertilt (2020) — has mostly focused on explaining per-capita unsecured consumer borrowing and default, borrowing and default over the life-cycle, the evolution of per-capita borrowing and default rates over time, and business cycle fluctuations in borrowing and default. The literature has not addressed the sources of heterogeneity in credit card contract terms.

As noted in the introduction, our findings point to ex-ante differences in discount factors and default costs. In an experimental setting, Meier and Sprenger (2012) tied discount factor heterogeneity to differences in credit scores: participants who acted more patiently in experiments tended to have higher credit scores. More recently, Chatterjee, Corbae, Dempsey, and Ríos-Rull (forthcoming) explore the role of privately observed discount factor heterogeneity in driving differences in credit scores across individuals.[3] In Athreya, Tam, and Young (2009) differences in survival probabilities by educational attainment introduce differences in effective discount factors. In Athreya, Tam, and Young (2012) there are differences in default costs across education types.

Raveendranathan (2020) and Herkenhoff and Raveendranathan (2021) pioneered the credit card model of unsecured consumer borrowing with credit card search (significant earlier unpublished contributions are Mateos-Planas and Ríos-Rull (2007) and Drozd and Nosal (2008)). They use a life-cycle model while we continue in the tradition of heterogeneous agent macro models that abstract from life-cycle features.

Motivated by high spreads and limited pass-through of changes in the card companies' cost of funds, Herkenhoff and Raveendranathan (2021) study the existence of monopoly power in the credit card industry. Dempsey and Ionescu (n.d.) also study spreads or borrowing premia in unsecured consumer loans. The former uses a credit card model while the latter uses the "supply curve of credit" model pioneered in Eaton and Gersovitz (1981) (and adopted in Chatterjee, Corbae, Nakajima, and Rios-Rull (2007) and Livshits, MacGee, and Tertilt (2007)). We employ the credit card model and assume directed search in the credit card market. Our focus is not on monopoly power per se but on understanding the underlying nature of heterogeneity among borrowers.

There is a large literature on the implications of heterogeneity for macroeconomics (see, for instance, Krueger, Mitman, and Perri (2016)) which gives great importance to the reasons why consumption differs across people and what this dispersion means for macroeconomic policies (Kaplan and Violante (2014)). This literature has found that the workhorse model with incomplete markets and plausible ex-post heterogeneity in incomes cannot easily account for the observed wealth inequality. But augmented with ex-ante heterogeneity in discount factors (and stochastically finite lives), it can do so — as shown in Krusell and Smith (1998) and Carroll, Slacalek, Tokuoka, and White (2017). More recently, Aguiar, Bils, and Boar (2021) argue that consumption behavior of high MPC people suggests they have high discount factors, as did an earlier study by Gelman (2021). We examine the implications for heterogeneity suggested by the patterns in credit

[3] If discount factors are privately observed, credit scores signal a person's type and are rationally used by lenders to regulate the amount of credit granted to an individual. If discount factors are publicly observed they will affect credit scores, but a credit score does not serve any signaling role.

card contracts and explore its implications for the distribution of marginal propensity to consume across the population.

3 Credit Card Terms, Usage, and Performance

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, it is natural to focus on new card originations as we expect contract terms to be tied to the perceived characteristics of customers most closely when a card is first issued.[4]

We begin with some basic facts about the terms of credit card contracts, the usage of cards, and their performance. Our data source is confidential supervisory data, maintained by the Board of Governors of the Federal Reserve System. The data tracks, at a monthly frequency, the terms, usage, and performance of credit card accounts issued by all U.S. bank holding companies with assets greater than $50 billion. The data covers about 80 percent of all U.S. credit card accounts in existence at a point in time.

The facts we present pertain to a 1-in-200 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. We chose our sample period with the following consideration in mind. The majority of credit cards accounts 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.[5]

Following the initial random selection, we trimmed the sample for outliers.[6] 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

[4] As time passes, the circumstances of the account holders might stray far from what was expected at origination.

[5] The Y-14M data set begins in mid 2012, but not all variables are reported/available for the first year or so.

[6] 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.

income with the estimated earnings processes commonly used in the calibration of macro models. Second, accounts with credit scores in the bottom 10 percent of the credit-score-at-origination distribution were removed. The terms and performance of these most risky accounts seem outside the explanatory scope of our model.[7]

The contract terms we focus on are non-promotional interest rate on revolving balances (the annual percentage rate or APR) and credit limits. We measure these terms not at origination when promotional terms are often active, but 18 months following origination when non-promotional terms are typically in place.

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: For instance, the leftmost dot in the top left panel is the mean APR on contracts with credit scores in the bottom 20 percent of the distribution of credit scores at origination. The plotted lines are simply lines of best fit through these dots.

There is systematic variation in contract terms with respect to score and income at origination. Individuals with low scores hold contracts with higher interest rates on average as do individuals with low income. While contract interest rates decline with scores and incomes, the variation is not pronounced: The mean interest rate in the top quintile of scores is only 3.3 percentage points lower than the mean interest rate in the bottom quintile and the difference in mean APRs between the top and bottom quintiles of income is only 1.5 percentage points.

In contrast, 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 score in the top quintile have a ratio that is, on average, about 2.7 times higher than 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, 1.7 times higher than individuals with incomes in the top quintile. This pattern arises because credit limits increase with income but increase less in proportionate terms.

Figure 2 displays the performance patterns of credit card accounts. We focus on two measures of performance. One is the frequency of 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

[7] 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.

Figure 1:
Variation in Contract Terms by Score & Income

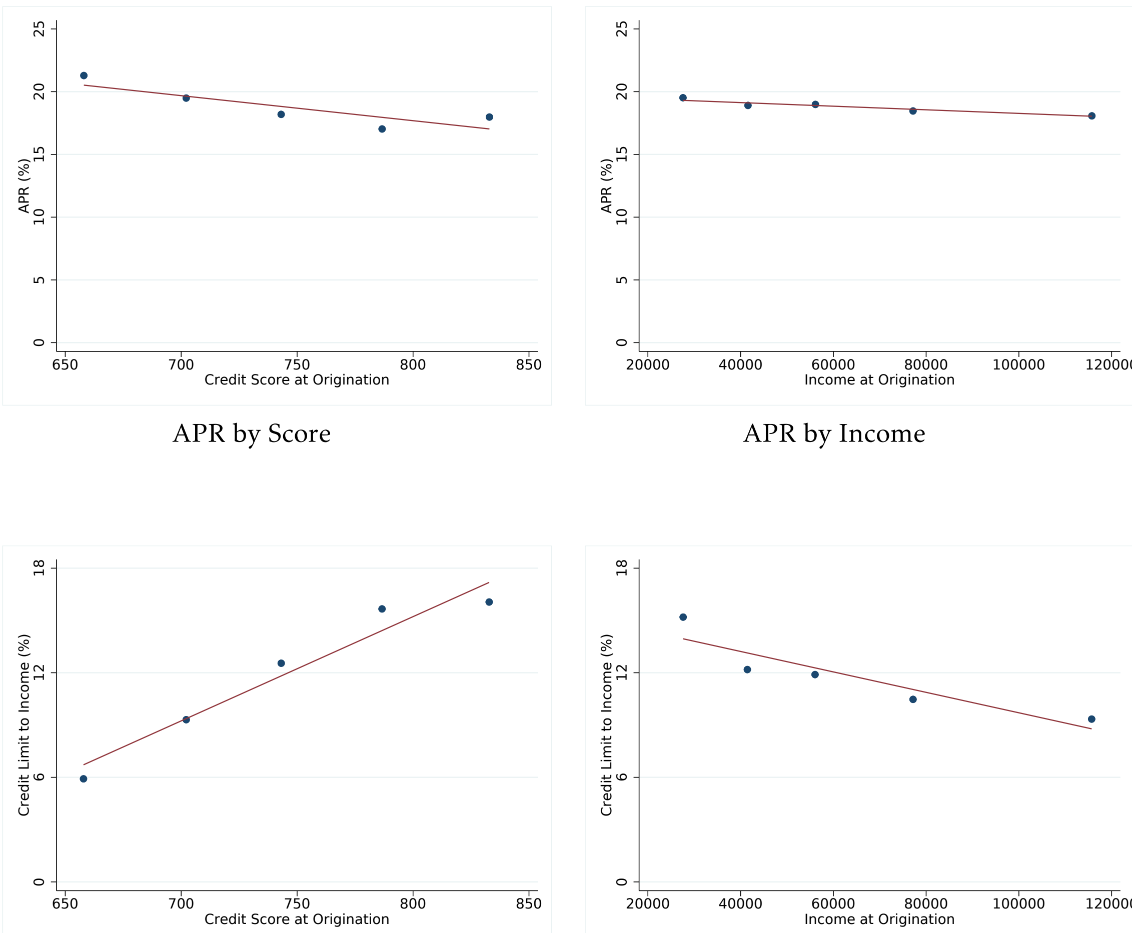

APR by Score

APR by Income

Credit Limit to Income Ratio by Score

Credit Limit to Income Ratio by Income

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

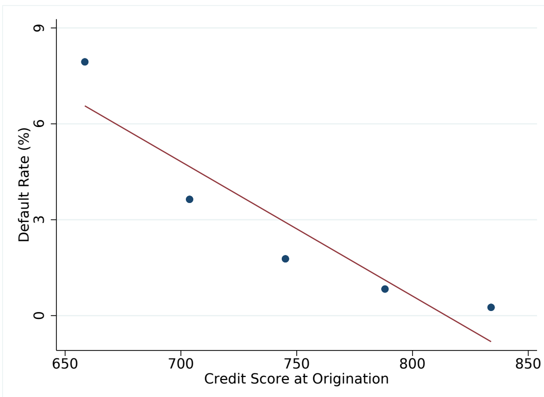

Default Frequency by Score

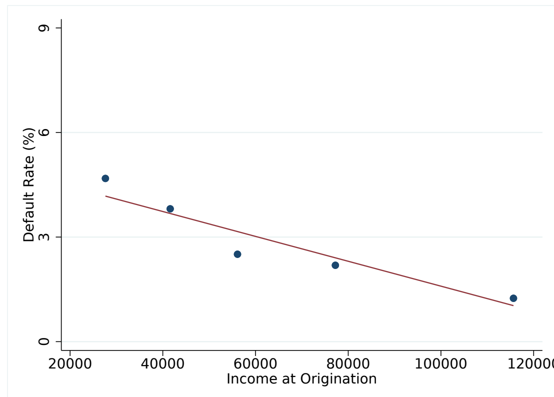

Default Frequency By Income

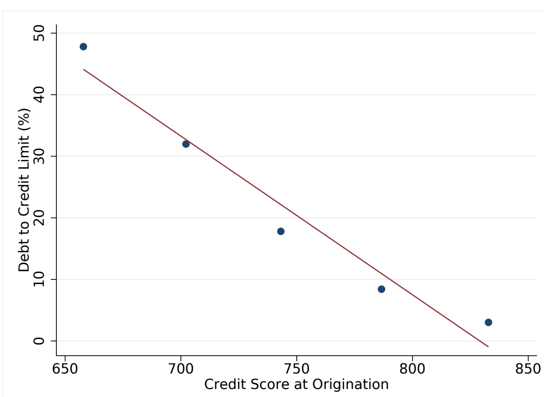

Debt to Credit Limit Ratio by Score

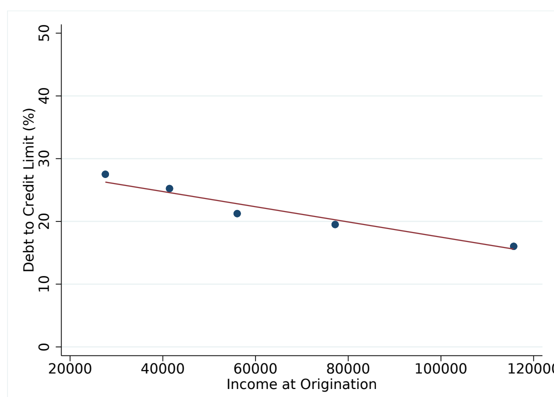

Debt to Credit Limit Ratio by Income

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.[8] The second measure is the utilization rate of the card for which we use the ratio of revolving debt to credit limit 18 months following origination.

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 7.9 percent while it is 0.3 percent for the top quintile. Default frequencies also decline with income but the variation is less pronounced. The default frequency in the bottom quintile of incomes at origination is 4.7 percent, while for the top quintile it is 1.2 percent.[9]

The bottom panels display the ratio of revolving balances to credit limits. Individuals in the bottom quintile of scores have a utilization rate of 48.5 percent while the top quintile has utilization rate of only 3.0 percent. The utilization rate also varies negatively with income, but much less so. Individuals in the bottom quintile of incomes have a mean utilization rate of 17.8 percent, while those at the top income quintile have a mean utilization rate of 16.2 percent.

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

4 Model

4.1 Environment

Time is discrete. There are $i \in \mathbb{I}$ types of people in the economy who differ in their discount factor and/or default costs. We assume equal measures of each type. Individuals have CRRA (per-period) utility function with a common intertemporal elasticity of substitution $1/\gamma$ and an individuals survive from period to the next with constant probability $\nu \in (0,1)$.

An individual's income is stochastic and given by $y+m$, where y is the persistent component of income and m is the transitory component. The persistent component is distributed $F(y'|y)$ and the transitory component is distributed uniformly on support that depends on the persistent component of earnings: $[-\lambda y, \lambda y]$, $\lambda > 0$. Thus we permit the transitory shock to be negative, meaning that an individual's income can temporarily fall below her persistent income level y.

[8] 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.

[9] A noteworthy fact is that the contribution of bankruptcy and debt waivers to defaults is relatively minor, with majority of defaults – 71 percent – stemming from delinquencies.

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 and we use $\omega = (\underline{a}, R)$ to denote the contract terms of a card. We use $\Omega \subset \mathbb{R}^2$ to denote the space of all possible contracts. An individual can hold at most one card and a card is forever associated with one set of contract terms.[10] 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 she does, she loses her card and (i) she cannot borrow or save in the period of default, (ii) she gives some fraction $0 < [1-\phi_y^i] \leq 1$ of her transitory income in excess of $-\lambda y$ to her creditors, and (iii) with probability $(1-\delta)$ she is shut out of the credit card market in the following period. And, conditional on being shut-out, she continues in that state with probability $(1-\delta)$. The fraction of income given to creditors upon default is permitted to depend on the person's type i and her persistent income level y.

A person who is not shut out of the credit card market but does not have a card searches for one with probability $\mu > 0$. Search happens at the start of a period before the transitory shock m is realized. Since there is complete information, she can only search for contracts that are being offered to people of her type i, her asset position a, and her persistent earnings y. But within this "market", there are potentially many submarkets offering different contract terms ω and card companies commit to honoring the terms of a contract if accepted. The probability of encountering a contract ω is given by $f^i(\omega; a, y)$ and this function is an equilibrium object. These assumptions makes our search environment one of directed search.

In order to generate active search markets, we assume that an individual can get separated from her card with probability ξ. 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 can be transferred to a prespecified contract $\bar{\omega}^i(y)$ that depends on her persistent income at the time of the separation shock and her type. If the contract is profitable for the card company, i.e., generates nonnegative expected profits, the balance transfer goes through; otherwise, the individual continues on with her existing card.

[10]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. Our assumption that terms never change sidesteps this institutional detail.

4.2 Individual's Decision Problem

Let $h^i(a,y,m)$ denote the value of an individual who is not shut-out of the card market but failed to get a card in the current period. This could be because she did not search or she searched but failed to make contact. In this situation, she can only save. Let $S^i(a,y)$ denote the *ex-ante* value of a type i person in state (a,y) who is without a card but searching for one. The qualifier *ex-ante* means that this is an expected value prior to the realization of the transitory income shock m. Then $h^i(\cdot)$ solves:

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

Let $x^i(a,y,m)$ denote the value of an individual who is without a card and is excluded from the credit card market following a default. Such a person is also limited to saving and $x^i(\cdot)$ solves:

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

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

$$v^i(y,m,a;\omega) = \max_{a'\geq \underline{a}} \left\{\frac{c^{1-\gamma}}{1-\gamma} + \nu\beta^i \mathbb{E}_{y'|y} W^i(a',y';\omega)\right\}$$
$$c = y + m + Pa - a', \text{ where}$$
$$P = \begin{cases} R \text{ if } a < 0 \\ R_f \text{ if } a \geq 0. \end{cases}$$

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 (optimal) asset and consumption decision rules.

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

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

In the event of default, the individual retains some fraction of her transitory income and pays the rest to her creditors. Since the support of m as $[-\lambda y, \lambda y]$, the individual pays zero if m is at its lower support and pays $2(1-\phi^i_y)\lambda y$ if it is at its upper support.

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

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

where $\xi^i(a,y)$ is the probability that an indebted person will have a balance transfer; it is the product of ξ and an indicator function that is 1 if the balance transfer is profitable and 0 otherwise. For a indebted individual this probability is either ξ or 0, given her state $\{a,y\}$. For future use, let $D^i(a,y,m;\omega)$ denote the default decision rule: $D^i(\cdot)$ is 1 if $v^i(a,y,m;\omega) < v^i_{\text{DEF}}(y,m)$ and 0 otherwise.

Finally, the recursion for $S^i(a,y)$ solves

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

The inner max chooses over ω for the best contract. While the domain of choice is indicated as Ω, we may assume without loss of generality that the choice is over only those contracts for which $f^i(a,y;\omega)$ is strictly positive. The outer max recognizes that the individual always has the option to not engage in search and get the value $H^i(a,y)$.

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 have access to funds at the risk-free interest rate 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,

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

If the individual is carrying a balance on the card and defaults, the card company gets $[1-\phi_y^i](m+\lambda y)$ and credit line is closed. If she does not default, the card company receives the (gross) interest payment $a \cdot R(\omega)$ minus the individual's new borrowing $A^i(a,y,m;\omega)$, if any. If the individual is not carrying any balances, the company "receives" the minus of new borrowings, 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 $\xi^i(a',y')$.[11] 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.

A card company chooses the profit-maximizing contract ω for people of type i in state (a,y) and, given profit associated with the best contract, chooses the measure of contracts to post at the cost of $\tau > 0$ per post. In making these choices, 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

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

[11] For $a' \geq 0$, $\xi(a',y') = \xi$ and for $a' < 0$, $\xi(a',y')$ is either ξ or 0 (as explained in the text).

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,

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

4.4 Contact Probabilities and Market Tightness

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

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

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

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

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

where

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

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

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

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

$$\begin{aligned} &S^{i*}(a,y) = \\ &\max\left\{\max_{\omega\in\Omega}\left\{f(\theta^{i*}(a,y;\omega))\cdot[V^{i*}(a,y;\omega)-H^{i*}(a,y)]+H^{i*}(a,y)\right\}, H^{i*}(a,y)\right\} \end{aligned} \tag{1}$$

$$q^{i*}(a,y;\omega) = \begin{cases} 1/\left[1+f^{-1}\left(\frac{S^{i*}(a,y)-H^{i*}(a,y)}{V^{i*}(a,y;\omega)-H^{i*}(a,y;)}\right)\right] & \text{if } V^{i*}(a,y;\omega) > H^{i*}(a,y) \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

$$q^{i*}(a,y;\omega)\Pi^{i*}(a,y;\omega) \leq \tau \text{ for all } \omega \in \Omega \tag{3}$$

$$\left[q^{i*}(a,y;\omega)\Pi^{i*}(a,y;\omega) - \tau\right]\theta^{i*}(a,y;\omega) = 0 \text{ for all } \omega \in \Omega \tag{4}$$

Equation (1) 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 $S^{i*}(a,y)$.

Equation (2) 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 $H^{i*}(a,y)$ don't attract any customers and, so, $q^{i*}(a,y;\omega)$ is zero (in effect, market tightness for such a contract is infinite).

Equation (3) asserts that all feasible contracts earn nonpositive expected net profits.

Equation (4) 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 (3) and (4) 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. 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 and are in state $\{a,y\}$.[12] The aggregate measure of posts in this (sub)market is then $E^{i*}(a,y;\omega^*) = B^{i*}(a,y;\omega^*) \cdot \theta^{i*}(a,y;\omega^*)$.[13]

4.6 Computation

For our quantitative analysis, we compute an approximate equilibrium of our 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 $\{i,y\}$ person with assets a, searches in the $\{i,a_k^i,y\}$ market, 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.

5 Quantitative Analysis: Preliminaries

5.1 Income Process

We calibrate the earnings process, we use the estimates provided in Floden and Lindé (2001). These authors estimate an earnings process for annual earnings given by $y(t)z(t)$, where y, the persistent component of earnings, is an AR1 process in logs and z, the transitory component, is i.i.d. with $\ln(z)$ having standard deviation σ_z. Their estimates are reported in Table 1.

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

[12]These include households that separated from their cards at the start of the current period and those who entered the period without a card.

[13]If these individual search for several distinct contracts in equilibrium then the measure of searchers in each submarket is indeterminate. But given some distribution of searchers, the measure of posts in each submarket can be determined in the same way.

Table 1:
Income Process Parameters

Parameter	Description	Value
ρ	Autocorrelation of log of persistent component	0.91
σ_ε^2	Var. of innovation to log of persistent component	0.043
σ_z^2	Var. of log transitory shock	0.042

Notes: Source: Floden and Lindé (2001), Table ? p. All parameters are at annual frequency.

distributed with support $[-\lambda y, \lambda y]$, given y. Note that this latter assumption is equivalent to assuming that z is uniformly distributed with support $[(1-\lambda), (1+\lambda)]$. Depending on the experiment, we discretize y into 3 or 5 levels and assume that it follows a first-order Markov process. We pick λ and the values of the Markov transition matrix such that the model earnings generate a ρ, σ_ε and σ_z at annual frequencies that matches what is reported in Table 1.

5.2 Parameters Set Independently

Table 2 displays the other parameters whose values are set independently. The (real) risk-free (gross) rate R_f was set at 1.01 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 cost of posting a contract, τ, is set to 0.001, a value low enough to ensure that a profitable credit card exists for all $\{i, a, y\}$ triples. The probability of search following a separation, μ, is set at 0.50: Delays in getting a new card creates a lock-in effect for existing cards which increases profits.

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 a year on average, the effective exclusion from the credit card market following a default is about 7 years.

The separation parameter, ξ was set to 0.02. We pin this down from the observed frequency, between June 2013 and December 2017, of accounts in good standing that disappear from our sample.

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

Table 2:
Parameters Set Independently

Parameter	Description	Value
λ	Bounds for transitory shock	0.67
R_f	Gross risk-free rate	1.0025
ν	Probability of survival	0.998
γ	Curvature of CRRA utility function	2.0
δ	Entry rate after default	1/28
τ	Cost of posting a contract	0.001
μ	Probability of searching for a card	0.50
ξ	Probability of exogenously separating from card	0.02
α	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.

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.

These parameter choices will be held fixed in the subsequent analyses.

5.3 Estimation Strategy

The parameters that will vary across the different economies studied below are the number of types, the β^i and ϕ^i for each type, and the sensitivity of default costs to y. For the latter, we assume that

$$\phi^i_k = \phi^i + \zeta\left[\frac{K+1}{2} - k\right],$$

where K is the number of income states (3 or 5) in the Markov chain and $k \in \{1,2,\ldots K\}$ indexes income states from lowest to highest. If ζ is estimated to be positive, lower income individuals of all types get to keep a larger portion of their income in default.

In order to address the facts discussed earlier, 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 a default within the next two years (i.e., 8 model periods). This probability can be computed recursively as described in Appendix C.

The model we build is of an individual. However, our data set follows credit card *accounts* and it is not possible to merge accounts belonging to the same individual. To take account of this, we scale up the credit-limit-to-income ratios shown in Figure 1 by a factor of 2.2, which is the average number of cards that individuals in the lowest quintile of credit scores possess, if they possess any card at all.[14] 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.

All estimation results presented in subsequent sections stem from minimizing the following objective function:

$$
\begin{aligned}
&w_1 \times \sum_{j=1}^{5} \Big(\text{Def Prob}_{\text{data, score}_j} - \text{Def Prob}_{\text{model, score}_j}\Big)^2 + \\
&w_1 \times \sum_{j=1}^{5} \Big(\text{Def Prob}_{\text{data, inc}_j} - \text{Def Prob}_{\text{model, inc}_j}\Big)^2 + \\
&w_2 \times \sum_{j=1}^{5} \Big(\text{APR}_{\text{data, score}_j} - \text{APR}_{\text{model, score}_j}\Big)^2 + \\
&w_2 \times \sum_{j=1}^{5} \Big(\text{APR}_{\text{data, inc}_j} - \text{APR}_{\text{model, inc}_j}\Big)^2 + \\
&w_2 \times \sum_{j=1}^{5} \Big(\text{Util}_{\text{data, score}_j} - \text{Util}_{\text{model, score}_j}\Big)^2 + \\
&w_2 \times \sum_{j=1}^{5} \Big(\text{Util}_{\text{data, inc}_j} - \text{Util}_{\text{model, inc}_j}\Big)^2 + \\
&w_3 \times \sum_{j=1}^{5} \Big(\text{CL-to-Y}_{\text{data, score}_j} - \text{CL-to-Y}_{\text{model, score}_j}\Big)^2 + \\
&w_3 \times \sum_{j=1}^{5} \Big(\text{CL-to-Y}_{\text{data, inc}_j} - \text{CL-to-Y}_{\text{model, inc}_j}\Big)^2 .
\end{aligned}
$$

We set $w_1 = 1000$, $w_2 = 5$ and $w_3 = 1$. Thus, the estimation strives to match the pattern in default frequencies. We do this because our goal is to explain the fundamental drivers of contract terms and an explanation that does not get the pattern in default frequencies right would not be credible. The other moment deviations are given the same weight, 5, except for the credit limit to income ratios which are given a weight of 1. The reason for down weighting the CL-to-Y ratio is simply

[14] This information comes from industry sources. We use 2.2 as the scaling factor rather than the average (which is 2.7) because we want to avoid overestimating the debt capacity of lowest credit score people.

that any given change in this ratio, say a change of 0.10, is less economically consequential than the same change in interest rates: Across quintiles, the former is roughly close to 1 (when income is measured at quarterly frequency) while the latter is close to 0.20. Details of the computation and estimation procedures are given in Appendix (to be added)

6 Can the Workhorse Model Explain These Patterns?

We start off our quantitative analysis with workhorse version in which there is a single type, i.e., all individuals have a common β and a common ϕ (and ζ). For this investigation, the income process in Table 1 is approximated by a Markov chain with 5 states.

Figure 3 plots average incomes in the income quintiles relative to the middle quintile for both the model and the data. There is a close fit. This is not a given since incomes in the model is determined by the income process in Table 1 which is estimated on data from a completely different source. Furthermore, the set of people looking for a new card is partly endogenous: Some searchers are looking for a new card after being shut out of the credit card market following a default. Also, among indebted people, the ones getting a new card is a a selected group, i.e., the ones for whom the balance transfer contract is profitable. That said, the concordance is helped by the fact that our data sample considers only individuals in 7th to 93rd percentiles of incomes-at-origination.

Figure 4 compares the default frequencies across income and score quintiles with the data. The model can replicate the default frequency by income quintiles fairly well. The pattern has a simple explanation in the model. A high-income person expects her income to decline in the future which induces her to borrow less or save. In effect, temporarily high income makes an individual patient and, therefore, she is less likely to default over a two year horizon. The opposite is true for a low-income individual who expects her future income to rise. She is more likely to borrow and, therefore, more likely to default (within a two year horizon).

The model cannot explain the pattern in default frequencies across score quintiles to the same degree of accuracy. It substantially underpredicts the default frequency for the bottom quintile and overpredicts it for top three quintiles. In the model, households with higher probability of default (equivalently, lower credit scores) are people with low incomes: they are more likely to be indebted (or become indebted) and more likely to default. Given a single dimension of heterogeneity (incomes), the model attempts to get the pattern right by making people very impatient

Figure 3

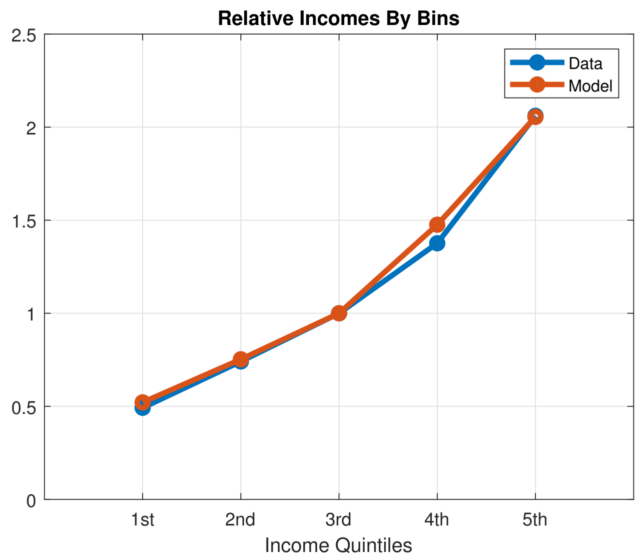

Figure 4

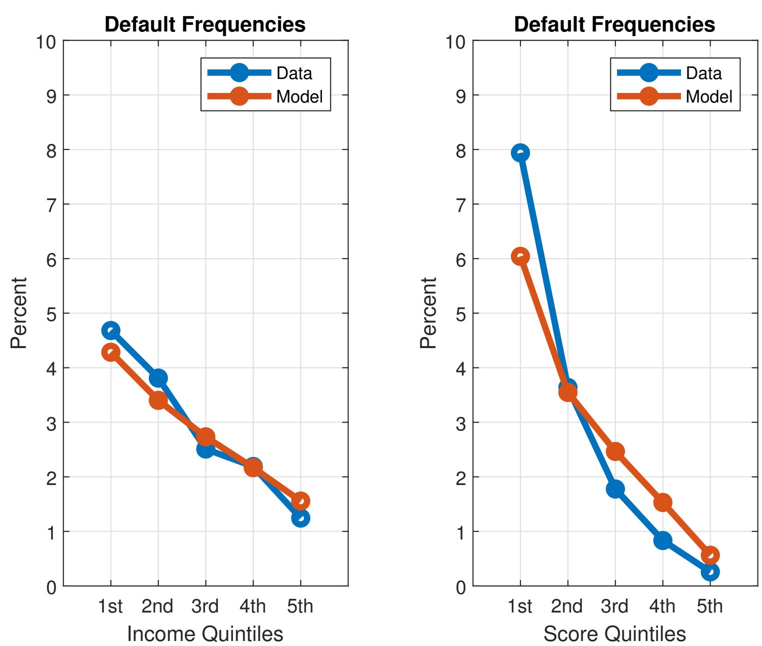

Figure 5

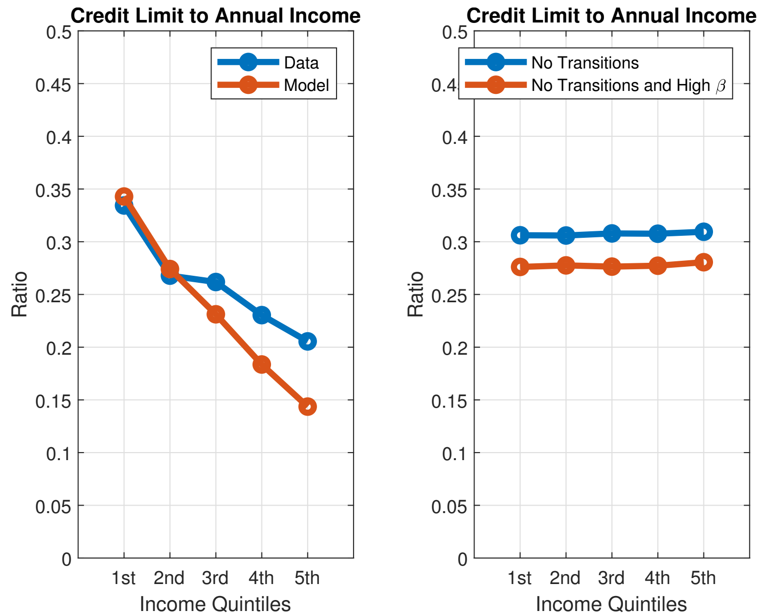

(the estimated value of β is 0.7153, annualized). This helps to get the default frequency in the bottom quintile closer to the observed value, but at the expense of overpredicting default frequencies for a majority of the sample.

We turn next to the relationship between income and credit-limit-to-income ratios. The left panel of Figure 5 shows that the model can capture the the decline in this ratio with income but overpredicts the drop for the top three income quintiles. Why does the credit limit to income ratio decline with income in the model? There are two reasons. First is the nature of the card contract which commits to a credit limit and an interest rate, regardless of future income states. To see why this matters, consider the case where there is no transition between income states, i.e., a person's persistent income is actually permanent. Since support of the (uniform) distribution of the m-shock is proportional to y, ϕ_y^i depends on the ranking of y rather than y itself, and the utility function is homogeneous (of degree $1-\sigma$), the individuals decision problem is homogeneous of degree 0 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 in our setup because of the fixed cost of posting contracts, but it is almost true. Consequently, the equilibrium of this model features credit card contracts that scale with y. The right panel of Figure 5 displays the rough constancy of the model-implied credit-limit-to-income ratios (labeled "No Transitions"). But with positive probability of transitions, a uniform credit limit *ratio* is likely to be very sub-optimal for card companies. If the credit limit-

Figure 6

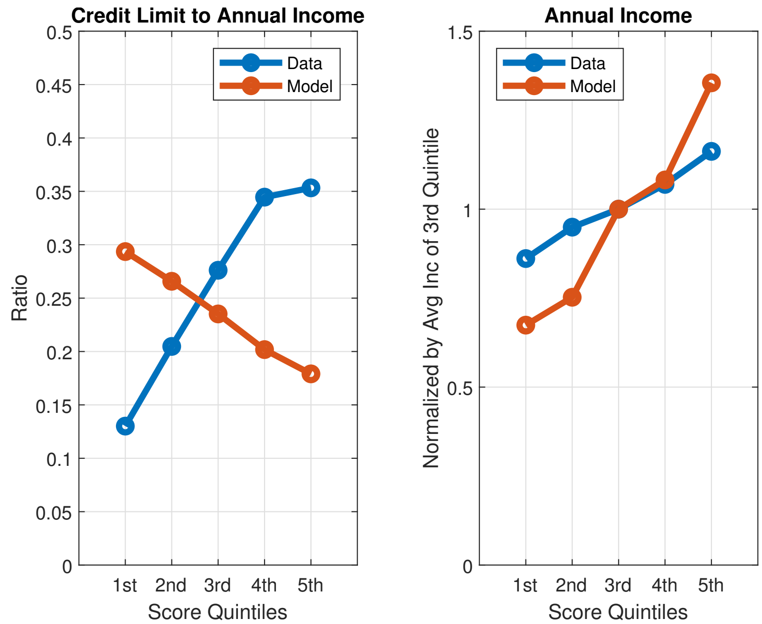

to-income ratio was identical across persistent income levels, an individual holding a credit card designed for a person with, say, double her income would have a strong incentive to default: Her expected default cost would be half as large and she might be tempted to borrow all the way to her large limit and default. In other words, 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 the credit limit of high-income individuals.

A second reason is that high-income individuals know their incomes can fall, which makes them act more patiently: They tend to save or borrow less. Consequently, they prefer cards with a lower credit limits and lower interest rates.[15] One way to see the role of patience on contract terms is to examine the impact of greater patience on credit limits. The left panel of Figure 5 also displays the credit limit ratios if the discount factor is raised in a model where persistent income levels are permanent. Observe that ratios are systematically lower. In sum, these two reasons work to make credit limits rise less than proportionately with income and, hence, for credit limit-to-income ratio to fall with income.

We now turn to the relationship of credit access to credit scores. The left panel of Figure 6 reports the model-implied relationship between the credit-limit-to-income ratio and scores with

[15]The situation here is similar to that in insurance: the good risks prefer contracts that have high deductibles and low premium.

the relationship in 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*. As explained above, the credit limit ratio is a decreasing function of y since card companies are constrained on how generous a credit limit they can offer to high-income individuals when there is a chance that an individual's income might fall in the future. If individuals can differ only in their incomes, high-income individuals will have low probability of default and high credit scores. The combination of these two features implies that in the model the credit limit-to-income ratio *must* decline with credit scores.

The most plausible explanation for this mismatch is that there are factors other than income that also affects credit scores and these factors are not strongly correlated with income. This conclusion is reinforced if we compare the average (relative) incomes in the model score bins with the data. In the model, there is a very sharp rise in average incomes across the score quintiles. In light of our discussion of how the default frequency is affected by a person's permanent component of income, this is to be expected: people with low scores and high default probability must be people with low y and those with high scores and low default probability must be people with high y. In contrast, the link between credit scores and income at origination is weaker in the data. Incomes do rise with score quintiles but the rise is not as steep: The average income of individuals in the highest score quintile is only 1.35 times the average income of people in the bottom quintile, whereas in the model this factor is 2.1.

To summarize: With income as the only source of heterogeneity, the model can match the relationships between income-at-origination and the frequency of default and between income-at-origination and the credit-limit-to-income ratios reasonably well. But it fails to match the observed relationships of these variables to credit score quintiles. The lesson is that there is some other dimension of heterogeneity that make people good or bad risks besides their incomes.

7 Expanding the Model to Have Type Heterogeneity

In the context of our model, the dimensions of heterogeneity that are most natural are discount factors and default costs. The former is a key determinant of default frequency (impatience leads to debt and the possibility of default) and the latter is a key determinant of debt capacity and, hence credit limits. Our framework allows types to vary by β and ϕ. Since the earnings process does not vary by type, types and incomes are uncorrelated.

In what follows, we explore two models: One in which both β and ϕ are allowed to vary across types — called the "Baseline" model — and the other in which only ϕ is allowed to vary across types — called the "Default Cost Only" model (henceforth the DCO model). The reason for presenting both sets of results is to show that while the Baseline model naturally gives a better explanation of the data, the Default Cost Only model is a surprisingly close runner up. Thus, the heterogeneity that seems to matter the most is heterogeneity in the costs of default.[16]

7.1 Parameter Estimates

We postulate 3 types of equal measure. The assumption of equal measure of types is not by itself restrictive since any joint distribution of discount factors and default costs can be approximated by appropriate choices of (β^i, ϕ^i) if N is chosen large enough. Of course, the fact that $N = 3$ is restrictive, but having more than 3 types makes the estimation of the model quite time consuming. It turns out that allowing ζ to be non-zero does not affect the estimation results much at all, so to cut down on computation time we set $\zeta = 0$.

The estimated parameters for the two models are given in Table 3. People are estimated to be quite impatient. In the Baseline model, the most impatient type (type 1) discount the future at a 14.7 percent quarterly rate, while the least impatient type (type 3) discount at a 4.4 percent quarterly rate. The costs of default are estimated to be positively related to patience: The most impatient types pay 41 percent of their transitory income to creditors upon default, whereas the least impatient type forfeit 105 percent. In the DCO model, people discount at an intermediate value of 9.2 percent quarterly rate and the variation in default costs has roughly the same range as in the Baseline model.

7.2 Model Fit

Figure 7 displays the closely targeted data moments and their model counterparts. The left panel shows that both models come quite close to explaining the pattern in default frequencies with respect to income quintiles and do so about equally well. The right panel shows that the Baseline model explains the pattern in default frequencies with respect to score quintiles almost perfectly.

[16]Heterogeneity in discount factors alone cannot explain the facts, the main difficulty being that it cannot account for the positive association between credit scores quintiles and credit limit to income ratios. The logic of the model is that, all else constant, *more* patient individuals want cards with *smaller* limits. And since more patient individuals are also the ones with higher credit scores, the equilibrium relationship between scores and credit-limit-to-income ratios remains negative as in the workhorse model.

Table 3:
Estimate of Type Parameters

Type	Baseline		Def Cost Only	
i	β^i	$(1-\phi^i)$	β	$(1-\phi^i)$
1	0.8717	0.4102	0.9156	0.3745
2	0.9469	0.9414	0.9156	0.7339
3	0.9580	1.0500	0.9156	1.0504

Notes:

Figure 7

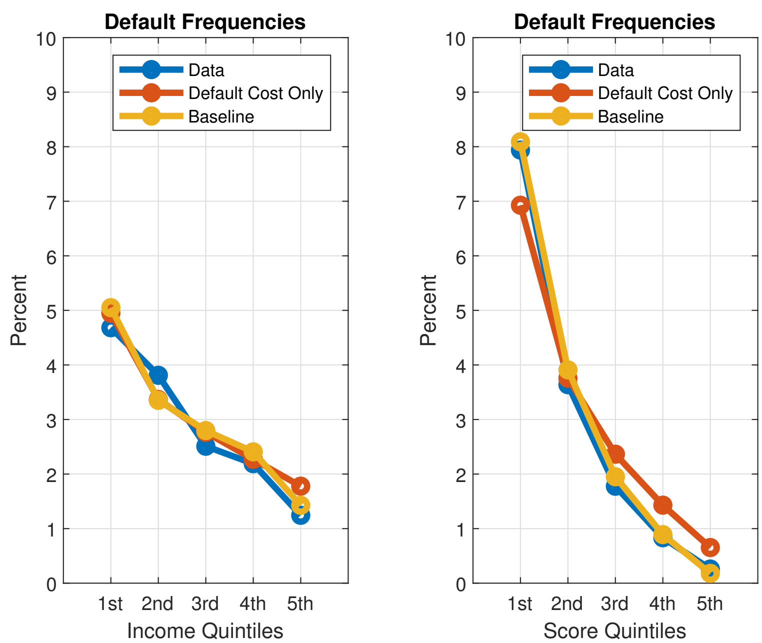

The Default Cost Only model underpredicts the default frequency in the bottom quintile and over predicts it for the top three quintile, a pattern reminiscent of the income heterogeneity only model.

Why does heterogeneity in discount factors help to fit the pattern with respect to credit score better? The heterogeneity in discount factor results in heterogeneity in default frequency. If some highest discount factor individuals are searching for and acquiring new cards, they will likely show up in the top quintiles and help reduce the default frequencies of those quintiles. Combined with the heterogeneity in default frequencies that come from differences in income (high-income individuals are less likely to default), discount factor heterogeneity helps produce the *sharply* declining pattern of default frequencies with score quintiles.

Figure 8

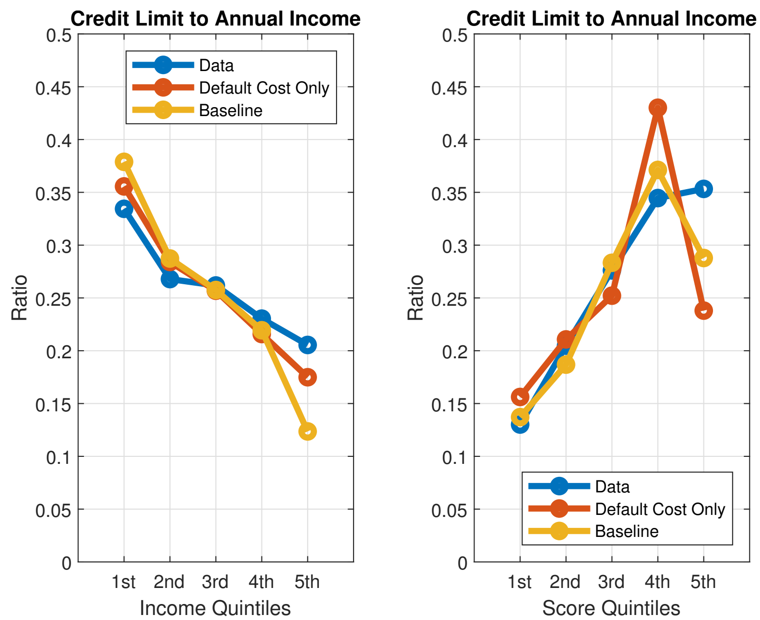

To verify this, the top panel of Table 4 reports the distribution of the three types across score quintiles among the population of people who acquire new cards in any period. The most impatient type (Type 1) is mostly concentrated in the bottom two quintiles and the least impatient type (Type 3) is mostly concentrated in the top two quintiles. Those with the middle level of patience are mostly concentrated in the top three quintiles.[17]

Figure 8 turns to credit-limit-to-income ratios. The left panel shows that both models can get the decline in this ratio with respect to income quintiles. As shown earlier, heterogeneity with respect to income alone can get this pattern but type heterogeneity significantly improves the fit. The DCO model does better as the over- and under-prediction for the bottom and top quintiles, respectively, is less than in the Baseline model.

Why does the Baseline perform worse than the DCO model? In the Baseline, the most impatient type has a substantially lower discount factor than households in the DCO model and the most patient type has a substantially higher discount factor. To accommodate the greater (lesser) desire of the most (least) impatient to borrow, the equilibrium contract features higher (lower) credit limit to income ratio for Type 1 (Type 3). Since there is not much difference in the repre-

[17]The table also reports the prevalence of the three types among the people who are getting new cards. While there is an equal measure of the three types in the economy, the most impatient type is slightly over-represented and the least impatient type is slightly under-presented. The reason is 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.

sentation of the three types across income bins — as shown for the Baseline model in the bottom panel of Table 4 — the average credit limit to income ratio for the bottom (top) quintile is higher (lower) than in the DCO model. Of course, the variation in discount factors was brought in to increase the fit between the model and data with respect to default frequencies. Evidently, this comes at the expense of a poorer fit with respect to the pattern in credit-limit-to-income ratio for income quintiles.

The left panel of Figure 4 shows the pattern with respect to score quintiles. Unlike the income heterogeneity only model, which failed to get the positive association, both models can account for the positive relationship between the credit-limit-to-income ratio and the first four quintiles of credit scores, with the Baseline model showing a closer fit. But both models predict a *decline* in the ratio for the top quintile, with the DCO model predicting a bigger drop than the Baseline model.

Why does the Baseline model get the association between credit-limit-to-income ratio and score quintiles mostly correct? Higher default costs generally support higher credit limits. Since more patient people are also the people with higher default costs, and more patient people are more represented in the higher score quintiles, the average credit-limit-to-income ratio is increasing in score quintiles. However, a countervailing force is that, all else the same, the more patient prefer cards with lower limits and lower interest rates as their need for credit is lower. For the top quintile, the countervailing force wins out and the credit-limit-to-income ratio actually declines (relative to the fourth quintile) rather than increasing slightly as in the data.

Table 4:
Type Distribution of New Originations

	I	II	III	IV	V	Fraction
Types	Score Quintiles					
1	0.48	0.33	0.14	0.05	0.00	0.39
2	0.04	0.17	0.29	0.24	0.26	0.32
3	0.00	0.05	0.18	0.36	0.40	0.28
	Income Quintiles					
1	0.23	0.19	0.18	0.19	0.21	0.39
2	0.20	0.20	0.20	0.20	0.20	0.32
3	0.15	0.21	0.23	0.22	0.18	0.28

Notes: Model outcomes.

7.3 Equilibrium Contracts

In the Baseline model, we permit card companies to distinguish between 6 different asset levels for each type, and we permit 3 (persistent) income levels. Since there are 3 types, the total number of distinct contracts offered is 54. Tables 5 displays the (gross) interest rate and the credit-limit-to-income ratio for each of these contracts.

Focusing on the contract offered to individuals with zero assets, the top panel shows that holding fixed type, card interest rate is increasing in income and, holding fixed income, it is decreasing in type, meaning that more patient types are offered lower interest rates. The bottom panel shows that holding fixed type, the credit-limit-to-income ratio is decreasing in income, and holding fixed income, it is increasing in type. Both panels show that holding fixed type and income, contract terms are generally insensitive to financial wealth, if it is positive. If terms change, higher financial wealth is generally associated with lower interest rates and higher credit limit ratios.

Holding fixed income, the pattern across types is intuitive: Since default costs and patience are both increasing in type, the former pushes credit limits to increase with types and the latter pushes interest rates to decrease with types.

Holding fixed type, the pattern with respect to income results from the same forces that were discussed for the model with a single type: Higher income increases credit limits but less than in

Table 5:
Equilibrium Contracts

Gross Interest Rates, Annualized

Type 1				Type 2				Type 3			
$a \downarrow y \rightarrow$	0.50	1.00	1.98	$a \downarrow y \rightarrow$	0.50	1.00	1.98	$a \downarrow y \rightarrow$	0.50	1.00	1.98
0	1.23	1.25	1.24	0	1.11	1.12	1.13	0	1.06	1.06	1.06
0.45	1.21	1.20	1.21	0.85	1.11	1.11	1.12	1.18	1.05	1.05	1.05
1.05	1.20	1.19	1.21	2.03	1.10	1.11	1.12	2.82	1.05	1.06	1.05
1.82	1.19	1.19	1.21	3.43	1.10	1.11	1.12	4.84	1.05	1.05	1.05
2.54	1.19	1.19	1.21	4.87	1.11	1.12	1.12	6.76	1.05	1.05	1.05
3.30	1.19	1.19	1.21	6.32	1.11	1.11	1.12	8.80	1.05	1.05	1.05

Credit-Limit-to-Income Ratios

Type 1				Type 2				Type 3			
$a \downarrow y \rightarrow$	0.50	1.00	1.98	$a \downarrow y \rightarrow$	0.50	1.00	1.98	$a \downarrow y \rightarrow$	0.50	1.00	1.98
0	0.57	0.36	0.22	0	1.82	1.08	0.57	0	2.93	1.68	0.86
0.45	0.62	0.38	0.23	0.85	1.88	1.07	0.57	1.18	2.96	1.64	0.83
1.05	0.63	0.38	0.22	2.03	1.89	1.06	0.56	2.82	2.96	1.64	0.81
1.82	0.64	0.38	0.22	3.43	1.90	1.06	0.56	4.84	2.98	1.59	0.79
2.54	0.64	0.38	0.22	4.87	1.94	1.06	0.55	6.76	2.98	1.58	0.78
3.30	0.65	0.37	0.22	6.32	1.94	1.04	0.54	8.80	2.98	1.53	0.78

Notes: a and y refer to asset and income levels at origination. In the model, credit limits are negative numbers since debt is a negative asset position. For ease of exposition, the credit limit ratios are reported as positive numbers in this table.

proportion to income and so the credit limit ratio declines with income. It is more surprising that interest rates *rise* with income, given type. This is because higher income individuals have higher limits and, so, when they default they do so on larger debts. This is reflected in the interest rate on the card. This point is discussed in greater detail in the next section.

The lack of sensitivity of contract terms to wealth is a consequence of the long-run behavior of consumers, given income and type, being independent of initial wealth levels (due to impatience). Since a card is long-duration contract, its performance is heavily influenced by long-run behavior of cardholders and, consequently, its terms are not very sensitive to initial wealth levels.

8 Behavior of Card Holders

The market arrangement in our model resembles an Aiyagari-style model with a borrowing constraint and idiosyncratic income shocks. However, there are also important differences. Since people can default on their debts, the borrowing rate exceeds the risk-free saving rate. And, the $(\underline{a}, R)$ pair that a type i has access to at the current time is history-dependent as it is a function of her (a, y) at the time she accepted the contract. This history dependence introduces a new layer of heterogeneity: Two individuals with the same $\{i, a, y\}$ could behave differently because they are operating with different credit card contracts.

8.1 Precautionary Savings and Utilization Rate of Credit Cards

Given the discount factors of each type and the borrowing interest rates they face, every type would want to borrow if there was no income uncertainty. By this we mean that the product of the highest gross borrowing interest a type can face and the type's discount factor is always strictly less than 1. Thus, if incomes were constant, every type would have an incentive to borrow. But uncertainty in incomes provides a reason to accumulate precautionary balances.

The top panel of Table 6 reports the average steady-state assets of different types with different persistent income levels, as a proportion of median income of the economy (which is 1). When persistent income level is at the lowest level, all types are, on average, are indebted. But for each type, the average asset position is increasing in y, as the utility cost of accumulating precautionary balances declines.

For a given persistent income level, the variation with respect to type depends on the income level. For the lowest income level, increase in type implies *more*, not less, indebtedness, on average. This is a consequence of the fact that, holding income constant, higher types are offered contracts with higher credit limits and lower interest rates as they have higher default costs and are more patient.[18] When the individual has the lowest income level, she expects her income to rise over time, on average, and so has a preference to borrow against the future. Higher types can indulge this preference more effectively and, so, they are more indebted.

In contrast, the average assets of the middle income individuals is not monotonic in type. For these individuals, the preference to borrow is weaker as they do not expect rising incomes, on average. Higher types are more patient and, all else constant, will accumulate more precautionary savings. But higher types are also offered contracts with a higher credit limit and lower interest rates (see Table 5) which works to blunt the precautionary savings effect. Going from Type 1 to Type 2, the precautionary savings effect dominates and the average savings of middle income individuals go up. But, going from Type 2 to Type 3, the borrowing effect dominates and average savings decline. This despite the fact that high income individuals expect their incomes to fall, on average.

The bottom panel of Table 6 reports the steady state average utilization rate on cards across (persistent) incomes 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 decreasing in income.

For a given type, the utilization rate is consistently (and strongly) decreasing in persistent income levels: As income rises, individuals borrow less frequently and when they do borrow they borrow smaller amounts.

For a given income level, the utilization rate is not monotonic. For the lowest income level, the utilization rate is rising in type, which mirrors the results for average assets shown in the top panel. For middle income individuals, the utilization rate varies in a way that also mirrors the behavior of average assets. What is more striking, though, is that the utilization does not vary much across types, despite substantial variation in the credit-limit-income ratio with types. For

[18] 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.

Table 6:
Average Assets and Utilization Rates

Income	Type 1	Type 2	Type 3
Average Assets			
0.50	−0.01	−0.30	−0.75
1.00	0.22	0.39	0.25
1.98	0.76	2.15	2.82
Average Utilization Rates			
0.50	0.36	0.48	0.59
1.00	0.20	0.18	0.24
1.98	0.11	0.05	0.05

Notes: Average assets of a cell is the mean assets of the cell expressed as a proportion of median income, which is 1. The Average Utilization Rate of a cell is the mean utilization rate of a cell, where individuals with positive assets have a utilization rate of zero.

high income individuals, the utilization rate is flat for Type 2 and Type 3 but more than double this rate for the Type 1 individual.

8.2 Default

In the event of a default, the cardholder's debt is erased, she pays $(1-\phi^i)$ of the excess of m over $-\lambda y$ to the card company, and she is excluded for some random length of time from the credit card market. Hence her consumption in the period of default is $y - \lambda y + \phi^i \cdot [m + \lambda y]$. If $m = -\lambda y$, the term in square brackets is zero and she does not pay anything to the card company; if $m = \lambda y$, she pays the maximum possible (given y) which is $(1-\phi^i)2\lambda y$; for intermediate values of m, her cost is somewhere in between.[19]

Consider a type 2 individual with low y who holds a card. If this individual has debt and draws a low m, she has can buffer her consumption against the low m by shedding her debt at a low cost, i.e., defaulting. On the other hand, if her debt is well below her card limit, she can also buffer her consumption against the low m by borrowing more. We would expect default to be more likely in situations where the second option is unavailable or its scope is limited because her utilization rate is close to 1 or at 1.

Figure 9:
Model Default Probability

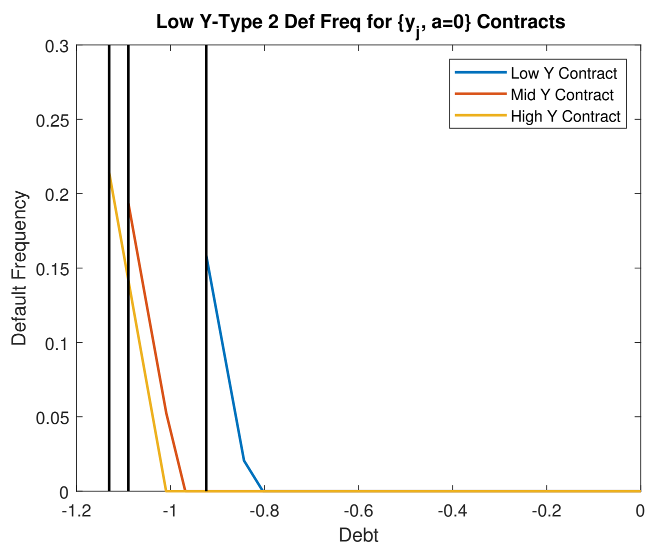

[19]We could model a fixed cost of default that depends on y but there are two drawbacks to this. The reason people default for low m is that the default cost is close to zero for low m. Once a fixed cost is added, it will make it harder induce high default rates without making individuals even more impatient. Another possibility is a utility cost of default that, for tractability, has an extreme value distribution. But this has the complication that it can distort borrowing behavior: some people will borrow just so they can get the potential utility benefit of default. Even if we ignore this, we will have to introduce some income cost of default anyway in order to account for recoveries which are important for delinquencies.

Figure 9 plots a low-y-type-2 individual's default frequency at different debt levels for three different contracts she might hold in equilibrium — contracts offered to type 2 individuals with low, medium or high income and little or no assets at origination.[20] The solid vertical lines are drawn at the credit limit of each of these three contracts. Observe that for each type of contract, default frequency is zero until the debt level gets close to the credit limit. Thus, the individual never defaults unless her utilization rate is close to 1. Evidently, if the individual can borrow on her card, she prefers that over defaulting. The plot would look similar if we focused on Type 1 or Type 3 individuals.

Looking across contracts, the default frequency is increasing in the credit limit but remains well below 1. That said, it is true that an individual's debt level *per se* is not informative about the likelihood of default: an individual with low y will never default on a debt level of 1 (the point -1 on the x-axis) if she has the mid or high y contract, but she may default on debt levels *lower* than 1 if she has the low y contract. What matters for default is the cardholder's proximity to the credit limit, i.e., her utilization rate.

Figure 10:
Model Default Probability Across Incomes

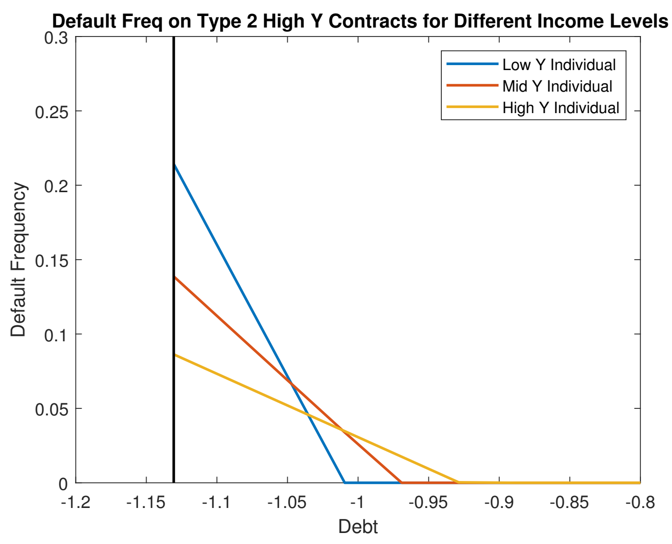

Figure 10 compares the default behavior a type 2 individual who holds a card offered to individuals with high income and zero assets when her income level is low, medium or high. Observe that the set of debt levels on which a low y individual shows positive probability of default is a

[20]As noted earlier, the contracts do not depend much on the level of assets at origination.

Table 7:
Share Defaulting By Income and Type

Income, y	Types			Sum
	1	2	3	
0.50	0.29	0.21	0.12	0.62
1.00	0.19	0.09	0.04	0.32
1.98	0.05	0.01	0.00	0.06
Sum	0.53	0.31	0.16	1.00

strict *subset* of the debt levels on which the mid y individuals shows a positive probability of default and that set, in turn, is a strict subset of the set for which a high y individual shows positive probability of default. In this sense, the low y individual is the most reluctant to default. On the other hand, default probabilities rise fastest for the low y individuals and the slowest for high y individuals.

What explains these patterns? A card with a limit meant for a high y individual is very valuable to the low y individual, since she is unlikely to get such generous terms on her new card subsequent to a default. This makes her reluctant to default. On the other hand, a given debt level is more onerous to service if (mean) income is low versus when it is high. On this account, a low y individual is more likely to default than a high y individual. The pattern of default probabilities shown in Figure 10 mixes these two forces. Similar patterns emerge if the graphs for type 1 or type 3 individuals are plotted.

Table 7 displays the incidence of default across income and types in steady state. Default is most prevalent among people with low income (last column) and among low types (last row). That said, the differences between types 2 and 3 are not pronounced for the middle and high income groups. In contrast, for each type, the incidence of default drops quickly with income.

9 Behavior of Card Companies

9.1 Spreads

A card company chooses the interest rate and credit limit on a contract offered to a person of a particular type, income and wealth level. Given these terms, it extends loans on demand until the contract expires because of a default or because the individual gives up the card to search for a better one. In choosing these contract terms, a card company takes as given the outside option of the individual which, in turn, depends on the contract terms being offered by other card companies and how quickly the individual can expect to get another offer.

In a directed search market, the market tightness adjusts to make the ex-ante value from a contract always equal to the individual's outside option. Given R, this equality defines a locus of (q, L) pairs that are equally attractive to the individual but generates varying levels of ex-ante profits for the card company. Starting from (q, L) combination that is on the locus, increasing L can increase ex-ante profits for two reasons: A higher L generates more revenues from borrowing and hence a higher Π and higher L will also increase the value the individual places on the card which then generates a higher q. However, increases in L eventually decrease Π because loss conditional on default rises with L: recall that at the point of default utilization rate on the card is at 1 or close to 1. Thus, for a given R there is a value of L that maximizes ex-ante profits while giving individuals their outside option. This is the content of the first-order optimality condition with respect to L.

This logic helps us understand why credit-limit-to-income ratios are increasing in types, as seen in Table 5. For card companies, the only cost of increasing L is the greater loss incurred in default. All else the same, the likelihood of default on a card with limit L will be lower for more patient types since their default costs are higher and, therefore, it is optimal to offer a higher credit limit to a more patient type.

Similarly, given L, there is a locus of (q, R) combinations that give the individual her outside option. Starting on the locus and increasing R can increase Π if the increase in R does not discourage borrowing too much. Of course, a higher R is welfare-reducing for borrowers which means that market tightness must fall to keep the individual indifferent with respect to her outside option. So, the increase in Π will be partially offset by a decrease in q. Eventually, though, increases in R will discourage borrowing and both Π and q will decline and, hence, ex-ante profits will decline as well. Thus given L, there will be a value of R that maximizes ex-ante profits. This is the content of the first-order condition with respect to R.

It is possible that for some market the best (L, R) combination implies negative expected net profits. In that case, the company will not offer any contracts in that market.

This first-order logic does not readily illuminate the pattern of interest rates we see in Table 5. To better understand these patterns, it is helpful to consider the following stylized two-period example. Imagine a card company has committed to a credit limit L and an interest rate R. Assume that the opportunity cost of its fund is the (gross) risk-free R_f. The card company expects the cardholder to borrow l in the first period. In the second period, there are two possibilities. In the good state of the world, which occurs with probability p, the card-holder repays Rl. In the bad state of the world, which occurs with probability $1-p$, the card-holder borrows all the way to the limit, defaults, and pays nothing on the defaulted debt. The expected (second-period) profits of the card company, Π, can be expressed as $\Pi = -lR_f + p \cdot Rl - (1-p)[L - Rl]$. This identity can be written as

$$[R - R_f]\frac{l}{L} = (1-p) + \frac{\Pi}{L}.$$

The l.h.s of this identity is the product of the contract spread $R - R_f$ (the excess of the contract rate over the risk-free rate) and the utilization rate l/L. The r.h.s. is sum of the probability of default $(1-p)$ and the profit rate expressed in units of loan commitment L. For some given competitively determined profit rate, this identity shows that the contract spread depends negatively on the utilization rate l/L and positively on the probability of default $(1-p)$ The negative dependence on the utilization arises because l is what is repaid with interest and lower l is the higher must the spread be to generate the same profit rate.

This logic explains a puzzling aspect of contract interest rates. For the moment assume that $\Pi = 0$. Then the contract spread that solves the above equation is the breakeven (zero-profit) contract spread. If the utilization rate l/L less than 1, the breakeven contract spread $[R - R_f]$ will *exceed* the probability of default. This divergence occurs because repayment occurs on small debts l but default occurs on large debts L. In the macro consumer default literature, the typical market arrangement is such that l is always equal to L and the utilization rate is 1. In this case, the breakeven contract spread is exactly equal to the default probability (assuming risk-neutral lenders, of course). But with a credit card contract, the card company must factor in the fact that repayments occur on debts that are likely to be smaller than the debts on which defaults occur. Since the same interest rate applies to all debts, the breakeven contract spread will need to exceed the default probability and the amount by which it needs to exceed it will be larger the smaller is the utilization rate (i.e., the smaller is the debts on which repayments occur). Figure 11 illustrates the gap between contract spreads and default probabilities for the contract offered to Type 1, low y individuals with zero assets.

Figure 11:
Spreads and Default Probabilities

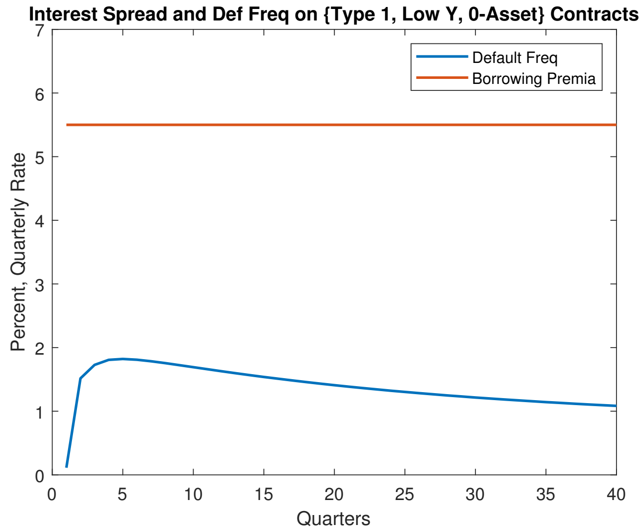

Of course, if there are monopoly profits, i.e., $\Pi/L > 0$, contract spreads will be higher still. Thus, the "loan commitment" nature of a credit card contract, as well as monopoly power, play a role in elevating credit card interest rates above default probabilities.

9.2 Contract Cash Flows and Profitability

How much of the contract spread results from monopoly power? Table 8 gives one measure of this. For each contract, it reports the excess, in proportionate terms, of the equilibrium contract interest rates over the interest rate that would give zero profits to the card company, holding the terms for all the other contracts fixed. For the bulk of the contracts, the difference is less than 5 percent, suggesting modest monopoly power. In a few cases, the difference is 10 percent or larger. This is true for the contracts offered to Type 1 individuals with zero assets. These individuals really wish to borrow, so their outside option of

We turn now to examining how the baseline model performs with respect to interest rate on credit card contracts. Figure 15 displays the APR on (newly originated) card by income and score quintiles for both the Baseline and DCO model. As noted earlier, interest rates decline with credit scores but the decline is not as sharp as the decline in default frequencies. The model counterparts for the bottom two quintiles are close to the data, but there is substantial discrepancies between model and data for the top three quintiles: model interest rates drop off much more steeply than

Table 8:
Excess of Interest Rate Over Zero-Profit Interest Rates

Type 1				Type 2				Type 3			
$a \downarrow y \rightarrow$	0.50	1.00	1.98	$a \downarrow y \rightarrow$	0.50	1.00	1.98	$a \downarrow y \rightarrow$	0.50	1.00	1.98
0	0.11	0.10	0.08	0	0.06	0.05	0.07	0	0.06	0.04	0.06
0.45	0.06	0.05	0.05	0.85	0.03	0.02	0.03	1.18	0.03	0.02	0.10
1.05	0.05	0.04	0.03	2.03	0.02	0.02	0.03	2.82	0.03	0.02	0.05
1.82	0.05	0.04	0.03	3.43	0.02	0.02	0.03	4.84	0.03	0.03	0.11
2.54	0.05	0.03	0.03	4.87	0.02	0.04	0.03	6.76	0.03	0.03	0.36
3.30	0.05	0.03	0.03	6.32	0.02	0.16	0.03	8.80	0.03	0.06	0.38

Notes: a and y refer to asset and income levels at origination. Each cell reports (Gross Contract Interest Rates - Gross Zero Profit Interest)/(Gross Contract Interest Rate -1).

Figure 12:

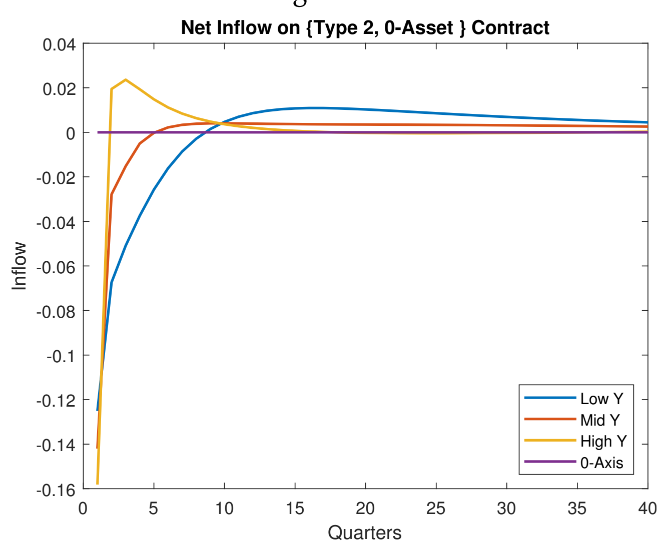

Notes: All contracts are for zero assets at origination.

Figure 13:

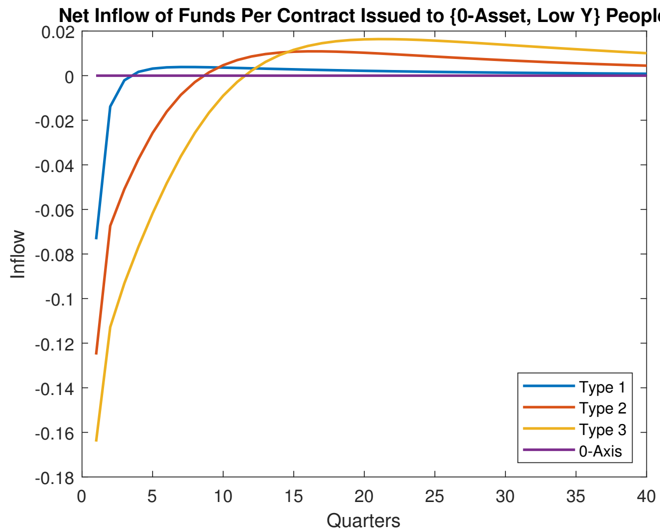

Notes: All contracts are for zero assets at origination.

in the data. This discrepancy is then reflected in utilization rates which are much higher in the model for the top 3 quintiles than in the data.

Figures 15 and 14 display the contract interest rates and utilization rates for the model and the data. The fit between the observed patterns and either model is quite imperfect. That said, both models do a reasonable job of accounting for the utilization rates with respect to income quintiles. With respect to score quintiles, the Baseline model can get the sharp decline in utilization rates, but underpredicts the utilization rate for the bottom two score quintiles and overpredicts for the next two. The performance of the DCO model is quite poor in this dimension.

Turning to contract interest rates, the Baseline generally underpredicts APR while the DCO over predicts it. One of the key differences between the two models is that Types 2 and 3 in the Baseline are substantially more patient people in the DCO model. Yet, the variation in utilization rates is closer to the data for the Baseline model. The model accomplishes this by lowering the APR.

However, it is worth noting that very high APRs are not outside the scope of credit card models of the type developed in this paper. This is clearly shown by the DCO model: the model implied APRs are even higher than in the data. This point is important because high APRs are often seen as a sign of monopoly power. While card companies do exercise monopoly power in our model,

Figure 14

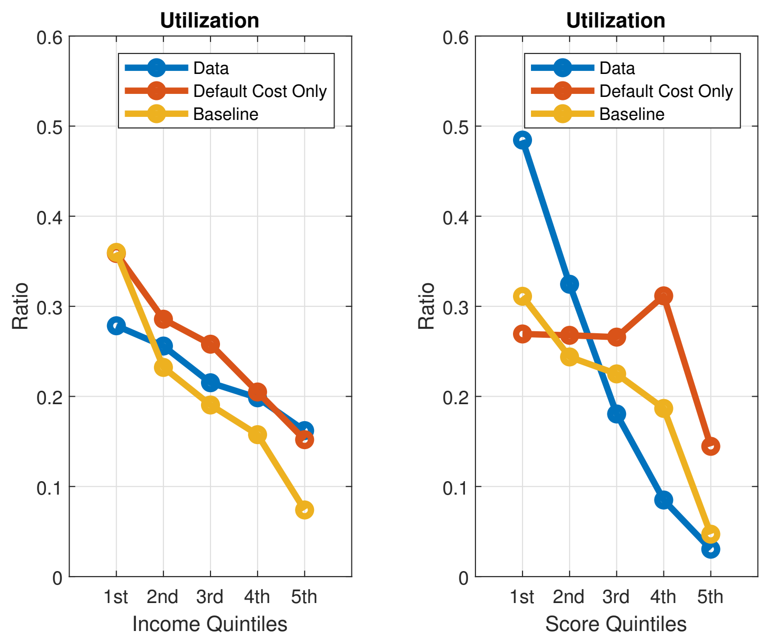

Figure 15

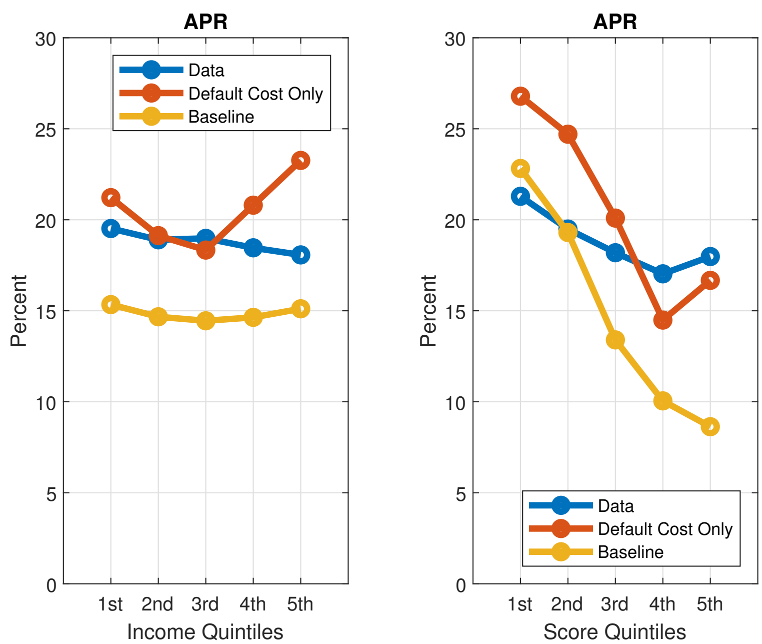

that power is constrained by ex-ante competition in the acquisition stage. Nevertheless, APRs can be quite high.

10 Macro Implications

(Preliminary) We now turn to the macro implications of the credit card industry.

10.1 Distribution of MPC

We examine the implications of credit card borrowing for the distribution of MPC. As is well known, the distribution of MPC has important implications for the effects of certain types of fiscal policies, such as fiscal transfers.

The top panel of Table 9 displays the mean MPC for various cells in the steady state of the Baseline model. For each income level, MPC falls as type increases. This pattern reflects the fact that higher types are more patient and, therefore, allocate a larger portion of a unit increase in wealth toward increasing future consumption via savings or debt reduction. For each type, MPC declines as income rises, as one would expect.

The next panel displays the same information for a model in which the credit card industry does not exist. In this model, people can only save as is typical of heterogeneous agent macro models. Comparing cell by cell, we see that the MPCs are higher in the no-borrowing model. This reflects the fact that without borrowing, people's consumption is more depressed and, so, the MPC is higher. In other words, the existence of the credit card industry is a force in favor of lowering MPCs.

The bottom two panels of Table 9 report MPCs separately for card holders and non card holders. The MPC of card holders is lower than the MPC of non card holders, as we would expect.

Table 9:
MPC

Income	Type 1	Type 2	Type 3
MPC in Baseline, Avg. 0.22			
0.50	0.40	0.22	0.19
1.00	0.36	0.17	0.13
1.98	0.32	0.11	0.08
MPC without a Credit Card Industry, Avg. 0.24			
0.50	0.43	0.28	0.24
1.00	0.37	0.19	0.15
1.98	0.33	0.12	0.08
MPC of Card Holders, Avg. 0.21			
0.50	0.38	0.20	0.17
1.00	0.35	0.16	0.13
1.98	0.32	0.11	0.05
MPC of Individuals Without Cards, Avg. 0.28			
0.50	0.47	0.36	0.34
1.00	0.41	0.22	0.15
1.98	0.35	0.11	0.07

Notes:

11 Conclusion

(To be added)

References

Aguiar, M., M. Bils, and C. Boar (2021): "Who Are The Hand-to-Mouth?," Unpublished.

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

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 Informationa nd Unsecured Credit," *American Economic Journal: Macroeconomics*, 4(3), 153–183.

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 (forthcoming): "A Quantitative Theory of the Credit Score," *Econometrica*.

Chatterjee, S., D. Corbae, M. Nakajima, and J.-V. Rios-Rull (2007): "A Quantitative Theory of Unsecured Consumer Credit with Risk of Default," *Econometrica*, 75(6), 1525–1589.

Dempsey, K., and F. Ionescu (n.d.): "Borrowing Premia in Unsecured Credit Markets," Unpublished.

den Haan, W. J., G. Ramey, and J. Watson (2000): "Job Destruction and Propogation of Shocks," *American Economic Review*, 90(1), 482–498.

Drozd, L., and J. Nosal (2008): "Competing for Customers: A Search Model of the Market for Unsecured Credit," Unpublished.

Eaton, J., and M. Gersovitz (1981): "Debt with Potential Repudiation: Theoretical and Empirical Analysis," *Review of Economic Studies*, 48(2), 289–309.

Exler, F., and M. Tertilt (2020): "Consumer Debt and Default: A Macro Perspective," Unpublished.

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

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

Herkenhoff, K. F., and G. Raveendranathan (2021): "Who Bears the Welfare Cost of Monopoly? The Case of the Credit Card Industry," Unpublished.

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

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

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

Livshits, I., J. C. MacGee, and M. Tertilt (2007): "Consumer Bankruptcy: A Fresh Start," *American Economic Review*, 97(1), 402–418.

Mateos-Planas, X., and J.-V. Ríos-Rull (2007): "Credit Lines," Unpublished.

Meier, S., and C. D. Sprenger (2012): "Time Discounting Predicts Creditworthiness," *Psychological Science*.

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

Appendix A Data Tables

Table 10:
Terms, Utilization and Performance

Description	I	II	III	IV	V
Mean Bin Scores	658	702	743	787	833
APR (in %)	21.29	19.49	18.19	17.03	17.98
Credit Limit to Income (in %)	5.92	9.31	12.55	15.66	16.06
Default Frequency (in %)	7.94	3.64	1.78	0.83	0.26
Debt to Credit Limit (in %)	48.45	32.45	18.05	8.5	3.04
Relative Bin Incomes	0.86	0.95	1.00	1.07	1.16
Mean Bin Incomes	27606	41448	56049	77180	115695
APR (in %)	19.52	18.90	18.98	18.46	18.07
Credit Limit to Income (in %)	15.20	12.18	11.90	10.47	9.34
Default Frequency (in %)	4.68	3.81	2.51	2.19	1.24
Debt to Credit Limit (in %)	27.84	25.59	21.52	19.85	16.21
Relative Bin Incomes	0.49	0.74	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.

Table 11:
Utilization at the Time of Default

Description	Overall	Delinquency	Bankruptcy & Debt Waiver
Util. Rate At time of Default (in %)	84.20	92.60	63.10
Frequency (in %)	2.89	2.06	0.83

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

Table 12:
Attrition Rate of Accounts

Description	Overall
Monthly Attrition Rate (in %)	0.80

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

Appendix B Model Tables

Table 13: **Model with One Type**

Description	I	II	III	IV	V
By Score Quintiles					
APR (in %)	13.77	13.85	14.00	14.37	13.44
Credit Limit to Income (in %)	29.36	26.58	23.52	20.18	17.91
Default Frequency (in %)	6.04	3.55	2.47	1.53	0.56
Debt to Credit Limit (in %)	40.11	31.63	25.01	18.03	10.87
Relative Bin Incomes	0.70	0.83	1.00	1.29	1.47
By Income Quintiles					
APR (in %)	13.64	13.51	13.39	14.22	14.67
Credit Limit to Income (in %)	34.29	27.41	23.21	18.36	14.36
Default Frequency (in %)	4.51	3.42	2.72	2.19	1.52
Debt to Credit Limit (in %)	41.41	31.07	24.33	17.34	11.51
Relative Bin Incomes	0.52	0.75	1.00	1.48	2.05

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.

Table 14: **Baseline Model**

Description	I	II	III	IV	V
By Score Quintiles					
APR (in %)	22.82	19.30	13.40	10.05	8.63
Credit Limit to Income (in %)	13.72	18.69	28.33	37.12	28.78
Default Frequency (in %)	8.09	3.91	1.95	0.89	0.18
Debt to Credit Limit (in %)	31.11	24.38	22.50	18.67	4.71
Relative Bin Incomes	0.62	0.95	1.00	0.95	1.22
By Income Quintiles					
APR (in %)	15.34	14.67	14.45	14.64	15.11
Credit Limit to Income (in %)	37.89	28.73	25.75	21.94	12.36
Default Frequency (in %)	5.05	3.35	2.80	2.41	1.43
Debt to Credit Limit (in %)	35.99	23.22	19.05	15.75	7.40
Relative Bin Incomes	0.50	0.88	1.00	1.28	1.98

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.

Table 15: **Model With Default Cost Heterogeneity Only**

Description	I	II	III	IV	V
By Score Quintiles					
APR (in %)	26.79	24.70	20.09	14.49	16.67
Credit Limit to Income (in %)	15.62	21.07	25.22	43.01	23.78
Default Frequency (in %)	6.93	3.76	2.36	1.43	0.65
Debt to Credit Limit (in %)	26.93	26.78	26.58	31.16	14.46
Relative Bin Incomes	0.58	0.85	1.00	0.96	1.26
By Income Quintiles					
APR (in %)	21.22	19.13	18.34	20.80	23.26
Credit Limit to Income (in %)	35.56	28.40	25.68	21.58	17.49
Default Frequency (in %)	4.95	3.37	2.77	2.27	1.78
Debt to Credit Limit (in %)	35.86	28.57	25.80	20.49	15.12
Relative Bin Incomes	0.50	0.86	1.00	1.49	1.98

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 C Computation of Credit Score

(Preliminary! Does not use the notation of the main text!) Let $j = 1, 2, 3, \ldots, 8$ denote the time period that is j periods in the future ($j = 1$ is the next period). We initialize the credit score with 1 at the end of 8 periods at all the nodes the individual might end up in, and find its current credit score with backward induction. The credit score when the person holds a credit card is given by: Denote probability of death with g

$$k^i_{cc}(a, y; \omega, t) = (1-g) * \begin{pmatrix} \zeta^i(a, y; \omega) * \begin{bmatrix} f * \mu * E_{m,y'} k^i_{cc}(a'(a, y, m; \widetilde{\omega}(a, y)), y'; \widetilde{\omega}(a, y), t+1) \\ +(1 - f * \mu) E_{m,y'} k^i_{ncg}(a'(a, y, m), y'; t+1) \end{bmatrix} \\ +\left(1 - \zeta^i(a, y; \omega)\right) * \left(E_{m,y'}\left[(1 - D(a, y, m)) * k^i_{cc}(a'(a, y, m), y'; \omega, t+1)\right]\right) \end{pmatrix}$$

$$\begin{aligned} k^i_{ncg}(a, y; t) &= g * 1 \\ &(1-g) * \begin{pmatrix} f * \mu * E_{m,y'} k^i_{cc}(a'(a, y, m; \widetilde{\omega}(a, y)), y'; \widetilde{\omega}(a, y), t+1) \\ +(1 - f * \mu) E_{m,y'} k^i_{ncg}(a'(a, y, m), y'; t+1) \end{pmatrix} \end{aligned}$$

$$\begin{aligned} k^i_{ncb}(a, y; t) &= g * 1 \\ &(1-g) * \begin{pmatrix} \delta * f * \mu * E_{m,y'} k^i_{cc}(a'(a, y, m; \widetilde{\omega}(a, y)), y'; \widetilde{\omega}(a, y), t+1) \\ +\delta * (1 - f * \mu) * E_{m,y'} k^i_{ncg}(a'(a, y, m), y'; t+1) \\ +(1-\delta) E_{m,y'} k^i_{ncb}(a'(a, y, m), y'; t+1) \end{pmatrix} \end{aligned}$$

Appendix D Finding the default threshold

(Preliminary! Does not use the notation of the main text!) When the individual defaults, the m-shock reduces to $c(m) = \underline{m} + s * (m - \underline{m})$.

$$c(m) = (1-s)\underline{m} + sm$$

Say, in the case of no default, the country chooses the asset level a_k between the m-thresholds $[m_k, m_{k+1}]$ Now we will need to compare the utility under default and no-default: Denote the utility under repayment with:

$$-\frac{1}{c(a, a_k) + m} + W$$

Denote the utility under default with

$$-\frac{1}{c_d + sm} + W_d$$

where $c_d = c(0,0) + (1-s)\underline{m}$ There are some conditions where there might be a dominant strategy. These are: (1) If $W > W_d$ and $c(a,a_k) + m_k > c_d + sm_k$ then for all $m \in [m_k, m_{k+1}]$ the individual chooses to not-default. (2) If $W_d > W$ and $c_d + sm_{k+1} > c(a,a_k) + m_{k+1}$ then for all $m \in [m_k, m_{k+1}]$ the individual chooses to default. (3) Else, we need to find the two m-thresholds where the two values are identical.

$$-\frac{1}{c(a,a_k)+m} + W = -\frac{1}{c_d + sm} + W_d$$

$$-(c_d + sm) + (W - W_d) * (c(a,a_k) + m) * (c_d + sm) = -(c(a,a_k) + m)$$

$$\frac{c(a,a_k) - c_d}{(W - W_d)} + \frac{(1-s)}{(W - W_d)} m + sm^2 + smc(a,a_k) + c_d m + c_d c(a,a_k) = 0$$

$$sm^2 + \left(sc(a,a_k) + c_d + \frac{(1-s)}{(W - W_d)} \right) m + c_d c(a,a_k) + \frac{c(a,a_k) - c_d}{(W - W_d)} = 0$$

As seen, this is a quadratic equation. If $W > W_d$, then the person prefers to default between these two thresholds, and otherwise does not default. If $W < W_d$, then the person prefers to not-default between the two points, otherwise default. this way, the behavior of the person between $[m_k, m_{k+1}]$ is found.