

Competing for Customers: A Search Model of the Market for Unsecured Credit*

Lukasz A. Drozd[†] Jaromir B. Nosal[‡]

February 2008

Abstract

This paper develops a theory of the market for unsecured credit with market incompleteness closely resembling the key features of the credit card market in the US. A friction of targeting credit card offers to specific customers is introduced and the implications of lowering the strength of this friction is studied. The results indicate that such change – motivated by the rapid progress in information technology during the 90s – is promising to quantitatively account for several observations occurring during this period: (i) growing availability and use of revolving lines of credit, (ii) growing indebtedness of households, (iii) rising filing rate for bankruptcy protection, (iv) rising debt discharged per statistical bankrupt, (v) falling cost of revolving credit measured by the interest rate premium over cost of funds, (vi) massive increase in the credit card solicitations.

*We thank V.V. Chari, Patrick Kehoe and Fabrizio Perri for valuable advice and encouragement. We are also grateful to Manuel Amador, Narayana Kotcherlakota, Erzo Luttmer, Salvador Navarro, and Victor Rios-Rull for their help and valuable comments. We acknowledge and thank for access to the computational resources of the Department of Economics, University of Wisconsin-Madison, and the University of Minnesota Supercomputing Institute. All errors are ours.

[†]University of Wisconsin-Madison, Department of Economics, 1180 Observatory Drive, SSC Bldg., Madison, WI 53706. Contact: ldrozd@ssc.wisc.edu.

[‡]Job Market Candidate: University of Minnesota, Department of Economics, 1035 Heller Hall, 271 19th Avenue South, Minneapolis, MN 55455. Contact: jaromir@econ.umn.edu.

1 Introduction

Over the last few decades we have witnessed significant changes in the market for unsecured consumer credit. There has been a massive increase in both the use and availability of unsecured credit through a new product: revolving lines of credit in the form of credit cards¹. At the same time, we have seen an increase in filing rates for personal bankruptcy protection and the volume of unsecured debt discharged through bankruptcy, both measured relative to income and relative to total unsecured debt held by the public (charge-off rate). Not surprisingly, these changes raised growing concerns in policy and academic circles about the desirability of bankruptcy protection, and questions about the underlying mechanism behind these changes. In this paper, we propose a theory of the unsecured credit market, and attribute these changes to progress in information technology modeled by a falling cost of screening and soliciting credit customers. We use our theory to perform a disciplined comparative statics exercise of changing this cost to quantitatively account for the rise of bankruptcy related statistics.

To accomplish this task, we build a fully dynamic equilibrium model of borrowing and lending in which costs of screening and soliciting customers are explicit, and both market and institutional arrangements closely resemble those seen in the US data. More specifically, banks in our model pay a fixed cost to make targeted loan offers to consumers with afore-chosen characteristics. The loan contracts specify a credit limit and an interest rates to which the bank is committed. The credit card contract is long-lasting, and lasts until the consumer decides to switch to a better deal, default or dies. Consumers in the model face idiosyncratic income/expense shocks, and use their credit lines or savings to smooth consumption. They also have an option to file for bankruptcy protection and default on their debt in the spirit of Chapter 7 of the US bankruptcy law. Access to the unsecured credit, as well as competitiveness of the offered contracts, are endogenous in the sense that in each instance of time consumers stochastically receive an endogenous

¹This product revolutionized consumer finances, by facilitating the use of credit to smooth consumption. Revolving credit lines are particularly suitable for this purpose, as the credit line is constantly available and no interest is charged when it is not being used. In contrast, personal loans typically feature high application fees (\$100-250), and the interest rate is accrued on the credited amount from day one. Credit is also no longer available after the debt has been repaid.

number of offers from banks and select the best one². The key tension of the model is generated by the fact that banks need to finance the fixed cost of acquiring a credit card customer from the revenue on a long-lasting revolving contract specifying a credit limit and an interest rate. The novel force this model introduces is the propagation mechanism through the implications of progress in informational technology for competition for customers and surplus extraction by banks.

Our theory suggests the following explanation of the impact of technological progress on the use of credit and bankruptcy related statistics. The falling cost of acquiring a customer creates a situation of a fiercer competition for a credit card customer among banks, and leads to lower surplus extraction on the offered credit card contract as well as shorter waiting time for the first and subsequent credit card offers. These effects significantly improve access to credit, which becomes available to a broader set of households and on better terms. In particular, access to credit is available on better terms earlier in the life-cycle of households. Since early in the life-cycle households determine the likelihood of their reliance on self-insurance versus bank-provided insurance (borrowing), the use of credit increases drastically. The reliance on bank provided insurance is hardwired to filing for bankruptcy protection, which similarly to other models completes the market by providing a discharge of debts in case of a long spell of negative income shocks. As a result, the bankruptcy filing rate increases in proportion to the increased reliance on borrowing. At the same time increased competition leads to a fall in the average interest rates (net of cost of funds) and expansion of credit limits. Due to the more generous credit limits, more debt is also discharged per statistical bankrupt.

Quantitatively, the model comes close to replicating the current bankruptcy figures in the US, and it demonstrates a potential to account for the changes we have observed during the last two decades. Specifically, for the cost of acquiring a new customer set equal to \$75 – which is consistent with the estimates of the cost of acquiring a new credit card customer for the most recent period³ 2000-2005 – our model predicts that the filings

²Our theory builds on the work by Burdett and Judd (1983). We extend their setup to a dynamic environment.

³The estimated cost of acquiring a new customer comes from Visa U.S.A., see Evans and Schmalensee (2001). In the model we evaluate the fixed cost relative to the average income of a household. When we assume that the average income is \$35000 net of taxes and social security, it implies that this cost

for bankruptcy protection (full discharge) should be around 3 per 1000 households⁴, charge-off rates should be around 7%, and the unsecured debt to income ratio should slightly exceeds 10%. The increase in the cost of acquiring new credit card customers needed to take the model back to 1990s, replicating the debt-to-income ratio of 5%, implies an almost two-fold decrease in the filing rate for bankruptcy protection, higher average interest rates on credit cards (net of charge-off rate and risk free cost of funds), tighter credit limits, less debt discharged per statistical bankrupt (relative to income), and a fall in credit card solicitations.

The main quantitative shortcoming of the theory turns out to be the underpredicted increase in the average debt discharged under bankruptcy between 1990–95 and 2000–05. More specifically, the model predicts that the amount of debt discharged per statistical bankrupt should increase by about 6%, whereas in the data it has increased by as much as⁵ 50%. As a result, even though the model has correct predictions about the growing indebtedness of the households and bankruptcy filing rate, it falls short in terms of generating a concurrent increase in the charge-off rate – which is given by the ratio of the total volume of debt discharged due to bankruptcy (filing rate \times debt discharged per bankrupt) to the total unsecured debt held by the public.

Since in our paper we do not model the deep reasons behind the frictions we introduce, it is important to emphasize the motivation behind our key assumptions. First, the evidence shows that directed solicitations are in fact a dominant method of acquiring new credit card customers. According to the estimates available in the literature, as much as 74% of new accounts are acquired by direct mail/phone solicitations initiated by banks. Furthermore, consistent with the key prediction of our model, these solicitations have been on the rise during the period under consideration⁶ and just like in the model, they have been accompanied by a falling response rate (see Figure 1). In addition, during

would be about \$75.

⁴Since we do not model expense shocks, which actually in the data account for a large fraction of bankruptcy filings, this number is smaller but close to what we see in the data. In PSID only 53% of reported bankruptcy are due to credit misuse or job loss (see Chatterjee et al. (2007)).

⁵See Livshits et al. (2007a), Table 1.

⁶The source of this data is Synovate. Synovate is a global market research company providing commercial data to business. One of their products is mail monitor, informing banks about number of credit card solicitations sent out by banks in a given year, and the response rates to these solicitations. See www.synovate.com.

this time the credit card industry witnessed a massive increase in balance transfers from one credit account to another. According to the estimates by Evans and Schmalensee (2005) (pp. 226) as much as 60% of solicited offers allow for transferring balances at the mere cost of providing information about existing credit cards. In 2002 balance transfers accounted for as much as 17% of total outstanding balances. Also, over the time period 1992-2002 balance transfers were growing at an average rate of 38% annually. Our restriction on the space of admissible contracts to these which specify a credit limit and an interest rate is supported by the data, characterizing the terms on the most important credit cards on the market. Evans and Schmalensee (2005) (pp. 219-223) based on the evidence provided by Visa U.S.A., report that finance charges are the major source of credit card companies revenues, with annual fees⁷ playing at best a minor role⁸.

The literature provides also some evidence that the cost of acquiring new accounts constitutes a significant portion of the overall costs of credit card companies, and that this cost has been falling. Evans and Schmalensee (2005) (pp. 224), based on data provided by Visa U.S.A., report that a large fraction of total costs excluding costs of funds and charge-offs can be attributed to application processing costs. Ausubel (1991) provides additional indirect evidence suggesting that acquiring a credit card customer is a costly activity. Ausubel studies the prices at which credit card accounts are traded between banks. Because purchasing an account is to some extent an alternative to acquiring a new account, the high price at which these accounts are traded indicates that such cost must be substantial. In addition, consistent with our story of technological progress leading to a lower costs of soliciting new customers, Mester (1997) reports that the time needed to process a credit application had fallen from as much as weeks to hours over the last two decades.

Another issue we abstract from in this paper are forms of unsecured credit other than revolving credit lines. In support of this assumption, according to the estimates provided by Livshits, MacGee and Tertilt (2007*a*) for the time period 1990-2005, non-revolving forms of unsecured borrowing played an insignificant role. First, the data shows that for that time period about 80% of the unsecured borrowing took the form of revolving lines

⁷Annual fee is charged for having the card (on monthly or annual basis).

⁸At most 3% of revenue.

of credit. Second, almost all of the increase in unsecured borrowing over that time is attributed to the growth of revolving debt (see Figure 2).

Related Literature

There is a growing theoretical literature studying personal bankruptcy in a macroeconomic setup. Chatterjee, Corbae, Nakajima and Rios-Rull (2007) provide a general equilibrium model of bankruptcy and show that it can be consistent with the level of US bankruptcy rate and debt to income ratio. Livshits, MacGee and Tertilt (2007a, 2007b) in a series of papers propose a model of personal bankruptcy and evaluate the potential of the model to account for the rise of personal bankruptcy in the US. They provide a comprehensive quantitative evaluation of several mechanisms using a life-cycle version of Eaton and Gersovitz (1981) model of borrowing and lending, and they conclude that there is no single explanation that would accomplish the task. Similar findings are reported also by Athreya (2004), who explores the potential of theories to account for the rise. The main difference between this literature and our model is the type of market incompleteness and matching frictions. In our model the contracts are long-lasting with a pre-specified credit limit and interest rate, whereas the literature focuses on contracts that are renegotiated on a period-by-period basis, and abstracts from matching frictions. Related papers that also model endogenous access to unsecured credit in the context of technological progress are Narajabad (2007) and Livshits, MacGee and Tertilt (2007b). Narajabad (2007) explicitly models improvement in informational technology as leading to a higher precision of signals about consumer risk types. This can induce a switch from the equilibrium in which types are pooled and credit is scarce to an equilibrium in which types are separated and credit is abundant (to the low risk types). Livshits, MacGee and Tertilt (2007b) propose a model in which high fixed cost leads to a similar pooling of credit card applications due to high fixed price of posting a contract. The key difference between this work and our model is that while our model assumes no private information and abstracts from sorting effects of better technology, it gives us potential for quantitative work.

In an independent work, Rios-Rull and Mateos-Planas (2007) consider a model of long-term relations between banks and consumers. Contracts are a specified credit limit and an interest rate. Consumers face a fixed utility cost in order to switch from their

current contract to the best available to them on the market. Banks face a fixed cost of issuing a credit line, and they can unilaterally terminate or change the contract. Our way of modeling competition differs in several ways. We assume that banks compete for customers not only by setting the terms of the contract, but also by actively soliciting the customers – which is the predominant way customers are acquired in the data. Consumers, face no switching cost once they are targeted by a bank, and switching is not possible otherwise. We assume bank’s commitment to the terms of the contract.

The rest of the paper is organized as follows. Section 2 presents the data we use to evaluate the model. Section 3 presents the model. Parameterization and results are described in Sections 4 and 5. Section 6 concludes.

2 Key Observations in the Data

The goal of this section is to document three major changes which occurred in the US credit market between the late 1980s and the recent period 2000-2005. We focus here solely on this subperiod because it was a time period with relatively stable bankruptcy regulations, no interest rate caps imposed by usury laws, and with the credit card system already well developed. Data from this section will later be used to evaluate the theory.

Growing Use of Revolving Credit The use of unsecured borrowing in the US, particularly in the form of bank credit cards, boomed in the 1990s. Figure 3, reproduced from Livshits, MacGee and Tertilt (2007*a*), presents the time series for total unsecured consumer credit outstanding relative to disposable income and total revolving consumer debt relative to disposable income. Two features of the data stand out. First, over the 1990s most of the increase in consumer borrowing is accounted for by the increase in the unsecured revolving credit. Second, revolving credit constituted most of the unsecured debt in the US - accounting for about 80% to 90% of total unsecured debt in the early 90s. According to these estimates, unsecured debt of US households rose from about 5 – 6% of disposable income in 1989 to about 9.2% in 1998 (data for a later time period is unavailable due to the fact that FRB discontinued the time series used to construct this estimate).

These general trends are also reflected in the use of credit cards by households (see Table 1). The Survey of Consumer Finances provides detailed characteristics of the use and ownership of credit cards by its respondents. The data shows that in 1989, about 56% of respondents had a bank card with a revolving feature, and 29% actually carried a balance on it. In contrast, in 2004, the fraction of respondents with a bank card rose to 71%, and nearly 40% of population reported that they carry a positive balance. Over the same time period credit limits measured relative to household income had more than doubled.

Falling Price of Credit Card Services In the aggregate data reported by the FRS Board of Governors, the average interest rate charged on credit card balances after adjusting for inflation fell slightly between 1990 and 2005, from an average 13.8% in time period 1990-1995 to an average 10.8% in time period 2000-2005. However, taking into account the soaring charge-off rate, which is actually part of the cost for the credit card issuer, the ex-post real interest rate net of this cost fell drastically. Figure 2 illustrates the two time series of the real interest rate charged on credit cards: (i) net of cost of funds that include the charge-off rate and yield on 3m T-bills, and (ii) net of cost of funds that include chargeoff rate and yield on Moody's seasoned AAA all industry corporate bonds. As we can see from the figure, both series exhibit a significant downward sloping trend suggesting that the interest rate premium charged on credit accounts fell drastically. Using different data and measurement, Evans and Schmalensee (2005) (pp. 234) arrive at a similar conclusion. Figure 2, and Tables 1, 2 summarize the key facts discussed above.

Table 1. Credit card use and credit limits

Year	1989	2004
Population with a bank card	56%	71%
Population with positive (revolving) balance on a card	29%	39%
Average credit limit to income	0.2	0.46

Soaring Bankruptcy Statistics Figure 4 presents Chapter 7 bankruptcy filings per 1000 adults⁹, and the charge-off rate¹⁰. We can see the filing rate nearly doubled during the 1990s, going from 3.08 in 1990 to 5.75 in 2004, and charge-off rate rose from 3.46% in 1990 to 4.86% in 2004.

A closer look at the related statistics reported by two studies of court files of bankrupts (Sullivan, Warren and Westbrook (2000) and Bermant and Flynn (1999)) suggests the rising charge-off rate comes from the fact that both the bankruptcy filing rate has almost doubled as well as the average debt defaulted on by a statistical bankrupt (Sullivan, Warren and Westbrook (2000) report that the average debt of a statistical bankrupt in 1997 was \$43,000, whereas in 1991 it was only about¹¹ \$26,000). As a result, even though the total debt held by the public doubled as well (denominator of charge-off rate), the charge-off rate actually went up over the period under consideration. The bankruptcy related statistics for the years 1991 and 1997 are summarized in Table 2.

Table 2. Debt and Bankruptcy Related Statistics

Year	1991	1997
Charge-off rate	4.5%	5.4%
Bankruptcy filings per 1000 adults	3.5	5.5
Average unsecured debt of bankrupts to income*	111%	160%
Average unsecured debt to income**	6%	9%

*Estimates from Warren et al. (2000) and Bermant et al. (1999).

**Estimates from Livshits et al. (2006).

3 The Model

Time is continuous and horizon infinite, $t \in (0, \infty)$. The economy is populated by a large number of ex-ante identical households, each facing a constant Poisson arrival rate ρ of death. Population size is stationary because upon each death, a new household is born

⁹Data on filings were taken from the American Bankruptcy Institute webpage. We count adult population aged 20-74, with the data taken from the Census Bureau.

¹⁰A bank's net charge-offs (gross charge-offs minus recoveries) divided by the average level of its loans outstanding.

¹¹Both measured in 1997 dollars.

to replace the old one.

Households face idiosyncratic income/expense shocks, derive utility from consumption, and use unsecured lines and savings to smooth consumption. They can only borrow from financial intermediaries called *banks* or save at a risk free return r^f . In every instance of time, households have an option to file for bankruptcy, which gives them a full discharge of debt in the spirit of Chapter 7 of the US bankruptcy code¹².

Banks are assumed to have an unlimited access to borrowing and lending at a constant risk free interest rate r^f . They also have access to a technology to make directed take-it-or-leave-it *credit card* offers to households at a fixed cost of screening and soliciting, χ . A credit card offer specifies an unsecured credit line L and interest rate R to which the bank is committed. Banks profit from the interest rate spread $R - r^f$ whenever borrowing takes place. We assume exclusivity of credit card offers by assuming that in any instance of time a household can only be matched with one bank.

The key feature of our model is the search friction of acquiring a new customer by a bank. Two features of the search and matching process are important. First, banks must actively search for a consumer in order to establish a new account, and this activity is *costly*. Second, when a bank solicits a customer, the number and kind of *other offers* that the consumer will be evaluating at the same time is stochastic. Thus, the bank faces a trade-off between profitability of the offered contract and the number of consumers that will respond positively to a given solicitation. Banks, facing a trade-off between the endogenous acceptance probability of a credit card offer and its ex-post profitability, randomize across offers that provide different levels of continuation utility to the household, in the spirit of Butters (1977) and Burdett and Judd (1983) price dispersion models.

3.1 Households

The economy is populated by a large number of ex-ante identical households, each facing a constant Poisson arrival rate ρ of death. Population size is stationary because upon

¹²Chapter 7 bankruptcy gives full discharge of debt provided that bankruptcy does not occur more often than every six years. This period length was changed in 2005 to eight years.

each death, a new household is born to replace the old one.

Households face idiosyncratic income/expense shocks, and maximize their life-time utility from consumption. To smooth their consumption path, they save in a state un-contingent asset or borrow from the banks. Households are born with no initial access to credit, but once they are targeted by a bank, they obtain one. Access to credit is exclusively in the form of a revolving credit lines with a pre-specified credit limit L and interest rate R .

The persistent component of the income of the household is assumed to be governed by a Poisson process with values $[y_1, \dots, y_n]$, subject to stochastic transition according to matrix $M_{\bar{y}}$ with an arrival rate η . In addition, the household is subject to transitory income/expense shocks that affect directly her income level. The volatility of this transitory component depends on the persistent level of income, i.e. $\sigma(\bar{y}) = \hat{\sigma}\bar{y}$. Given the state of the persistent component of income at date t , \bar{y}_t the *cumulative income* process Y_t evolves according to¹³

$$dY_t = \bar{y}_t dt + \sigma(\bar{y}_t) dz_t.$$

The asset evolution equation is then

$$da = radt + dY_t - c_t dt = (ra + \bar{y} - c)dt + \sigma(\bar{y})dz, \quad (1)$$

where

$$\begin{aligned} r &= R \text{ for } a \leq 0, \\ r &= r^f \text{ for } a > 0. \end{aligned} \quad (2)$$

The domain of assets is determined by the credit limit L ,

$$a \geq -L, \quad (3)$$

and it is understood that whenever the net assets fall below this limit, the household

¹³A similar formulation of output is used for example in Sannikov (2007). Discretized approximation of this process is: $y_t \Delta t = \bar{y} \Delta t + \sigma \varepsilon \Delta t^{\frac{1}{2}}$, where $\varepsilon \sim N(0, 1)$.

must ‘involuntarily’ default. It is important to stress that the non-negativity constraint imposed on consumption ($c \geq 0$) implies that whenever the net assets are on the bound, $a = -L$, this ‘involuntary’ default necessarily occurs. It follows from a general property of a continuous time process implying that within any fixed interval of time the Brownian motion term dz necessarily pushes net assets a below the prescribed limit L . The above considerations are summarized in Lemma 1.

Lemma 1 *Suppose $a = -L$. For any Δt ‘involuntary’ default occurs with probability one, as a must fall below the limit.*

Proof. Assume $c = 0$ to maximize the probability of not defaulting. See Corollary 9.2.7 (page 183), Oksendal (1998)¹⁴. ■

To simplify the problem, we do not model here explicitly what happens after default. Instead, we assume an exogenous continuation value for the household given $V^D(\bar{y})$, and assume that this value depends on the persistent component of income \bar{y} only. In particular, in our parameterization, we set $V^D(\bar{y})$ equal to the value of a household that has $L = 0$ and cannot obtain access to credit less a fixed utility cost of default φ .

The initial state of each household is $(a_0, \bar{y}, 0, r^f)$, where \bar{y} is drawn from the stationary distribution. Given (1), and the initial state (a_0, \bar{y}, L, R) , the objective of the household is to maximize the expected present discounted lifetime utility,

$$E\left[\int_0^T e^{-\beta t} u(c_t) dt + e^{-\beta T} W_{exit}\right] \quad (4)$$

where T denotes the stopping time related to the first occurrence of the following events:

- (i) The household decides to file for bankruptcy, in which case the exit value W_{exit} is given by $V^D(\bar{y})$.
- (ii) The household gets a new credit card offer, with the exit value determined by the continuation value implied by the best new credit card at hand.

¹⁴Intuitively this property follows from a basic property of Brownian motion having the variance proportional to the square root of the underlying interval of time, \sqrt{dt} , rather than the interval of time itself, dt .

- (iii) The household dies, in which case the exit value W_{exit} is zero.
- (iv) An shock to persistent level of income arrives, in which case the exit value is the value associated with initial state $(L_T, R_T, a_T, M_y \bar{y})$

The household maximizes the above expression by choosing the policy for consumption $c(a, \bar{y}, L, R)$, and the stopping time related to default (which is an option whenever $a > -L$). The maximized value of the above program is denoted by $V(a, \bar{y}, L, R)$.

3.1.1 Recursive Formulation of the Household Problem

To represent the household problem recursively, define first the default set D as the set of values of asset levels for which the household will choose to default or will be forced to default (i.e. $(-\infty, -L] \subset D$).

The Bellman equation for the household is given by:

$$\beta V(a, \bar{y}, L, R) = \max_{c \geq 0} \left\{ u(c) + \frac{E(dV)}{dt} \right\} \text{ for } a \notin D \quad (5)$$

subject to: i) the law of motion for assets

$$da = (ra + \bar{y} - c)dt + \sigma(\bar{y})dz,$$

where

$$\begin{aligned} r &= R \text{ for } a \leq 0, \\ r &= r^f \text{ for } a > 0, \end{aligned}$$

ii) value matching condition for the state of default ($a \in D$)

$$V(a; \bar{y}, L, R) = V^D(\bar{y}), \quad (6)$$

iii) smooth pasting condition for the state of default for $a > -L$ ($a \in D \setminus (-\infty, -L]$)

$$V_a(a; \bar{y}, L, R) = 0, \quad (7)$$

and iv) Ito's generator for $\frac{E(dV)}{dt}$, featuring: standard Ito drift terms, the endogenous probability of receiving a new credit card offer, and Poisson arrival rates of the major life-time risks (income shock and death)

$$\begin{aligned} \frac{E(dV)}{dt} = & V_a(a, \bar{y}, L, R)da + \frac{1}{2}\sigma^2 V_{aa}(a, \bar{y}, L, R) + \\ & + (\text{a new credit offer}) \lambda(a, \bar{y}, L, R)(N(a, \bar{y}; L, R) - V(a, \bar{y}, L, R)) \\ & + (\text{income shock}) \eta(M_{\bar{y}} \cdot V(a, \bar{y}', L, R) - V(a, \bar{y}, L, R)) \\ & + (\text{death}) \rho(-V(a, \bar{y}, L, R)), \end{aligned} \quad (8)$$

where N denotes the continuation utility starting from a new draw of a credit card offer and λ is the arrival rate of such event. The next sections describe how these objects are determined in equilibrium, and characterize their properties.

In the solution of the model, we guess and verify that the only time default occurs is actually at $a = -L$, so that $D = (-\infty, -L]$. In the Appendix, we derive simple sufficient conditions for this to hold that are satisfied for our parameter choices.

3.2 Banks

There is a unit measure of banks which provide access to unsecured borrowing to the households. Banks are assumed to have unlimited access to borrowing and lending technology at a constant risk free interest rate r^f . They can make directed take-it-or-leave-it *credit card* offers to households at a fixed cost of targeting an offer χ . Banks choose the number of offers to send out to each household type, and select the terms (credit limit L' and interest rate R'). Households accept or reject the offers made by the banks. Banks profit from the interest rate spread $R - r^f$ whenever borrowing takes place. Exclusivity of credit card offers is assumed, in the sense that in any instance of time a household can only be matched with one bank.

To define the banks problem formally, we first need to note that every credit card (L', R') the bank may offer to a customer of type $\theta = (a, \bar{y}, L, R)$ will imply some continuation utility $v = V(a, \bar{y}, L', R')$ to this customer. Since the acceptance of the offer by the customer depends only on this value, bank should deliver this value in a profit

maximizing way.

Assuming that an actual offer made by a bank takes the form of a lottery on a discrete grid of credit limits and interest rates $(R_i, L_j)_{i \in I, j \in J}$,

$$\mathcal{C} = \{x \in R^{\#I \times \#J} : \sum_{i \in I, j \in J} x_{i,j} = 1\}, \quad (9)$$

the first stage of bank's optimization can thus be formalized as a choice of a feasible lottery $x \in \mathcal{C}$ to maximize the profits subject to the constraint of delivering a fixed continuation value v to the household of type a, \bar{y} :

$$\Pi(v|a, \bar{y}) \equiv \max_{x \in \mathcal{C}} \sum_{i,j} x_{ij} \pi(a, \bar{y}, L_i, R_j), \quad (10)$$

subject to

$$v = \sum_{i,j} x_{ij} V(a, \bar{y}, L_i, R_j),$$

where the ex-post profit from a customer are defined by:

$$r^f \pi(a, \bar{y}, L, R) = -(R - r^f) \min(a, 0) + \frac{E(d\pi(a, \bar{y}, L, R))}{dt}, \quad (11)$$

with Ito's generator given by

$$\begin{aligned} \frac{E(d\pi(a, \bar{y}, L, R))}{dt} &= \pi_a(\bar{y} + ra - c(\theta)) + \frac{\sigma^2}{2} \pi_{aa} + \\ &\quad - \lambda(a, \bar{y}; L, R) \pi(a, \bar{y}, L, R) \\ &\quad \text{(balance transfer)} \\ &\quad + \rho(\min\{a, 0\} - \pi(a, \bar{y}, L, R)) \\ &\quad \text{(death)} \\ &\quad + \eta(M_{\bar{y}} \pi(a, \bar{y}', L, R) - \pi(a, \bar{y}, L, R)), \\ &\quad \text{(income shock)}, \end{aligned}$$

the law of motion for assets given by

$$da = (\bar{y} + ra - c(a, \bar{y}, L, R))dt + \sigma(\bar{y})dz,$$

and the initial condition determined by the loss due to bankruptcy protection law

$$\pi(-L, \bar{y}, L, R) = -L.$$

It is easy to establish the following properties of Π that will become important in the next section.

Lemma 2 *Assuming risk free rate is on the grid of available interest rates, $\Pi(\cdot|\theta)$ has the following properties: (i) it is a continuous and weakly concave function of v , (ii) it is defined on a closed interval, (iii) for high enough v , $\Pi(v|\theta) < \chi$, (iv) if V and π are continuous functions of assets a , Π is also continuous w.r.t. a .*

Proof. See Appendix. ■

In the second stage, given the Π schedule, the bank needs to choose the intensity of sending offers to each type and the type of offers made. Before proceeding with the analysis of this second stage of bank optimization, we will first need to formalize how matching takes place in our model.

3.3 Matching and Equilibrium Solicitation Strategy of the Banks

The central assumption how matching takes place in our model is that new credit card offers arrive to the households at an exogenous Poisson rate k (e.g. monthly), and upon the arrival, the actual number of offers is drawn from a power distribution with a parameter $p(\theta) \in [0, 1)$, i.e. the household receives at least one new credit card offer with probability $p(\theta)$, at least two offers with probability $p^2(\theta)$, at least three with probability $p^3(\theta)$ etc. The household can then review all these offers, and choose the best one among them.

This type of matching gives rise to contract dispersion of the type described in Burdett and Judd (1983). Just like in their paper, the consumer gets a random number of price

observations, and the only equilibrium is in mixed strategies – in our case the banks will mix over the continuation utilities granted to the consumer¹⁵ (or equivalently, over the underlying contracts that deliver these values). Denoting the mixed strategy of the bank over the continuation values for type θ by $F(v|\theta)$, we will now show how these objects of matching are uniquely pinned down in equilibrium, and how they are consistent with the objective of profit maximization by the banks. (Note that making an offer to any household type θ costs the bank χ , and profit maximization objective requires that the bank makes an infinite number of offers whenever the expected profits are above this cost, and makes no offers otherwise.)

Since the consumer always chooses the best of her offers, the acceptance probability of an offer delivering continuation utility v is not always one. Denote the acceptance probability of an offer delivering v to consumer type θ by $\mathcal{P}(v|\theta)$, and the distribution of the best offer at hand conditional on having at least one new offer by $\mathcal{F}(v|\theta)$. These object can be easily computed¹⁶ given $F(v|\theta)$ and $p(\theta)$:

$$\begin{aligned}\mathcal{F}(v|\theta) &= \frac{(1-p(\theta))F(v|\theta)}{1-p(\theta)F(v|\theta)}, \\ \mathcal{P}(v|\theta) &= \frac{1-p(\theta)}{1-p(\theta)F(v|\theta)}.\end{aligned}\tag{12}$$

What the bank's profit maximization objective imposes, is that in equilibrium all possible offers a bank can make at best cover the fixed cost of making such offer – as otherwise the bank would want to make infinite amount of offers. In other words, after factoring in the endogenous acceptance probability $\mathcal{P}(v|\theta)$ of an offer delivering

¹⁵Imagine all banks offering the same v and making positive profit. A deviation offering $v + \varepsilon$, and capturing the whole market gives strictly higher profit. If all banks offer v that is zero profit, offering $v - \varepsilon$ and getting the customers with only one offer is a profitable deviation.

¹⁶

$$\begin{aligned}\mathcal{F}(v|\theta) &= \frac{1}{p(\theta)}[(1-p)pF + (1-p)p^2F^2 + (1-p)p^3F^3 + \dots] \\ \mathcal{P}(v|\theta) &= (1-p) + p(1-p)F + p^2(1-p)F^2 + \dots\end{aligned}$$

continuation utility v to the household, in equilibrium, we must require

$$\mathcal{P}(v|\theta)\Pi(v|\theta) - \chi \leq 0, \quad (13)$$

with equality whenever an offer is actually made.

We now proceed to show that the above condition imposes enough structure to uniquely pin down the key objects of matching $p(\theta)$ and $F(v|\theta)$. For this purpose, consider the following two cases that can arise. First, consider the case in which for any continuation utility v higher than the current value $V(\theta)$, we have $\Pi(v|\theta) \leq \chi$. In other words, even if the acceptance of the new offer is certain, the bank is losing money (weakly) on any feasible offer it can make. In such case, consistency with the objective of profit maximization by the bank, requires that in equilibrium no offer is made to such household type, and consequently we must have $p(\theta) = 0$. Second, consider the case when for some $v \geq V(\theta)$ the ex-post profits are positive, i.e. $\Pi(v|\theta) > \chi$ for at least one v from the domain. Clearly, if the acceptance probability of such offer was certain, now the bank would want to make an infinite number of offers. Thus, we must require that in this case $p(\theta) > 0$. Now, denote $\underline{v} = \max(V(\theta), \operatorname{argmax}_v \Pi(v|\theta))$, and note that by property (i),(ii) in Lemma 2 \underline{v} is well defined, by properties (i), (ii), (iii) for all $v \geq \underline{v}$ $\Pi(v|\theta)$ is a decreasing function of v , and by properties (i), (iii) there exists a value \bar{v} such that

$$\Pi(\bar{v}|\theta) = \chi. \quad (14)$$

Clearly, the only offers consistent with the objective of profit maximization by the bank are the ones from the interval

$$\mathcal{V}(\theta) = [\underline{v}, \bar{v}]. \quad (15)$$

This is because all offers below \underline{v} are on the increasing portion of $\Pi(\cdot|v)$, and are dominated by those above \underline{v} , and offers above \bar{v} can not cover the fixed cost χ even if the acceptance probability is one. With these properties in hand, we are now ready to construct the distribution function $F(v|\theta)$ on $\mathcal{V}(\theta)$, and find the value of $p(\theta) \in [0, 1)$.

First, note that at the lower bound \underline{v} , we have $F(\underline{v}|\theta) = 0$, and therefore by (13) and

(12)

$$p(\theta) = 1 - \frac{\chi}{\Pi(\underline{v}, a, \bar{y})}. \quad (16)$$

Thus, by the fact that $+\infty > \Pi(\underline{v}, a, \bar{y}) > \chi$, p is uniquely pinned down, and $p(\theta) \in [0, 1)$. But, given $p(\theta)$, and the domain $\mathcal{V}(\theta)$, the distribution $F(v|\theta)$ is now also pinned down by zero profit condition:

$$\frac{1 - p(\theta)}{1 - p(\theta) F(v|\theta)} \Pi(v|\theta) - \chi = 0. \quad (17)$$

By the properties of $\Pi(\underline{v}, a, \bar{y})$ established in Lemma 2, $F(v|\theta)$ is increasing, continuous, and satisfies $F(\underline{v}|\theta) = 0$ and $F(\bar{v}|\theta) = 1$. Thus, F is, in fact, a distribution function, and by arguments implicit in the construction, no other function can satisfy our requirements¹⁷. We summarize the above result into a formal proposition stated below.

Proposition 1 *There exists $p(\theta) \in [0, 1)$, and whenever $p(\theta) > 0$, there exists a unique and continuous distribution function $F(v|\theta)$ defined on a closed interval $\mathcal{V}(\theta)$, such that the equilibrium zero profit condition*

$$\mathcal{P}(v|\theta)\Pi(v|\theta) - \chi \leq 0 \text{ with equality if } v \in \mathcal{V}(\theta) \quad (18)$$

is satisfied, and the implied solicitation strategy of the banks is consistent with the objective of profit maximization.

Proof. See in text above. ■

Finally, we are now ready to explicitly define the expected continuation utility $N(\theta)$ of the household following the arrival event of at least one new credit card offer

$$N(\theta) = \int_{\mathcal{V}(\theta)} v \mathcal{F}(dv|\theta), \quad (19)$$

¹⁷Implicit in our analysis is the argument that excludes the possibility of a mass point in F . It is instructive to make this argument explicit. So, suppose there is a mass point at v_0 . Then, there must be a profitable deviation of making an offer that delivers $v_0 + \varepsilon$ (ε arbitrarily small) to the consumer. Clearly, the acceptance probability of $v_0 + \varepsilon$ is strictly higher by an amount independent of the choice of ε (due to the mass point). By continuity of Π , the loss in profits is of order ε . Since the zero profit condition holds at v_0 , this contradicts that F is consistent with the objective of profit maximization by the banks.

and the Poisson rate at which banks lose existing customers due to a successful solicitation by another bank

$$\lambda(\theta) = kp(\theta). \tag{20}$$

With these objects defined, we are ready to close the model by formally defining the equilibrium:

Definition 1 *The equilibrium in this economy is: value function of the household V , the policy function c , the profit functions Π, π , the distribution of new offers F , and probability p , such that: (i) the policy function c is admissible¹⁸, (ii) the value functions V and π are of class C^2 , (iii) both the policy function and the value functions V solve (5), (iv) the value functions π and V satisfy the transversality conditions $\lim_{t \rightarrow \infty} E[e^{-r^f t} V(a, \bar{y}, L, R)] = 0$, $\lim_{t \rightarrow \infty} E[e^{-r^f t} \pi(a, \bar{y}, L, R)] = 0$, (v) profit function π solves the Bellman equation, (vi) profit function Π is consistent with (10), (vii) the distribution function F and the probability p satisfy the zero profit condition given by (13), and are consistent with the profit maximization objective of the bank.*

4 Parameterization

In this section we describe how we choose functional forms and parameter values. The baseline period in the model is chosen to be one year, i.e. $\Delta t = 1$ is interpreted in the model as 1 year. The instantaneous utility function is CRRA,

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma},$$

with the coefficient of risk aversion $\gamma = 2$. We set the risk-free rate at $r^f = 2.5\%$ – the average return on T-bills over the postwar period (see Athreya (2004)). We choose the discount factor β of 6% to reflect the risk-free rate of 2.5% augmented by the expected postwar growth of real GDP in the US of 3.5%– which we embed here into the discount factor rather than explicitly model. The arrival rate of death is chosen so that the average

¹⁸ Satisfy the Lipschitz condition.

life span of our households is 50 years: i.e. individuals enter the economy at age 22 and exit at age 72. This gives the number $\rho = 0.02$.

The continuation value following default decision $V^D(\bar{y})$ is set equal to the continuation utility of an agent who does not receive offers and has $L = 0$, less a one time assumed utility cost of bankruptcy φ chosen to match the aggregate debt to income ratio in year 2004.

To calibrated our income/expense process, we start-off with the standard formulation of the annual income process taken from Guvenen (2005)

$$\begin{aligned}\log y &= \log z + \log \varepsilon, \\ \log(z) &= \phi \log(z_{-1}) + \eta,\end{aligned}$$

where z_t is the persistent component with $\phi = 0.988$ and i.i.d. innovations $\eta \sim N(0, \sigma_\eta^2)$, $\sigma_\eta^2 = 0.015$, and $\varepsilon \sim N(0, \sigma_\varepsilon^2)$ is a transitory i.i.d. component with variance $\sigma_\varepsilon^2 = 0.061$. (A similar formulation is widely used in the literature, see for example Storesletten, Telmer and Yaron (2004).) We map the above income process onto our Brownian motion component, in place of the transitory part $\log \varepsilon$, and a two-state Markov chain \bar{y} in place of the persistent AR(1) component¹⁹ $\log(z)$. We think of the transitory part as the lowest bound on the temporary income/expense shocks that we should have in our model. In the benchmark case, we thus set $\hat{\sigma} = 0.3$, slightly above the value implied by $\sigma_\varepsilon = 0.25$.

The cost of soliciting and screening χ is calibrated so that in the recent period 2000-2005 the cost of a *successful* credit card application (factoring the acceptance probability of an offer) is about \$100 (when measured relative to the median income of a household in our model economy, which is identified with the median US household income of \$35000 net of taxes and Social Security). The cost of acquiring a new customer is close to the one given in Evans and Schmalensee (2005) (estimate by Visa U.S.A.). In our quantitative exercise, we increase this cost to reduce the debt-to-income ratio of a statistical household to the level consistent with the data in 1990.

¹⁹We use the Tauchen method to map the AR(1) z -process into a symmetric Markov chain. The value of \bar{y} in each state of our Markov chain is set to match the average values of the AR(1) process when computed conditional on being above and below the mean.

Finally, we arbitrarily set the initial asset holdings of the household a_0 equal to 10% of the average persistent level of income.

The list of the chosen parameter values is summarized in Table 3 below.

Table 3. Model parameters

Parameter		Value
Discount factor	β	0.06
Risk free rate	r^f	2.5%
Arrival rate of death	ρ	0.02
Cost of soliciting and screening	χ	0.001
Income process	\bar{y}_l	0.6
	\bar{y}_h	2.13
	σ	$\{0.3\bar{y}_l, 0.3\bar{y}_h\}$
	η	1
	M_y	$\begin{bmatrix} 0.95 & 0.05 \\ 0.05 & 0.95 \end{bmatrix}$
Stigma	φ	13.5

5 Findings (preliminary and incomplete)

This section presents preliminary results from our model economy for the parameterization described in the previous section. The aggregate statistics are presented in Table 4 below.

Implications for the most recent period Overall, the model matches well several dimensions of the recent data: filing rate for bankruptcy protection per 1000 of households (excluding reasons due to divorce, unplanned pregnancy and medical bills that account for about half of filings), high charge-off rate, high fraction of population with credit card and revolving balance, and debt-to-income ratio of households (served as our calibration target). It overpredicts by a factor of two the debt-to-income ratio of a statistical bankrupt, and the average credit limit to income of a statistical card-holder. The latter discrepancy is rather hard to interpret, because in the model there is no cost of

accepting a credit card offer. In other words, households in the model necessarily must have a maximal potential credit limit when they are granted a credit card – even if they do not intend to use it. In the data, since there is a cost of accepting an offer or holding an outstanding credit card, it is reasonable to expect that households have lower credit limits relative to what they could possibly have.

The most serious shortcoming of the model in our opinion is the level of debt discharged for a statistical bankrupt relative to her income. In the model it is twice as high as it is in the data. As a direct consequence of this discrepancy, the model overpredicts the level of the charge-off rate, which is twice as high as in the data. (Note that the denominator of the charge-off rate – total debt of the public – has been implicitly pinned down in calibration to match the debt-to-income ratio in the data).

Table 4. Data and Model Results

Statistic	Model		Data	
	High χ	Low χ	1990 – 95	2000 – 05
Bankruptcy filing rate	1.6	3.23	3.08	5.75
Charge-off Rate	8.3%	7.1%	3.5%	4.9%
Debt/Income	4%	10%	5%	10%
Debt/Income of bankrupts	3.48	3.68	1.11	1.60 ²⁰
Average credit limit/income	2.68	2.79	0.2	0.46
Interest rate premium*	12.4%	9.0%	8.5%	5.6%
Fraction of population with a c-card	24%	46%	56%	71%
Fraction of cc w/ revolving balance	31%	35%	51%	55%
Number of solicitations (per year)	0.82	3.0	8.5	26
Response rate to solicitations	38%	27%	2.2%	0.54%
Average duration of cc relation (years)	0.82	0.59	-	-

*Net of charge-offs and $r^f = 2.5\%$ in the model, (T-bills in the data).

Results from the comparative statics experiment The results of our comparative statics experiment look somewhat promising, even though quantitatively the model does

²⁰This number is the 1997 estimate.

not match the data exactly. As we can see from the same table, the implied change in the filing rate for bankruptcy protection is close to the one observed in the data. The expansion of credit limits relative to income is also close to the one observed in the data. The change in the interest rate premium (average interest rate on credit cards less risk-free rate and charge-off rate), even in levels a bit off, falls similarly to the data. The major shortcoming seems to be the behavior of the charge-off rate, which in the model goes down while in the data it goes up. This failure can be explained as follows. In the model, the filing rate almost doubles – just like in the data – but the debt-discharged per statistical bankrupt increases by about 6%, which is much less than the almost 50% increase in the data. Since the numerator of the charge-off rate is given by the product of the filing rate times the average debt discharged per statistical bankrupt, and since the denominator more than doubles as we calibrate the model to replicate the change in the debt-to-income ratio, the charge-off rate must fall. (In the data, since debt discharged per statistical bankrupt increases by 50%, the charge-off rate goes up despite the twofold increase in total debt held by the public.)

An additional check on our story is the increase in the credit card solicitations. Our model has direct predictions about it, and in our comparative statics exercise, credit card solicitations intensify by a factor of 4 times. This increase is reasonably close to what we observe in the data (see Figure 1). However, the level of solicitations is off by an order of magnitude with respect to the data, and so is the response rate.

Overall, we conclude that our mechanism is promising in accounting for a significant portion of the observed changes in the US unsecured market over the last two decades, but quantitatively falls short in some dimensions.

Discussion of the key mechanism The key mechanism of the model is driven by the increased use and reliance on bank insurance versus self-insurance. Preliminary quantitative experiments with the model (not reported yet) reveal that the way continuation utility is delivered to the households through the equilibrium contracts does not play any significant role for the result following a change in χ . It is the number and frequency of the new credit card offers which is playing a critical role. We design a quantitative experiment to show that this is true. In this experiment we use the contracts from low

χ economy, and lower solicitation intensity in consistency with the economy with high χ . It turns out that we obtain very similar quantitative results to the benchmark model, which led us to this conclusion. (to be completed).

The force we identify as critical, suggests the following interpretation of how χ affects the key bankruptcy statistics. The increased frequency of solicitations has two effects on the households' decision to accumulate assets over the life-cycle. First, the lower costs of soliciting and screening implies a shorter waiting time for the first credit card offer (after the household is born). Second, once this offer arrives, due to the increased competition for a customer, the surplus extraction by banks is lower – i.e. the first contract is more favorable to the customer. Third, the household climbs up much faster towards better contracts – an effect which comes in our model from a form of history dependence: the better contract a household has, the higher is its reservation utility for accepting a new contract offer, and thus for the new offer to be competitive, it has to give more favorable terms to the household²¹. As a result of these changes, the attractiveness of borrowing increases when χ is low, and so does the use of credit, and bankruptcy filings.

Figures 5 through 7 illustrate how these effects show up in the implied cross-sectional characteristics of households in our model economy. As we can see from Figure 5, households accumulate less assets over the life-cycle when χ is low. (The fraction of households with revolving balance also significantly expands.) These findings show that for all age groups there is a greater reliance on bank-provided insurance versus self-insurance. In addition, Figures 6-7 show that the average contract L, R when χ is low features slacker credit limits and lower interest rates. This is driven by the increased competition for customer, which implies that agents overall have access to credit on better terms over their life-cycle. (We observe similar effects when we condition on the state of the households. To be completed.)

(Sections on dispersion of credit limits and interest rates in model and data, and propagation mechanism through competition are to be added.)

²¹This feature of the model can be loosely related to the idea that nowadays it is much easier to quickly build credit history to get access to borrowing on good terms.

6 Conclusions

In this paper we argued that in the 90s we have witnessed quite a revolution in consumer finances: the growing availability and use of revolving lines of credit. We have argued that this product made it more affordable for consumers to use unsecured credit to smooth consumption. What have attributed to this change was the fall in the ‘targeting/screening’ friction that precluded banks from offering this product on a large scale in the proceeding period. By abstracting from the possibility of consumers paying the screening cost up-front, we have shown that the fall in this friction may be critical role to account for the recently observed changes in the US market for unsecured credit.

Some important questions still remain. Quantitatively, our theory turned out to overpredict the level of indebtedness of a statistical bankrupt, and underpredict the change of this statistic in the comparative statics exercise. This failure resulted in a counterfactual implication of the model for the behavior of the charge-off rate – which in the data went up over 90s, but fell in the model. Our model also turned out to overpredict the level of credit limits to income. We conjecture that a more careful modeling of the state of default might help partially address this puzzle, and plan to extend our model in this direction. In the future revisions of the paper, we also plan to examine the welfare consequences of the bankruptcy law changes introduced in 2005 (limiting bankruptcy protection only to those households whole income falls below the median).

References

- Athreya, Kartik. 2001. “The Growth of Unsecured Credit: Are We Better Off?” *Federal Reserve Bank of Richmond Economic Quarterly* 87/3.
- Athreya, Kartik. 2002. “Welfare Implications of the Bankruptcy Reform Act of 1999.” *Journal of Monetary Economics* 49:1567–1595.
- Athreya, Kartik. 2004. “Shame As It Ever Was: Stigma and Personal Bankruptcy.” *Federal Reserve Bank of Richmond Economic Quarterly* 90.

- Athreya, Kartik. 2006. “Fresh Start or Head Start? Uniform Bankruptcy Exemptions and Welfare.” *Journal of Economic Dynamics and Control* 30:2051–2079.
- Ausubel, Lawrence M. 1991. “The Failure of Competition in the Credit Card Market.” *American Economic Review* 81:50–81.
- Bermant, G. and E. Flynn. 1999. “Incomes, Debts and Repayment Capacities of Recently Discharged Chapter 7 Debtors.” *American Bankruptcy Insitute, manuscript* .
- Burdett, Kenneth and Kenneth Judd. 1983. “Equilibrium Price Dispersion.” *Econometrica* pp. 955–69.
- Burdett, Kenneth and Melvyn Coles. 2000. “Equilibrium Wage-Tenure Contracts.” *Econometrica* 71:1377–1404.
- Butters, Gerard. 1977. “Equilibrium Distributions of Sales and Advertising Prices.” *Review of Economic Studies* 44:465–491.
- Chatterjee, Satyajit, Dean Corbae, Makoto Nakajima and Jose-Victor Rios-Rull. 2007. “A Quantitative Theory of Unsecured Credit with Risk of Default.” *forthcoming, Econometrica* .
- Eaton, Jonathan and Mark Gersovitz. 1981. “Debt with Potential Repudiation: Theoretical and Empirical Analysis.” *The Review of Economic Studies* 48:289–309.
- Evans, David S. and Richard Schmalensee. 2005. *Paying with Plastic: The Digital Revolution in Buying and Borrowing*. The MIT Press.
- Gross, David B. and Nicholas S. Souleles. 2002. “An Empirical Analysis of Personal Bankruptcy and Delinquency.” *Review of Financial Studies* 15:319–347.
- Kushner, Harold J. and Paul Dupuis. 2001. *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer-Verlag.
- Livshits, Igor, James MacGee and Michele Tertilt. 2007a. “Accounting for the Rise in Consumer Bankruptcies.” *Manuscript* .

- Livshits, Igor, James MacGee and Michele Tertilt. 2007b. “Costly Contracts and Consumer Credit.” *manuscript* .
- Mester, Loretta J. 1997. “What’s the Point of Credit Scoring?” *Federal Reserve Bank of Philadelphia Business Review* .
- Narajabad, Borghan N. 2007. “Information Technology and the Rise of Household Bankruptcy.” *manuscript* .
- Oksendal, Bernt K. 1998. *Stochastic Differential Equations: An Introduction with Applications*. Springer.
- Rios-Rull, J.V. and X. Mateos-Planas. 2007. “Credit Lines.” *University of Minnesota, manuscript* .
- Storesletten, Kjetil, Christopher I. Telmer and Amir Yaron. 2004. “Consumption and Risk Sharing over the Lifecycle.” *Journal of Monetary Economics* 51:609–633.
- Sullivan, T.A., E. Warren and J.L. Westbrook. 2000. *The Fragile Middle Class*. Yale University Press.

Appendix

Omitted proofs and derivations

Condition for default occurring only on the bound, i.e. at $a = -L$ This section derives a sufficient condition for default occurring only at the maximum asset level $a = -L$. Consistent with our parameterization, we restrict the persistent level of income to take two values: \bar{y}_l and \bar{y}_h .

We will compare two options for the household in state $(a_t, \bar{y}_{ti}, L_t, R_t)$, $a_t > -L_t$. The first one is default, with the corresponding value equal to $V^D(\bar{y}_{ti})$. The second one is the following plan. Instead of defaulting at period t_0 , set high enough consumption equal to $c = C \in R_+$, and default either after $\hat{\Delta}$ or when first leave the set $a \geq -L$ (whichever occurs first). Denote by $\Delta(C)$ the stopping time associated with default. Without loss

of generality of the argument below, assume that if a new credit card offer is received by the consumer, the consumer declines such offer. Clearly, this assumption is without loss of generality as accepting the new offer is an option, and can only improve things. Note also that $\Delta(C)$ has the following properties *on each fixed sample path*: (i) starting from C high enough, $\Delta(C)$ is strictly decreasing w.r.t. C , and (ii) $\Delta(C) \rightarrow_{C \rightarrow \infty} 0$, as it is *only* bounded below by 0. (For simplicity, we incorporate death probability into the discount factor.)

Along any sample paths the following flow is received by the household from the second plan: (i) utility from consuming C until stopping occurs,

$$\int_0^{\Delta(C)} e^{-(\rho+\beta)t} u(C) dt,$$

(ii) continuation value from autarky in case a switch of \bar{y} occurs before stopping²²,

$$\eta \Delta(C) e^{-(\beta+\rho)\Delta(C)} M \cdot V^D(\bar{y}') + o(\Delta(C)),$$

(iii) continuation value from autarky in case a switch of \bar{y} does not occur before stopping,

$$(1 - \eta \Delta(C)) e^{-(\beta+\rho)\Delta(C)} V^D(\bar{y}) + o(\Delta(C))$$

Ignoring residual terms of order higher than $\Delta(C)$, for every sample path we must have

$$\begin{aligned} & \int_0^{\Delta(C)} e^{-(\rho+\beta)t} u(C) dt + \eta \Delta(C) e^{-(\beta+\rho)\Delta(C)} M \cdot V^D(\bar{y}') + (1 - \eta \Delta(C)) e^{-(\beta+\rho)\Delta(C)} V^D(\bar{y}) \\ & > \Delta(C) [u(C) + \eta M \cdot (V^D(\bar{y}') - V^D(\bar{y}))] + e^{-(\beta+\rho)\Delta(C)} V^D(\bar{y}) \text{ if } (u(C) \text{ negative}) \\ & > \Delta(C) [e^{-(\rho+\beta)\hat{\Delta}} u(C) + \eta M \cdot (V^D(\bar{y}') - V^D(\bar{y}))] + e^{-(\beta+\rho)\Delta(C)} V^D(\bar{y}) \text{ if } (u(C) \text{ positive}) \end{aligned}$$

Note that for our CRRA formulation, the utilities and values are negative. A sufficient

²²The term $o(\Delta)$ denotes a residual that is by an order of magnitude smaller than $\Delta(C)$ and can be ignored when $\hat{\Delta}$ is small enough.

condition for the consumer to prefer to wait with default is

$$\begin{aligned} & \Delta(C)[u(C) + \eta M \cdot (V^D(\bar{y}') - V^D(\bar{y}))] + e^{-(\beta+\rho)\Delta(C)}V^D(\bar{y}) \\ & > \Delta(C)(\beta + \rho)V^D(\bar{y}) + e^{-(\beta+\rho)\Delta(C)}V^D(\bar{y}) \end{aligned}$$

which can be rewritten as

$$u(C) - (\beta + \rho)V^D(\bar{y}) > -\eta M \cdot (V^D(\bar{y}') - V^D(\bar{y}_{t_0})) + \varepsilon, \text{ some } \varepsilon > 0.$$

Since the utility function we assume is CRRA, for C big enough, $u(C)$ can be arbitrarily small, and so a sufficient condition that ensures that the above inequality is satisfied is

$$-(\beta + \rho)V^D(\bar{y}) > -\eta M \cdot (V^D(\bar{y}') - V^D(\bar{y})). \quad (21)$$

Two cases are possible.

1. $\bar{y} = \bar{y}_l$. Then the term on the right is negative as long as

$$V^D(y_h) > V^D(y_l) \quad (22)$$

and the inequality is satisfied.

2. $\bar{y} = \bar{y}_h$. Denote the diagonal terms of M by ϕ (the persistence of the Markov chain). The right hand side is

$$\begin{aligned} & -\eta(\phi V^D(\bar{y}_h) + (1 - \phi)V^D(\bar{y}_l) - V^D(\bar{y}_h)) \\ & = -\eta(1 - \phi)[V^D(\bar{y}_l) - V^D(\bar{y}_h)], \end{aligned}$$

and (21) becomes

$$-(\beta + \rho)V^D(\bar{y}_h) > -\eta(1 - \phi)[V^D(\bar{y}_l) - V^D(\bar{y}_h)].$$

Assuming $V^D(y_h) > V^D(y_l)$, it is equivalent to

$$(\beta + \rho + \eta(1 - \phi))|V^D(\bar{y}_h)| > \eta(1 - \phi)|V^D(\bar{y}_l)|. \quad (23)$$

In our numerical solution, we check that equations (22) and (23) are satisfied, which is a sufficient condition for default to occur only on the bound, i.e. at $a = -L$.

Proof of Lemma 2. (i) and (ii): These properties are implied by the fact that lotteries are allowed on the discrete grid of contracts.

(iii): Consider a contract featuring $R = r^f$ and the highest credit limit $L = L_{\max}$ possible. Clearly, in this case the continuation utility $v_{\max} = V(a, \bar{y}, L_{\max}, r^f)$ to the household is the highest possible. Thus, this contract (or lottery that puts all the weight on this contract) must solve (10) for $v = v_{\max}$. But since $R = r^f$, the banks do not make any profits, and thus $\Pi(v_{\max}|\theta) < \chi$ (χ is assumed strictly positive).

(iv) This property follows from the Maximum Theorem. A change in a is nothing else but a change affecting the weights in the objective function and the constraint set – both changing continuously by assumption. By the Maximum Theorem the implied value at optimum, Π , changes continuously as well. ■

Numerical solution algorithm

We discretize the continuous state space and compute the equilibrium of the model using an implementation of the Message Passing Interface (MPI-2) in Fortran 90 on a Beowulf cluster of 100 Processors of class Xeon 5300 2.66MHz. The household's problem and bank's problem is solved using the finite difference Markov chain approximation method described in Kushner and Dupuis (2001). The following steps describe the procedure we follow to compute an approximate equilibrium of our model:

- Step 1. **Discretization of the state space.** Here, we demonstrate how we apply Markov chain discretization method for the household's problem. Because bank's problem is tackled by an exact analogy, we skip it. The discretized grid obtained in this step is important, as it is inherited by the proceeding steps of our numerical procedure. Given the process for assets a

$$da = (\bar{y} + ra - c) dt + \sigma dz,$$

where

$$\begin{aligned} r &= R, \text{ for } a \leq 0, \\ r &= r^f, \text{ for } a \geq 0, \end{aligned}$$

define

$$b(a, c) \equiv \bar{y} + ra - c.$$

For simplicity, ignore Poisson jump terms (see comment at the end how to incorporate these terms). By Ito's Lemma, the following HJB representation of the household's asset accumulation problem is valid:

$$\begin{aligned} \beta V(a) &= \max_c \left[V_a(a) b(a, c) + \frac{1}{2} \sigma^2 V_{aa}(a) + U(c) \right], \quad a \in [-L, \bar{a}] \\ V(a) &= V^D(a), \quad a = -L \end{aligned}$$

To arrive at the finite difference Markov approximation of the above, the following approximations are used:

$$\begin{aligned} V_a(a) &= \frac{V(a+h) - V(a)}{h} \text{ if } b(x, c) \geq 0, \\ V_a(a) &= \frac{V(a) - V(a-h)}{h} \text{ if } b(x, c) < 0, \\ V_{aa}(a) &= \frac{V(a+h) + V(a-h) - 2V(a)}{h^2}, \end{aligned}$$

To simplify notation, define

$$\begin{aligned} b^+ &= \max[b, 0], \\ b^- &= \max[-b, 0], \end{aligned}$$

and note that

$$|b| = b^+ + b^-.$$

Also, denote the optimal policy by $c(a)$, and then use the above to obtain

$$\begin{aligned}\beta V(a) &= \frac{V(a+h) - V(a)}{h} b(a, c(a))^+ - \frac{V(a) - V(a-h)}{h} b(a, c(a))^- \\ &\quad + \frac{1}{2} \sigma^2 \frac{V(a+h) + V(a-h) - 2V(a)}{h^2} + U(c(a)),\end{aligned}$$

and thus

$$\begin{aligned}V(a) \frac{\beta h^2}{h |b(a, c(a))| + \sigma^2} &= V(a+h) \left[\frac{hb(a, c(a))^+ + \frac{1}{2}\sigma^2}{h |b(a, c(a))| + \sigma^2} \right] \\ &\quad + V(a-h) \left[\frac{\frac{1}{2}\sigma^2 + hb(a, c(a))^-}{h |b(a, c(a))| + \sigma^2} \right] \\ &\quad - V(a) + U(c(a)) \frac{h^2}{h |b(a, c(a))| + \sigma^2}\end{aligned}$$

Identify the underlying time interval with the expression

$$\Delta t(x, c) \equiv \frac{h^2}{h |b(x, c(x))| + \sigma^2},$$

and rewrite the above as

$$\begin{aligned}V(a) (1 + \beta \Delta t) &= V(a+h) \frac{\frac{1}{2}\sigma^2 + hb(a, c(a))^+}{h |b(a, c(a))| + \sigma^2} + \\ &\quad + V(a-h) \frac{\frac{1}{2}\sigma^2 + hb(a, c(a))^-}{h |b(a, c(a))| + \sigma^2} \\ &\quad + U(c(a)) \Delta t.\end{aligned}$$

To tackle discounting, divide by $1 + \Delta t \beta$, and using the fact that $e^{\beta \Delta t} \approx 1 + \beta \Delta t$, and $\frac{\Delta t}{1 + \beta \Delta t} \approx \Delta t$, observe that

$$\begin{aligned}V(a) &= e^{-\beta \Delta t(a, c(a))} V(a+h) \frac{\frac{1}{2}\sigma^2 + hb(a, c(a))^+}{h |b(a, c(a))| + \sigma^2} \\ &\quad + e^{-\beta \Delta t(a, c(a))} V(a-h) \frac{\frac{1}{2}\sigma^2 + hb(a, c(a))^-}{h |b(a, c(a))| + \sigma^2} \\ &\quad + \Delta t(a, c(a)) U(c(a)).\end{aligned}$$

Finally, define Markov transition probabilities

$$\begin{aligned} p(a, a+h|c(a)) &= \frac{\frac{1}{2}\sigma^2 + hb(a, c(a))^+}{h|b(a, c(a))| + \sigma^2}, \\ p(a, a-h|c(a)) &= \frac{\frac{1}{2}\sigma^2 + hb(a, c(a))^-}{h|b(a, c(a))| + \sigma^2}. \end{aligned}$$

and rewrite the above problem in a more compact form

$$\begin{aligned} V(a) &= e^{-\beta\Delta t(a, c(a))} p(a, a+h|c(a)) V(a+h, c(a)) + \\ &+ e^{-\beta\Delta t(a, c(a))} p(a, a-h|c(a)) W(a-h, c(a)) \\ &+ \Delta t(a, c(a)) U(a, c(a)). \end{aligned}$$

The expression above a readily applicable equation to proceed with a value function iteration scheme:

$$V(a) = \max_c \left[\begin{array}{l} e^{-\beta\Delta t(a, c)} p(a, a+h|c) V(a+h) \\ + e^{-\beta\Delta t(a, c)} p(a, a-h|c) V(a-h) \\ + \Delta t(a, c) U(c) \end{array} \right]. \quad (24)$$

Poisson jump events can be incorporated in the above reasoning by expanding the term

$$V(a-h) e^{-\rho\Delta t} p_- + V(a+h) e^{-\rho\Delta t} p_+$$

into

$$(1 - \eta\Delta t) [V(a-h) e^{-\rho\Delta t} p_- + V(a+h) e^{-\rho\Delta t} p_+] + e^{-\rho\Delta t} \eta\Delta t V^+,$$

where V^+ is the value function after the Poisson event occurs. See Kushner and Dupuis (2001) for more details.

Step 2. **Efficient frontier of contracts.** We assume some initial guess for $\lambda = kp$, which is $\lambda_0 = 0$. Using the approach outlined in step 1, for all combinations of $(L, R) \in \mathcal{C}$

from a discrete grid of values, we obtain the set of tuples

$$\{V(a, \bar{y}, L, R), \pi(a, \bar{y}, L, R)\}_{(L,R) \in \mathcal{C}}.$$

We then take the convex hull of these tuples, and compute the efficient frontier of the convex hull. (Note that each tuple is an actually independent problem to solve, which implies huge efficiency gains from parallel computing.)

- Step 3. **Arrival rate of new offers.** Let L, R denote the old (existing) contract a consumer has at hand. For each $(L, R) \in \mathcal{C}$, from the zero profit condition (13) we update λ , i.e. $\lambda'(a, \bar{y}, L, R) = k(1 - \frac{\chi}{\Pi(\underline{v}, a, \bar{y})})$.
- Step 4. **Distribution of the best contract.** We approximate the distribution F using (13), and solving it on a grid of 50 different values of v from the interval $\mathcal{V}(\theta) = [\underline{v}, \bar{v}]$. We use a simplicial interpolation to interpolate F continuously for the remaining values on $[\underline{v}, \bar{v}]$.
- Step 5. **Update.** We update the guess for λ , i.e. we compute $\lambda_i = 0.9\lambda_{i-1} + 0.1\lambda'$, and repeat steps 2-5 starting from λ_i instead of λ_0 . We iterate until convergence.
- Step 6. **Verification.** We verify the solution by increasing the precision of our grid and solving again. Since we have not established that the equilibrium is unique, we examine numerically whether the solution is invariant to the initial conditions, i.e. if it remains unchanged when we start solving the model from different initial guesses for λ_0 .

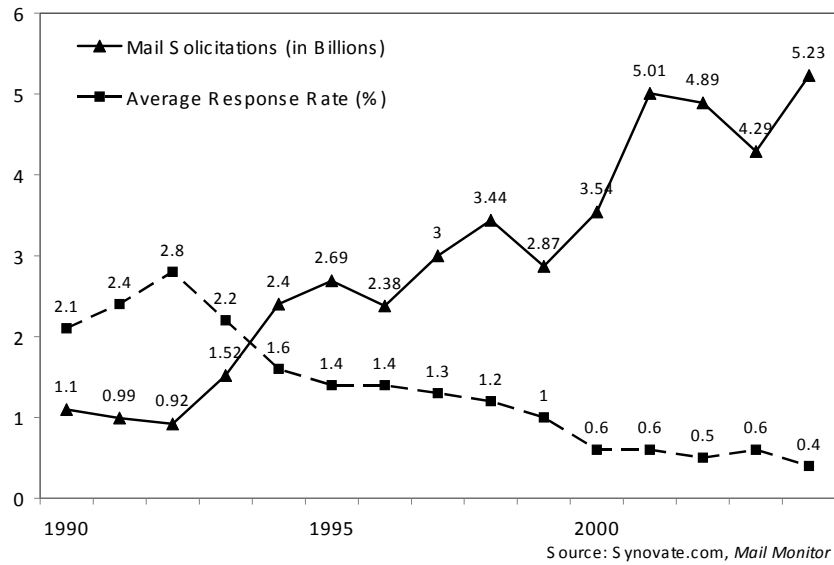


Figure 1: Mail Solicitations with Credit Card Offers in the US, 1990-2005.

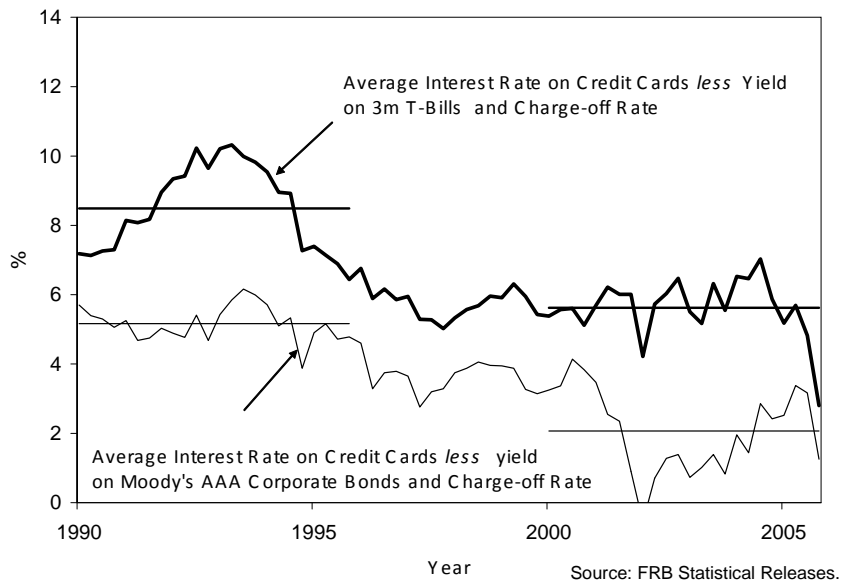


Figure 2: Credit Card Interest Rate Premium in the US, 1990-2005: Interest Rate on Credit Cards Less Cost of Funds.

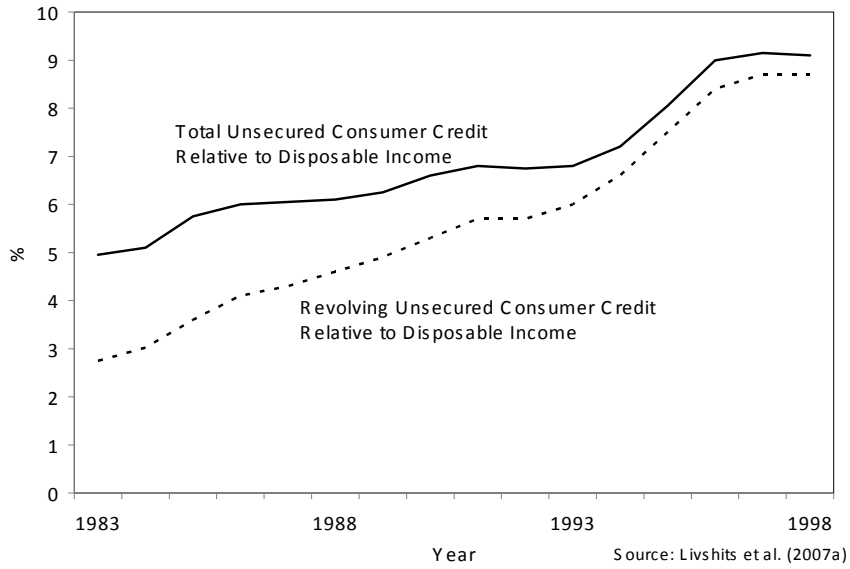


Figure 3: Total Unsecured Consumer Credit versus Revolving Consumer Credit in the US, 1983-1998.

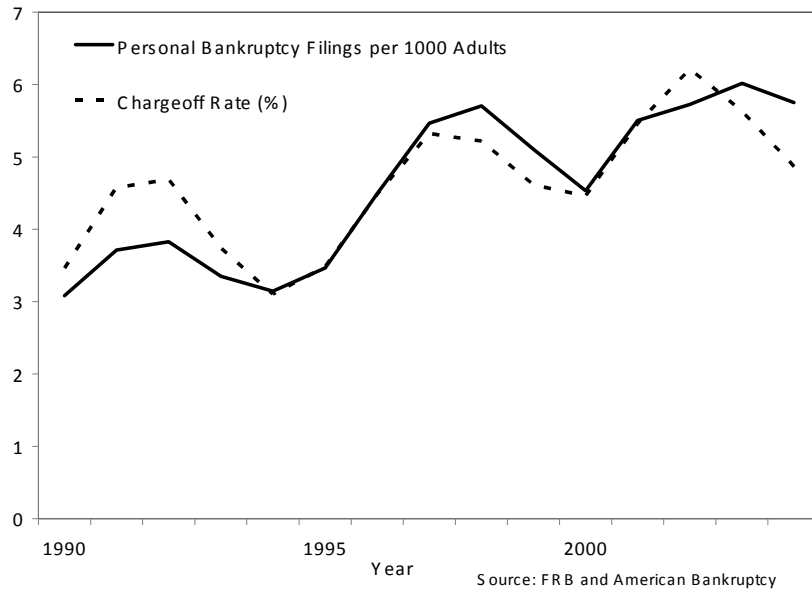


Figure 4: Bankruptcy Filing Rate and Charge-off Rate in the US, 1990-2004.

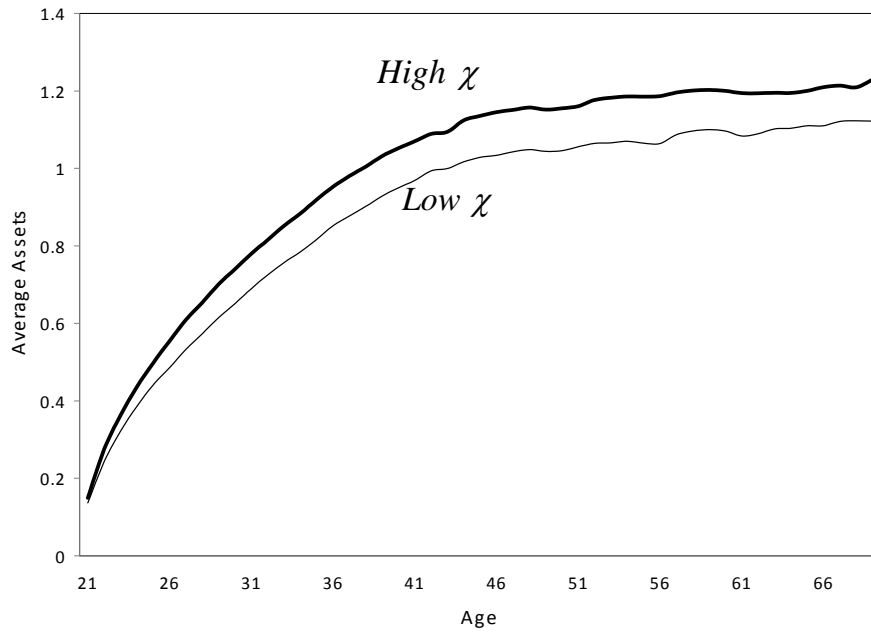


Figure 5: Average assets by age.

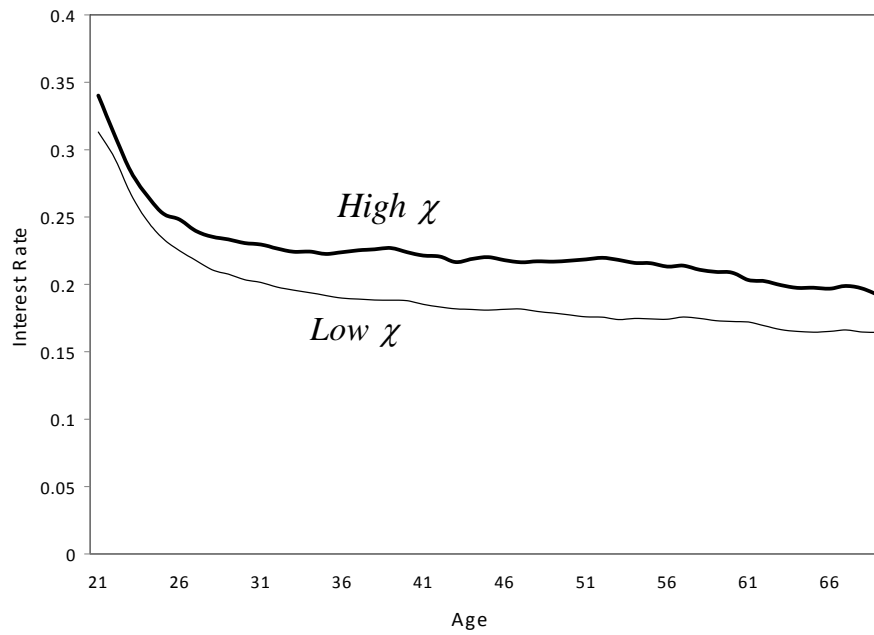


Figure 6: Average interest rate by age.

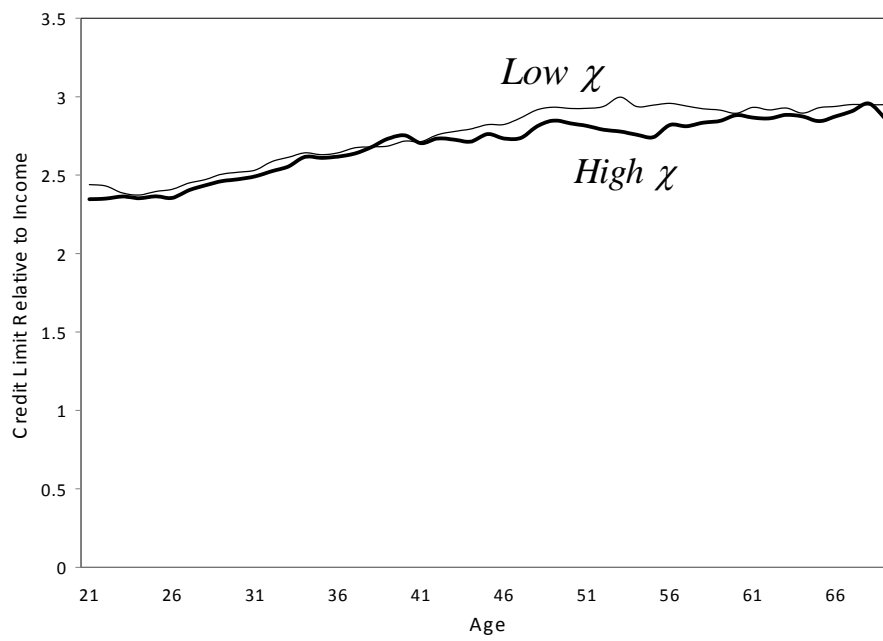


Figure 7: Average credit limit by age.