The Evolution of Unsecured Lending Terms and Information Technology

Andrew Davis
Acadia University

Jiseob Kim
Korea Development Institute

January 27, 2015*

Abstract

We consider two stylized changes in the cross-sectional distribution of credit card contracts over time: the increasing variance in interest rates and the increasing variance in credit limits. This pattern holds even when conditioning on several traditional observables, notably income. Various corroborating evidence suggests that these changes reflect a financial sector with increasing access to and sophistication in using a wider set of consumer information. Within a life-cycle framework, we show that a model of unsecured lending using contracts reflecting modern credit cards can qualitatively capture these two main trends along with several other moments relating to the evolution of lending terms following an increase in the information set available to the financial sector. The spread of information technology generates an increase in household welfare accruing primarily to younger households.

*Andrew Davis: andrew.davis@acadiau.ca. Jiseob Kim: jiseob83@gmail.com. The authors thank defense, seminar, and workshop participants at the University of Rochester, Korea Development Institute, the Korean Economic Association Economic Joint Conference, Korea University, Finance Canada, and the Bank of Korea for helpful comments and advice. All errors and oversights are our own.

1 Introduction

Unsecured credit, primarily exemplified by credit cards, has long since become a common feature of daily life in America and many other countries, with the 2010 edition of the Survey of Consumer Finances showing that some 65% of American households hold at least one credit card. While for many credit cards are effectively only a means of payment, rather than an actual instrument for borrowing, for many others the interest rate they pay on credit card debt and their credit limit are important variables in their economic decision-making. Empirical evidence for this is presented in Gross and Souleles (2002), who find that exogenous changes in a household's credit card interest rate or credit limit generate significant impacts on consumption/savings behavior.

Thus, given that credit access and credit pricing influences consumer actions, it is natural to wonder what determines how much consumers pay and how much they can borrow in the unsecured credit market. In economic literature, interest rates are often determined through a zero-profit condition that takes into account default probabilities and the costs of default, while borrowing capacity is typically determined as some function of what can be pledged as collateral. However, in the case of credit cards and more generally other forms of unsecured debt, collateral is not required as security against the value of the loan. Rather, the lender is relying on other incentives faced by the consumer not to default (instead of forfeiting collateral), such as future exclusion from financial markets or the cost of bankruptcy proceedings. This implies the necessity of credit limits (credit rationing) for unsecured debt. Absent such limits, many would simply borrow large amounts and then immediately default. In turn, this implies that credit limits should be based on individual characteristics given optimizing lenders, just as are interest rates, and that credit limits and interest rates cannot be determined separately from each other.

In this paper, we first document how credit limits and interest rates have evolved over time in the credit card market, relying on data from the Survey of Consumer Finances. First, we update the facts on how the variance of interest rates charged on credit cards is increasing over the length of our sample. Second, we examine a similar pattern in terms of credit limits, the cross-sectional variance of which has also increased over time. In particular, we believe that the empirical information presented on credit limits is new to the literature.

We also discuss how the determinants of credit limits have changed over time, reflecting changes in the information available to lenders, or changes in how they use the information available to them in setting contract terms. We find that a much wider set of information has become strongly correlated with the terms of unsecured credit contracts. While at the beginning of our sample in the 1980s, income was the only strong predictor of an individual's credit terms, in the 2010s information including homeownership status, educational attainment and credit history has become much more relevant. Not only does this reflect the usage of more information, but the usage of information more strongly linked to default - seemingly attempting to measure permanent income and unemployment risk.

These stylized facts motivate our model, which relies on imperfect information in the financial industry. We assume lenders cannot know everything about a prospective borrower; they can only decide on how much credit to extend and what interest rate to charge based on what they can observe about any given individual and their expectations given what they know about the population. Competition in the financial sector and household preferences determine the actual terms of the contract struck with any given household. We do not pursue an optimal contracting approach here, rather we restrict the financial sector to offering contracts solely in terms of an interest rate and credit limit in an attempt to mimic the reality of the modern unsecured lending market.

For our quantitative exercise, we examine a world where the information available to lenders improves in quality. We find that this informational change generates model outcomes that qualitatively match several trends observed in the data, particularly the increasing cross-sectional variances in rates and limits. Intuitively, with relatively poor information about the riskiness of borrowers, lenders initially operate in an equilibrium that is more 'pooled', where the best they can do is aggregate risks together and set common interest rates and credit limits. As information improves, lenders can increasingly extend credit to each individual based upon their individual characteristics, generating increases in the variance of credit limits and interest rates. The improved information available to lenders alters the composition of the pool of households which decide to borrow, driving out households revealed to be severe credit risks, but drawing in secure households which couldn't previously borrow at favorable terms that reflected their true financial position.

We conclude our numerical exercise with some welfare calculations. We find that the change from the low-information baseline to the high-information alternative generates a small positive increase in social welfare. More interestingly, not all households benefit from the change in information technology. With more accurate credit pricing, households with a high risk of expenditure shocks are often driven out from the lending marketplace, to their detriment, while the rest of society enjoys superior terms on their loans. The removal of this insurance through cross-subsidization, however, is dominated by the efficiency gains from better pricing. We believe this phenomena of households being driven out of the unsecured credit market could be coupled with the growth of the payday loans industry in future work.

Empirically, other related papers include Ausubel (1999), who runs a series of randomized trials demonstrating the existence of informational imperfections in the credit card market. Karlan and Zinman (2009) also conduct similar randomized trials experiments. Dey and Mumy (2009) compute a series of cross-sectional regressions with credit card interest rates and credit limits as the endogenous variables in a reduced-form framework, though they do not consider changes over time. We are not aware of any empirical papers which examine credit limit data to the extent we have.

Theoretically and computationally, Livshits, MacGee and Tertilt (2011) is perhaps the most similar paper to our work. They also consider a model of lenders with imperfect information on potential borrowers, with fixed costs of contracting resembling those in this paper. We extend on their work primarily through our empirical work on credit limits, and the inclusion of that data in disciplining the model and computing results. Our model also incorporates a full dynamic life-cycle instead of their two-period framework, and both these factors drive somewhat different quantitative results. Also closely related, Athreya, Tam and Young (2012) consider a series of data moments and rationalize them, as we do, through an informational mechanism, but their model generates distinctly different welfare predictions.

Beyond these papers, Mateos-Planas (2011) presents a model of the joint determination of the credit limit and the interest rate under imperfect information which is somewhat related to our. However, the model presented has to make some unrealistic assumptions concerning how loan characteristics are determined that we can sidestep. Chatterjee, Corbae and Rios-Rull (2011) work in a model with the same broad characteristics, but their empirical exercises do not deal with either interest rates or credit limits. Narajabad (2012) links changes in the information lenders have about borrowers to increases in consumer default, but assumes a fixed interest rate on all lending. Drozd and Serrano-Padial (2013) also deal with information technology and unsecured lending, but discuss a story dealing with improved debt recovery following default rather than increased information about borrowers. Sanchez (2012) also discusses a similar informational story to ours, but takes a much different approach to contracting.

Several other papers also consider similar models though with different goals than our work. Drozd and Nosal (2011) discuss a model with long-lasting relationships between borrowers and lenders, much as we do, but are not concerned with examining trends in the data over time. Guler (2010) studies changes in information in the mortgage market rather than the credit market, also with a life-cycle component, finding similar welfare results as we do.

The remainder of the paper is organized as follows. Section 2 discusses several empirical findings from the SCF data. Section 3 presents our formal model of unsecured consumer lending. Section 4 covers our calibration and main quantitative results, Section 5 conducts a brief welfare analysis, and the paper subsequently concludes.

2 Data

The primary data source for this paper is the Survey of Consumer Finances (SCF), as collected by the US Federal Reserve. The SCF was initially administered in 1983,[1] and every three years following. However, the survey has generally expanded over time, and only becomes useful to us in 1989 when detailed questions about credit card holdings were introduced. Further, it was not until 1995 that all survey questions we use were included in the survey. Therefore our sample consists of up to eight series of surveys, culminating in 2010, with 4000-5000 households included in each edition of the survey. While this intermittent sampling is limiting in some ways, we feel the trends we present are sufficiently robust for the data we have to suffice. There is a small panel element in the SCF, when the 2007 panel was reinterviewed in 2009, but we do not exploit that information in this work.

In particular, we draw upon the information the SCF contains on both credit card credit limits and interest rates. The SCF also contains data on multiple credit card holdings, other sources of debt, various measures dealing with credit history, interest rates on other loans, and many typical 'characteristics' questions.

The first fact we document, and the one most examined by prior literature, is the increased cross-sectional variance in credit card interest rates. In Figure 1, we plot both the cross-sectional standard deviation in quoted interest rates and also the interquartile range of the distribution.[2] The standard deviation across the population of credit card holders

[1]There was an orphan 1962-63 survey as well, which we do not use.

[2]We plot nominal rates, as these numbers may be more familiar. Relative to inflation, credit card interest

has risen steadily from 4.3 in 1995 (the first year for which relevant data exists in the SCF) to 6.1 in 2010. A similar pattern emerges if we instead use interquartile range instead of variance, and in general all the data we examine are robust to different presentations unless otherwise noted.

Figure 1: Variance measures, credit card interest rate

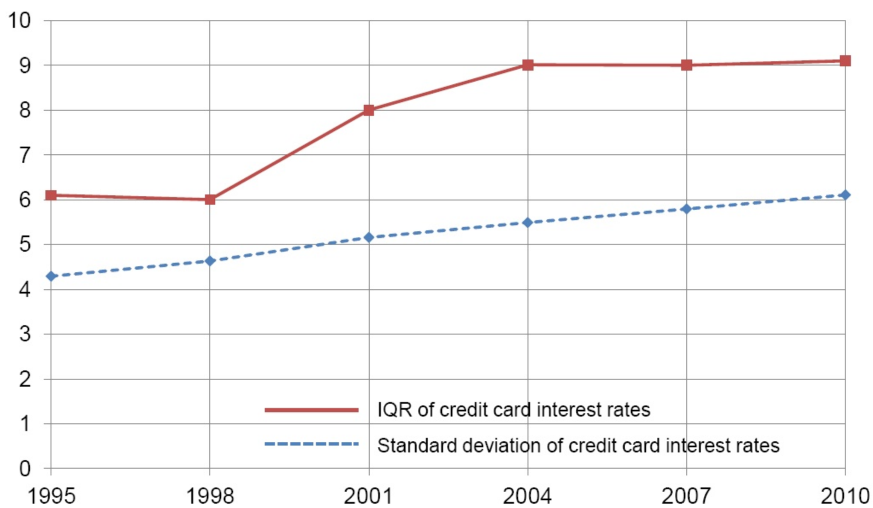

As a brief aside, we also present the mean and median of quoted credit card interest rates in Figure 2. There's no significant trend, though we do observe a sharp decline in rates in 2004 and a rebound that approaches average 1990s rates in 2010. We suspect this might be related to a general decrease in lending standards before the Great Recession and the well-documented subsequent tightening thereafter, but the level of rates is not something we directly seek to address in this work.

In terms of credit limits, we would ideally like to measure the credit extended by lenders against a borrower's future ability and willingness to repay. Changes in this statistic would perfectly reflect changes in the information available to lenders. However, since neither

rates are sufficiently high so that real and nominal measures are approximately equal.

Figure 2: Level measures, credit card interest rate

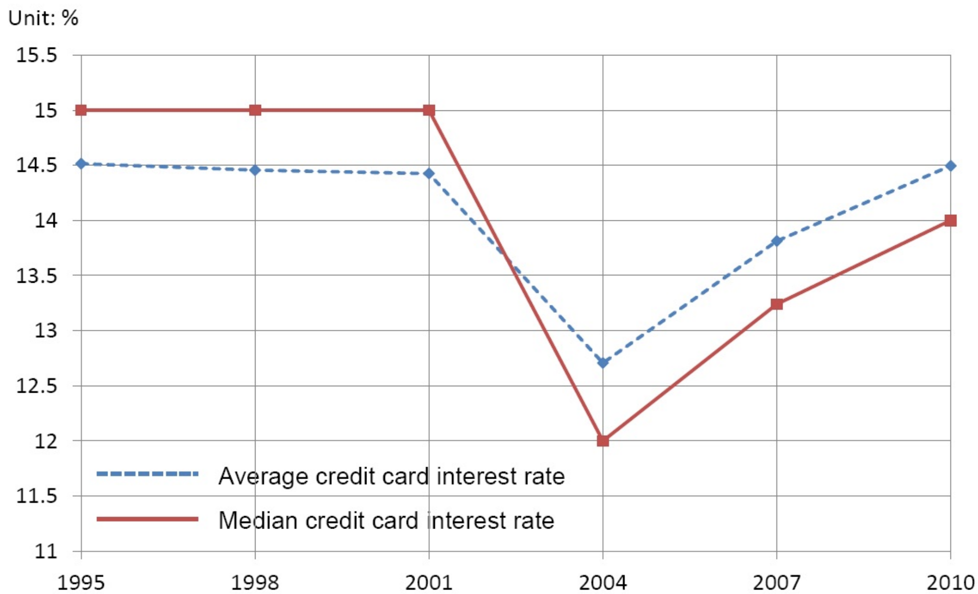

we nor lenders can observe the future, we instead use credit limits relative to income as our main statistic. We believe that at the beginning of our sample, in 1989, lenders could observe income. This is supported by a large number of individuals reporting that their credit applications were denied due to insufficient income.[3] If we suspect that over time lenders are accumulating access to more information, or are simply developing more capability to use and deploy information about the true riskiness of any individual, the variance in this statistic should increase as lenders base their decisions on other pieces of information more correlated with future default risk. This pattern is precisely what we observe, save in the recession year of 2010, when there's plenty of evidence that lenders had tightened standards considerably (see Figure 3 and 4).

Also note that it is not simply the variance that's increasing with respect to credit limits. Turning to the levels, it's unsurprising that the real credit limit has increased across households. Given that incomes and wealth are increasing over time, credit limits would

[3]This information is self-reported by households in the SCF.

Figure 3: Standard deviation, credit limit to income ratio

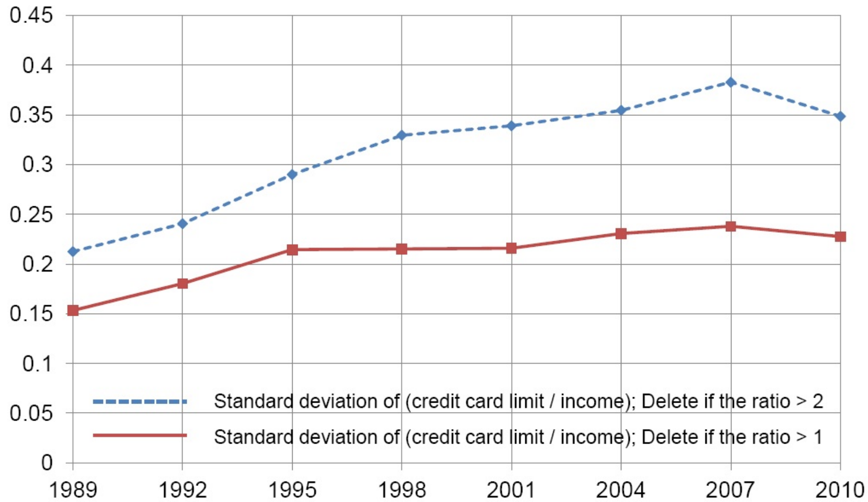

naturally increase absent any changes in financial markets. The trend persists, however, even when we condition on income, and both series are plotted in Figure 5. While our paper is primarily concerned with cross-sectional heterogeneity, we will address this strong trend in our quantitative results as well.

This series of figures we have discussed capture the key trends we want to explain in this paper, and the story we will propose to explain them revolves around informational developments in the financial sector. However, there are numerous stories aside from our informational mechanism one could tell to rationalize, at least initially, these movements. These could include increasing income inequality (which we partially address by conditioning the credit limit changes on income) or demographic changes.[4] Consequently, we now document a variety of corroborating evidence to support our proposed mechanism.

To start, we can examine the correlation between credit limits and observables. While at the beginning of our sample, current income is most correlated (out of the relevant vari-

[4]Drozd and Serrano-Padial (2013), mentioned earlier, discusses changes in the effectiveness of debt collection operations.

Figure 4: Interquartile range, credit limit to income ratio

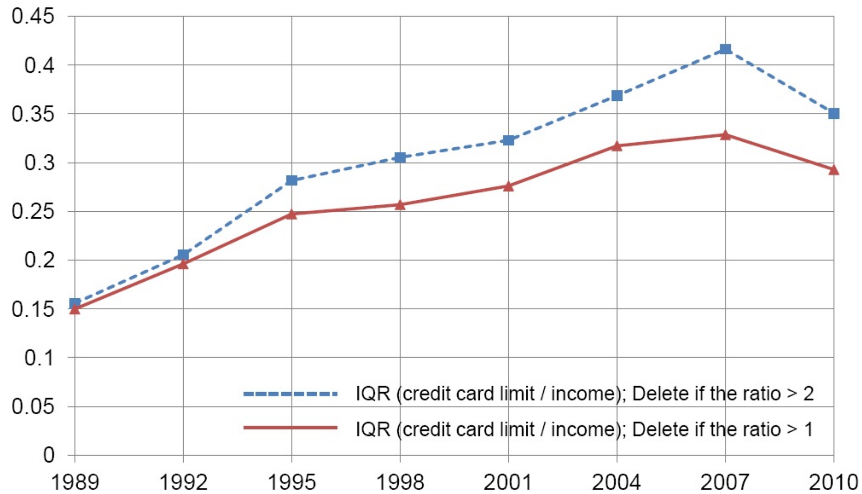

ables within the SCF) with credit limits, by 2010 other characteristics such as educational attainment, credit history and homeownership status have all become much more correlated with credit limits than current income. We demonstrate this through a simple regression, run separately for each edition of the SCF in our sample period, the results from which are presented in Table 1, specified as follows:[5]

$$\begin{aligned} log(creditlimit) &= \beta_0 + \beta_1 log(income) + \beta_2 I(creditrejected) \\ &\quad + \beta_3 I(homeowner) + \beta_4 age + \beta_5 age^2 + \beta_6 I(educ) + \beta_7 I(\#cards) + \epsilon \qquad (1) \end{aligned}$$

Several distinct patterns stand out from this series of estimates. First, having been previously rejected for new credit has become much more strongly linked with one's credit limit. We take this as evidence on the rise of credit bureaus and more generally the availability of more information about an individual's credit history. This is reflected in self-reported

[5]Our 'educ' variable is a dummy based on no high school/high school/some college/college, and $I(\#cards)$ reflects dummy variables for households with 1,2,3,.. household credit cards.

Table 1: Regression results across SCF waves[6]

	1989 SCF	1992 SCF	1995 SCF	1998 SCF	2001 SCF	2004 SCF	2007 SCF	2010 SCF
log(income)	0.279***	0.225***	0.231***	0.235***	0.267***	0.257***	0.263***	0.304***
	(0.018)	(0.014)	(0.015)	(0.013)	(0.013)	(0.012)	(0.011)	(0.012)
I{credit turned down}	-0.101*	-0.212***	-0.323***	-0.508***	-0.588***	-0.624***	-0.684***	-0.543***
	(0.058)	(0.044)	(0.045)	(0.052)	(0.057)	(0.063)	(0.066)	(0.048)
I{own house}	0.054	0.170***	0.142***	0.382***	0.419***	0.471***	0.564***	0.536***
	(0.055)	(0.041)	(0.043)	(0.051)	(0.049)	(0.053)	(0.057)	(0.047)
age	0.026***	0.028***	0.012*	0.030***	0.022***	0.033***	0.035***	0.030***
	(0.009)	(0.007)	(0.006)	(0.007)	(0.007)	(0.007)	(0.007)	(0.006)
age^2	-0.0002**	-0.0002***	-0.00002	-0.0002***	-0.0001*	-0.0002***	-0.0002***	-0.0002***
	(8.6E-5)	(6.1E-5)	(6.1E-5)	(6.5E-5)	(6.4E-5)	(6.7E-5)	(6.5E-5)	(5.5E-5)
I{high school}	0.052	0.034	0.05	0.155*	0.319***	0.231**	0.055	0.248***
	(0.07)	(0.06)	(0.063)	(0.086)	(0.078)	(0.095)	(0.09)	(0.08)
I{some college}	0.074	0.097	0.179***	0.251***	0.429***	0.265***	0.253***	0.304***
	(0.079)	(0.071)	(0.066)	(0.089)	(0.083)	(0.1)	(0.094)	(0.083)
I{college or more}	0.237***	0.223***	0.274***	0.385***	0.611***	0.500***	0.416***	0.638***
	(0.076)	(0.058)	(0.062)	(0.083)	(0.076)	(0.093)	(0.088)	(0.077)

Stars indicate statistical significance based on the standard error in parentheses. * represents significance at the 10% level, ** at the 5% level, and *** at the 1% level.

Figure 5: Average credit limit, credit limit to income

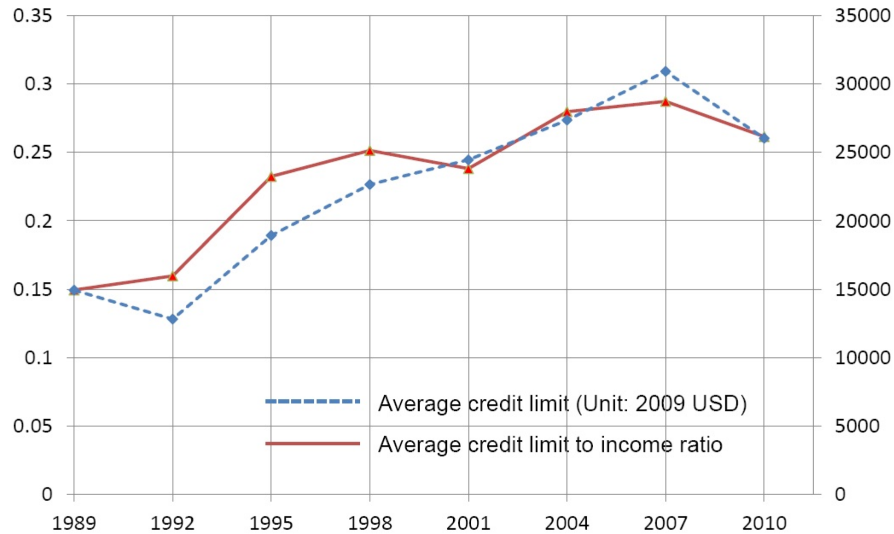

data within the SCF on why consumers were rejected for new credit: the share of rejections attributed to credit bureaus has risen by an order of magnitude from 1989 (3%) to 2010 (27%).[7]

Education has also become increasingly correlated with credit limits over the sample period, though the mechanism of action is not quite as clear or direct, at least from our perspective. We theorize that educational attainment could capture a component of permanent income, and consequently is related to repayment ability in the future. Further, given how strongly education and unemployment are linked,[8] education may be more indicative of the probability of future default than current income. Similarly to education, whether a household is a homeowner or not has also become increasingly predictive of credit limits over time. While the financial sector could likely have observed homeownership status at the beginning of our sample, it's quite possible that the usage of this information over time

[7]Given the self-reported nature of these reasons, we cannot identify if this really reflects a significant change in utilization of credit bureaus - perhaps the financial sector today finds it easier or more convenient to claim credit history instead of specific components of credit risk, such as income.

[8]Riddell and Song (2011).

has become more sophisticated, resulting in the increasingly significant correlation.

Another way to provide support for our proposed explanation is to look at the correlation between how borrowers are evaluated by different parts of the financial sector. In particular, we note that the correlation between an individual's credit limit and the interest rate paid on their mortgage decreases significantly over time (see Figure 6).[9] That is, as an individual is judged less risky by unsecured lenders (credit limit increases) they are also judged as less risky by mortgage lenders (interest rate decreases). This might reflect greater informational sharing between financial institutions, possibly stemming from centralized credit agencies, or perhaps more simply that all financial institutions are obtaining increasingly better information about borrowers - effectively, the signals each is observing are less noisy, and consequently more correlated.

Figure 6: Correlation between mortgage rates and (log) credit limit

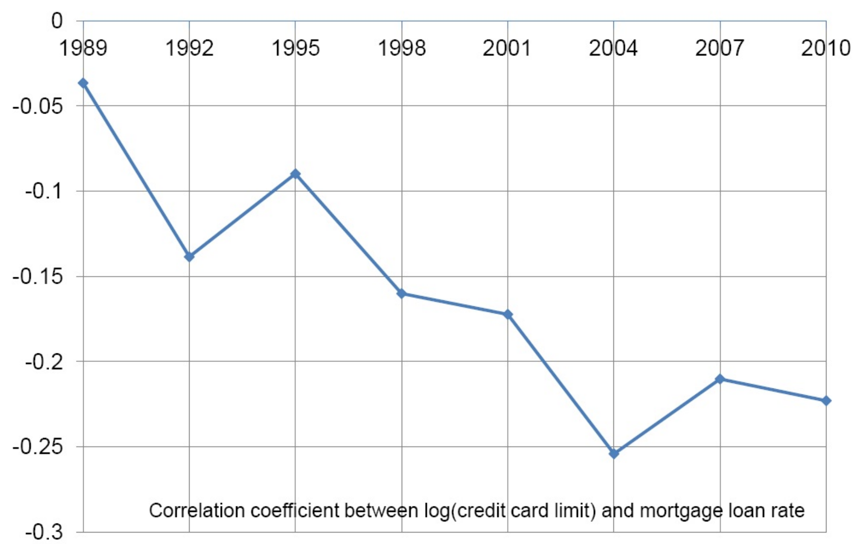

All together, we take this collection of facts and hypothesize that changes in the information available to banks about borrower quality has driven these trends.[10] We stress that we

[9]This correlation is calculated for each edition of the SCF by using only mortgages that were originated since the prior SCF. For example, the calculation uses mortgages originated in 2008-2010 for the 2010 SCF.

[10]Further regressions involving the mortgage rate yielded very similar results compared to the prior re-

are not making strong claims about these pieces of supportive evidence, but are only taking them as motivation for our model mechanism.

The following model thus presents an environment where optimizing banks make interest rate and credit limit decisions based on incomplete information about borrowers, where the environment allows us to vary the information available to lenders over time to see if these empirical findings can be rationalized quantitatively.

3 Model

We now construct a model of the consumer credit card market in order to evaluate our proposed theory - that the positive trends in the cross-sectional variance of credit card interest rates and credit limits, which we take as representative of unsecured debt in general, are a result of increased information available to the financial sector. In turn, this allows for more individualized contracts. The basic intuition here is that with more information, the economy moves from more of a 'pooling' equilibrium to a 'separating' one, and the variance measures thus increase. In describing the model, we define everything formally in terms of the baseline model, that is the model in which the financial sector is relatively ignorant.

3.1 Basic Environment

Time is discrete. The economy is comprised of a continuum of households and a financial sector. Households are endowed with a stochastic income stream, z, a stochastic i.i.d. expenditure shock, θ, have the option to borrow - denote borrowing as a - but not to save, forcing $a \leq 0$, and have a distinct life-cycle, aging and eventually dying after T periods. There is no retirement or youth in the model - households die at the end of their working lives and newborn households immediately enter the labor market.

gression specification.

A household i, of age j, has their income, y_{ij}, evolve via the following process:

$$\ln(y_{ij}) = \ln(\bar{y}_j) + \ln(z_{ij}) \tag{2}$$
$$\ln(z_{ij}) = \rho \ln(z_{ij-1}) + \varepsilon_{ij}$$

where $\bar{y}_j$ is the average income of age j households, z_{ij} is a persistent innovation to income, and ε_{ij} is an i.i.d shock.

Households of any age, j, have the opportunity to seek credit from the financial sector. Credit is unsecured, as are credit cards, and a credit contract consists of a credit limit, L, and a one-period price of credit, q (which implicitly defines the interest rate). Financial intermediaries can observe households' state variables, except for the expenditure shock θ. However, they know the distribution of θ over the population. These values, (L, q), will be jointly pinned down by a competitive lending sector and consumer maximization over different types of contracts that are consistent with driving out profits based on the expectations of firms given their knowledge.

Agents, armed with their credit contract, decide whether or not to default on their obligations and then make their corresponding borrowing choices for the next period. This timing assumption that lenders must extend credit terms to households before borrowing decisions are made will play a key role in our model. If households made their borrowing choice simultaneously with the determination of their contract terms, then firms could infer the expenditure shock from the household's preferences over different zero-profit contracts. We feel this timing reflects the real decisions faced by households - credit cards work in exactly this manner, with the terms being decided upon before borrowing occurs.

Throughout, we prohibit agents from saving for a very basic reason - we wish to only consider those agents who are borrowing. In the data, we of course observe many wealthy households holding credit cards. However, it is unlikely the terms on those cards accurately represent the cost of unsecured credit. Such households are probably only using their credit cards as a method of payment, and not as an actual borrowing vehicle. If we extended the model to include personal savings, very little would change, since in calculating all the relevant moments we investigate, households with significant positive savings would have to be dropped.

This assumption is not new to the literature. Chatterjee and Eyigungor (2012) impose this same restriction, both for tractability and clarity, though in terms of a different context - sovereign borrowing. Further, households in the data that do carry positive credit card balances have minimal assets. In the 2010 wave of the SCF, the median such household has net financial assets of \$370. Consequently, we view the prohibition on saving as a weak assumption, if not outright preferred to the alternative.

We later extend the model to give the financial sector information about expenditure realizations, representing increases in information technology. Hence, the timing assumption just discussed is crucial - otherwise the firms would in effect not learn anything new from this change. Carrying out this experiment quantitatively and comparing the results to the discussed data will form our main results.

3.2 Lenders

Suppose a household with observable debt level a and without a current credit contract approaches the financial sector for an unsecured loan. The lender needs to offer terms, (L, q) to the household, but cannot reasonably do so without knowing something about the household's financial situation. Screening the household, drawing up the new credit contract and so forth costs χ, and tells the lender what the household's current income is. Age, j, is always observable. It does not, however, reveal any information about the household's expenditure shock.

Given that the credit contract has two parameters, a credit limit and an interest rate,

(L,q), one zero-profit condition is insufficient to pin down terms, as there are potentially many possible contracts that satisfy zero profit. The contract offered then is the one that maximizes household utility, conditional on the bank's expectations over the unobserved expenditure shock. Alternatively, one can interpret this as the household choosing between the portfolio of contracts that satisfy zero-profit, where zero-profit is conditional on the bank's expectations over the population.

That is, when $j \leq T-2$, a bank solves the following problem:

$$\max_{L,q} E_\theta \left[V_j^R (L,q,a,z,\theta)\right] \tag{3}$$

$$\begin{aligned}
s.t.\quad \pi_j (L,q,a,z) &= E_\theta \left[\begin{array}{c} I\left\{V_j^D (z,\theta) > V_j^R (L,q,a,z,\theta)\right\} \cdot 0 \\ +I\left\{V_j^D (z,\theta) \leq V_j^R (L,q,a,z,\theta)\right\} a_j' (L,q,a,z,\theta) \end{array} \right] q \\
&- (1 - D_j (L,q,a,z))\, a \\
&+ \frac{1}{1+r_f} E_{\theta,z'|z} \left[\begin{array}{c} I\left\{V_j^D (z,\theta) > V_j^R (L,q,a,z,\theta)\right\} \cdot 0 \\ +I\left\{V_j^D (z,\theta) \leq V_j^R (L,q,a,z,\theta)\right\} \times \\ \left\{ \begin{array}{c} I\left\{0 < \pi_{j+1}\left(L,q,a_j',z'\right) < \chi\right\} \pi_{j+1}\left(L,q,a_j',z'\right) \\ +I\left\{\pi_{j+1}\left(L,q,a_j',z'\right) \leq 0\right\} 0 \\ +I\left\{\pi_{j+1}\left(L,q,a_j',z'\right) \geq \chi\right\} \chi \end{array} \right\} \end{array} \right] \\
&= \chi
\end{aligned}$$

Here, $V_j^R(.)$ is the household's value function when it repays its debt, and $\pi_j(.)$ is the bank's net profit function. As discussed, a zero profit-condition is insufficient to pin down a two-parameter contract. We assume that the consumer effectively chooses the contract that it prefers over all zero-profit contracts. This is not quite true, however, because the financial sector cannot observe the expenditure process in the baseline model. Hence, the bank offers a contract which the consumer would prefer in expectation over the expenditure process.

Let $f_L(a, z, j)$ and $f_q(a, z, j)$ be optimal solution of the banks' problem.

The contribution to profit in the present period (the first term of π_j) is broken into pieces depending on whether or not the consumer defaults (V^D). In the first term, inside the expectation, there is an outflow of resources from extending new loans at price q. If the consumer defaults ($V^D > V^R$), then they are excluded from credit markets and do not receive loans. If the consumer stays current on their debt, then they take a new loan, a', and they repay their old loan, a, in the second term, where D the probability of default. If default occurs, no repayment happens and thus that term vanishes.

Continuing the breakdown of the lender's problem, the bank discounts future profits at the risk-free rate. These profits depend on whether or not the contract terms are updated in the future and whether default occurs. If default does not occur and if the expected profit, from the firm's perspective, becomes negative or rises above χ, then the contract gets reset. Otherwise, the same (L, q) are carried forward to the next period. Again, the zero-profit condition is only zero-profit with respect to the menu cost upon renegotiation of terms - the bank can make some positive flow profits off an existing relationship since the consumer can't do better by shopping around.

The zero-profit condition therefore depends on the current period's income, the cost of a potential default, and the value of the future relationship to the bank. The cost of origination pins down the ongoing flow profit the financial sector can obtain in the future. This implies that unmodeled factors that create profits in the financial sector may also be captured in χ, instead of treating it as a strict menu cost. These could include costs for an individual to move to a new financial institution or a bank's ability to make positive profits due to imperfect competition.[11] This aligns with our previous discussion on the problem. The consumer is unwilling to tolerate the bank making a significant profit - they realize that they

[11] The existence of a nontrivial degree of imperfect competition is argued for in Calem and Mester (1995) and elsewhere.

can get a better deal by searching around and would move to a new credit provider if flow profits are too high. Since the firm wants to maintain the relationship, as the expected profit is slightly positive, they offer better terms to the individual in order to keep their business. That is, as long as the expected profit to the bank is less than χ, the consumer cannot do better by going to a competitor, since that will incur the cost and then the zero-profit terms offered by that competitor would be dominated by the household's current credit terms.

Mechanically, as the consumer ages, the lender's problem needs to be modified to reflect the incentives imposed by the end of life. When $j = T-1$, the bank makes the final loan of the household's life cycle (as no firm would lend in the final period, since repayment can never occur):

$$\max_{L,q} E_\theta \left[V^R_{T-1}(L,q,a,z,\theta) \right] \tag{4}$$

$$\begin{aligned} s.t. \quad \pi_{T-1}(L,q,a,z) = {} & E_\theta \begin{bmatrix} I\left\{V^D_{T-1}(z,\theta) > V^R_{T-1}(L,q,a,z,\theta)\right\} \cdot 0 \\ +I\left\{V^D_{T-1}(z,\theta) \le V^R_{T-1}(L,q,a,z,\theta)\right\} a'_{T-1}(L,q,a,z,\theta) \end{bmatrix} q \\ & - \left(1 - D_{T-1}(L,q,a,z)\right) a \\ & + \frac{1}{1+r_f} E_{\theta,z'|z} \begin{bmatrix} I\left\{V^D_{T-1}(z,\theta) > V^R_{T-1}(L,q,a,z,\theta)\right\} \cdot 0 \\ +I\left\{V^D_{T-1}(z,\theta) \le V^R_{T-1}(L,q,a,z,\theta)\right\} \pi_T\left(a'_{T-1}, z'\right) \end{bmatrix} \\ & = \chi \end{aligned}$$

Here, things have simplified somewhat, reflecting the fact that future credit terms are of no concern, since no lending will happen next period. This simplifies the continuation value of the relationship considerably with respect to the general problem, but otherwise things remain unchanged.

Finally, when $j = T$, the bank's profit simplifies again to and is given by:

$$\pi_T (a, z) = E_\theta \left[\begin{array}{c} I \{\bar{y}_T z + a - \theta \geq \underline{c}\} |a| \\ +I \{\bar{y}_T z + a - \theta < \underline{c}\} (\bar{y}_T z + a - \theta - \underline{c}) \end{array} \right] \tag{5}$$

This terminal condition reflects how much repayment the bank can expect on any outstanding debt. We kill off the option to default, insisting that everyone repay all debt possible before dying up to the consumption, $\underline{c}$, reflecting social services and welfare, which is required to keep the problem well behaved. One might ask why the lending sector would permit anyone to enter the final period holding sufficient debt to the point where full repayment might not be possible, but imposing stricter lending terms in prior periods has a cost, and tightening lending standards to the point where full repayment is guaranteed may not be optimal. Also note that, consistent with US law, there is no intergenerational transfer of debt - one cannot pursue newborn households for payments of debts left by passed agents.[12]

3.3 Households

For any household of age $j \leq T - 2$, each period they face the decision to default, V^D, or not to default (repay), V^R, solving the following problem:

$$V_j (L, q, a, z, \theta) = \max \left\{ V_j^D (z, \theta), V_j^R (L, q, a, z, \theta) \right\} \tag{6}$$

[12] We briefly considered a world where debts, were passed on to freshly minted households, but we saw no major quantitative changes.

where the value of default is given by the following program:

$$V_j^D(z,\theta) = u(c) + \beta E_{\theta',z'|z}\left[\gamma V_{j+1}(L',q',a'=0,z',\theta') + (1-\gamma)V_{j+1}^D(z',\theta')\right] \tag{7}$$

$$s.t.\ \ c = \left\{ \begin{array}{ccc} (1-\kappa)\bar{y}_j z - \theta & if & (1-\kappa)\bar{y}_j z - \theta \geq \underline{c} \\ \underline{c} & if & (1-\kappa)\bar{y}_j z - \theta < \underline{c} \end{array} \right\}$$

$$L' = f_L(a'=0, z', j+1)$$

$$q' = f_q(a'=0, z', j+1)$$

The structure of this program can be broken down in parts. The value of default depends upon today's consumption and a discounted continuation value. Once a household defaults, it is excluded from financial markets by the banking sector. With some probability, γ, the household will be forgiven next period and be allowed to borrow, with any defaulted-upon debts forgotten upon, i.e. $a' = 0$. Otherwise, the household remains in financial autarky, and consequently does not have (L, q, a) as state variables. The subscript j reflects the aging of the household.

Turning to the budget constraint, the defaulting agent simply consumes everything when it is without access to financial markets, inclusive of κ which reflects costs of default. These costs include attempts to repossess or reclaim the defaulted-upon debt, restrictions on what can be owned and purchased due to the bankruptcy process, etc. We also impose a minimum consumption value, $\underline{c}$, reflecting basic government benefits and social services, that a household enjoys even if their assets do not permit them to finance it.

The functions L' and q' capture that a rational household realizes borrowing decisions today influence credit terms in subsequent periods. f_L and f_q denote the agreed credit contract coming out of the financial sector's optimization for a household of a given type. The value of default thus depends on the credit terms that can be obtained upon readmittance to the financial sector. For a defaulted household, however, these functions are effectively trivial since readmittance to the financial sector always happens with $a' = 0$ after the period of exclusion is served.

When the household instead chooses to stay current on its debt, it solves:

$$V_j^R(L,q,a,z,\theta) = \max_{a'} u(c) + \beta E_{\theta',z'|z}\left[\begin{array}{c}\phi V_{j+1}(L',q',a',z',\theta') \\ +(1-\phi)V_{j+1}(L,q,a',z',\theta')\end{array}\right] \tag{8}$$

$$\begin{aligned} s.t.\quad c + a'q &= \bar{y}_j z + a - \theta \\ -L &\le a' \le 0 \\ L' &= f_L(a',z',j+1) \\ q' &= f_q(a',z',j+1) \\ \phi &= 1 - I\{0 < \pi_{j+1}(L,q,a',z') < \chi\} \end{aligned}$$

Again, the value function for the household depends on consumption and the household's continuation value. The household which does not default has certain credit terms, (L,q), a certain asset position, a, and a certain income-expenditure type captured by z and θ. Here, the continuation value is a function of whether or not their credit contract will be renegotiated next period, the probability of which we denote as ϕ. ϕ is a function of the household's state variables, but we suppress that notation for simplicity. In the next period, the credit contract will be updated if the expected value of the contract to the financial sector either generates losses or if the contract is generating too much profit - in which case the bank must renegotiate or risk losing the customer.

Households also have the opportunity to borrow, or maximize over a', in an effort to smooth consumption in the face of variable income and expenditure shocks. This leads to a very standard-looking budget constraint. Conversely to the default case, however, the laws of motion that govern the evolution of potential future credit terms are not trivial - the household's debt choice influences whether or not they can expect new credit terms in the next period and what those terms will be if so.

As the end of the life cycle approaches, the household's problem changes somewhat, reflecting the fact that the financial sector is aware that this is the final period in which they can extend credit and expect repayment.

When $j = T-1$, the household solves the following problem, again facing the choice to default or not:

$$V_{T-1}\left(L,q,a,z,\theta\right) = \max\left\{V_{T-1}^{D}\left(z,\theta\right), V_{T-1}^{R}\left(L,q,a,z,\theta\right)\right\} \tag{9}$$

When the household defaults, it solves the following problem:

$$V_{T-1}^{D}\left(z,\theta\right) = u\left(c\right) + \beta E_{\theta',z'|z}\left[\gamma V_{T}\left(a'=0,z',\theta'\right) + \left(1-\gamma\right)V_{T}^{D}\left(z',\theta'\right)\right] \tag{10}$$

$$s.t.\quad c = \left\{\begin{array}{ccc} \left(1-\kappa\right)\bar{y}_{T-1}z-\theta & if & \left(1-\kappa\right)\bar{y}_{T-1}z-\theta \geq \underline{c} \\ \underline{c} & if & \left(1-\kappa\right)\bar{y}_{T-1}z-\theta < \underline{c} \end{array}\right\}$$

Effectively, all that is changing here is that the financial sector is respecting certain terminal conditions imposed by the end of the life cycle. Since no lender is willing to lend in the final period, a household which defaults at age $T-1$ has slightly different continuation values, reflecting the fact it cannot expect any credit next period, simplifying the problem.

When the age $T-1$ household repays, it solves the following problem:

$$V_{T-1}^{R}\left(L,q,a,z,\theta\right) = \max_{a'} u\left(c\right) + \beta E_{\theta',z'|z}\left[V_{T}\left(a',z',\theta'\right)\right] \tag{11}$$

$$s.t.\quad c + a'q = \bar{y}_{T-1}z + a - \theta$$

$$-L \leq a' \leq 0$$

Again, things simplify somewhat as the household ages. The continuation value collapses to a single term without any uncertainty about future credit terms

Things simplify even more when $j = T$, and the household solves the following problem:

$$V_T(a, z, \theta) = V_T^R(a, z, \theta) \tag{12}$$

At age-T, there is no further opportunity to borrow. Hence, the default option is absent at age-T.

$$V_T^D(z, \theta) = \left\{ \begin{array}{ccc} u((1-\kappa)\bar{y}_T z - \theta) & if & (1-\kappa)\bar{y}_T z - \theta \geq \underline{c} \\ u(\underline{c}) & if & (1-\kappa)\bar{y}_T z - \theta < \underline{c} \end{array} \right\} \tag{13}$$

While this seems trivial, note that the value of default is needed to define V_{T-1}^D. The value of repayment is then:

$$V_T^R(a, z, \theta) = \left\{ \begin{array}{ccc} u(\bar{y}_T z + a - \theta) & if & \bar{y}_T z + a - \theta \geq \underline{c} \\ u(\underline{c}) & if & \bar{y}_T z + a - \theta < \underline{c} \end{array} \right\} \tag{14}$$

When $\bar{y}_T z + a - \theta < \underline{c}$, the remaining debt is eaten as a loss by the financial sector, and is not transmitted to new household by assumption as discussed previously.[13] Conversely, when $\bar{y}_T z + a - \theta \geq \underline{c}$, the household repays in full. When $\bar{y}_T z + a - \theta < \underline{c}$, along with $\bar{y}_T z - \theta - \underline{c} \geq 0$, the bank recovers $\bar{y}_T z - \theta - \underline{c}$. When $\bar{y}_T z + a - \theta < \underline{c}$, along with $\bar{y}_T z - \theta - \underline{c} < 0$, the debt is sufficiently large that the bank recovers nothing and books the full loss of the outstanding debt. With saving prohibited as well, the age-1 household's initial asset/debt position is always zero.

3.4 Informational Experiment

The quasi-experiment we carry out in our quantitative work models an improvement in information technology corresponding to our motivation for this paper. Initially, given the timing of the model, and reflecting the reality of the credit card market, lenders could not

[13]We assume that the expenditure shock, θ, has first priority in the repayment order.

observe the expenditure shock faced by households. We now assume that lenders can observe each household's realization of θ, reflecting their increased knowledge about households.

This extra information changes the financial sector's problem slightly. Previously, the financial sector was forced to take an expectation over θ when solving its optimization problem. Now, since lenders know θ for each household, the expectation vanishes and banks can compute the relevant quantities directly.

A somewhat more sophisticated option here would be to instead give the firm a signal about the expenditure instead of revealing it outright, but such a twist only complicates the framework without any significant enhancement of our story.

4 Results

As just discussed, the experiment that we model is based around increasing the information available to the financial sector, reflecting improvements in information technology over the duration of the SCF sample. Consequently, our baseline model calibration corresponds to the beginning of our sample, in 1989, where banks can only observe household income and debt.

Most of our calibrated parameters are taken from the literature and only one is calibrated to match moments from the data. Our calibration, in full, is presented in Table 2, and discussion of most of the moments follows.

A time period in the model is taken as two years, and the working lifespan (parametrized by T) is taken as 42 years or 21 periods. The yearly discount rate, β, is set at 0.96, with a corresponding risk-free rate, r_f, of 0.04. We also take consumers as having constant relative risk aversion over consumption, with a risk aversion parameter of $\sigma = 2$, that is:

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

Table 2: Main Calibration

Parameter	Description	Value	Target
T	number working years	21	24-65, inclusive
$\bar{y}_j$	common income component	-	average income from SCF
z_{ij}	individual income component	-	STY (2004)[a]
θ_1	low expenditure shock	0.264	26.4% income, Livshits et al. (2007)[b]
ν_1	Pr(low expenditure shock)	0.047	Livshits et al. (2007)
θ_2	high expenditure shock	0.8218	82.2% income, Livshits et al. (2007)
ν_1	Pr(high expenditure shock)	0.0031	Livshits et al. (2007)
γ	Pr(market reinstatement|default)	0.33	6 years exclusion
κ	cost of default	0.355	Livshits et al. (2007)
β	annual discount rate	0.96	Athreya and Neelakantan (2011)
σ	risk aversion	2	Athreya and Neelakantan (2011)
r_f	annual risk-free rate	0.04	Livshits et al. (2007)
χ	contract adjustment cost	0.055	credit card holding rate
$\underline{c}$	minimum consumption	0.181	average subsidies[c]

[a] Storesletten, Telmer and Yaron.

[b] Note average model income is 1.069, not 1.

[c] We take this number extrapolating from the average household food stamp benefit and then multiplying by the fraction of spending food stamps attempt to fill as per federal guidelines. Effectively, our minimum income is what an average household subsisting exclusively on welfare benefits consumes. The finalized value of 0.181 represents 17% of average income. Specifically, we calculate it based on the size of the average household, the national average individual food stamp allowance, and federal guidelines for what fraction of total consumption is expected to be covered by food stamps for a household reliant on welfare.

For the aggregate component of income, $\bar{y}_j$, we calibrate to match average income over time in the SCF data, and for the stochastic process in incomes facing individuals, we adopt the process proposed by and widely adopted from Storesletten, Telmer and Yaron (2004).

Within the model, the expenditure process is the primary driver of default. We take the process proposed in Livshits, MacGee and Tertilt (2007), which is well-suited to our model as they also deal with a framework involving unsecured consumer debt. The expenditure process can take three states - zero, a relatively small expense, and a large expense that often induces default. We also take a cost of default from Livshits et al., which captures that default may involve a good faith repayment, legal costs, etc. This is in contrast to γ, which controls how forgiving the financial sector is, i.e. how long households can expect to remain excluded from financial markets after defaulting. We set γ to imply an average exclusion period of six years, reflecting US law, which states that households cannot re-file for Chapter 7 bankruptcy within this period. Thus, six years reflects how long until households are 'forgiven'. We also need to fix the level of minimum consumption, which we take based on US welfare benefits.[14]

We also assume that a newborn household has basic initial credit terms.[15] The 1989 SCF shows that households aged 24-25 hold credit cards with credit limits that average 10% of their income and an average interest rate of 14.5%. We match both these figures.

The final parameter we have to calibrate is χ, or the frictional cost of switching and renegotiating debt terms. We calibrate this to match the average credit card holding rate.[16] Further, a number of model households possess nontrivial credit card contracts but do not actually borrow intertemporally, and so simply calibrating χ to match the fraction of debt-holding households in the data would not be appropriate. These households would potentially

[14] See the calibration table.

[15] This assumption is made out of computational necessity.

[16] A better moment to match may be the percentage of households which carry a balance on their credit card, which is somewhat lower, but the phrasing of the relevant SCF question introduces some uncertainty as to what is being reported, making it unsuitable.

Table 3: Selected model moments from the baseline experiment

	Data	No information	Full Information
SD(credit limit / 2-year income)	$0.11 \rightarrow 0.17$	0.09	0.22
SD(annual interest rate)	$0.04 \rightarrow 0.06$	0.04	0.05
Avg(credit limit / 2-year income)	$0.09 \rightarrow 0.14$	0.05	0.07
Avg(credit limit)/Avg(2-year income)	$0.07 \rightarrow 0.13$	0.04	0.06
Avg(annual interest rate)	$15\% \rightarrow 14\%$	15%	10%
Card holding rate	$56\% \rightarrow 65\%$	56%	61%

borrow based on the realization of their expediture shock, however, and thus credit terms matter to them.

Having calibrated the model to our baseline date of 1989 (Table 2), we then simulate what happens when we change the information set of the financial sector. As discussed previously, this mechanically happens by letting firms observe the household expenditure shock, whereas in the baseline model they could not. Simulating this exercise yields the following results in Table 3.

Our model qualitatively captures the movements in several key moments in the sample. Unsurprisingly, since we only calibrate one parameter to match any of these moments (the contract reset-'stickiness' parameter), we cannot expect to precisely capture the levels in many variables of interest.

In particular, our results do well with respect to the three main points identified in the data section. Most importantly, the model predicts increases in the cross-sectional variances of credit limits and interest rates that closely corresponds to the magnitude of the change observed in the data. The predicted changes are 0.13 and 0.01, respectively, compared to 0.06 and 0.02 observed (the first two lines in the table). Our model also predicts movements in the level of credit limits that roughly replicates what we observed in the SCF. This suggests to us that our informational explanation is plausible.

We also plot life-cycle profiles for the average credit limit and its standard deviation, along with the standard deviation of credit interest rates in Figures 7, 8, 9. One reason that we adopted a life-cycle framework is a concern over potential demographic effects. Given the length of our sample and the well-known aging of the American population, a potential explanation for some of the observed trends might revolve around demographic changes. If the variance in the population age increased, credit trends could follow, and the life-cycle framework allows us to evaluate this theory.

Figure 7: Life cycle: model and data, average credit limit to income ratio

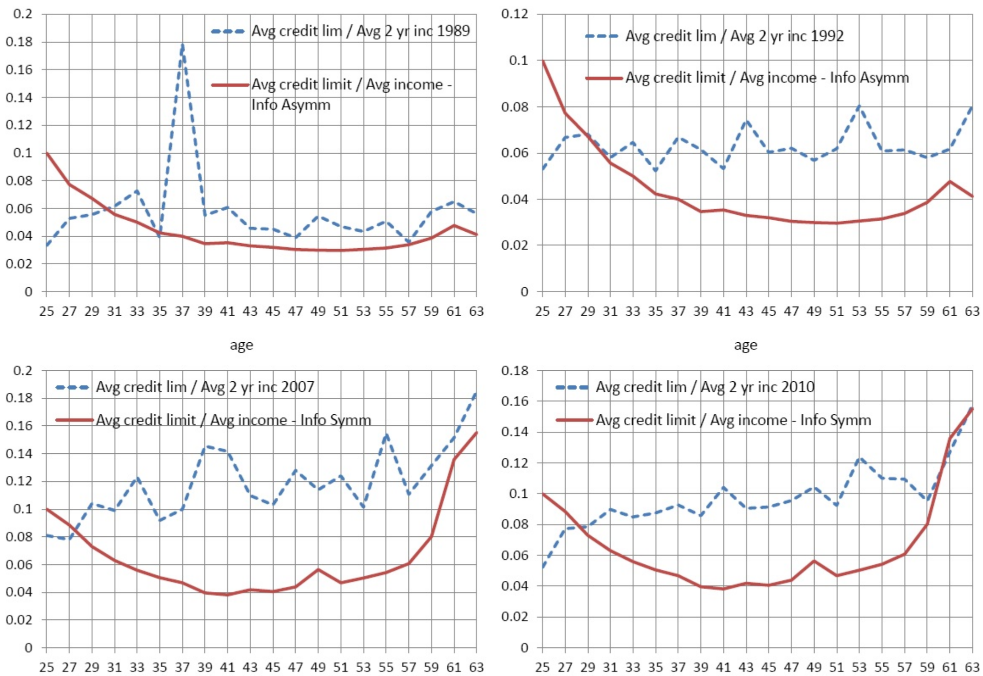

The main noticeable change that occurs in terms of the life-cycle profiles following the model experiment is increased variance in borrowing capacity near the end of the life-cycle. This pattern also shows up in the data - the average credit limit jumps at the age of 60 in the 2007 and 2010 SCF, referring back to Figure 7. However, this is not a demographic effect. Looking at the 1989 life-cycle pattern, no such increase is evident among older workers, and consequently we argue our observed trends are not a demographic effect. If the 1989 life-cycle variance profile had been maintained, any demographic shift would not have driven a trend

Figure 8: Life cycle: model and data, standard deviation, credit limit to income ratio

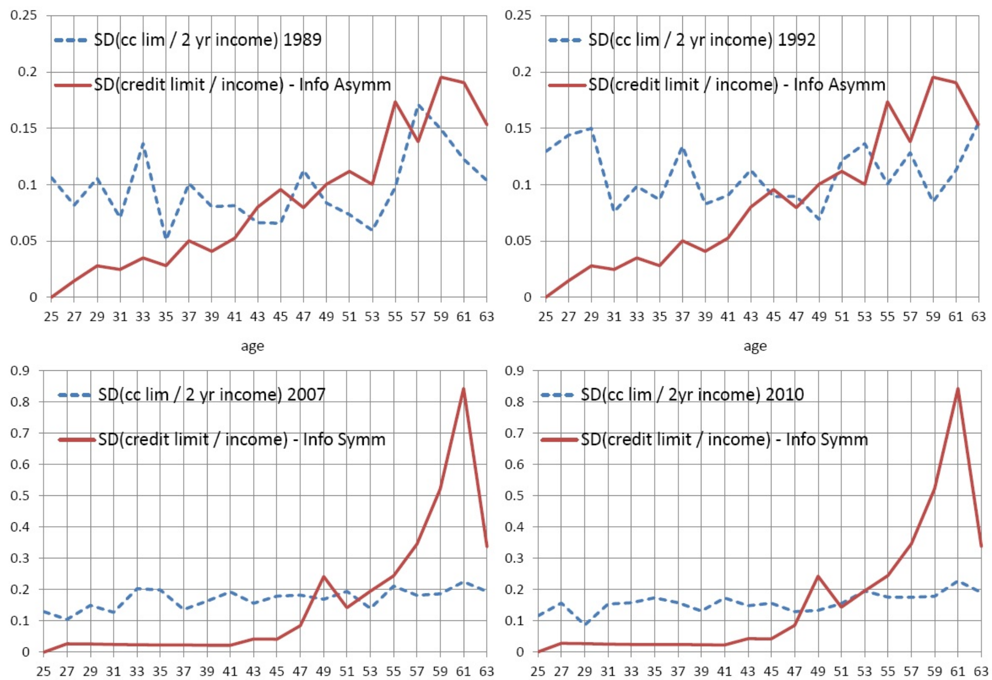

in credit limits. Indeed, the fact that the data profile and the model profile are shifting in a similar fashion lends extra support to our informational story. The model does not perform as well in terms of credit limit variance over the life cycle, but the data suggests (Figure 8) there is zero evidence that demographics are driving any of our motivating trends.

One qualitative issue with our model predictions is with respect to the credit card holding rate. Within our sample, there's a notable rise in the fraction of the population who possess at least one credit card, and similarly we observe an increase in the number of people who carry a positive balance on their credit card. Similarly, within the model, an increase in information available to the financial sector predicts more borrowers. The intuition for this result is with more individualized pricing of debt, some households who didn't want to borrow at high interest rates to smooth their consumption (as opposed to those households financing significant expenditure shocks via borrowing) are now revealed to the financial

Figure 9: Life cycle: model and data, standard deviation, credit interest rate

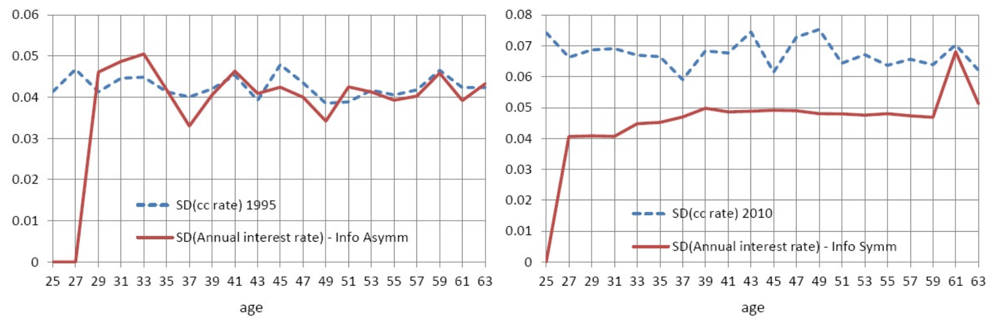

sector as relatively safe, avail themselves of lower interest rates, and enter the unsecured debt market.

However, there's a little more to this story. In Figure 10, we plot the observed credit card holding rates (according to the SCF). Notice that the rate peaks roughly at the turn of the century and thereafter exhibits a robust, sustained decline. In 1989, at the beginning of our sample, we consider it likely that credit cards had not fully saturated the market and were still in the process of complete adoption by society. The increase in holding rates at the beginning of our sample is thus perhaps not indicative of the real changes in the technology sector, and the increase in holding rates in the model is then at odds with the data if the increase over the whole sample period is simply driven by technology adoption with the decline in later years representing actual changes in the financial sector.[17]

A consequence of this hypothesis shows up in another unsecured lending industry: the payday loans industry. Within our model, we cap the maximum interest rate at 100% over a two year period, which we take as a rough approximation of credit card lending norms. Given that in our quantitative experiment the increase in credit card holding rates is driven primarily by low-interest-rate households starting to borrow, an increase that overwhelms the households that are driven out due to revelation of their expenditure type, it is natural that

[17]Lyons (2005) captures the first half of this story empirically, but her sample ends too early to observe the following decline.

Figure 10: Credit card holding rates

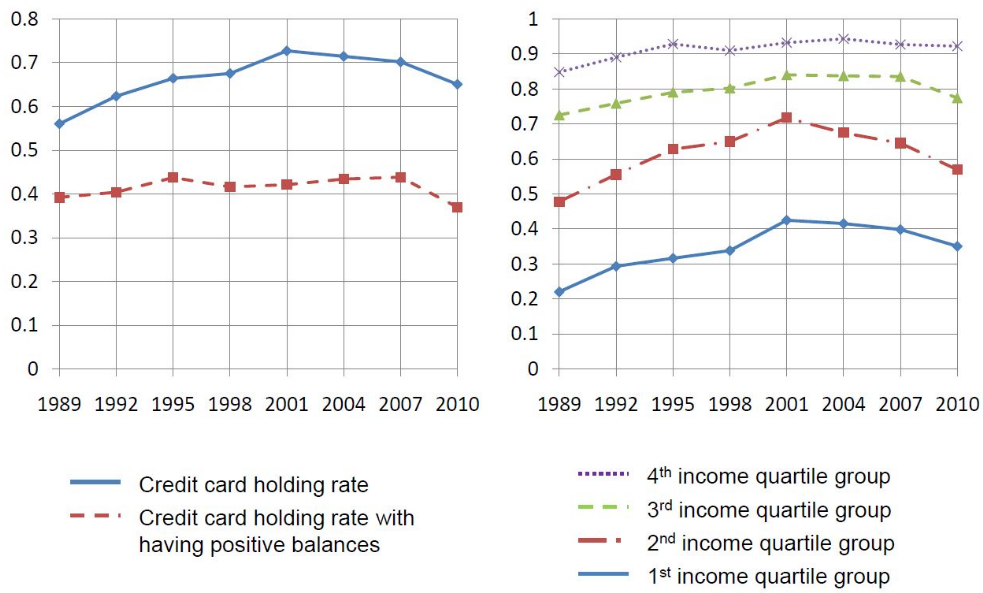

an alternative financial sector that charges higher rates might emerge for these households that can no longer borrow in the traditional financial sector. The growth in the payday lending sector has been muted in the United States, due to regulation, but has been grown rapidly in other countries, particularly the UK. We believe that further investigation into high-interest lending may be potentially informative, but it is not an avenue we pursue in this paper.

Another minor issue with our results is that our model cannot take into account all details of credit card pricing processes. For example, while some credit card contracts are really only a function of their interest rate and credit limit, many others offer extensive and complicated reward schemes, annual fees or other stipulations, and so on. Since the model misses such credit pricing features, it seems natural that our predicted moments cannot perfectly match the data, though the model performs well at least in qualitative manner.

One alternate way to model the change in the financial sector might be to, instead of modifying the informational structure and calibrating χ to a certain moment as occurs in

Table 4: Selected model moments from an alternative experiment

	No info ($\chi = 0.055$)	Full Info ($\chi = 0.055$)	No info ($\chi = 0.05$)	Full Info ($\chi = 0.05$)
SD(credit limit / 2-yr income)	0.09	0.22	0.22	0.58
SD(annual interest rate)	0.04	0.05	0.04	0.05
Avg(credit limit / 2-yr income)	0.05	0.07	0.07	0.14
Avg(credit limit)/Avg(2-yr income)	0.04	0.06	0.06	0.12
Avg(annual interest rate)	15%	10%	15%	10%
Card holding rate	56%	61%	51%	61%

our base model, use a lower value of χ and maintain the information asymmetry between banks and households. We feel like our informational story better fits the narrative[18] than a change in the cost of arranging a contract. As shown in Table 4, our brief examinations of this alternative approach suggest that the majority of the qualitative changes that result from lowering χ also point in the right direction. As the cost of arranging a contract decreases, the variance of credit limits and interest rates goes up. Furthermore, the credit limit to income ratio also increases. As the contract friction decreases, financial intermediaries could pricing credit terms with low costs, which leads to diversified credit terms.

Given the similarity of the results under both experiments, we continue to prefer, in line with the literature, our baseline experiment. However, the alternative modelling choice of how IT progress influences contracts does not pose any problem to the story we wish to tell.

5 Welfare and Transition Analysis

A natural question arising out of the prior exercises is, given that our informational story seems a compelling explanation for the observed moments, what are the welfare consequences for society following this development in information technology? In order to seriously treat the calculation of welfare, we first deal with characterizing the transition between the two

[18]The literature seems to agree on this, with several papers taking an informational approach, while we are not aware of any that drive their results through a change in what is effectively a menu cost.

steady-states.

The transition occurs as follows. Initially, we begin in the limited information economy, where firms are unable to observe the household's expenditure shock, θ. As before, this is calibrated to represent the 1989 economy, which is the earliest set of data we have. We model the informational change as happening in one period, allowing lenders to now observe household types. Figure 11 shows the transition path resulting from this change in the financial sector's information set in terms of various key statistics, where the red squares correspond to the model generated data and the blue diamonds the observed data.[19]

When lenders can observe household types, the credit card holding rate increases slightly between steady states. However, initially, the financial sector's increased information set allows them to eject bad credit risks from the pool of unsecured borrowers. Over time, however, the logic of the previously discussed steady-state comparisons dominates as the economy approaches equilibrium, where the lower interest rates and higher credit limits following from the decreased risks involved in lending (since less uncertainty about default with more information) lead more households to find unsecured borrowing worthwhile.

Following the same logic, and as also seen in the data, total borrowing increases with the decline in the price of borrowing and the relaxing of constraints on the quantity of borrowing. Our measures of dispersion - the standard deviations in terms of both credit limits and interest rates - both increase, as in the data and model discussed in Section 4. With lenders able to discern more information about households, more accurate pricing occurs. Before, when the financial sector could not observe expenditure shocks, lenders had no way to differentiate households that were identical except in terms of θ. As they now can, the optimal contracts are now different, increasing aggregate dispersion. Effectively, there is less pooling and more separation in terms of loan contracts.

[19]The data on aggregate outstanding loans is computed from total outstanding revolving loans. Loans in 1989 are normalized to 1.

Figure 11: Transition from asymmetric to symmetric information economy

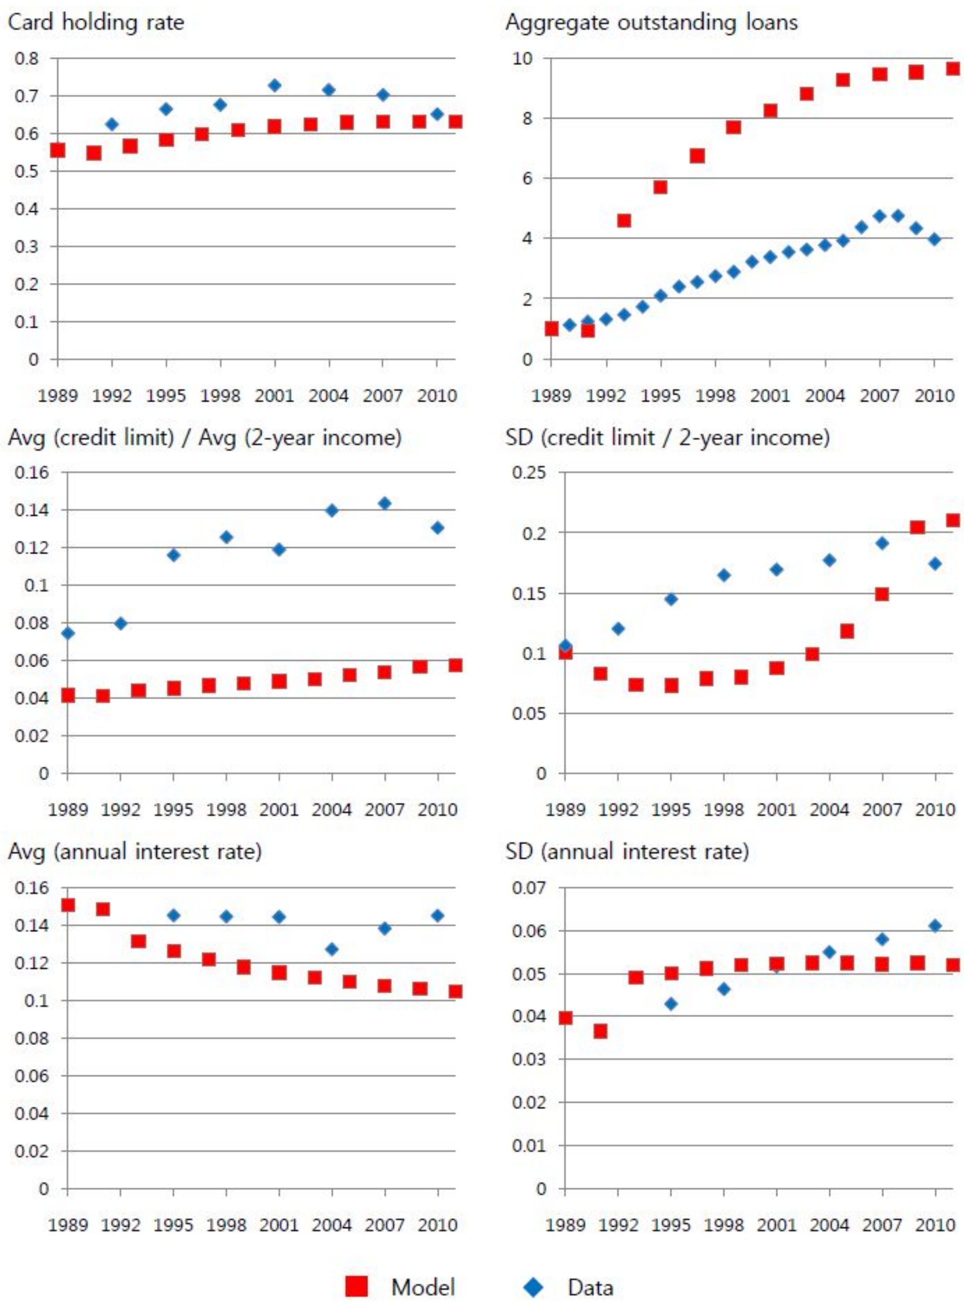

In sum, our calculated model transition captures the observed trends in the data in terms of our six most important moments. We find this both reassuring in terms of supporting our proposed mechanism, and in terms of our ability to accurately estimate the resulting welfare effects from the change.

5.1 Welfare Analysis

We now put our model to use by examining how these informational changes affect household welfare. Our calculation of welfare is inclusive of the transition, and not simply a comparison of the two steady-states. As discussed, households begin from the low-information economy steady-state and the introduction of the information is modelled as occurring between period 0 and period 1.

To calculate welfare changes, we first need to convert each household's welfare into consumption units. Let $V_j^i(\Delta)$ be an age-j household's lifetime utility value for state (i, Δ), where $i \in \{D, R\}$, and $\Delta = (L, q, a, z, \theta)$. Then, we can convert a given lifetime utility value into a corresponding permanent periodic consumption $(\tilde{c})$ as follows:

$$V_j^i(\Delta) = \sum_{t=0}^{T-j} \beta^t \frac{(c_{t+j})^{1-\sigma}}{1-\sigma} = \sum_{t=0}^{T-j} \beta^t \frac{\left(\tilde{c}_j^i(\Delta)\right)^{1-\sigma}}{1-\sigma}$$

Utility-equivalent consumption can then be defined as:

$$\tilde{c}_j^i(\Delta) = \left[\frac{V_j^i(\Delta)(1-\sigma)(1-\beta)}{1-\beta^{T-j+1}}\right]^{\frac{1}{1-\sigma}}$$

And aggregate household welfare from a change in the informational regime as:

$$W = \frac{\int \tilde{c}_j^i(\Delta)_{SI} - \tilde{c}_j^i(\Delta)_{AI}\, d\Psi_{AI}}{\int \tilde{c}_j^i(\Delta)_{AI}\, d\Psi_{AI}}$$

where $\tilde{c}_j^i(\Delta)_{SI}$ is the household's (constant) periodic consumption after switching to the symmetric information economy and $\tilde{c}_j^i(\Delta)_{AI}$ is consumption in the asymmetric information economy for a household in state (i, j, Δ). Here, Ψ_{AI} is the invariant distribution of households in the limited information economy and W is the change in aggregate household welfare when moving from the limited to the full information economy, inclusive of the transition. Hence, if W is larger than 0, net household welfare improves from the increased access to information in the financial sector.

Under our calibration, W is 0.06%. That is, when households move from the asymmetric to the symmetric information world, aggregate household welfare increases 0.06% in terms of lifetime consumption. This is not uniform across types, though. Households ending up with high θ shocks prefer the (baseline) asymmetric information economy, where they have relatively easy access to credit markets as the financial sector cannot identify their true financial position, leading to their cross-subsidization by low θ households. Given the rarity of the high-expense shock, though, 93% of households (strictly) prefer the symmetric information world to the asymmetric information economy. In the symmetric information economy, these households enjoy (on average) higher credit limits and lower interest rates, reflecting the greater efficiency of the credit market. As this is the experience of the vast majority of households, it is unsurprising that welfare increases, along with our natural priors that the use of information to achieve superior pricing should improve things on net.[20]

To stress that the gains are unequally distributed, consider that while the net welfare gain is only 0.06%, newborn households find their welfare improved by 0.51% following the informational change.[21] Following the model change, young households enjoy higher credit limits along with a lower borrowing interest rate over their entire life cycle. Further, the opportunities to smooth consumption are greatest early in life. Calculating the welfare gains

[20]Other papers that work with welfare calculations in related settings, for example Athreya, Tam and Young (2012), also find an overall increase in welfare from an informational improvement but with some possibility that already-born households become worse off.

[21]Welfare changes are calculated for each age-j household as follows:

$$W_j = \frac{\int \tilde{c}^i_j \left(\Delta\right)_{SI} - \tilde{c}^i_j \left(\Delta\right)_{AI} d\Psi_{AI}(j)}{\int \tilde{c}^i_j \left(\Delta\right)_{AI} d\Psi_{AI}(j)}$$

where $\Psi_{AI}(j)$ is the measure of such households before the transition.

by age group shows that the gains are strictly decreasing with age. While it may be intuitive to think that the increasing information held by lenders would hurt younger households the most, the ability to access accurate pricing is most valuable to young households who do not encounter severe adverse shocks.

6 Conclusion

Using data from the Survey of Consumer Finances, we update the literature and discuss how the cross-sectional heterogeneity in unsecured lending contracts, measured using interest rates and credit limits of household credit cards, has been steadily increasing over the SCF sample period from 1989 to present. This increasing heterogeneity remains even when conditioning on traditional observables in the financial sector, such as income. We also present a variety of evidence that suggests this pattern reflects the increasing accumulation and deployment of consumer information by the financial sector, e.g. the increasing correlation of educational attainment and homeownership with credit terms.

We also go beyond the current quantitative literature by computing a realistic model of unsecured borrowing within a life-cycle framework where consumers borrow according to contracts modelled on modern credit cards against expenditure shocks in order to smooth consumption. These contracts are characterized by a credit limit and interest rate, and are endogenously 'sticky' in the sense that lenders do not update the terms of the credit contract every period. The lending terms are jointly pinned down by a zero-profit condition and consumer preference over zero-profit contracts. Within our model, we show that increasing the information set available to lenders generates model predictions that qualitatively and quantitatively match the key trends in interest rate and credit limit variances, among other moments. Our model experiment also captures movements in the life-cycle profiles of these statistics and potentially applies to the growth of the payday loans industry. The improved deployment of information by the financial sector generates a small gain in total welfare that accrues primarily to young households.

References

Kartik Athreya and Urvi Neelakantan. The cost of unanticipated household financial shocks: Two examples. *FRB Richmond Economic Quarterly*, 97(4):431–450, 2011.

Kartik Athreya, Xuan S. Tam, and Eric R. Young. A quantitative theory of information and unsecured credit. *American Economic Journal: Macroeconomics*, 4(3), 2012.

Lawrence M Ausubel. Adverse selection in the credit card market. Technical report, 1999.

Sayajit Chatterjee, Dean Corbae, and José-Víctor Ríos-Rull. A theory of credit scoring and competitive pricing of default risk. working paper, July 2011.

Satyajit Chatterjee and Burcu Eyigungor. Maturity, indebtedness, and default risk. *American Economic Review*, 102(6):2674–99, 2012.

Paul S Calem and Loretta J Mester. Consumer behavior and the stickiness of credit-card interest rates. *The American Economic Review*, 85(5):1327–1336, 1995.

Shubhasis Dey and Gene Elwood Mumy. Determinants of borrowing limits on credit cards. working paper, November 2009.

Lukasz A. Drozd and Jaromir B. Nosal. Competing for customers: A search model of the market for unsecured credit. working paper, February 2011.

Lukasz A. Drozd and Ricardo Serrano-Padial. Modeling the credit card revolution: The role of debt collection and informal bankruptcy. working paper, April 2013.

David Gross and Nicholas Souleles. Do liquidity constraints and interest rates matter for consumer behaviour? evidence from credit card data. *Quarterly Journal of Economics*, 117(1), 2002.

Bulent Guler. Innovations in information technology and the mortgage market. *University of Texas at Austin mimeo*, 2010.

Dean Karlan and Jonathan Zinman. Observing unobservables: Identifying information asymmetries with a consumer credit field experiment. *Econometrica*, 77(6), 2009.

Igor Livshits, James MacGee, and Michele Tertilt. Consumer bankruptcy: A fresh start. *The American Economic Review*, 97(1):402–418, 2007.

Igor Livshits, James MacGee, and Michèle Tertilt. Costly contracts and consumer credit. NBER working paper 17448, September 2011.

Angela C. Lyons. How credit access has changed over time for U.S. households. *Journal of Consumer Affairs*, 37(2), 2005.

Xavier Mateos-Planas. A model of credit limits and bankruptcy. working paper, June 2011.

Borghan N. Narajabad. Information technology and the rise of household bankruptcy. working paper, April 2012.

W Craig Riddell and Xueda Song. The impact of education on unemployment incidence and re-employment success: Evidence from the us labour market. *Labour Economics*, 18(4):453–463, 2011.

Juan M. Sánchez. The it revolution and the unsecured credit market. working paper, March 2012.

Kjetil Storesletten, Christopher I Telmer, and Amir Yaron. Consumption and risk sharing over the life cycle. *Journal of Monetary Economics*, 51(3):609–633, 2004.