Modeling the Credit Card Revolution: The Role of Debt Collection and Informal Bankruptcy

Lukasz A. Drozd and Ricardo Serrano-Padial*

April 24, 2013

ABSTRACT

In the data, most consumer defaults on unsecured credit are informal and the lending industry devotes significant resources to debt collection. We develop a new theory of credit card lending that takes these two features into account. The two key elements of our model are moral hazard and costly state verification that relies on the use of information technology. We show that the model gives rise to a novel channel through which IT progress can affect outcomes in the credit markets, and argue that this aspect is potentially critical to understand the trends associated with the rapid expansion of credit card borrowing in the 1980s and over the 1990s. Independently, the mechanism of the model helps reconcile the high levels of defaults and indebtedness observed in the US data.

JEL: D1,D8,G2
Keywords: credit cards, consumer credit, unsecured credit, revolving credit, informal bankruptcy, debt collection, moral hazard

---

*Drozd: Finance Department, The Wharton School of the University of Pennsylvania, email: ldrozd@wharton.upenn.edu; Drozd is also a visiting scholar at the Payment Cards Center of the Philadelphia Fed: The views expressed here are those of the authors and do not necessarily represent the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System; Serrano-Padial: Department of Economics, University of Wisconsin - Madison, email: rserrano@ssc.wisc.edu. Drozd acknowledges generous support provided by the Vanguard Research Fellowship and the Cynthia and Bennet Golub Endowed Faculty Scholar Award. Serrano-Padial acknowledges generous support provided by the Graduate School Research Committee Award. We thank Satyajit Chatterjee, Robert Hunt, Scott Richard, and Amir Yaron for valuable comments. All remaining mistakes are ours.

1 Introduction

Existing theories of consumer bankruptcy on unsecured debt abstract from the costs of enforcing repayment from defaulting borrowers. The underlying assumption is that consumers always file for personal bankruptcy protection in court, and thus eligibility for debt discharge is determined and enforced by the court system rather than by the lending industry.

Recent evidence suggests that this view is grossly at odds with the data. This is because at least half of credit card debt defaulted on in the US takes the form of an 'informal' bankruptcy, implying that many consumers default without ever formally filing in court.[1] Furthermore, the lending industry actually devotes significant resources to debt collection.[2]

Our interpretation of this evidence is that it suggests the presence of costs assumed by the lending industry to implement credit contracts that admit default. As we show, such costs can have important ramifications on the ex-ante pricing of default risk, which we explore here in the context of the rapidly expanding credit card lending in the 1980s and over the 1990s. The literature has generally attributed this expansion to the widespread adoption of information technology by the lending industry (e.g. information sharing through the credit bureaus, automated risk scoring models, etc.).[3] In this context, our paper identifies a novel channel that allows to endogenously link the IT-implied improvements in the quality of information to the growing default exposure assumed by the US credit card industry over this time period.

To develop this theme, we build on the costly state verification literature (e.g. Townsend (1979)), and propose a model in which the lending industry must monitor defaulting borrowers

[1] Using a panel of 50,831 pre-approved credit card recipients tracked in time, Dawsey and Ausubel (2004) show that as much as half of discharged credit card debt is not attributable to formal bankruptcy filings thereafter. Our work is very much inspired by their findings. Various recent industry surveys are also consistent with this view. For example, 1999 Annual Bankruptcy Survey by Visa U.S.A. Inc. finds that two-thirds of credit card loans that are charged-off as uncollectible are not attributable to bankruptcy filings.

[2] According to the BLS, employment in the third-party debt collection industry is about 150,000 people. PricewaterhouseCoopers estimates that the industry as a whole directly or indirectly supports employment between 300,000 and 420,000 jobs (see "Value of Third-Part Debt Collection to the US Economy in 2004: Survey and Analysis," by PricewatehouseCoopers, prepared for ACA International). According to an ACA-commissioned survey (Ernst &Young "The Impact of Third-Party Debt Collection on the National and State Economies," ACA International 2012), credit card debt account for 20% of all debt collected by the third-party debt collection industry. For comparison, the total size of the US police force stands at about 700,000 officers across all agencies. See the appendix for an overview of the consumer debt collection process. See Hunt (2007) for an overview of the debt collection industry in the US.

[3] For an overview of the key changes in the US credit card market, see Evans and Schmalensee (2005) and White (2007).

in order to extend credit contracts that admit default in equilibrium. Monitoring is necessary because default in our model involves *asymmetric information* that leads to *moral hazard*. Asymmetric information is brought about by the fact that default is largely informal and thus does not automatically reveal the true solvency status of the defaulting borrowers to lenders.[4] Moral hazard arises because all debtors are generally prone to defaulting whenever informal debt forgiveness seems overly likely to them. Since the monitoring cost required to sustain repayment by the ex-post 'solvent' borrowers crucially depends on the probability of default by the 'insolvent' borrowers, this cost affects the ex-ante pricing of default risk. Importantly, its impact is *endogenously* linked to the state of information technology that lenders can use to maximize the effectiveness of their monitoring effort on borrowers' incentives. This mechanism allows us to relate the predictions of our model to the trends in the data.

The central implication of our model is that information technology, by enhancing the effectiveness of monitoring, influences the risk composition of aggregate debt through prices. As a result, an IT-driven expansion of credit card borrowing in the model naturally leads to an increase in the average riskiness of the outstanding debt. This particular aspect of the 'revolving revolution' is widely debated in the literature (see, for example, White, 2007). While it is obvious how cheaper and more precise information might have increased the availability of credit by lowering its price, it is far from clear why this expansion also resulted in an almost twofold increase in the average default exposure of credit card debt. Most quantitative models fall short of accounting for this phenomenon, as we discuss below.[5]

Figure 1 illustrates our main quantitative findings. As we can see, our model is consistent with three basic observations in the data:[6] 1) expansion of revolving credit in the form of pre-

[4]Our model requires that informal default must be *marginally* more attractive than formal bankruptcy for the majority of both insolvent and solvent borrowers. The aforementioned evidence suggests this must be the case, as most borrowers in the data choose this option. If formal bankruptcy is marginally more costly than informal debt forgiveness, such feature is naturally implied by the fact that the option to file formally is retained by the borrowers.

[5]For a thorough analysis of the performance of the standard models in light of the trends in the US unsecured credit market, see also Livshits, MacGee and Tertilt (2010).

[6]To estimate the trend line we used the time period 1985-2005. The earlier time period is intentionally omitted due to the binding usury laws. Usuary laws have been gradually eliminated in the US by the 1978 Supreme Court ruling, Marquette National Bank of Minneapolis vs. First of Omaha Service Corp. This ruling let credit card issuers apply nationally usury laws from the state in which they were headquartered. To attract credit card companies, some states dropped their usury laws. The increase in credit card lending in the early

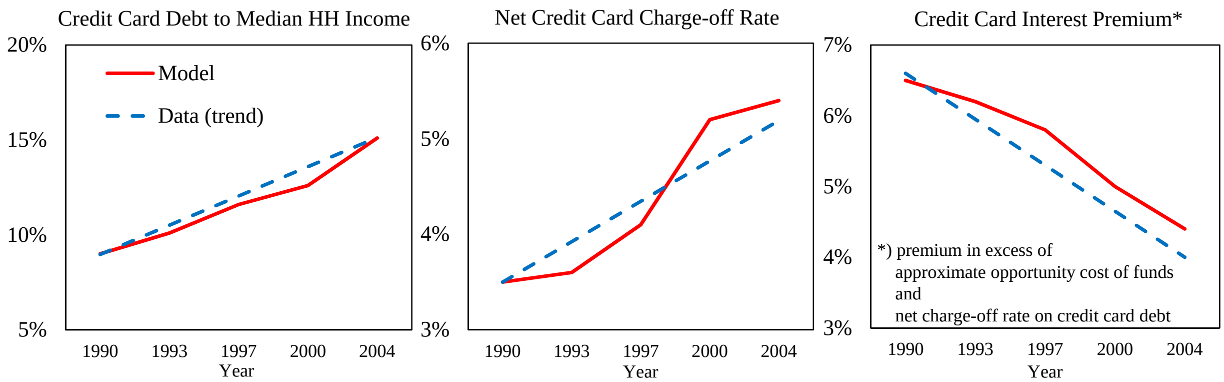

Figure 1: The Rise of Revolving Credit in the US

committed credit lines (credit cards), 2) a decline in interest premia on credit card debt, and 3) an increase in the average default exposure of credit card debt (as measured by the fraction of debt defaulted on). Using our model, we fully account for all these observations by raising the precision of information that lenders have about their delinquent borrowers by a factor of 3, and by assuming a simultaneously declining transaction cost of making credit card funds available to the consumers by about 20%. The change in the transaction cost is independently calibrated using the evidence on the excess productivity growth within the banking industry relative to the other sectors of the economy (as reported by Berger (2003)). In addition, as discussed at the end of our paper, our approach independently enhances the performance of the existing models in their ability to account for the high levels of defaults and indebtedness seen in the data.[7]

In a nutshell, our model paints the following picture of the observed changes in the US unsecured credit market. In the late 80s, the IT technology employed by the industry was underdeveloped, which made it difficult for the lending industry to optimize on the cost of enforcing repayment from the moral hazard-prone debtors. This is because risky contracts could only be sustained during this time period using costly "carpet" monitoring of all de-

80s was likely affected by this decision. In 2005 a major bankruptcy reform affected the credit card market, which we abstract from by focusing exclusively on the time period before 2005.

[7]Existing models imply a tension between the sustainability of defaultable debt and the attractiveness of the default option. In our model, repayment of defaultable debt is sustained through monitoring, which is focused on the solvent borrowers by definition. Hence, the attractiveness of the default option among the insolvent or distressed borrowers can be independently maintained. This allows our model to easily reconcile a high default rate in the data with a relatively high indebtedness.

faulting borrowers. As a result, credit limits were tight to prevent costly defaults. However, over the 90s, while learning how to assess credit risk in general, the industry also learned how to economize on unnecessary and wasteful monitoring of truly distressed and insolvent borrowers – whose default is actually constraint efficient from the ex-ante point of view. This effect lead the industry to assume greater default exposure, as the data shows. By studying the impact of information on the (ex-post) moral hazard problem faced by lenders, our paper thus uncovers a novel channel by which IT progress can shape outcomes in the credit markets. In contrast, most of the literature has focused on the link between information precision and adverse selection, i.e., lenders' ability to assess credit risk ex ante. As we discuss below, the predictions on default exposure from these models is unclear.

Finally, consistent with the key assumption of our model, the technological improvements in the collection industry largely paralleled those in the credit industry. For instance, TransUnion – the same company that collects credit history data and helps lenders to assess credit risk ex-ante – advertises its parallel product marketed to debt collectors this way: "*TransUnion combines data (...) to provide valuable insight into each account so you can better define consumers' willingness and ability to repay (...).*" More concretely, Portfolio Recovery Associates Inc – one of the major companies involved collecting charged-off credit card debt in the US – reports in its 10-K form filed with the SEC that, from 1998 to 2004, its effectiveness, as measured by cash spent on collection efforts relative to cash collected, has increased at least by a factor of two.

Our model is also consistent with other suggestive pieces of evidence. For example, Hynes (2006) studies court data and finds that the growth of garnishment court orders was far slower than the growth of defaults. To the extent that garnishment orders can be interpreted as an imperfect measure of monitoring, this phenomenon is predicted by our theory. Finally, using recent TransUnion micro-level data, Fedaseyeu (2013) studies the relation between tightness of debt collection laws at the state level and the supply of revolving credit. Consistent with our model, he finds that stricter regulations of debt collection agencies across states are associated with tighter supply of revolving credit, while having no effect on secured borrowing.

Related Literature To the best of our knowledge, no other study systematically relates information technology to the effectiveness of debt collection. However, viewed more broadly, our paper is related to a number of recent contributions in the literature. These include the adverse selection models of IT-driven credit expansion by Narajabad (2012), Athreya, Tam and Young (2008), and Sanchez (2012); papers on informal bankruptcy by Chatterjee (2010), Athreya et al. (2012), Benjamin and Mateos-Planas (2012), and White (1998); and other work on the effects of information technology on credit pricing, such as Drozd and Nosal (2007) and Livshits, MacGee and Tertilt (2011).

The first group of papers views events of the 90s as a transition from a pooling to a separating equilibrium. In the latter equilibrium, due to improvements in credit scoring techniques, low-risk types manage to separate themselves from high-risk types, and consequently borrow more and default more. However, at the same time, high risk types borrow less and default less. Compared to our approach, existing models of adverse selection do not account for the increased default exposure of debt in the data.[8] In addition, adverse selection models also involve an arguably counterfactual reallocation of credit toward the less risky segments of the credit market – to the extent that risk characteristics are correlated with observables such as income, education, wealth etc. See the discussion in White (2007).

The second group of papers focuses on the choice of formal versus informal bankruptcy, which we do not pursue here. These papers view informal bankruptcy as a renegotiation tool that can add more flexibility to the system. By justifying the attractiveness of the informal default option to borrowers under distress, the findings of this literature generally complement our work by highlighting the importance of informal bankruptcy.

The final set of papers are those that develop models exhibiting fixed costs of extending credit or designing contracts, such as Drozd and Nosal (2007) and Livshits, MacGee and Tertilt

[8]In terms of specific papers, Athreya, Tam and Young (2008) report (in Table 3) that the predicted change of discharged debt to income remains unchanged, while it doubles in the data. Given debt in the data increased faster than income, we infer from this statistic that their model fails to generate an increase in average riskiness of debt. A similar exercise is also performed by Sanchez (2012). However, in his comparative statics exercise the riskiness of the income process changes as well (raising the demand for risky credit). The reported experiments suggest that the model still falls short of accounting for both the level and the change in the charge-off rate (the former by a factor of 10). The predictions of the model by Sanchez (2012) are additionally difficult to relate to the charge-off rate in the data. This is because the model is calibrated to *net* debt rather then *gross* debt. In the data, the calculation of the charge-off rate in the denominator involves gross debt (total outstanding debt), which is larger than the net position by at least a factor of ten.

(2011). Both papers emphasize the extensive margin of credit card lending. In Drozd and Nosal (2007) the cost can be interpreted as an account acquisition cost, or ex-ante credit risk assessment costs to reveal borrower's initial state. Their model fails to account for the rising charge-off rate in the data. The fixed cost in Livshits, MacGee and Tertilt (2011) emphasizes contract design costs in the presence of adverse selection. Their cost thus implies increasing returns from lending (on the extensive margin). In contrast to our paper, a higher precision of information in Livshits, MacGee and Tertilt (2011) *lowers* the default exposure of debt in equilibrium. IT progress manifested through a lower cost of designing risky contracts or a fall in transaction cost can generate effects consistent with the data, although their quantitative implications remain to be explored.[9]

2 Environment

We begin with a two-period model, which we study analytically. In later sections, we extend this setup to a multi-period life-cycle environment and explore its quantitative implications.

2.1 Economic Agents

The model economy is populated by a continuum of consumers and a finite number of credit card lenders. Consumers live for one period that is composed of two sub-periods. Their objective is to smooth consumption across the two sub-periods by borrowing from lenders. They enter the period with some pre-existing exogenous stock of unsecured debt $B > 0$, and their income is fixed at $Y > 0$. In the second sub-period they are subject to a random realization of a binary distress shock $d \in \{0, 1\}$ of size $E > 0$ (e.g. job loss, divorce, unwanted pregnancy,

[9] Specifically, the setup involves an assumption that designing a risk-free contract is costless. As a result, an exogenous decline in the design cost leads to a change in the relative price of the two contracts, which raises the share of risky contracts in equilibrium. In our model a similar effect arises endogenously as a result of improvements in the precision of information. In addition, Livshits, MacGee and Tertilt (2011) make an interesting point that in the presence of a fixed cost of designing contracts, a decline in the proportional transaction cost or even the risk free interest rate can raise the share of risky contracts. This effect is associated with increasing returns in the lending industry that their setup implies. Specifically, in their model a larger number of revolvers reduces the design cost *per borrower*. Given the assumption that risk-free contracts are costless, this similarly affects the relative price of contracts. Since their model is not quantitative, the empirical relevance of this channel remains to be explored.

or medical bills). This shock hits with probability $0 < p < 1/2$ and is the only source of uncertainty in the two-period model (in our full model Y is stochastic and B is endogenous). Lenders have deep pockets and extend credit lines to consumers at the beginning of the period. Their cost of funds is exogenous and normalized to zero. Credit lines are characterized by an interest rate $R < 1$ and credit limit $L \geq 0$, and they are committed to borrowers as of the beginning of the period. Consumers can accept at most one credit line. It is assumed that there are perfect rental markets of consumer durables to justify our focus on unsecured credit.

2.2 Lenders

Lenders are Bertrand competitors. Hence, they offer a contract that maximizes the ex-ante expected indirect utility of their potential customers subject to a zero profit condition (in expectation). Since credit lines are accepted and committed to borrowers before the realization of the distress shock d (and signal), lenders are potentially exposed to the risk of default. In such a case, credit lines offer insurance against the distress shock, which borrowers typically find desirable in our setup.

2.2.1 Asymmetric Information

Information between borrowers and lenders is *asymmetric* upon default, which is brought about by the fact that lenders do *not* directly observe the distress shock d. Instead, they observe a noisy signal $s \in \{0, 1\}$ with exogenous precision $0 \leq \pi \leq 1$. The parameter is assumed to fully summarize *the state of information technology* available to lenders. In the presence of a formal option of bankruptcy, the crucial assumption here is that distressed agents marginally prefer informal default to formal default (which we do not model for simplicity).[10]

2.2.2 Monitoring Technology

Asymmetric information creates *moral hazard* among the non-distressed borrowers. This is because such borrowers may choose to default *strategically* in expectation of informal debt forgiveness by lenders. To deal with the moral hazard problem, lenders are equipped with a *monitoring technology*. Monitoring cost is λ and monitoring results in repayment by a

[10]See also the comment in footnote 4.

monitored non-distressed consumer, and it is ineffective in the case of distressed consumers. This assumption is justified by the fact that distressed borrowers are assumed to be insolvent and thus 'monitoring proof'; alternatively, they subsequently file for (formal) bankruptcy protection. At the same time, non-distressed borrowers are assumed to possess assets that a bankruptcy court could seize, and their propensity to default *formally* is significantly lower than their propensity to default *informally*. They may also be ineligible for formal bankruptcy, as is the case under the current law.[11]

To keep our model simple, we assume that lenders can ex-ante commit to a monitoring probability distribution: $0 \leq P(s) \leq 1$, which depends on the signal $s = 0, 1$. Such an approach parsimoniously implements the efficient contract as far as monitoring is concerned, and conveniently abstracts from any institutional characteristics of the debt collection industry. Formally, in the beginning of the period lenders choose a credit line contract $K = (R, L)$ and a monitoring strategy $P(s)$ to maximize the ex-ante expected utility of consumers:

$$\max_{K,P} V(K, P) \tag{1}$$

subject to ex-ante zero profit condition

$$\mathbb{E}\Pi(I, K, P) \geq \lambda \sum_{I=(d,s)} \delta(I, K, P) P(s) Pr(I),$$

In the above problem, $\mathbb{E}\Pi(I, K, P)$ is ex-ante profit of the lenders from a customer pool with normalized measure one (gross of monitoring cost), $I = (d, s)$ is the interim state of the consumer, λ is monitoring cost (per borrower or measure one of borrowers), and $\delta(I, K, P)$ describes the consumer's default decision function (defined formally in the next section), which equals one in case of default and zero otherwise. The interim profit function is given by

$$\Pi(I, K, P) = \begin{cases} 2\rho(K, b(I, K, P)) & \text{if } \delta(I, K, P) = 0, \\ -L + L(1 + \bar{R})(1 - d)P(s) & \text{if } \delta(I, K, P) = 1, \end{cases} \tag{2}$$

[11] After 2005 bankruptcy protection involves income means testing. Borrowers who do not qualify must partially repay their debt.

where $b(I, K, P)$ is the consumer borrowing and $\rho(K, b) \equiv R \max\{b, 0\}/2$ is interest income of the lender (collected in each sub-period when the consumer does not default). The function conveys the basic idea that default is costly for the lender because it allows the consumer to discharge L (and any thus far assessed interest), while revenue is derived from the interest payments when the consumer chooses not to default. Monitoring reverts any non-distressed defaulting consumer back to repayment by recouping $L(1+\bar{R})$, where $\bar{R}$ is an exogenous penalty interest paid whenever collection is successful.[12]

2.3 Consumers

Consumers choose consumption in the first sub-period c, consumption in the second sub-period c', borrowing within the period b, and the (informal) default decision $\delta \in \{0, 1\}$. For simplicity, these choices are made at the interim stage, that is, after the signal s and the distress shocks d are observed. However, the consumer still faces the residual uncertainty associated with being monitored, which is denoted by $m \in \{0, 1\}$. Clearly, this residual uncertainty only matters in the case of default.

Formally, given K and P, the consumer chooses the default decision δ to solve

$$V(K, P) \equiv \mathbb{E} \max_{\delta \in \{0,1\}} [(1 - \delta)N(I, K, P) + \delta D(I, K, P)]. \tag{3}$$

where $\mathbb{E}$ is the ex-ante expectation operator, $N(\cdot)$ is the indirect utility function associated with repayment and $D(\cdot)$ is the indirect utility function associated with defaulting.

Under repayment, the consumer chooses b, c, c' to solve

$$N(I, K) \equiv \max_{b \leq L} U(c, c') \tag{4}$$

[12] Although we consider here a deterministic case and generally assume $\bar{R}$ is non-negative, there is nothing that precludes $\bar{R}$ to be negative (as long as it is not too negative). $\bar{R}$ can also be stochastic. Hence, our model could incorporate the idea of partial debt forgiveness.

subject to

$$
\begin{aligned}
c &= Y - B + b - \rho(K, b) \\
c' &= Y - b - dE - \rho(K, b).
\end{aligned}
$$

In the above problem, U is the utility function satisfying the usual set of assumptions. The budget constraint states that the consumer can borrow in the first sub-period up to the credit limit L. In such a case, in the second sub-period, the consumer pays back the principal, and pays interest on borrowing ρ in each sub-period.

To define the indirect utility function for the case of default, we assume that consumers cash out the credit line in the first sub-period.[13] In addition, they also default on a fraction ϕ of their distress shock, which is discharged regardless of whether they are monitored. Consistent with our earlier discussion, defaulting *distressed* consumers incur a pecuniary cost of defaulting equal to θY. They can always fully discharge their debt. *Non-distressed* defaulting consumers can discharge their debt only if they are not monitored. If they are monitored, they revert back to repayment, pay back the principal L, a penalty interest $\bar{R}L$, and incur a smaller pecuniary cost associated with the state of temporary delinquency $\underline{\theta}Y < \theta Y$.

Formally, the defaulting consumers choose b, c, c' to solve

$$
D(I, K, P) \equiv \max_{b \leq 0} \mathbb{E}_I U(c, c') \tag{5}
$$

subject to

$$
\begin{aligned}
c &= Y - B + L + b \\
c' &= (1 - \theta)Y - (1 - \phi)dE - b - m(1 - d)[((\underline{\theta} - \theta)Y + L(1 + \bar{R}))],
\end{aligned}
$$

where the expectation operator $\mathbb{E}_I$ integrates over the unknown realization of the ex-post monitoring decision of lenders $m \in \{0, 1\}$, and the term in the square bracket captures the impact of monitoring on the consumer budget constraint.

[13] This assumption is consistent with the data, see Herkenhoff (2012).

Finally, to make our model non-trivial, we maintain two assumptions to ensure that: 1) non-distressed consumers do not *always* want to default, and 2) in the case of completely uninformative signals, the excess penalty interest rate associated with debt collection over the interest normally paid under repayment is insufficient to cover the monitoring cost needed to identify a non-distressed defaulting borrower. These assumptions are formalized below.

Assumption 1. *$\underline{\theta}Y + \bar{R}L$ is sufficiently high to ensure that a non-distressed consumer does not default when monitored with certainty.*[14]

Assumption 2. *For any feasible contract $K = (R, L)$, we assume $\lambda > (1-p)(\bar{R}L - \rho(K, b^*))$, where b^* is optimal borrowing associated with repayment (given K).*

Optimal default decision of consumers and default exposure of credit lines. We next turn to the characterization of the policy function that governs the default decision in our model. This result is summarized in the proposition below. It implies that there are both risky and risk-free contracts in this environment. Furthermore, sufficiently high credit lines always bundle a potentially welfare enhancing transfer of resources from the state of no distress toward the state of distress.[15] In the rest of the paper, we refer to default by a distressed consumer as *non-strategic*, and to default by a non-distressed consumer as *strategic*. Unless otherwise noted, all proofs are relegated to the Appendix.

Proposition 1. *For any feasible $K = (R, L)$ and $P(s)$, the default decision of consumers is consistent with the following set of cutoff rules:*

1. *There exists a cutoff $L_{\min}(d = 1, R) > 0$, continuous and strictly decreasing w.r.t. R, such that a distressed borrower defaults (repays) if $L >$ $(<)$ $L_{\min}(d = 1, R)$, regardless of P.*
2. *There exists a cutoff $L_{\min}(d = 0, R) > 0$, continuous and strictly decreasing w.r.t. R, such that a non-distressed borrower repays if $L < L_{\min}(d = 0)$, regardless of P.*

[14] This can be globally assured by assuming $\underline{\theta}Y + \bar{R}L > \max_R Rb^*(R, L)$, where $b^*(R)$ is borrowing associated with repayment given R, L.

[15] When intertemporal smoothing needs are particularly high, it is possible that the transfer associated with the equilibrium contract is inefficient (too high) and the presence of a default option lowers welfare. This arises rarely in our parameterized model and overall the presence of default option is welfare enhancing.

3. *If $L > L_{\min}(d=0)$, there exists a cutoff $\bar{P}(R,L) \in (0,1)$, continuous and increasing in R, such that a non-distressed borrower repays (defaults) when $P(s) > (<)\, \bar{P}(R,L)$, and she is indifferent between defaulting or not for $P(s) = \bar{P}(R,L)$. Furthermore $\bar{P}(R,L)$ is independent of information precision π.*

In the case of zero-profit contracts, we denote $L_{\min}(d) \equiv L_{\min}(d,R)$. If no agent defaults when indifferent between defaulting or not, then $L_{\min}(\cdot)$ is independent of information precision π, with $L_{\min}(d=1) = \theta Y - \phi E$.

The above result allows us to distinguish between the following three classes of contracts in our model.

Definition 1. *We refer to a contract as a:*

i) *risk-free contract, if $L < L_{\min}(d=1)$,*
ii) *non-monitored insurance contract, if $L \in [L_{\min}(d=1), L_{\min}(d=0))$,*
iii) *monitored insurance contract, if $L > L_{\min}(d=0,R)$.*

Finally, by *equilibrium* in this economy we mean a collection of indirect utility functions $V(\cdot), N(\cdot), D(\cdot)$ and decision functions $\delta(\cdot), b(\cdot), K(\cdot), P(\cdot)$ that are consistent with the definitions and optimization problems stated above.

2.4 Characterization of Equilibrium

The goal of this section is to characterize the impact of information precision π on the risk composition of debt. To accomplish this task, we first characterize the optimal monitoring strategies that can arise in equilibrium to sustain any insurance contracts, and then discuss the pricing implications of our model.

However, before we proceed with the analysis, we introduce a technical assumption that helps us isolate the effect of information precision from an uninteresting effect of ex-ante segmentation of consumers into types unrelated to information. This assumption allows us to eliminate potential exotic equilibria in which lenders might want to use the signal as a segmentation device, without any regard for their informational content.[16] Under reasonable conditions this problem never arises, and it is best to assume it away to keep the paper focused.

Assumption 3. *Suppose the signal is completely uninformative. Then, lenders monitor to prevent strategic default from occurring within both signal types or none of them. Formally, if* $\pi = 0$ *then* $P(s) \geq \bar{P}$ *for all* s *or* $P(s) < \bar{P}$ *for all* s*.*

2.4.1 Optimal Monitoring Strategy

In this section, we demonstrate that only two types of monitoring strategies will used by lenders to sustain monitored insurance contracts under Bertrand competition: i) *full monitoring*, and ii) *selective monitoring*. Under full monitoring, lenders simply ignore the signal, and uniformly monitor all defaulting borrowers up to the point at which strategic default is fully prevented (i.e. non-distressed consumers are indifferent between defaulting and repaying). Under selective monitoring, lenders prevent strategic default only in the case of signal of no distress, while defaulting consumers associated with the signal of distress are not monitored enough to prevent them from defaulting strategically. Selective monitoring might also involve some monitoring of the distressed signal types. A *necessary* condition for this to be the case is that the ex-post yield from monitoring is strictly positive. To make clear how our results depend on the precision of information, throughout the paper we write $P_\pi(.)$ instead of $P(.)$ to indicate all cases in which monitoring strategy actually directly depends on the precision of signal π.

Proposition 2. *Monitored insurance contracts consistent with* (1) *are exclusively supported by one of the following two monitoring strategies:*

i) full monitoring: $P(s) = \bar{P}(R, L)$*,* $s = 0, 1$*, or*

ii) selective monitoring: $P(0) = \bar{P}(R, L)$ *and* $0 \leq P_\pi(1) < \bar{P}(R, L)$*,*

[16]For example, consider the following setup: the utility function is Leontief across the sub-periods, but across states the agent is risk neutral (e.g. $U(c, c') = U(G(c, c'))$, where G is Leontief and U is linear). In such a case, the consumer does not care about consumption smoothing across states, and interest rate is non-distortionary across sub-periods. If monitoring cost is sufficiently high, the equilibrium may involve monitoring to prevent default only within a random subset of borrowers identified by some uninformative signal (e.g. first letter of the last name). This is because credit lines are still useful in this case to relax borrowing constraints, and if the cost of strategic default is small enough, it is optimal to economize on monitoring costs by setting $P(s) \geq \bar{P}$ and $P(1 - s) < \bar{P}$, with s being used just as a tool to segment the market independently of its information content. Since these forces are difficult to compare analytically, we directly assume away the possibility of such strange equilibria.

where $\bar{P}(R, L)$ is uniquely defined by $N(d = 0, s = 0, K) = D(d = 0, s = 0, K, \bar{P})$. Furthermore, if $\pi > \pi^* \equiv 1 - \frac{\lambda}{L(1+\bar{R})(1-p)}$, then $P_\pi(1) = 0$.

Corollary 1. *Monitored insurance contracts consistent with* (1) *involve:*

i) no strategic default in the case of full monitoring,

ii) istrategic default by borrowers associated with signal $s = 1$ and no strategic default by borrowers associated with signal $s = 0$ in the case of selective monitoring.

The above corollary additionally establishes that at the monitoring intensity that makes non-distressed borrowers indifferent between defaulting or not, lenders prevent default. That is, lenders can and will raise the monitoring infinitesimally to eliminate strategic default.

Having characterized the set of possible monitoring strategies, we next discuss how lenders price defaultable debt.

2.4.2 Equilibrium Pricing of Defaultable Debt

Pricing of defaultable debt is more complex in our model due to the use of credit line contracts. This is because such contracts involve a varying rate of utilization, which depends on the interest rate. Nevertheless, partial characterization is still possible by decomposing the zero profit interest rate into a *monitoring premium* and a *default premium,* and separating the endogenous *utilization rate* in the form of a scaling factor.

The first proposition provides a decomposition of the zero profit interest rate for the class of *fully* monitored insurance contracts. It is derived from the corresponding zero profit condition associated to a credit line of size L given Proposition 2 and Corollary 1. This zero profit condition is given by

$$R\mathbb{E}b(L, R) - pL - p\bar{P}(R, L)\lambda = 0, \tag{6}$$

where $\mathbb{E}b(L, R)$ is the expected debt level associated with repayment. The remaining terms capture the the expected losses attributable to defaults and the overall cost of monitoring. The proof of the proposition directly follows from the equation above and is therefore omitted.

Proposition 3. *Given a monitored insurance contract, the zero profit interest rate under* ***full monitoring*** *can be decomposed into a monitoring premium $\mathcal{M}$ and a default premium $\mathcal{D}$ as* *follows:*

$$R = (\mathcal{D} + \mathcal{M}) \times \frac{L}{\mathbb{E}b(L,R)},$$

where

$$\mathcal{D} = p \text{ and } \mathcal{M} = p\frac{\lambda}{L}\bar{P}(R,L),$$

and $\mathbb{E}b(L,R)/L$ *is the expected utilization rate associated with repayment. Furthermore,* $\mathcal{M} = 0$ *and* $\mathcal{D} = p$ *in the case of non-monitored insurance contracts, and* $\mathcal{M} = \mathcal{D} = 0$ *in the case of risk-free contracts.*

The above result is intuitive. First, under full monitoring, only distressed consumers default, and the probability of such occurrence is p. To break even, lenders must be compensated for bearing the default risk, which we refer to as *default premium.* Second, lenders must be also compensated for the expected cost of monitoring. This cost is referred to as *monitoring premium*, and it is given by $p\lambda\bar{P}/L$. To obtain the final effect of these premia on the interest rate charged in equilibrium, the formula scales them depending on the expected utilization rate of the credit line. This adjustment is necessary because the borrowers default on the entire credit line, while interest is assessed only when borrowing actually takes place.

The next proposition derives an analogous result for *selectively* monitored contracts.

Proposition 4. *Given a monitored insurance contract, the zero profit interest rate under* ***selective monitoring*** *can be decomposed into a monitoring premium* $\mathcal{M}$ *and a default premium* $\mathcal{D}$ *as follows:*

$$R = (\mathcal{D} + \mathcal{M}) \times \frac{L}{\mathbb{E}b(L,R)},$$

where

$$\begin{aligned}
\mathcal{D} &= p + \underbrace{p(1-p)(1-\pi)}_{\text{“strategic” default premium}} \\
\mathcal{M} &= p(1-p)(1-\pi)\left(\left(\bar{P}(R,L) + \frac{P_\pi(1)}{(1-p)(1-\pi)}\right)\frac{\lambda}{L} - \left(1+\bar{R}\right)P_\pi(1)\right).
\end{aligned}$$

Furthermore, as $\pi \to 1$, $\mathcal{D}$ *and* $\mathcal{M}$ *converge to* p *and* 0, *respectively.*

As we can see, the default premium involves an additional cost associated with strategic default of non-distressed borrowers under the signal of no distress. Thus, the default premium is in this case strictly higher than under full monitoring, and the difference is a function of the precision of information. At the same time, a smaller mass of agents is monitored, and the monitoring premium is generally lower, although the difference depends on the precision of information. Moreover, the monitoring premium is reduced by the expected recovery rate – which takes place in case lenders decide to monitor distressed signal types with some intensity ($P_\pi(1)$). Under Assumption 2, expected recoveries net of monitoring costs do not fully offset the losses implied by strategic default.

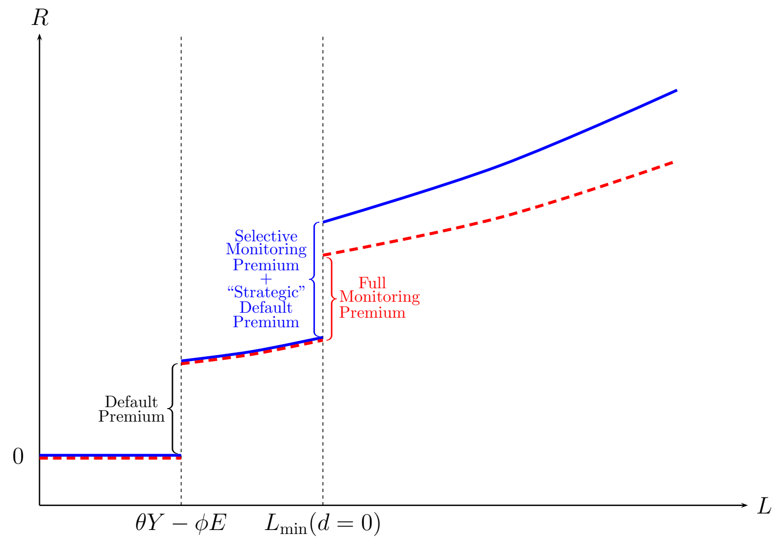

Figure 2: Equilibrium Interest Rate Schedule

Figure 2 graphically illustrates the resulting pricing schedule. By Proposition 1, we know that some contracts do not require monitoring at all, and some of them are risk-free. This makes the pricing schedule discontinuous. The figure shows the case in which selective monitoring is more expensive, but this need not be the case. Furthermore, note that selective monitoring may prevail in equilibrium even if it is associated with a higher interest rate. This is because Bertrand competitors internalize the fact that strategic default also raises utility.

2.4.3 Effects of IT Progress on Prices and Risk Composition of Debt

We next turn to the comparative statics exercise. Before we begin, we introduce an additional assumption to gain analytic tractability. This assumption makes the utility function quadratic across the sub-periods. By introducing this assumption, we effectively restrict attention to intertemporal borrowing policy functions that are linear. The assumption can thus be interpreted as an approximation of a more general case.[17]

Assumption 4. *Let the utility function be of the form $u(G(c, c'))$, where*

$$G(c, c') = c + c' - \mu(c - c')^2, \quad \mu > 0,$$

and let u be any utility function satisfying the usual set of assumptions (e.g. CRRA).

Our first result completes the characterization of the selective monitoring strategy. Namely, we establish that $P_\pi(1)$ singled out in Proposition 2 is monotonically decreasing in the precision of information π.

Proposition 5. *Given any zero profit selectively monitored insurance credit line, $P_\pi(1)$ is monotonically decreasing w.r.t. π and $\lim_{\pi \uparrow \pi^*} P_\pi(1) = 0$, where π^* is defined in Proposition 2.*

The above results allow us to establish that, as the precision of information improves, the optimal monitoring strategy that sustains any fixed credit line necessarily switches from *full* to *selective* monitoring. This is because pricing under selective monitoring is a function of the precision of information, while under full monitoring it is not. Thus, to make sure that the precision of information matters in our model, we must establish that selective monitoring prevails when the signal is sufficiently precise.

Proposition 6. *For each $L > L_{\min}(d = 0)$ there exists $\bar{\pi}(L) < 1$ such that for all $\pi > \bar{\pi}(L)$ the preferred zero profit contract under selective monitoring yields a strictly higher utility than the preferred zero profit contract under full monitoring.*

The proof directly follows from Propositions 3, 4, and 5.

[17] Quadratic intertemporal utility establishes a sufficient condition to guarantee that the following intuitive result holds: a decline in the interest burden ($Rb(R)$) for any fixed L and P by some infinitesimal amount due to a drop in R leads to a higher gain in indirect utility N the higher the initial value of R is. Our numerical results suggest that this property applies very broadly, but we are not able to establish it for a generic utility function. It can also be established for log utility.

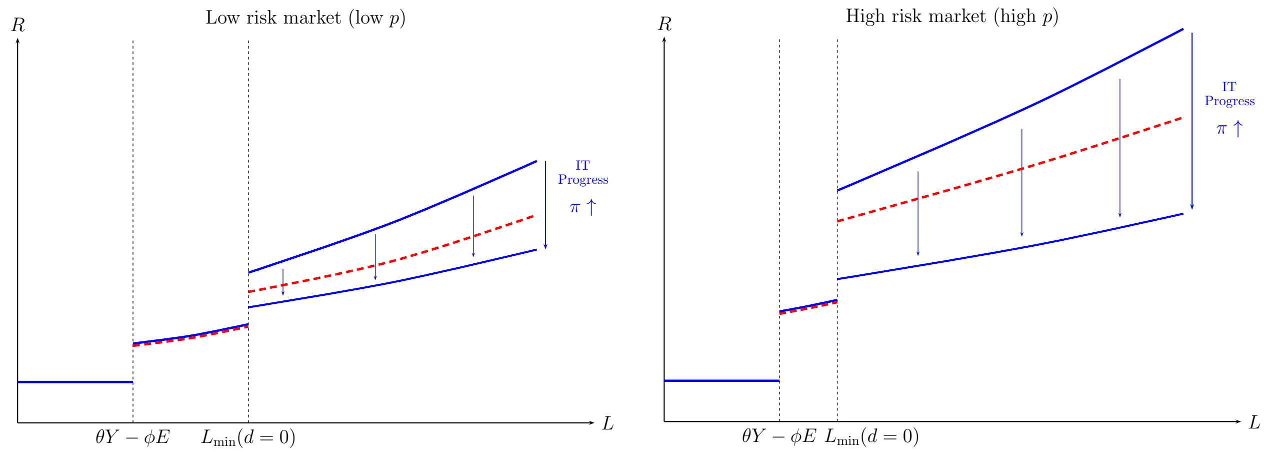

Figure 3: IT Progress and Equilibrium Interest Rates

As illustrated in Figure 3, the results that we have established so far imply that selectively monitored insurance contracts become relatively *more* attractive relative to risk-free or non-monitored insurance contracts (note that, by Proposition 1, the range of L that can be sustained without monitoring remains unchanged). Moreover, above a certain level of signal precision, selective monitoring surely prevails. Consequently, our results imply that more precise information is consistent with a larger share of risky contracts in the contract pool, as such contracts will appeal to a broader range of borrowers in the space of B and Y due to their lower price. The exact quantitative effect of this mechanism will be studied in the next section. Non-monitored insurance contracts can arise as well, but in our full model they are dominated by either risk-free contracts or monitored insurance contracts.

Furthermore, as it is illustrated in Figure 3, Proposition 4 shows that the higher the risk of distress is (p), the more sensitive prices should be to the precision of information.[18] As a result, the model implies that the riskier a given segment of the consumer market is, as implied by p, the more the price of risky credit contracts declines relative to the price of the risk-free contracts as the precision of information improves. This result is qualitatively consistent with the idea of "democratization of credit", as discussed in Johnson (2005) and White (2007). In our quantitative model we follow a parsimonious approach of not incorporating any ex-ante

[18]Up to an ambiguous effect of p on $P_\pi(1)$, which matters only in the case of $\pi < \pi^*$.

heterogeneity in risk types. To the extent that an additional source of ex-ante heterogeneity in risk types is a feature of the US data, our results might only understate the quantitative potential of our model.

3 Quantitative Analysis

In this section, we extend our baseline setup to a quantitative T-period life-cycle model ($T = 27$ in our calibration). We use this model to demonstrate that our mechanism can quantitatively account for the US data. At the end of this section we also provide a comparison of our model to the standard theory.

3.1 Extended Life-cycle Model

The dynamic aspects of the model are fairly standard and similar to Livshits, MacGee and Tertilt (2010). The previously presented static model describes the environment within each period, which is also divided into two sub-periods for the sake of consistency with the above setup. However, the consumer now borrower across the periods and B is endogenous. Furthermore, in order to reasonably capture income risk in the data, Y is allowed to be stochastic. Since contracts last only one-period, dynamic aspects only affect the consumer problem, which we discuss below.

Persistent income shocks. Income in period t is given by $Y_t = e_t z_t$ and is governed by a Markov process characterized by a transition matrix $P_t(z|z_{-1})$, and an age-dependent deterministic component e_t. The transition matrix is identical from period 1 through $T - N$, and then again from period $T - N$ to period T. The first set of periods is identified with working-age, while the second correspond to retirement.[19] During retirement, it is assumed that there is no income uncertainty, and so $P_t(z|z) = 1$, and zero otherwise.

Endogenous intertemporal borrowing. The consumer can carry debt across the periods,

[19]Retirement income is determined to yield a replacement rate of 55%, as assumed by Livshits, MacGee and Tertilt (2010). It is given by a weighted average of realized $T-N$ income of the agent (20%), and the average income in the economy (35%).

implying a modified second sub-period budget constraint given by

$$c' = Y + B' - b - \rho(K, b),$$

where $B' \leq L$ is the consumer choice of next period's debt level.

Intertemporal preferences and discounting. The utility function is assumed to be CES,[20] and it is discounted at rate β. We follow the standard approach of adjusting the units of consumption to account for the varying family size over the life-cycle. Specifically, consumption is scaled by a factor $1/s_t$ in any given period. This gives rise to a hump-shaped consumption profile over the life-cycle consistent with the data. The adjustment factors are taken from Livshits, MacGee and Tertilt (2010).

Life-cycle optimization. Agents in the model live for T-periods and solve a dynamic life-cycle optimization problem. This optimization problem determines the evolution of debt B across the periods. For example, the interim value from not defaulting is given by

$$V_t^{\delta=0}(.) = \max_{b,B'}\{U(c/s_t, c'/s_t) + \beta V_{t+1}(B', Y')\}. \tag{7}$$

The remaining value functions are defined analogously.

Formal bankruptcy. In the quantitative model, we explicitly allow both distressed and non-distressed borrowers to file for formal bankruptcy. This feature plays a role only in the case of non-distressed consumers, and places an upper bound on the maximum credit limit that is sustainable in equilibrium. It is introduced to facilitate calibration. Under formal bankruptcy, we assume that debt can always be discharged, although at a higher pecuniary cost denoted by $\overline{\theta} > \theta$ (also applies to non-distressed agents). Formal bankruptcy is never used by a non-distressed consumer in our environment, but it can be used by a monitored distressed consumer (after she finds out she is monitored). This, however, arises rarely.[21]

[20] The utility function is assumed to be of the form $U(c, c') = u(G(c, c'))$. We assume u is CES, and assume that within the period G is a quadratic function given by: $G(c, c') = c + \beta c' - \mu(c - c')^2$. We select μ to be consistent with CES utility (specifically, we consider a median income consumer, and fit the function around the point $1 = c = \beta c'$ for $R = 0$.

[21] It would be cumbersome but possible to extend our model so that both types of default actually coexist in equilibrium. Since such models already exist, we do not consider such extension here. The simplest extension would entail a random dis-utility associated with monitoring. In such a case, all distressed borrowers who

Temporary exclusion to autarky. We incorporate the standard feature of excluding a borrower to autarky for one period after default of any kind.[22]

Transaction costs. To facilitate calibration, we introduce a transaction cost of making credit card funds available to consumers. This cost is denoted by τ and lenders incur it in proportion to consumer borrowing b.

3.2 Parameterization

Model parameters are calibrated to account for the trend values of various moments in the US data for 2004. The comparative statics exercise is designed to take the model back to the 1990s. For more details how we obtain trend values, see data appendix. Period length is 2 years, implying each sub-period is 1 year long.

Income process. We start from the usual annual AR(1) process for income taken from Livshits, MacGee and Tertilt (2010) (the original AR(1) specification is of "RIP" type, and features a single income profile). In addition, we assume that any income drop above 25% is attributable to a distress shock, and place it under the distress shock (discussed below). We convert the residual to obtain a biannual Markov process using the Tauchen method. We start our simulation from the ergodic distribution of the fitted Markov process.

Distress shocks. As mentioned above, the distress shock soaks up all income shocks that imply an income drop of 25% or more within the original AR(1) process that we used to obtain the income process. We augment this shock by including three major lifetime expense shocks singled out by Livshits, MacGee and Tertilt (2010): medical bills, the cost of an unwanted pregnancy, and the cost of divorce. We use their estimated values, although appropriately adjusted to obtain a single biannual distress shock. This procedure gives $E = .4$ (40% of median annual household income), and the shock hits with 10% biannual frequency. We consider medical bills to be the only shock that can be directly defaulted on, and consequently

are monitored and face a relatively high utility cost associated with monitoring would decide to file. In anticipation of monitoring, some of them would file before monitoring takes place. The model would generate an imperfect but positive correlation between formal and informal bankruptcy.

[22] In autarky the agent can save but cannot borrow. At an arbitrary penalty interest rate of 30%, she can roll-over a fraction ϕ of the distress shock. In the data, delinquency status stays on record for 7 years, which is somewhat longer than in our model.

set $\phi = .24$. We should emphasize that this approach crucially departs from the usual practice of treating the distress shock as almost fully defaultable. In contrast to the literature, in our model only a small fraction of the shock is actually accounted for by medical bills, which results in a largely non-defaultable distress shock.

Discount factor. We set the discount factor to match the level of credit card debt to median household income of 15%, which corresponds to our estimated trend value for 2004. We chose median household income as the base because we consider it a better measure than the mean.[23]

Transaction costs. The cost of bank funds and the saving rate are both normalized to zero. However, the use of credit lines involves an exogenous transaction cost τ. We set $\tau = .12$ (biannual) to match the trend-implied annual interest premium on revolving credit card accounts of 4.0% in 2004 (as reported by FRB). This premium is defined by the difference between the average interest rate on revolving credit card accounts assessing interest, and an approximate measure of the opportunity cost of credit card funds (measured by 5-ytm yield on Treasuries),[24] and the aggregate net charge-off rate on the credit card debt as reported by the FRB.[25]

Risk aversion. We assume a standard relative risk aversion of 2.

Pecuniary costs of defaulting. We set $\theta = .32$ to match the trend value of the charge-off rate of 5.1% for 2004. We set $\overline{\theta}$ to match the average credit card debt defaulted on per statistical bankruptcy filer relative to her income of 93% (as reported by Sullivan, Westbrook and Warren (2001), and extrapolated using a linear trend regression to obtain the trend value for 2004). Due to data limitations, this calibration target pertains to formal rather than informal bankruptcy filers, and the latter would be more appropriate for our model. Nonetheless, our model is not particularly sensitive to this target and could accommodate a lower number. Finally, we set $\bar{R}$ to assure $\bar{R}L = \rho(K, b)$ (where b is borrowing associated with repayment) and assume $\underline{\theta}$ is

[23] The increase in the per capita income over the sample period is almost fully accounted for by the top income groups (mostly the top 1% of earners). In our view, this makes the average income ill-suited to study the evolution of the unsecured credit market that serves mostly the rest of the population. Calibrating our model to lower debt to income would make our job only easier.

[24] The trend is almost identical for bond maturities 1-5 (or 1 year LIBOR), although the levels are obviously higher when shorter maturities are used instead. We chose 5ytm to match the maturity of credit card accounts in the data, and to avoid the volatility associated with the short-term rates.

[25] The net charge-off rate is the fraction of credit card debt discharged by lenders (180 days or less) net of recoveries. As such, it is an aggregate measure of the idiosyncratic default risk premium borne by lenders.

zero so that $\bar{P} = 1$. This simplification only matters for the interpretation of the calibrated value of the monitoring cost, and is otherwise innocuous.

Information technology and monitoring. We arbitrarily set the precision of signal to $\pi = .5$, and assume a mean-preserving spread of 50% around this number to capture IT progress. We choose the monitoring cost so that the transition to selective monitoring is centered around the middle of the signal precision interval, obtaining $\lambda = .3$. This value implies that monitoring cost is about 30% of annual median household income (per monitored borrower). This is fairly high, although not unreasonably high. It can be lowered to about 20% by assuming noisier information. If lowered further, our results gradually weaken.

Comparative statics exercise. Our comparative statics exercise involves two technology parameters: the precision of information π and the transaction cost τ. We vary these parameters to account for the trends in the data. To discipline the change in τ, we note that the average productivity growth within the banking industry outpaced the rest of the economy by about 22% from 1990 to 2004, as reported by Berger (2003). Consistent with this data, we lower the value of τ by 20% from the base value of .15 to .12. Regarding the precision of information π, as already pointed out, we choose a mean-preserving spread around 50% precision to ensure a transition from full to selective monitoring over the 1990s, resulting in $\pi_{90s} = .25$ and $\pi_{00s} = .75$.

3.3 Quantitative Findings

Table 1 reports our key findings. Figure 1 illustrates the implied trends, and compares them to the data. In what follows next, we first discuss the implications of our model for trends. We also compare our model to the standard model, emulated using our framework to facilitate a consistent comparison. The results of this comparison are summarized in Table 2.

3.3.1 Quantitative Effects of IT Progress

As is clear from Figure 1 and rows 1,2 and 5 of Table 1, our model fully accounts for the key trends seen in the US data. Importantly, all of these effects can be traced back to technological progress within the credit card industry (in the case of τ, in excess of the technological progress in the rest of the economy).

Table 1: Quantitative Results from the Benchmark Model

	Data[a]	Model	Data	Model	Decomposition		
	00s	00s	90s	90s	$\tau = \tau_{90s}$	$\pi = \pi_{90s}$	
					$\pi = \pi_{00s}$	$\tau = \tau_{00s}$	$\tau = \tau_{\text{fit}}$
(in % unless otherwise noted)	(1)	(2)	(3)	(4)	(5)	(6)	(7)
1) CC Debt to Median HH Income	15.1	15.1	9.0	9.0	11.2	13.9	15.1
2) Net CC Charge-off Rate	5.3	5.4	3.5	3.5	5.5	4.1	4.1
3) Defaults (households per 1000)	-	10.8	-	4.5	9.0	7.5	7.9
- fraction monitored ($m = 1$)	-	18	-	30	17	31	32
- fraction strategic ($m = 0, d = 0, s = 1$)	-	19	-	0.0	19	0	0
4) Frequency of insurance contracts[b]	-	36.6	-	21.4	35.7	31.3	31.1
- fraction fully monitored	-	1	-	100	0	100	100
- fraction selectively monitored	-	99	-	0	100	0	0
5) CC Discharge to Income of Defaulters[c]	93	89	48	74	82	80	82
6) Mean CC Interest Premium[d]	4.0	4.4	6.6	6.5	6.1	5.3	4.6
7) Fraction of Revolvers in Population	40	53	27	49	50	57	57

All values in % unless otherwise noted.

[a] Data corresponds to trend values for 1990 and 2004 (due to major bankruptcy reform we do not consider here years after 2004). This procedure allows us to abstract from business cycle fluctuations in the underlying variables. Linear trends are estimated using time series from 1985 to 2004, whenever possible. Mean unsecured debt defaulted on per statistical (formal) bankruptcy filer, as default data is only available for formal filings. Linear trends fit all time series reasonably well. See data appendix for more details and a list of sources.

[b] As a fraction of the total number of revolving contracts ($b > 0$).

[c] Reported data value pertain to total unsecured discharged debt to income of formal bankruptcy filers.

[d] See previous footnote; interest rate on revolving consumer credit card accounts assessing interest, less the charge-off rate and the opportunity cost of funds (our preferred measure of the opportunity cost of cc-funds is the yield on 5-ytm US Treasuries).

Columns 5-7 of Table 1 decompose the contribution of π and τ. Specifically, column 5 reports the contribution of an increase in π, by fixing τ at the level from the 90s, while the next column considers an analogous exercise with τ. Finally, column 7 aims at fitting the change of indebtedness over this time period by changing τ alone ($\tau_{\text{fit}} = .11$). Together, these exercises demonstrate that, while τ is important to account for the growing indebtedness of the household sector, its contribution to the rise in the default exposure of credit card debt is modest and falls short of the data. As already discussed, this result is intuitive: the most powerful channel that raises default exposure is the one that affects the relative price of risky contracts to risk-free contracts. Furthermore, the comparison of the last two columns reveals that the effect of τ is non-linear. This is because it relies on the effect of "bunching" at the discontinuity point of the pricing schedule.

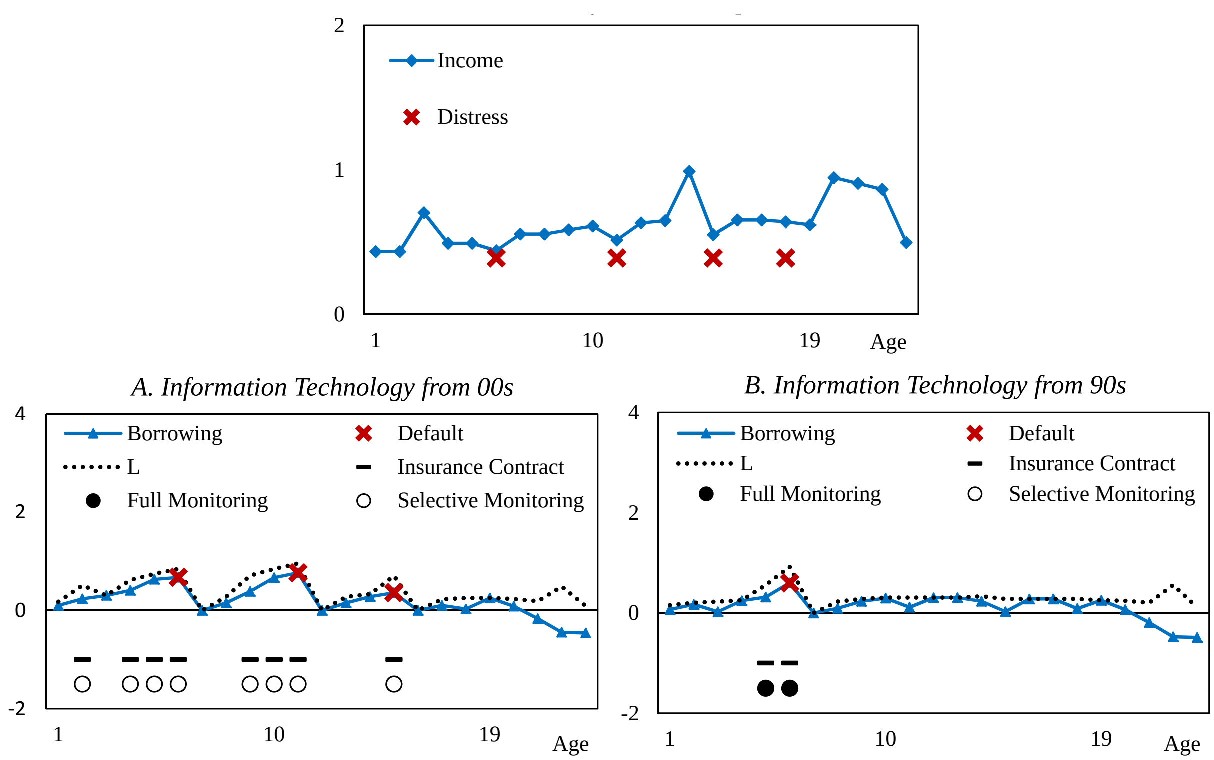

Figure 4: Life-Cycle of a Severely Distressed Borrower

Simulated time-series from the model confirm the intuition behind the results discussed earlier in the paper. For example, consider the simulation illustrated in Figure 4. The figure presents a life-cycle profile of a highly distressed consumer. As illustrated in the top panel, this consumer has low income over her entire life, and suffers from several distress shocks. Bottom panels compare the simulation of this consumer assuming the equilibrium from the 00s and the 90s, respectively. Bars below each plot indicate when a given contract involves insurance. The circle indicates whether a given contract is sustained using full or selective monitoring strategy. As we can see, while the consumer defaults three times in the 00s, she only defaults once in the 90s. The other defaults are simply eliminated through tighter credit limits. Furthermore, the only default that occurs in the the 90s is associated with a lower discharge. This is because contracts prior to the event also feature tighter credit limits, making it more difficult for the borrower to accumulate debt prior to defaulting.

The 3rd row in Table 1 shows that our model implies an annual default rate of about 11 per 1000 households. This is about twice as high as the observed formal bankruptcy filling rate in the data, which is not surprising. Given that we have calibrated our model to match the charge-off rate from the data, and also required that the model matches the average discharge per bankruptcy filer in the data, the default rate generated by our model should be about twice as high as the formal default rate in the data. Had we calibrated our model to a lower target, default rate would be correspondingly higher.

Row 4 of the table shows that the share of insurance contracts increases along the transition, and monitoring intensity per insurance contract declines. These effects confirm the intuition that follows from the static model. The prediction is also qualitatively consistent with the data, as discussed earlier in the paper.

Row 5 of the table shows that our model falls short of accounting for the changes on the extensive margin. However, by incorporating some degree of ex-ante heterogeneity in p, our model could account for these changes as well. As noted earlier, such departure would only reinforce our results.

3.3.2 Comparison to the Standard Model

Before we conclude, it is instructive to compare some of the implications of our model to the standard theory of consumer bankruptcy.[26] To this end, we emulate the standard theory within our framework. Specifically, we shut down the access to a monitoring technology and assume

[26] For example, see Chatterjee et al. (2007) or Livshits, MacGee and Tertilt (2006).

$\underline{\theta} = \theta = \overline{\theta}$. Similarly to the benchmark model, we choose the value of the pecuniary cost θ to match the charge-off rate in the data, and attempt to find the value of the discount factor β to match debt to median household income. (Our version of the standard model features one period credit lines instead of one-period loan contracts. Since maturity of contracts is matched, this does not matter for the predictions of the model.)

Table 2 reports the results from two versions of the standard model. The first version reported in the table is labeled "$\phi = .25$" , and it shares with our benchmark model the same values for the other parameters. The second version, labeled "$\phi = 1$" , departs from the benchmark model in a single dimension by assuming a fully defaultable distress shock ($\phi = 1$).

The reason why we include two calibrations is because the standard model cannot simultaneously match the level of indebtedness and the net charge-off rate (for any value of β, θ), and both calibrations are of interest given the existing literature. Specifically, when required to match the charge-off rate, the standard model falls short in terms of indebtedness by at least an order of magnitude. This is because the model can only match the charge-off rate by making the consumer default very frequently on small amounts, which is only consistent with a low average level of indebtedness. This version of the standard model represents one of the early approaches to relate the theory to the data. However, under this interpretation one should be aware that the charge-off rate in the model and in the data are two different objects and cannot be readily compared. This is because the calculation of the charge-off rate in the data involves gross debt (in the denominator), while the model has only prediction for net debt.[27]

The second calibration resorts to the usual assumption of a fully defaultable distress shock ($\phi = 1$) to match the gross positions in the data. This version of the standard theory is in the spirit of the approach favored by Livshits, MacGee and Tertilt (2010). The key strength of this particular calibration is that the model can generate the kind of gross positions that make sense from the perspective of the lending industry – arguably offering a better framework to think about the supply side of the market. Nevertheless, despite this extreme assumption,

[27]To justify such modeling approach researchers have used net asset positions rather than gross positions. For an example of such modeling approach, see Chatterjee (2010). It is also used by Sanchez (2012) and Athreya, Tam and Young (2008).

Table 2: Comparison to the Standard Model[a]

	Benchmark Model[b]		Standard Model[c]			
	00s	90s	$\phi = .25$		$\phi = 1$	
			00s	90s	00s	90s
(in % unless otherwise noted)	(1)	(2)	(3)	(4)	(5)	(6)
1) CC Debt to Median HH Income	15.1	9.0	1.9	1.5	15.1	12.5
2) Net CC Charge-off Rate	5.4	3.5	5.2*	5.3*	5.2	5.3
3) Defaults (households per 1000)	10.8	4.5	34.8	34.6	34.5	33.7
4) Frequency of insurance contracts	36.6	21.4	88.0	89.1	95.7	98.5
5) CC Discharge to Income of Defaulters	88.7	73.9	2.1	1.8	18.3	16.3
6) Mean CC Interest Premium	4.4	6.5	6.0	7.7	6.5	8.1
7) Fraction of Revolvers in Population	53.3	48.5	10.6	5.4	21.8	18.9

[a] Notes to previous table apply.

[b] Benchmark model corresponds to base calibration of our model (see previous table).

[c] Standard model corresponds to an emulated version of the standard model using our framework (monitoring technology and $\underline{\theta} = \theta = \overline{\theta}$); the columns present two calibrations: Case $\phi = .25$ features partially defaultable distress shocks as in the benchmark model and case $\phi = 1$ assumes that distress shocks can be fully discharged.

[b] If debt is interpreted as unsecured debt net of assets, charge-off rate from the model is not comparable to the data. See explanation in text.

matching the data still turned out difficult. As is clear from row 3, even in the best case scenario the model over-predicts the default rate by about a factor three. The reason is that the model simply cannot match the targeted discharge per bankruptcy filer (see row 4).[28]

On the basis of the above results, we conclude that our mechanism also allows to improve the performance of existing models in terms of levels. The explanation why this is the case is related to the fact that the standard model involves a inherent tension between debt sustainability and attractiveness of default. More specifically, on one hand a severe punishment for defaulting is required to sustain defaultable debt in equilibrium, but on the other hand the same punishment discourages default. Our model naturally resolves this tension by incorporating endogenous monitoring that targets only solvent defaulters. As a result, high indebtedness can be sustained per solvent borrower, yet the attractiveness of default can be maintained in the case of insolvent (distressed) borrowers.

Finally, to complete our analysis, we consider the effect of IT progress in the standard model. Since the model features full information, we can only evaluate the effect of IT progress manifested through the transaction cost τ. As we can see, the charge-off rate falls, and the result highlights the non-linearity associated with τ. We observe the same property in our model. For example, when we lower τ within the parameterization matching the end of the sample (00s), the charge-off rate similarly declines (for example, see the decomposition in Table 1). Intuitively, this is because the transaction cost does not change the relative attractiveness of risk-free contracts relative to risky contracts, and only increases the fraction of the population that uses credit. Whenever risk-free contracts are extended in equilibrium sufficiently often, the transition from saving to borrowing may in some cases raise the share of insurance contracts. However, this effect is weak and non-linear. Based on these results, we conclude that a fall in the transaction cost τ is not a promising approach to match the data. This result is consistent with the findings reported by Livshits, MacGee and Tertilt (2010).

[28] In the original paper, Livshits, MacGee and Tertilt (2010) manage to match the charge-off rate and the level of discharged debt per statistical bankruptcy filer much better than we do here. This likely happens for three reasons. First, the maturity of loans is 3 years in their case, and it is only 2 years in our case. Second, their calibration uses the average income rather than median household income (which gives debt to income target that is about half of what we have). Third, they use ‘forced partial’ repayment as a punishment for defaulting, which helps sustain more debt, and raise the amount of debt defaulted on.

4 Conclusion

Existing theories of consumer bankruptcy rule out the option of informal bankruptcy by assumption, and abstract from costs of debt collection. Here we argue that this assumption is not only at odds with the data, but it can drastically change the predictions regarding the impact of IT progress on the ex-ante pricing of unsecured credit. In particular, we show that an endogenous link to IT progress operating through this margin is potent enough to account for all the facts underlying the IT-driven expansion of credit card borrowing in the 1980s and over the 1990s. Independently, our approach can help account for the high levels of default rate and credit card debt seen in the recent data.

Appendix

A1. Overview of the Debt Collection Process in the US

In this appendix, we briefly describe the key debt collection process in the US. We focus on credit card debt, but similar methods are generally used in the case of other forms of unsecured credit.

Life-cycle of delinquent debt. Delinquent debt is first handled by the original creditor. As far as credit card debt is concerned, the first contact with the borrower is made right after 30 days of delinquency. The effort intensifies every 30 days. Between 120 and 180 days, credit card debt enters a pre-charge-off stage. the debtor might be offered a settlement deal. After 180 days, in the case of credit cards, debt must be charged off by the original creditor (in case of credit card), meaning that the account is no longer listed as an account receivable on the creditor's books, and its value is charged against the creditor's reserves for losses. Any payment on the charged-off debt obtained later in the collection process is then treated as income.[29]

[29] See Federal Trade Commission, "Collecting Consumer Debts: The Challenges of Change", A Workshop Report 2009, FTC, and references therein. Herkenhoff (2012) reports cure rates of delinquent debt. The data pertain to the time period after the 2005 bankruptcy reform. According to this evidence, about 45% of accounts that are 30 days delinquent become current next quarter. The recovery rate falls by slightly less

After discharge, unless the consumer has filed for bankruptcy, debt is sent to collection, typically first in-house, and then it is likely to be sold to a third-party debt buyer. The collection effort extends until the statute of limitations on debt expires, which varies from 3 to 7 years depending on the state (unless a judgment is obtained beforehand).

In-house collections. Discharged debt is first placed in collection in-house. In the case of in-house collections, the process is handled either by an in-house collection department or more typically by a third-party credit collection agency that works with the original creditor on a contractual basis. From 1996 onward, in-house collection is *practically* limited to 36 months due to tax regulations.[30] Generally, in-house collection uses the same technology as third-party collection, although methods may differ simply because collection takes place earlier.

Third-party collections. Uncollected debt by the original creditor, especially after 1996 (see footnote 30), is typically sold to third-party debt buyers.[31] As already mentioned, by selling discharged debt, the original creditors can avoid premature tax-related closure of the debt collection process while cashing out on the residual value of debt immediately. In fact, tax treatment of uncollected debt might have been a decisive in driving the explosive growth of the third-party collection industry in the mid to late 90s.[32]

Price of discharged debt. The price of discharged debt crucially depends on the number of previous collection attempts and on the age of debt. Average prices of debt portfolios seem

than half for every 30 days up to 120+ days.

[30]From 1996 onward, if creditors discharge debt, after 36 months of bona fide collection efforts, they are required to file a 1099 cancelation of debt form with the IRS. In practice, only ongoing litigation or packaging of debt for resale is a valid exception to this rule (see 6 CFR Ch. I (4112 Edition), 1.6050P-22). Interestingly, third-party debt collection agencies were excluded from this regulation until about 2003. In 2004-2006, IRS filed a lawsuit to enforce the rule. While the court sided with the IRS, due to the inability of the credit collection industry to file proper 1099 forms, the IRS suspended the application of the 36-month rule to third-party debt collectors (for more details, see amendment RIN 1545-BH99 of 1.6050P-1 26 CFR 1 in Federal Register Volume 73, Number 218 (Monday, November 10, 2008), Rules and Regulations, pp. 66539-66541). In reality, due to the statute of limitations applying to IRS audits, which is 3 or 6 years depending on the amount, cancelation of debt handled by third-party collectors is not taxed (see, for example, David Scott Stewart and Carla Annette Stewart v. Commissioner of Internal Revenue Service Docket No. 10374-11S, United States Tax Court).

[31]For example, in Figure 4, Hunt (2007) compares the volume of credit card debt charged-off to the volume of debt sold to third-party debt buyers. The data comes from Nielsen report. Back of the envelope calculation suggests that, in the more recent period, most debt is eventually sold.

[32]See Hunt (2007) more information about this industry.

to hover around few cents on a dollar, while pre-charge-off 'fresh' discharged debt can be priced around 20 cents on a dollar.[33]

Collection methods. To collect debt, credit collection agencies today rely on sophisticated statistical models that are based on databases of past collections, the credit history of borrowers, and other supplementary sources. These methods were not available in the past on such a wide scale. According to anecdotal evidence, the IT progress within the debt collection industry very much mirrored the IT progress in the lending industry.[34]

In a fraction of cases, collection process may involve litigation. By entering the legal path, collectors usually seek a judgment in state courts, allowing them to garnish debtors' wages, seize bank accounts or place a lien on debtor's property. Judgments may also be obtained to extend the option of collecting unpaid debt beyond the expiration of the statute of limitations. Since judgments can be eliminated through a formal filing for bankruptcy, they generally do not diminish the attractiveness of the informal default option relative to a formal bankruptcy filing.[35]

According to a study of 1999 credit records by Avery, Calem and Canner (2003) about a third of consumers had one judgement on their record in 1999, and another third had two (judgments stay on record for 7 years). However, very few (15.8%) of these judgements have been paid, which suggests that this tool might be predominately used to gain leverage against debtors or extend the statute of limitations on uncollected debt.

Lawsuits potentially carry substantial risk for debt collectors. This is because they entail

[33] Our estimates are based on the data reported in 10-K forms by one of the major credit card debt collection agency: Portfolio Recovery Associates, Inc. The numbers are consistent with those reported by another large agency, Encore Capital Group, Inc.

[34] For example, a major debt collection company, Encore Capital Group, Inc., in its Annual Report (10-K filed with SEC) describes the company's methods as follows: "*We pursue collection activities on only a fraction of the accounts we purchase, through one or more of our collection channels. The channel identification process is analogous to a decision tree where we first differentiate those consumers who we believe are unable to pay from those who we believe are able to pay.*" In the same report, Encore Capital also states: "*We have assembled a team of statisticians, business analysis and software programmers that has developed proprietary valuations models, software and other business systems that guide our portfolio purchases and collection efforts. (...) Our valuations are derived in large part from information accumulated on approximately 4.8 million accounts acquired since mid-2000.*". Portfolio Recovery Associates describes its methods in a similar manner, mentioning two econometric models that they use to price portfolios and select appropriate action. They also report significant gains in collection effectiveness over the recent years.

[35] According to Hynes (2006), judgments are rather difficult to enforce. See details in the paper.

substantial upfront administrative costs, and these costs are sunk in cases in which a delinquent debtor is found insolvent or chooses to file for formal bankruptcy protection thereafter. Not surprisingly, cash collected through litigation accounts for only a quarter of the cash collected by major debt buyers, such as Portfolio Recovery Associates, Inc. and Encore Capital Group, Inc. In addition, about 50% of such recoveries appear to be absorbed by the legal costs and fees associated with litigation.[36] This aspect of the data underscores the asymmetry of information associated with debt collection and the importance of information technology to properly identify 'solvent' debtors prior to litigation. Anecdotal evidence suggests that technology is used to precisely address this problem.[37]

A2. Omitted Proofs

Proof of Proposition 1. We begin by noting several properties of our model: i) indirect utility function D (defined in (5)) is decreasing in P (strictly for $d = 0$), and indirect utility function N (defined in (4)) is independent from P; ii) N is (strictly) decreasing in R, and D is independent from R, implying that the incentives to default monotonically increase with R for any $L > 0$; iii) default eventually takes place for sufficiently high L, given any fixed $P < 1$; iv) for L sufficiently close to 0, any $R < 1$ and also any P, the consumer strictly prefers not to default; v) D and N are continuous in R and P. These properties follow from the unambiguous and continuous effect of P and R on the size of the consumer budget set. The proof is trivial and is omitted.

Part 1): Fix $R = 0$. Compare problems (4) and (5) assuming $d = 1$: they share the same objective function, and the first sub-period budget constraint is identical (notice the difference in the constraint on b across (4) and (5)). As a result, the decision to default boils down to the comparison of the second sub-period 'net' resources between repayment and default. The former are given by $Y - E$, while the latter are given by $(1-\theta)Y - (1-\phi)E + L$. Thus, the consumer always defaults when $L > \phi E - \theta Y$, regardless of R by ii) (This fact will be used to establish the last part of the proposition). Given iv), there exists a strictly positive cut-off

[36] 2006 income statements of Encore Capital Group, Inc.

[37] See, for example, "Technology Brings a Kinder, Gentler Process to Collections", by Louis Barney, in Card and Payments, October 2005.

value of L at which the consumer switches her decision (given fixed R and P). We refer to this value as $L_{min}(d = 1, R)$, and note that the cutoff must be decreasing and continuous w.r.t. R by ii) and v).

Part 2): Assume the worst case scenario of $P = 0$ and fix $R = 0$. Using the same reasoning as in part 1), if net resources under repayment, Y, are higher than under *non-monitored* default, $(1 - \theta)Y + L$, a non-distressed consumer will certainly choose not to default. By ii) and iv) this establishes the existence of a (positive) cutoff. We refer to it as $L_{\min}(d = 0, R)$. By i), ii) and v) it is continuous and decreasing in R.

To prove part 3), we note from i) that the default decision of a non-distressed consumer must be weakly decreasing in the underlying monitoring probability; that is, if a non-distressed consumer decides not to default for $P(s) = \hat{P}$, she will not default for any $P(s) > \hat{P}$. Furthermore, we note that for $P(s) = 1$ the consumer will choose not to default (by Assumption 1), and for $P(s) = 0$ she will prefer default given the definition of $L_{min}(d = 0, R)$. Accordingly, there must exist some $\bar{P} < 1$ – which generally depends on contract terms (R, L) – such that a non-distressed agent defaults if $P(K, s) < \bar{P}$, and does not default when $P(K, s) > \bar{P}$. Continuity of the indifference point w.r.t. R follows from v). It is decreasing w.r.t. R by i) and ii). $\bar{P}(R, L)$ being independent of π is trivial since precision does not enter the consumer problem at all.

To prove that $L_{\min}(d = 1) = \phi E - \theta Y$ note that, when charged $R = 0$, distressed consumers are indifferent between defaulting or not at $L = L_{\min}(d = 1)$ while non-distressed consumer strictly prefer not to default. Hence, if indifferent consumers do not default, lenders face no default losses and can feasibly offer $R = 0$. Finally, $L_{\min}(d = 0)$ being independent of π trivially follows from the fact that at $L_{\min}(d = 0)$, monitoring does not affect consumer indirect utility for any realization of d since non-distressed agents are indifferent between defaulting or not. Thus, if they do not default as assumed, lenders will not monitor, regardless of precision, and their zero profit condition does not depend on π. □

Proof of Proposition 2 and Corollary 1. Fix $L > L_{\min}(d = 0)$, implying that non-distressed consumers will want to default as long as they expect to be monitored with sufficiently low probability (see Proposition 1). By Assumption 2, in order for lenders to break even, it must be the case that $P(s) \geq \bar{P}(R,L)$ for at least one signal realization, where $\bar{P}(R,L)$ stands for the cutoff identified in part 3) of Proposition 1. Hence, there are three possible scenarios that can arise in equilibrium: (i) $P(s) \geq \bar{P}(R,L)$ for $s = 0,1$, (ii) $P(0) \geq \bar{P}(R,L)$ and $P(1) < \bar{P}(R,L)$; and (iii) $P(0) < \bar{P}(R,L)$ and $P(1) \geq \bar{P}(R,L)$. We next show that (i) and (ii) must satisfy the conditions of the proposition, and we rule out (iii) under the stated assumption.

Case (i): Since monitoring is costly, lenders should commit to the lowest monitoring probability to sustain any targeted level of default. Thus, if equilibrium is of type (i), it must be that $P(s) = \bar{P}(R,L)$ for all $s \in \{0,1\}$. This is because in the case of $P(s) > \bar{P}(R,L)$, a lender has a profitable deviation. In particular, he can offer a contract with the same credit limit L and a slightly lower $P(s)$ without changing agents' default choices. By doing so, the lender incurs in lower monitoring costs and thus can charge a lower R, raising agents' utility while earning strictly positive profits. Note that $\bar{P}(R,L)$ declines as R is lowered by Proposition 1, which only reinforces the argument. Furthermore, none of the non-distressed agents default under i). To see why, note that if the mass of non-distressed defaulters is positive under some signal realization, a lender could increase the associated monitoring probability infinitesimally, driving such fraction to zero and unambiguously increasing profits. At the same time, we know that defaulting non-distressed consumers under s are indifferent between defaulting and not defaulting, and so their utility does not go down with the increase in $P(s)$. Utility of distressed agents or non-distressed agents who do not default remains the same as well. Since this deviation implies left-over resources for lenders, interest rate can be lowered, raising ex-ante utility – recall that $\bar{P}(R,L)$ is increasing w.r.t. R.

Case (ii): Using a similar argument as in case i), we note that $P(0) = \bar{P}(R,L)$ and so non-distressed agents under $s = 0$ do not default in equilibrium. To show that $P(1) = 0$ if $\pi \geq \pi^*$, notice that, since all agents default under $s = 1$ if $P(1) < \bar{P}(R,L)$ and borrowing policy under repayment is independent of P, for fixed contract terms, the marginal increase in profits due to a infinitesimal increase in $P(1)$ satisfies

$$\frac{\partial \mathbb{E}\Pi(I,K,P)}{\partial P(1)} = Pr(s=1)\left(Pr(d=0|s=1)L(1+\bar{R}) - \lambda\right) \leq p\left((1-p)(1-\pi)L(1+\bar{R}) - \lambda\right)$$

for all $P(1) < \bar{P}(R,L)$. It is easy to check that $\frac{\partial \mathbb{E}_{K,P}\Pi(I,K,P)}{\partial P(1)} \leq 0$ when $\pi \geq \pi^*$. As a result, lowering monitoring probabilities all the way down to $P(1) = 0$ with R, L unchanged is profit feasible and increases consumer welfare, since consumers prefer lower monitoring probabilities. Furthermore, due to the lower cost of monitoring under $s = 1$, lenders can offer a (weakly) lower interest rate R, which also lowers monitoring costs under $s = 0$, since $\bar{P}$ is decreasing in R.

Case (iii): Next, we rule out type iii) equilibria. By way of contradiction, assume $P(0) < \bar{P}(R, L)$ and $P(1) = \bar{P}(R, L)$, with $K = (R, L)$ such that (1) is satisfied. This implies that, in equilibrium, all non-distressed agents default whenever the signal (correctly) indicates no distress. At the same time, none of them defaults when $s = 1$ by the same argument used in (i). To see the contradiction, note that that whenever $P(1) = \bar{P}(R, L)$, we must also have $P(0) = \bar{P}(R, L)$. This is certainly the case when $\pi = 0$ by Assumption 3. Furthermore, monitoring costs associated with $P(0) = \bar{P}(R, L)$ are strictly decreasing in π, as they are given by

$$Pr(s = 0)Pr(1|0)\bar{P}(R, L)\lambda = p(1 - p)(1 - \pi)\bar{P}(R, L)\lambda.$$

At the same time, monitoring costs are independent of π for any fixed $P(0) < \bar{P}(R, L)$, as all agents default under $s = 0$. Accordingly, if setting $P(0) = \bar{P}(R, L)$ at $\pi = 0$ for a zero profit contract (R, L) is preferred to having another zero profit contract associated with a lower monitoring probability $P(0)$, the same must be true for any $\pi > 0$. This is because monitoring costs are lower and ex-ante utility can be increased by lowering the interest rate (again, recall that $\bar{P}(R, L)$ is decreasing w.r.t. R).

The last part of the proposition is trivial: by Proposition 1, if $L < L_{min}(d = 1)$ either only distressed consumers default or nobody defaults. In such cases, lenders optimally set monitoring probabilities to zero. Obviously, if nobody defaults monitoring costs are zero regardless of monitoring probabilities so lenders are indifferent between any P.

Corollary 1 directly follows form the above. □

Proof of Proposition 4. Since all distressed agents and non-distressed agents with $s = 1$ default, and a fraction $P_{\pi}(1)$ of non-distressed agents are reverted back to repayment under $s = 1$, the zero profit condition under SM is given by[38]

$$0 = R\mathbb{E}b(L,R) - [p + pPr(0|1)]\, L - \left[(1-p)Pr(1|0)\bar{P}(R,L) + pP(1)\right] \lambda + pPr(0|1)L(1+\bar{R})P_{\pi}(1).$$

To derive the expression stated in the proposition, we substitute $Pr(0|1) = (1-p)(1-\pi)$, $Pr(1|0) = p(1-\pi)$, solve for R and rearrange terms. The last part follows form the fact that, as signal precision is increased, $Pr(1|0)$ and $Pr(0|1)$ monotonically decrease to zero, and $P(1) = 0$ for all $\pi > \pi^*(< 1)$ by Proposition 2. □

Proof of Proposition 5. Fix credit limit L and signal precision π_L, and let $P^*(0) = \bar{P}$ and $0 \leq P^*(1) < \bar{P}$ be the monitoring probabilities under the best selectively monitored contract with credit limit L. Also, let $R(\pi_L, P^*)$ denote the associated zero profit interest rate. Suppose the precision of information improves to some level $\pi_H = \pi_L + \varepsilon$, for some infinitesimally small ε and that P^{**} denotes the monitoring probabilities of the best selectively monitored contract with L under π_H. We need to prove that $P^{**}(0) = \bar{P}$ and $P^{**}(1) \leq P^*(1)$, with strict inequality if $P^*(1) > 0$. Below, we first focus on establishing $P^{**}(1) \leq P^*(1)$, and comment at the end why a strict inequality actually applies.

By way of contradiction, suppose that $P^{**}(1) > P^*(1)$. Let V_π denote consumers' ex ante utility under the precision of information π. Our goal is to show that

$$V_{\pi_H}(L, R(\pi_H, P^{**}), P^{**}) \geq V_{\pi_H}(L, R(\pi_H, P^*), P^*)$$

implies

$$V_{\pi_L}(L, R(\pi_L, P^{**}), P^{**}) > V_{\pi_L}(L, R(\pi_L, P^*), P^*).$$

We first focus on the case that L is never binding and deal with the case of binding L at the end of the proof. Recall that, by Corollary 1, none of the non-distressed consumers default under $s = 0$ and all of the non-distressed consumers default under $s = 1$.

Let x denote the additional revenue collected by lenders under π_H by changing the mon-

[38]In the above expression, the second term in the RHS reflects *gross* default losses, given by the mass of distressed agents p, and the mass of non-distressed agents with $s = 1$, which is given by $pPr(0|1)$. The third term reflects the costs associated with monitoring. Finally, the last term captures expected recovered debt from defaulting non-distressed borrowers.

itoring from P^* to P^{**}. By assumption, the collected revenue must be strictly higher than the additional monitoring costs, otherwise the ex-ante utility of consumers would be lower under P^{**} than under P^*. Note that, while the additional (overall) monitoring costs under $s = 1$ implied by the deviation are independent of the precision of information (the mass of monitored agents is equal to p in both cases), the increase in revenue collected by deviating from P^* to P^{**} is actually strictly higher under π_L than under π_H. This is because the mass of non-distressed agents under signal $s = 1$ is higher when π is lower, while the revenue collected from each monitored agent, given by $L(1 + \bar{R})$, is independent of π and R. Specifically, the extra revenue collected by deviating from P^* to P^{**} under π_L is fx, where $f = \frac{Pr(0|1;\pi_L)}{Pr(0|1;\pi_H)} > 1$, where $Pr(y|z;\pi)$ the probability that $d = y$ conditional on signal $s = y$ when precision is given by π.[39]

Next, suppose the collected resources under π_L are redistributed equally to all repaying consumers by appropriately lowering the interest rate R (to ensure that the zero profit condition holds). Our goal is to show that, if such redistribution justifies the increase in monitoring probability under π_H relative to P^*, it must justify a similar increase under π_L. For now, we consider $\bar{P}(L, R)$ unchanged, and comment at the end why its change works to our favor.

To this end, we first note that the indirect utility of distressed agents is independent of P and R, and so the change in ex-ante utility due to the increase in monitoring probabilities and the corresponding reduction in interest rates under π is determined by:

$$\begin{aligned}\Delta V_\pi &= V_\pi(L, R(\pi, P^{**}), P^{**}) - V_\pi(L, R(\pi, P^*), P^*) \\ &= pPr(0|1;\pi)\Delta D_\pi(0|1) + (1-p)Pr(0|0;\pi)\Delta N_\pi(0|0),\end{aligned}$$

where $\Delta D_\pi(0|1)$ and $\Delta N_\pi(0|0)$ respectively denote the change in the indirect utility of a non-distressed defaulting consumer under signal $s = 1$ and a repaying consumer under signal $s = 0$. Importantly, the drop in utility of a non-distressed defaulting consumer, $\Delta D_\pi(0|1)$, is the same under π_L and π_H. This follows from the fact that for a non-distressed defaulting consumer the only relevant variable is the monitoring probability P, and these probabilities are identical

[39]Note that the extra (gross) revenue collected due to an increase in monitoring probability $\Delta P(1)$ is given by $Pr(s = 1)Pr(0|1)L(1 + \bar{R})\Delta P(1)$.

after the deviation in both cases (equal to P^{**}). Hence, we can re-write the change in the ex-ante utility in the case of π_L as follows

$$\begin{aligned}\Delta V_{\pi_L} &= pPr(0|1;\pi_L)\Delta D_{\pi_L}(0|1) + (1-p)Pr(0|0;\pi_L)\Delta N_{\pi_L}(0|0) \\ &= f\left(pPr(0|1;\pi_H)\Delta D_{\pi_H}(0|1) + (1-p)Pr(0|0;\pi_H)\frac{g}{f}\Delta N_{\pi_L}(0|0)\right),\end{aligned}$$

where $g = \frac{Pr(0|0;\pi_L)}{Pr(0|0;\pi_H)} < 1$ (since $\pi_H > \pi_L$).

The above expression implies that, by establishing

$$\frac{g}{f}\Delta N_{\pi_L}(0|0) > \Delta N_{\pi_H}(0|0), \tag{A1}$$

we establish $\Delta V_{\pi_L} > f\Delta V_{\pi_H}$. This is enough to establish the contradiction because we would have proven that $\Delta V_{\pi_H} \geq 0$ implies $\Delta V_{\pi_L} > 0$.

The see why the inequality in (A1) must hold, note the following. Under π_L, all non-distressed consumers under $s = 0$ receive a transfer of resources equal to fx. However, comparing to the case of high precision π_H, their mass is lower by a factor $g < 1$ (by definition of g). Hence, comparing to the transfer of resources under the high precision of information π_H, which in *per capita* terms (per non-distressed repaying borrower) is x, the *per capita* transfer of resources is actually higher in this case and equal to $\frac{f}{g}x$ (recall $f > 1$). Furthermore, since the default rate is higher under π_L, the zero profit condition implies that $R(\pi_L, P^{**}) > R(\pi_H, P^{**})$ and, from Lemma 1 below, we know that the marginal utility from an *identical* transfer is generally higher under π_L than it is under π_H. This finishes the proof that (A1) holds, as the increase in utility more than compensates the differences in masses from the ex-ante point of view. Finally, for the same reason, we note that the change in $\bar{P}(L, R)$ is higher under π_L due a larger impact on $N_{\pi_L}(0|0)$ (and no impact on $D_\pi(1|0)$). Moreover, the drop in the overall monitoring cost due to a change $\bar{P}(L, R)$ is also higher when signal is more noisy (as there are more defaulting consumers to monitor). This implies that our initial assumption of fixing $\bar{P}(L, R)$ is wlog, as relaxing it would only work to our favor.

When credit limits are binding it is straightforward to show that under Assumption 4 the increase in utility due to an increase in resources of x is the same under both π_L and π_H so the above reasoning still holds. Here note that when credit limit binds for R it binds for lower R as well. Strictly inequality can be established analogously. One simply needs to show that, since an infinitesimal deviation to a higher monitoring probability under π_H makes the consumer strictly worse off, an infinitesimal deviation to a lower monitoring probability must raise the ex-ante utility under π_L.

Finally, the fact that $P_\pi(1)$ is discontinuous at $\bar{P}$ directly follows from an argument similar to the one we used in part i) of Proposition 2. $\square$

Lemma 1. *Suppose Assumption 4 holds. Assume $R < R_{\max}$, where $R_{\max} \equiv \arg\max_R\{Rb(R)\}$. Consider an infinitesimal reduction of the interest rate R such that $\Delta(Rb(R)) = -x$. Then, the total gain in utility N under repayment of a non-distressed borrower is greater the higher the initial level of R is.*

Proof. First, assume L does not bind. By Assumption 4, $U(c, c')$ is given by $u(c + c' - \mu(c - c')^2)$, where $\mu > 0$ is a consumption smoothing parameter and u is any concave utility function (e.g. CRRA). In such a case, the policy function for borrowing is linear, and given by $b(R) = \left(B - \frac{R}{4\mu}\right)/2$. Interest rate has a distortionary effect on the intertemporal margin in this environment, causing a deadweight loss that can be explicitly calculated $\mu(c - c')^2 = \frac{R^2}{16\mu}$. This deadweight loss is higher the higher the interest rate R is. Furthermore, the total interest revenue collected from the agents is given by a quadratic function, $Rb(R) = \left(BR - \frac{R^2}{4\mu}\right)/2$. and it is concave and single peaked w.r.t. R at $R_{\max} = 2\mu B$.

Using the above properties, assuming borrowing constraints do not bind and the consumer is, in fact, a borrower, we can define:

$$\Delta N = u(2Y - B - Rb(R) - x - \frac{R'^2}{16\mu}) - u(2Y - B - Rb(R) - \frac{R^2}{16\mu}),$$

Where R' denotes the new interest rate (after the transfer x). This expression immediately follows from the assumed redistribution scheme in the lemma. Our goal is to show that $\Delta N(R)$ is increasing w.r.t. R. To see why this must be the case, note from

$$\frac{d(Rb(R))}{dR} = (B - \frac{R}{2\mu})/2$$

that the sensitivity of $Rb(R)$ is the lowest the closer the initial R is to its peak, $R_{\max}$. By the Fundamental Theorem of Calculus, we observe

$$\Delta(Rb(R)) = -\int_{R'}^{R} (B - \frac{R}{2\mu})/2dR,$$

and thus a higher R is generally associated with a larger drop of R as measured by $(R - R')$ to accomplish the assumed reduction in interest burden by x. Now, since the term $\frac{R'^2}{16\mu}$ is quadratic in the above expression of ΔN, the drop in R is highest when R is higher, and $2Y - B - Rb(R) - \frac{R^2}{16\mu}$ is lower to begin with when R is higher, it must be the case that ΔN is higher whenever $R^{**} > R^*$.

Finally, note that establishing this property in the case of binding credit limit is trivial (in a weak sense). If L binds for R it binds for $R' < R$. When L binds the interest rate distortion is absent, and the change in consumption in each sub-period is identical as the interest rate is reduced (see the budget constraint). $\square$

Proof of Proposition 6. The result follows directly from Propositions 3-5. To see why, fix an L associated with the monitored insurance region. By Proposition 3 the full monitoring interest rate associated to L is constant with respect to precision; while it is easy to check that the selective monitoring rate given by Proposition 4 is monotonically decreasing in π since $P_\pi(1)$ is decreasing in π by Proposition 5. Furthermore, the default premia of both rates are equal at $\pi = 1$ while the monitoring premium is lower under selective monitoring than under full monitoring for fixed R. Thus, since $\bar{P}$ is increasing in R, there exists a cutoff $\bar{\pi} < 1$ such that R is lower under selective monitoring for precisions higher than $\bar{\pi}$. To complete the proof, note that if a selectively monitored contract exhibits lower interest rates than a fully monitored contract with the same credit limit then consumers strictly prefer the former to the latter. This is because utility under repayment is higher at lower R and, in addition, selective monitoring involves lower monitoring probabilities than full monitoring for fixed R, implying higher utility of defaulting non-distressed agents. Finally, notice that, if there is a selectively monitored contract providing higher utility than the preferred full monitoring contract at precision π, by the above argument this must also be the case at higher precision. $\square$

A3. Data Sources and Notes

1. Data for net credit card charge-offs, revolving consumer credit (here referred to as credit card debt), and interest rate on revolving consumer credit are taken from the Federal Reserve Board Statistical Releases, G.19, historical data.
2. Median household income is taken from the US Census Bureau.
3. Credit card interest rate data on accounts assessing interest are available only from 1994 (Federal Reserve Board, G.19). Previous discontinued series reported interest rates based on a survey of credit card products, not actual rates. As an input to our regression, we used data from 1994 and subtracted the cost of funds measured by 5-ytm Treasury yields, and estimated the trend. The reported value for 1990 is the value implied by the regression.
4. Unsecured debt discharged to income per bankrupt is taken from Sullivan, Westbrook and Warren (2001), who conducted an extensive survey of formal bankruptcy filers in 1981, 1991, and 1997 (see Table 4.2, 4.3). We are not aware of any corresponding data reporting the average amount defaulted on informally, and we used this number as an imperfect proxy. Given that we lack any other data, we have used these three values to obtain a trend line and extrapolated the value for 2004 from this trend. While they are only 3 observations here, the trend line fits perfectly.

References

Athreya, Kartik, Xuan S. Tam and Eric R. Young. 2008. “A Quantitative Theory of Information and Unsecured Credit.” *American Economic Journal: Macroeconomics* 4(3):153–83.

Athreya, Kartik, Xuan S. Tam, Juan M. Sanchez and Eric R. Young. 2012. “Bankruptcy and Delinquency in a Model of Unsecured Debt.” *Federal Reserve of St. Louis Working Paper* 042B.

Avery, Robert B., Paul S. Calem and Glenn B. Canner. 2003. “An Overview of Consumer Data and Credit Reporting.” *Federal Reserve Bulletin* 2:47–73.

Benjamin, David and Xavier Mateos-Planas. 2012. "Formal and Informal Default in Consumer Credit." unpublished manuscript, Queen Mary University of London.

Berger, Allen N. 2003. "The Economic Effects of Technological Progress: Evidence from the Banking Industry." *Journal of Money, Credit, and Banking* 35(2):pp. 141–176.

Chatterjee, Satyajit. 2010. "An Equilibrium Model of the Timing of Bankruptcy Filings." 2010 Meeting Papers 1282, Society for Economic Dynamics.

Chatterjee, Satyajit, Dean Corbae, Makoto Nakajima and Jose-Victor Rios-Rull. 2007. "A Quantitative Theory of Unsecured Credit with Risk of Default." *Econometrica* 75(6):1525–1589.

Dawsey, Amanda E. and Lawrence M. Ausubel. 2004. "Informal Bankruptcy." unpublished manuscript, University of Maryland.

Drozd, Lukasz A. and Jaromir B. Nosal. 2007. "Competing for Customers: A Search Model of the Market for Unsecured Credit." unpublished manuscript, Wharton School of the Univeristy of Pennsylvania.

Evans, David S. and Richard Schmalensee. 2005. *Paying with Plastic: The Digital Revolution in Buying and Borrowing*. The MIT Press.

Fedaseyeu, Viktar. 2013. "Debt Collection Agencies and the Supply of Consumer Credit." unpublished manuscript, Bocconi University.

Herkenhoff, Kyle F. 2012. "Informal Unemployment Insurance and Labor Market Dynamics." *Federal Reserve Bank of St. Louis Working Papers* 2012-057A.

Hunt, Robert M. 2007. "Collecting Consumer Debt in America." *Business Review* Q2:11–24.

Hynes, Richard M. 2006. "Bankruptcy and State Collections: The Case of the Missing Garnishments." *Cornell Law Review* 91:603–652.

Johnson, Kathleen W. 2005. "Recent Developments in the Credit Card Market and the Financial Obligations Ratio." *Federal Reserve Bulletin* Q3.

Livshits, Igor, James MacGee and Michele Tertilt. 2006. “Consumer Bankruptcy: A Fresh Start.” *American Economic Review* 97(1):402–418.

Livshits, Igor, James MacGee and Michele Tertilt. 2010. “Accounting for the Rise in Consumer Bankruptcies.” *AEJ Macroeconomics* 2(2):165–193.

Livshits, Igor, James MacGee and Michele Tertilt. 2011. “Costly Contracts and Consumer Credit.” unpublished manuscript, University of Western Ontario.

Narajabad, Borghan N. 2012. “Information Technology and the Rise of Household Bankruptcy.” *Review of Economic Dynamics* 15:526–550.

Sanchez, Juan M. 2012. “The IT Revolution and the Unsecured Credit Market.” *Federal Reserve of St. Louis Working Paper* 022B.

Sullivan, Teresa A., Jay Westbrook and Elizabeth Warren. 2001. *The Fragile Middle Class: Americans in Debt.* Yale University Press.

Townsend, Robert. 1979. “Optimal Contracts and Competitive Markets with Costly State Verification.” *Journal of Economic Theory* 21(2):265–293.

White, Michelle J. 1998. “Why Don’t More Households File for Bankruptcy?” *Journal of Law, Economics, and Organization* 14(2):pp. 205–231.

White, Michelle J. 2007. “Bankruptcy Reform and Credit Cards.” *Journal of Economic Perspectives* 21(4):pp. 175–199.