

Information Design in Consumer Credit Markets *

Laura Blattner
Stanford University

Jacob Hartwig
Chicago Booth

Scott Nelson
Chicago Booth

November, 2022

Abstract

Over 30m US adults do not use formal consumer credit. How many of these are inefficiently excluded because they lack a credit history or have a poor credit score? We develop a framework to characterize the efficiency-maximizing system of credit histories and credit scoring, subject to the constraints imposed by the severity of adverse selection, and by the ability of credit histories to predict future risk. We find US consumer credit features a moderate amount of adverse selection and persistent consumer types. This adverse selection generates substantial welfare loss: a majority of today's non-borrowers would be first-best efficient to lend to. Credit reporting helps alleviate the costs of adverse selection, with the current US system recovering roughly two-thirds of the welfare that would be lost in a no-credit-reporting counterfactual, relative to a full-information first-best. We find that requiring histories to be shorter – or to forget past default sooner – would induce some market unraveling but also would help non-borrowing consumers escape the “no history trap.”

JEL codes: D43, D82, G28, G51, L14, L51

*Financial support from the University of Chicago Booth School of Business is gratefully acknowledged. Any errors or omissions are the responsibility of the authors. The results in this paper were calculated (or derived) based on credit data provided by TransUnion, a global information solutions company, through a relationship with the Kilts Center for Marketing at the University of Chicago Booth School of Business. Thanks to Art Middlebrooks and Heather McGuire at Kilts for their help and support. TransUnion (the data provider) has the right to review the research before dissemination to ensure it accurately describes TransUnion data, does not disclose confidential information, and does not contain material it deems to be misleading or false regarding TransUnion, TransUnion's partners, affiliates or customer base, or the consumer lending industry.

1 Introduction

Like people, societies have memories. Credit registries and criminal records are leading examples. Perhaps unlike people, societies also choose *what* to remember and for how long.

In this paper, we examine what a credit registry should remember, for how long, and why. We build a quantitative model to show how the optimal design of a credit registry depends on economic primitives, including consumers' risk exposure and the information asymmetries that exist in the absence of credit records. Estimating our model on the US consumer credit market, we characterize in equilibrium the distributional and efficiency effects of several potential redesigns of US consumer credit reporting.

The development of credit registries has long been associated with both greater credit availability and economic growth (Jappelli and Pagano, 2002; Djankov et al., 2007; Brown et al., 2009). However, there is wide diversity in how different countries remember credit histories (Djankov et al., 2007; Jappelli et al., 2000). Sometimes loan repayment is remembered longer than default, and sometimes repayment is not remembered at all; some countries' credit histories last only months, and others' last ten or more years. There is also little guidance from economics as to which of these systems is optimal for whom. Perhaps reflecting this lack of guidance, the US requirement that credit records expunge loan defaults after seven years appears Biblical, as in, "At the end of every seven years you shall grant a release of debts...".

To study the equilibrium consequences of different credit reporting systems, we build a quantitative model that clarifies how credit memory interacts with several primitives of an economy. Consumers in the economy have private information about their demand for credit and their likelihood of repayment. These demand and risk characteristics – what we refer to as two dimensions of consumers' types – change over time as shocks arrive. The model shows how the effects of credit reporting depend on the extent of information asymmetries about these types, the correlation between demand and risk, and how persistent these types are relative to the shocks that consumers face.

We show using calibrated examples that, depending on the primitives of an economy, different systems of credit reporting can be welfare enhancing, can have little effect, and in fact can be welfare reducing, relative to an environment with no credit reporting. Moreover, the type of credit memory that maximizes welfare in one economy may not be optimal in another.

Turning to the US consumer credit market, we estimate the primitives of our model on a large panel of US consumer credit report data, leveraging both the long panel dimension and quasi-experimental variation in credit histories generated from the US' requirements of expunging (most) loan defaults from a consumer's record after seven years. We find that the US features a moderate degree of adverse selection and persistent consumer types – two features of an economy that make credit reporting particularly welfare-enhancing. Indeed, we estimate that the current

US credit reporting system recovers roughly two-thirds of the welfare that would be lost in a no-credit-reporting counterfactual, relative to a first-best full-information counterfactual. Nevertheless, there are also potential gains from redesigns of US credit reporting. While our current baseline results are suggestive of these gains, in preliminary work we estimate that making credit history in the US more coarse would generate substantial gains for consumers with no credit history or consumers with long-ago defaults, with only modest overall efficiency costs.

The stakes are high in credit history design. Over 30m adults, conservatively estimated, do not currently participate in US consumer credit markets at the prices available to them. Our model suggests that the vast majority of these would be efficient to lend to, but that their pricing exceeds both the cost of lending to them and their willingness to pay for credit, because their credit histories pool them with other, higher-risk types. We also find, in estimates from the period after the Great Financial Crisis, that the number of inefficiently *included* consumers is modest: few consumers are accessing credit at a price that is “too low” from an efficiency perspective.

To be clear about the focus of our paper, our primary emphasis is on credit histories and credit reporting, as opposed to credit scoring. The former is the collection and regularization of borrower data for use in underwriting; the latter is the work of translating these data into predicted default risk (credit scores). Our model does have a notion of credit scores that are closely tied to equilibrium prices at each credit history, as we discuss in Section 2.4, because the predicted default rates that correspond to credit scores in our model are monotone in equilibrium prices.

Prior work on the design of credit reporting systems has largely come in two veins: one vein is primarily theoretical, for example studying the implications of coarse information (Bhaskar and Thomas, 2019), finite memory (Elul and Gottardi, 2015; Kovbasyuk and Spagnolo, 2018), or manipulable signals (Ball, 2021) in equilibria with credit histories;¹ a second vein is primarily empirical, studying in reduced-form or in partial-equilibrium the responses of consumers or lenders to changes in credit market information (Musto, 2004; Bos and Nakamura, 2014; Bos et al., 2018; Liberman et al., 2019; DeFusco et al., 2021; Jansen et al., 2022). In part our aim is to bridge these two parts of the literature, developing a quantitative equilibrium model that can be estimated using variation similar to what is highlighted in some of this prior empirical work.²

We also relate to structural and reduced-form work that either quantifies or studies the consequences of asymmetric information in credit and insurance markets (see Einav et al. (2021)

¹More broadly, theoretical work related to credit scoring includes the literature on general problems of information design and Bayesian persuasion (Bergemann and Morris, 2019; Kamenica, 2019; Kamenica and Gentzkow, 2011), a long line of papers characterizing the role of information in Akerlof economies (Levin, 2001; Garcia and Tsur, 2021; Bar-Isaac et al., 2021) or in insurance markets (Garcia and Tsur, 2021). Pagano and Jappelli (1993) also study endogenous information sharing among lenders in an adverse-selection setting.

²This bridging follows in some extent after Einav et al. (2012) and Einav et al. (2013b), though our focus on the design of credit reporting systems differs from theirs. Similar structural work has also studied the design of risk scores in the health insurance context (Bundorf et al., 2012; Handel, 2013; Vatter, 2021; Vatter and Marone, 2021).

for a review). One closely related paper in this literature is [Crawford et al. \(2018\)](#), who quantify the roles of both adverse selection and imperfect competition in determining the supply of small-business credit; while our focus is instead on a perfectly competitive model, our modeling of asymmetric information in many ways builds on theirs while adding dynamics in consumer types. These modeling approaches also trace back to the insights of [Chiappori and Salanie \(2000\)](#) and [Einav et al. \(2010\)](#), and to the empirical approach of [Cohen and Einav \(2007\)](#). In recent work, [Agarwal et al. \(2020\)](#) study a structural adverse-selection model that also incorporates loan search. In the consumer credit card market, [Ausubel \(1999\)](#) and [Agarwal et al. \(2010\)](#) use randomized-controlled trials to find evidence of adverse selection on price.³

Our paper also complements recent developments in the quantitative bankruptcy literature following the influential papers by [Chatterjee et al. \(2007\)](#) and [Livshits et al. \(2007\)](#). Whereas much of the earlier papers in this literature focused on unsecured consumer lending with perfect information where consumer’s risk is priced on a granular level, [Chatterjee et al. \(2020\)](#) and [Exler et al. \(2020\)](#) allow for type pooling and adverse selection, a central feature of unsecured credit markets. In contrast with these papers, abstracting from the consumption-savings problem and default as a choice which individuals optimize over, allows us to study the persistence of and correlation between consumers’ credit risk and demand as key primitives of consumer credit markets, and their interplay with different information structures. Other papers in this literature with either explicit or implicit notions of a credit score include [Athreya et al. \(2012\)](#), [Drozd and Serrano-Padial \(2017\)](#), and [Kovrijnykh et al. \(2019\)](#).

There also is of course a rich literature on optimal price discrimination and its welfare effects ([Varian, 1985](#); [Schmalensee, 1981](#); [Aguirre et al., 2010](#); [Bergemann et al., 2015](#); [Bonatti and Cisternas, 2020](#); [Dube and Misra, 2022](#)). While a key difference between our setting and many studies of price discrimination is the role of adverse selection, this literature also anticipates the important role of equilibrium consequences to any changes in information structure. For example, [Dube and Misra \(2022\)](#) note, “This empirical finding that consumer surplus is non-monotonic in the degree of consumer information suggests that any regulation of consumer data might need to consider carefully the welfare implications caused by downstream decisions based on such data.”

2 Model

This model aims to study the selection of consumers into and out of the consumer credit market in order to analyze alternative information structures. Based on the idea that the absence of collateral makes unsecured credit most sensitive to consumer credit information, the model focuses on the

³More broadly, this paper also contributes to the literature using tools from industrial organization to study consumer financial markets ([Allen et al., 2014, 2019](#); [Benetton, 2021](#); [Bhattacharya et al., 2019](#); [Buchak et al., 2020](#); [Clark et al., 2021](#); [Cuesta and Sepulveda, 2021](#); [Egan et al., 2017](#); [Robles-Garcia, 2020](#); [Galenianos and Gavazza, 2019a](#); [Jiang, 2020](#)).

pricing and pooling of consumers across different segments of the unsecured borrowing market. We describe our choice to focus on the unsecured borrowing margin in more detail in section 4.

2.1 Consumer Types and Histories

There is a continuum of consumers with mass 1 and time is discrete. Consumers are differentiated by their demand for borrowing and their riskiness of defaulting. The demand and riskiness characteristics of consumers is captured in their joint type $\theta = (\theta_d, \theta_r) \in \Theta = \mathcal{D} \times \mathcal{R} \in \mathbb{R}^2$.

Consumer type transition: Consumers' types change over time. The type transition probability $q(\theta'|\theta)$ is Markov, as is standard in the dynamic discrete choice literature. The distribution of consumers over types is denoted as $q(\theta)$.

In particular, types evolve according to a bivariate normal AR(1) process with persistence λ and potentially correlated innovation vector $u = (u_d, u_r)$:

$$\begin{bmatrix} \theta'_d \\ \theta'_r \end{bmatrix} = \begin{bmatrix} \lambda_d & 0 \\ 0 & \lambda_r \end{bmatrix} \begin{bmatrix} \theta_d \\ \theta_r \end{bmatrix} + u, \quad u \sim N \left(\begin{bmatrix} \mu_d \\ \mu_r \end{bmatrix}, \begin{bmatrix} \sigma_d^2 & \rho\sigma_d\sigma_r \\ \rho\sigma_d\sigma_r & \sigma_r^2 \end{bmatrix} \right) \quad (1)$$

As is standard, we discretize this process in our empirical work. The number of discretized demand and risk types is denoted as $N_d = |\mathcal{D}|$ and $N_r = |\mathcal{R}|$, and we denote the transition matrix across discrete types implied by the Markov probabilities $q(\theta'|\theta)$ as M_θ . Instead of relying on a Cholesky transformation of separate univariate processes to induce the shock correlation as commonly described in research on consumer income processes, we directly discretize the joint process using a bivariate version of the [Tauchen \(1986\)](#) method to discretize AR(1) type process ([Terry and Knotek, 2011](#)). This results not only in a more efficient discretization but also ensures that the shock correlation is reflected in the transition probabilities and not just in the type values.

We use $N_d = 5$ demand types and $N_r = 7$ risk types in our discretization. Instead of the equally-spaced grid points used as a default in the Tauchen method, we use equiprobable grid points that result in a more efficient discretization of the bivariate process; in particular, we are able to use a smaller number of discrete types when using an equiprobable discretization without introducing a “mechanical” pooling of consumers at the same discretized type.⁴

Consumer histories: In each period, consumers can choose to borrow b_t in the unsecured credit market. Based on their borrowing choice and their risk type, consumers can become delinquent or default d_t . At any point in time this gives rise to a borrowing and default outcome (b_t, d_t) and across time a history for each consumer. This consumer history $h \in H$ is assumed to be observable

⁴For intuition, consider the extreme case of only one type ($N_d = 1$ and $N_r = 1$); such a discretized model features full pooling of all (continuously distributed) consumer types because this continuous distribution is approximated with only one type. In unreported results, we find that only 5 discrete types are needed to generate similar equilibria as we observe with any higher number of discrete types, when using the equiprobable discretization, while closer to 10 discrete types are needed when using the equally-spaced distribution, because the equally-spaced discretization has a densely-populated middle type and very thinly-populated tail types.

for N_h time periods. In general, the space of all histories H given N_h is then given as the product space $H_N = \times_{t=\{1,\dots,N_h\}}(b_t, d_t)$.

2.2 Consumer borrowing choices

Consumers with a history h face a price $p(h)$ and derive a flow utility $u(b, \theta_d)$ from borrowing b that is dependent on the consumer type θ_d . Furthermore, each consumer i is subject to taste shocks $\epsilon_{i,b}$ across borrowing options b . Consistent with what is recorded in typical US credit histories, we focus on the extensive margin of borrowing and specify b_t as either 0 and 1. Given that consumer demand is taken as a primitive of the consumer, we assume that consumers' demand type is simply expressed as the logistic-transformed flow utility the consumers derive from borrowing

$$u(b, \theta_d) = \begin{cases} 0 & \text{if } b_t = 0 \\ \frac{\exp(\theta_d)}{1 + \exp(\theta_d)} & \text{if } b_t = 1 \end{cases}$$

where we normalize the utility from not borrowing to zero. The borrowing choice of a consumer i with history h and type θ conditional on price $p(h)$ and having received a vector of tastes shock $\epsilon_i = \{\epsilon_{i,b}\}$ is denoted as $b(\theta, h|p(h), \epsilon_i)$. Following the [Rust \(1994\)](#) and [Aguirregabiria and Mira \(2010\)](#), we model the consumer's intertemporal trade-off as a dynamic discrete choice problem, so that borrowing b solves,

$$\max_{b \in B} u(b, \theta_d) - p(h) \times b + \beta E[V(h', \theta')|\theta, h, b] + \epsilon_{\theta, h, b} = \max_{b \in B} v(\theta, h, b) + \epsilon_{i, b} = V(\theta, h) \quad (2)$$

where $V(\theta, h)$ is the beginning of period value function, β is the consumer's discount factor (here calibrated to 0.99), and $v(\theta, h, b)$ is the choice-specific value function. That is to say consumers solve an intertemporal optimization problem in which they internalize how their borrowing choice affects their transition into future history states and the resulting future prices they will face.

Given this intertemporal optimisation problem of the consumer and the assumption that taste shocks are drawn from a Gumbel distribution, the choice probabilities of a consumer of type θ with history h facing price $p(h)$ are then given in the familiar logit form:

$$P(b = k|p(h), \theta) = \frac{\exp(v(\theta, h, b = k))}{\sum_{j \in \{0,1\}} \exp(v(\theta, h, b = j))}, \quad \forall k \in \{0, 1\} \quad (3)$$

2.3 Consumer default

We assume there is no strategic default behavior and the probability of default is governed by consumers' risk type θ_r and their borrowing choice b_t . This modeling choice rules out a specific notion of moral hazard in which borrowing terms (e.g., prices) affect the probability default even

after conditioning on risk type.⁵ Instead, there is an alternative notion of moral hazard that affects borrowing probabilities through the reputational costs of defaulting that we capture here. We believe this is closest to the empirical evidence on moral hazard in consumer lending markets. We discuss the role of moral hazard in our model and its grounding in the empirical literature in more detail in section 3.1.

Although we focus on the borrowing decision of unsecured credit, we model explicitly consumer default on both secured \tilde{d}_t^s and unsecured credit \tilde{d}_t^u . Following the industry practice of distinguishing mild delinquencies from severe delinquencies and their importance for lenders inference about consumers' riskiness, we postulate a fairly detailed default process accounting for the severity of default on both secured and unsecured borrowing. More specifically, we define a consumer's default realization $\{\tilde{d}_t^k\}_{k \in \{s,u\}}$ to be equal to 2 if a consumer has any severe delinquency (i.e., 90 or more days past due), equal to 1 if the consumer has no severe but at least one mild delinquency (i.e., less than 90 days past due), and zero otherwise. Finally, in line with typical credit bureau variables, the default realization recorded in a consumer's history d_t captures the most severe delinquency by aggregating delinquencies across secured \tilde{d}_t^s and unsecured borrowing \tilde{d}_t^u

$$d_t = \max(\tilde{d}_t^s, \tilde{d}_t^u)$$

where $(\tilde{d}_t^s, \tilde{d}_t^u)$ are themselves drawn from non-parametric distributions specific to each risk type θ_r . Further details are provided in the appendix B.1.

2.4 The lender's problem

Lenders provide loans to consumers at cost and observe the (borrowing and default) history h of each consumer without observing the consumer type θ . Thus, prices are equal to the expected cost and are implicitly defined by

$$0 = E_\theta[b(\theta, h|p(h))(p(h) - c(\theta_r))|h], \quad \forall h \tag{4}$$

where $E_\theta[\cdot|h]$ is the expectation over types θ for a given history h and $b(\theta, h|p(h))$ is the optimal consumer borrowing choice given the price vector $p(h)$ while $c(\theta_r)$ is cost to the lender generated by a consumer of risk type θ_r . We assume that the cost $c(\theta_r)$ to the lender generated by a consumer

⁵Without an intensive borrowing margin, we naturally can't capture notions of moral hazard that postulate default increases with borrowing amounts. In the context of lending markets, the moral hazard notion of default increasing with prices is particularly hard to identify separately from adverse selection. Whereas in insurance markets risk exists outside the insurance contract, in lending markets the possibility of a risk realization is created by the lending contract itself. Thus, there is a notion of moral hazard in insurance markets that refers to the causal effect of the (insurance) contract on the risk probability that does not exist in the context of lending markets (Einav et al., 2021). Moreover, given a set of fixed contract terms, a correlation between the probability to accept a contract and the risk probability might be generated by either moral hazard or by adverse selection in the insurance context whereas it can only be generated by adverse selection in the lending market context.

of risk type θ_r is equal to the probability of him becoming severely delinquent on his unsecured borrowing $\omega^{s,u}(\theta_r)$, as defined in Sections 2.3 and B.1. There is no cross subsidization across consumers with different histories.

The prices that satisfy equation 4 also provide a notion of credit *scores*, as opposed to credit histories, in our model. These prices are monotone in expected default rates among consumers who choose to borrow at each history, as commercial credit scores are.⁶

2.5 Timing

The precise timing of the model is as follows. Within each period, first consumer types are realized, following the consumer type process in equation 1. Lenders then post zero-profit prices based on consumers' credit histories, following equation 4. Consumers make borrowing choices, as in equation 3, defaults are realized, following the process in section 2.3, and finally histories are updated using the new borrowing and default realizations for each consumer. This timing is illustrated in Figure 1.

2.6 Equilibrium

An equilibrium in this consumer credit market is defined by the triple $(b^*(\theta, h), q^*(\theta, h), p^*(h))$ satisfying the following three conditions:

1. **Individually-optimal borrowing choices:** Given prices $p^*(h)$, the borrowing choices $b^*(\theta, h)$ solve the intertemporal consumer borrowing problem for each type and history:

$$b^*(\theta, h) = \arg \max_{b \in B} u(b, \theta_d) - p^*(h) \times b + \beta E[V(h', \theta') | \theta, h, b], \quad \forall (h, \theta) \in H \times \Theta \quad (5)$$

2. **Stationary history-type distribution:** Given consumers' borrowing choices $b^*(\theta, h)$, the distribution of consumers over histories and types $q^*(\theta, h)$ is the stationary distribution associated with borrowing choices $b^*(\theta, h)$.
3. **Zero-profit pricing:** Given the borrowing choices $b^*(\theta, h)$ and the stationary history-type distribution $q^*(\theta, h)$, prices $p^*(h)$ are such that

$$0 = p^*(h) - \int b(\theta, h | p^*(h)) c(\theta_r) dQ(\theta | h), \quad \forall h \quad (6)$$

where $q(\theta | h) = \frac{q(\theta, h)}{q(h)}$ is the distribution over types for a given history h and $b(\theta, h | p(h))$ is the optimal consumer borrowing choice given the price vector $p(h)$.

⁶Precisely speaking, commercial credit scores are typically an affine transformation of the predicted log odds of default given any credit history, conditional on consumers having chosen to borrow at then-prevailing prices in the credit scorer's training data (Thomas, 2009).

In appendix A, we show using Kakutani’s theorem that an equilibrium exists under fairly general conditions.

3 Key Forces in the Model

We illustrate some of the key economic forces captured by the model. Two of these map closely onto parameters in the consumer type process in equation 1: adverse selection, which is governed by the parameter ρ that governs the correlation between demand and risk across consumers (Chiappori and Salanie, 2000; Einav et al., 2010; Mahoney and Weyl, 2017; Crawford et al., 2018); and the informativeness of credit history, which is governed by the autoregressive coefficients λ .

To show how these two forces interact to shape the equilibrium effects of credit reporting, we introduce a framework of four “welfare groups” that we will refer to throughout the paper. These four groups are essentially a two-by-two view of consumer types in equilibrium: whether a consumer participates or not in the credit market, and whether or not it is efficient for them to participate. We refer to these four groups as efficient inclusions, efficient exclusions, inefficient inclusions, and inefficient exclusions.

These four groups can also be plotted in four quadrants of the consumer type space. Panel (a) of Figure 2 illustrates these quadrants in the case of fully pooled pricing. Normalizing demand (on the y-axis) and cost (on the x-axis) such that efficient trades are those above the 45-degree line, we can then draw a horizontal price line (normalized in demand units) such that consumers borrow if and only if their demand is above price. These two lines – the horizontal price line and the 45-degree line – then divide consumer types into four welfare groups.

Panel (b) of the same figure adapts this framework for heterogeneous pricing based on credit histories h . We subtract prices $p(h)$ from both axes, such that consumers are translated towards the origin (and parallel to the 45-degree line) by a distance proportional to the price they face to borrow. As shown in the figure, a single horizontal line (at $y = 0$) then separates borrowers from non-borrowers in this price-translated demand-cost space.⁷

Different credit reporting regimes induce different price translations for consumers across this demand-cost space. An efficient credit reporting system will have two features: it will generate high enough prices for consumers whom it is inefficient to lend to (i.e., those for whom cost exceeds demand) in order to induce those consumers to not borrow, or equivalently, to translate them into the efficient-exclusion group; and it will generate low enough prices for consumers whom it is efficient to lend to so as to allow them to borrow, or equivalently, to avoid translating them into the inefficient-exclusion group. Because prices are equal to average costs among partic-

⁷A similar yet different graphical analysis appears in Bundorf et al. (2012). In that analysis, medical risk scores (analogous to credit scores in our context) are plotted on the x-axis, and both demand and risk are then projected on these scores. Our approach of instead using price-translated (i.e., score-translated) demand-cost axes allows us to conveniently study different credit reporting regimes on the same axes.

ipants in each history, this requires a credit reporting system to endogenously generate pools of consumers at each history such that average-cost pricing does not price out consumers who would borrow efficiently.

It is intuitive that adverse selection (as captured by ρ) and the informativeness of credit histories (as captured by λ) help determine the efficiency effects of credit reporting. First, note that the greater is ρ , the more that average cost and hence prices will be “pulled upward” in the demand-cost space of Figure 2, panel (b): this can be seen by noting that price must be equal to the average cost of consumers who fall above the price line, and as ρ increases, consumers above the price line will be shifted rightwards towards higher costs. Second, note that as λ increases, types will be more persistent and hence credit histories will be more informative about future default. The less informative are histories, the less feasible it is for credit reporting to induce a different allocation than under pooled pricing.

We formalize this intuition by illustrating what share of consumers fall into each of the four welfare groups in equilibrium, in four calibrated scenarios in the model corresponding to high and low values of correlation ρ , and to high and low values of persistence λ . These calibrated scenarios are plotted in the four panels of Figure 3. Within each panel, we show the share of consumers who fall in each welfare group in three different information environments: (A) a no-information environment with no credit histories and full pooling (homogeneous pricing), (B) an environment with credit reporting, and (C) a full-information environment.⁸ As the figure shows, the share of consumers who are inefficiently included or excluded is relatively low when demand-risk correlation (ρ) is low, as is the case in panels (b) and (d), regardless of the information environment. This illustrates how the welfare effects of credit reporting are modest when adverse selection, as captured by ρ , is small. In contrast, when ρ is higher, the share of consumers who are inefficiently included or inefficiently excluded in the no-information environment is higher. This is shown in environment (A) of panels (a) and (c). In these scenarios with relatively high adverse selection, there is a potential role for credit reporting to generate a more efficient allocation. However, as a comparison of panels (a) and (c) shows, credit reporting only does so when type persistence is also high.

Figure 4d shows total surplus, rather than the share of consumers in each welfare group, for these same four scenarios. As before, each panel shows one of the four $\rho - \lambda$ calibrated scenarios, and the three information environments of (A) no information, (B) credit reporting, and (C) full information are shown within each panel. Patterns in total surplus are consistent with those seen in welfare group shares: there is considerable welfare loss in the no-information environment from adverse selection when ρ is high, and credit scoring significantly reduces this welfare loss only when type persistence λ is also high.

⁸For details on how we solve for equilibria in each of these environments, see Section 8.

3.1 Reputational Incentives and Moral Hazard

The prior subsection emphasized the role of credit reporting in addressing adverse selection. Next, we discuss how credit reporting interacts with moral hazard, both in our model and in view of external evidence.

Our model departs somewhat from a longstanding literature that emphasizes credit reporting's role in preventing moral hazard, in the sense of disciplining borrowers to repay their loans. For example, [Brown and Zehnder \(2007\)](#) argue, "In credit markets dominated by short term interactions (e.g., consumer credit markets when borrower mobility is high), borrowers may only be motivated to repay if they know that, due to credit reporting, their current behavior is observable by other lenders." This argument of course builds on a long line of research showing how the threat of future exclusion from credit markets can be important in disciplining borrowers to repay (e.g., [Kehoe and Levine \(1993\)](#); [Boot and Thakor \(1994\)](#); [Bond and Krishnamurthy \(2004\)](#); [Brown et al. \(2004\)](#)).

These well-established arguments are perhaps in tension with institutional details in the US *consumer* credit market. Consumers who default, even those with especially severe defaults such as bankruptcy filings, are not excluded from future borrowing ([Han and Li, 2011](#)), and the effects of default on future loan pricing appear modest. Moreover, credit scores' ability to penalize borrowers for default operate only through *future* costs, whereas evidence on the determinants of consumer default decisions suggests that consumers are mostly sensitive to current liquidity, rather than the long-term costs and benefits of repayment or default, when choosing whether to repay ([Ganong and Noel, 2020a,b](#); [Indarte, 2021](#)). These findings suggest that the forces that discipline borrowers to repay are more likely to operate at short time horizons than credit reporting's incentive effects do. These near-term costs also appear to be substantial and numerous: the US has relatively aggressive debt collection practices that do not depend on the credit reporting industry for information about defaulted debts ([Hunt, 2007](#)); the social stigma of default is seen as substantial ([Gross and Souleles, 2002](#)); and defaulting on a loan also entails considerable fixed costs ([Livshits et al., 2010](#)).

Overall, the importance of these short-run costs support the notion that, especially in our setting where we study counterfactuals such as the shortening or lengthening of credit report memory, it is sensible to model default as arising stochastically, and not dependently on the particular credit reporting regime.

A separate form of moral hazard is whether changes in loan pricing affect default. While *some* level of price increase certainly would make any consumer unable to repay their loan, evidence in [Nelson \(2022\)](#) shows that for the types of credit products we focus on, principally unsecured loans such as credit cards, the extent of price variation typically seen in the market has negligible effects on default. Similar evidence supporting this price-invariance of default on credit cards is

shown in [Castellanos et al. \(2018\)](#).

Finally, it is worth noting there is a third notion of moral hazard that our model does capture directly: the possibility that consumers with high default risk borrow more in the first place when credit reporting has less of an effect on their access to low future pricing after defaults. We quantify the strength of these reputational incentives in [Section 8](#) and, consistent with the arguments above, find these long-run incentive effects are modest.

4 Data and Target Moments

We study the US consumer credit market using a large credit bureau dataset. This consumer credit panel (CCP) is a longitudinal, nationally representative panel of de-identified credit records from TransUnion. These data are derived from a 10% random sample of all people in TransUnion’s own database as of July 2000. Each subsequent month of data takes this original sample and (i) adds a 10% random sample of people that newly-appeared in that month (ii) removes deceased people. Each month of the CCP data contains for each sampled consumer the full credit record, including his credit score, credit inquiries, public records (e.g. tax liens, civil judgments and bankruptcies) and collections, as well as credit limits, balances, payments and detailed delinquency states for each tradeline.

These credit bureau records capture each consumer’s borrowing and default history which form the input to the commonly used proprietary risk-scores, like FICO or Vantage score. Credit scores are ubiquitous in the US consumer credit market and are typically used for underwriting a broad array of credit products from mortgages and car loans to credit cards and installment loans. As such, credit scoring models are built to predict future credit performance (likelihood that a consumer will be more 90 or more days past due on any of his credit payment obligations over the next 24 months). Thus, to estimate our structural model we map these rich credit report data into lower-dimensional, discretized histories captured in our model.

An important part of this exercise consists in approximating the consumer pooling effectively implemented in the US consumer credit market through consumers’ credit scores while maintaining the tractability in the estimation and counterfactual analysis of our structural model. Our empirical and modeling approach is based on two observations in the data: First, participation rates in secured credit products are high throughout the consumer risk distribution. For instance, [Appendix Figure 1](#) shows how at any point in time most individuals in the US hold a auto loan, have some student debt, or pay down a mortgage. Indeed, 85% of new-car purchases are typically financed by an auto loan ([Cross et al., 2019](#)). Moreover, loan pricing on secured credit is generally less dependent on credit scores than is the pricing of unsecured credit (e.g., [Argyle et al. \(2020\)](#); [Nelson \(2022\)](#); [Heilbron \(2022\)](#)). We believe secured borrowing can therefore be understood to provide signals about the riskiness of consumers while the secured borrowing decision itself is

largely independent of consumers credit score and credit history. This motivates our focus on the unsecured borrowing margin, as further defined below, while treating secured borrowing as a continuous stream of signals about the riskiness of consumers.

Second, the overall segmentation of consumers by credit scores is driven by consumers' borrowing and default behavior over periods of several years, while credit utilization rates affect credit scores only in the very short term. This insight drives our decision to summarize borrowing and default behavior in two multi-year periods at which past intensive margin utilization decisions are inconsequential for consumers credit score. This lines up well with our modeling choice to center our attention at the extensive borrowing margin.

We only treat credit card debt and consumer installment loans as unsecured borrowing. Moreover, focusing only on the extensive margin of unsecured borrowing allows us to both sidestep the complexities of the non-linear pricing problem of lenders, that need to set both credit limits and interest rates, and reflect well the actual industry practices in which lenders can only price the extensive borrowing margin without setting interest rates conditional on borrowing levels.⁹

To summarize the modeling choices we make in the face of the empirical insights just discussed, we synthesize the dynamics in consumer credit history at multi-year frequencies by treating secured borrowing as a signal stream of consumers' riskiness while only modeling the extensive margin of unsecured borrowing. In doing so, we stay close to the institutional details in capturing the most important information sources of consumers' risk used to price consumers while analyzing the borrowing margin most sensitive to lenders' information set without needing to disentangle any confounding roles of collateral present in secured borrowing.

Given these modeling choices it is straightforward to construct consumer borrowing and default histories in the data that map tightly into the model. In line with credit records reflecting at most seven years of consumers' past, consumers' credit histories in our model consist of the past two periods of three and half years each. That is to say, a period in the model is equal to three and a half years in the data. At any point, a consumer is said to be borrowing B unsecured if his combined credit card balances and other unsecured debt exceeds a low minimum threshold ($=1$) or has no unsecured borrowing otherwise ($=0$).¹⁰ Similarly, a consumer's default state D is based on her most severe delinquency status over the past three and half years, where we refer to zero delinquencies as $D = 0$, encode any mild delinquency without a severe delinquency as $D = 1$ and define $D = 2$ if the consumer has any severe default.

A consumer history is consequently of the form $[B_1, B_2 | D_1, D_2]$ where the most recent period borrowing B and default D of the consumer is denoted with the subscript 1 and the period two periods ago with the subscript 2. For instance, a consumer with credit history $[1, 0 | 1, 2]$ was bor-

⁹The analysis of the two margins (price and borrowing limits) involved in credit card pricing is challenging as reflected in the fact that the literature has not converged on a generally accepted modeling approach.

¹⁰We experiment with different thresholds but found that reasonable variation had little impact on the constructed model moments and thus estimates. Our results reported below are based on a minimum balance threshold of \$750.

rowing unsecured ($B_1 = 1$) and had some mild delinquency ($D_1 = 1$) in the most recent period but was not borrowing unsecured ($B_2 = 0$) while nevertheless having realized a severe delinquency ($D_2 = 2$) on his secured borrowing. Implicit in this history representation, is the aggregation of the consumer’s delinquencies across both secured and unsecured borrowing products. As already discussed in subsection 2.3, this is in line with the coarseness of typical credit bureau variables.

Having constructed borrowing and default histories for each consumer in the data such that they have a direct counterpart in the model, makes it straightforward to estimate conditional choice probabilities based on our consumer credit panel data. Conditional on each history we calculate the probability that a consumer with that history chooses to borrow and becomes mild or severely delinquent on his secured and unsecured borrowing. These conditional choice probabilities that we take from the data are shown in the left panel in Figure 7, and these play a key role as target moments in our estimation.

5 Estimation

We estimate the model using a full solution method for dynamic discrete choice models as originally developed by Rust (1987): that is, for each candidate parameter vector we solve for the full equilibrium of the model. The full solution method allows to estimate our model in spite of its rich unobserved Markov heterogeneity (where the unobserved demand and risk types evolve over time) and our embedding of this dynamic discrete choice model within a competitive lending market in which banks need to solve for optimal prices in the face of adversely selected borrower pools. As such, we are less concerned about the sensitivity of parameter estimates to measurement error in the choice probabilities, a more common concern in conditional choice probability (CCP) estimators in the tradition of Hotz and Miller (1993) and their recent extensions to dynamic games (Aguirregabiria and Mira, 2007) and unobserved heterogeneity (Arcidiacono and Miller, 2011).¹¹

We follow much of this literature in assuming that the data have been generated by only one (Markov perfect) equilibrium. So, even though our model potentially has multiple equilibria, we do not need to specify an equilibrium selection mechanism to estimate the model because the correct equilibrium will be identified from the conditional choice probabilities in the data.

The estimation algorithm is described in more detail in appendix section B.3 and we outline the estimation approach here only in broad terms. For a given parameter vector, we find the price vector that satisfies lenders’ optimality condition for each consumer history given how consumers sort into borrowing and default histories based on the prices they face. We search globally across the parameter space to minimize the distance between empirical choice probabilities and

¹¹For a review of dynamic discrete choice models see Keane and Wolpin (2009) and Aguirregabiria and Mira (2010) as well as Arcidiacono and Ellickson (2011) specifically for conditional choice probability estimators and Aguirregabiria (2021) for dynamic games.

the model implied choice probabilities. Thus, our estimation approach is a simulated method of moments estimator in the spirit of [McFadden \(1989\)](#). Computationally, this full solution approach is made feasible by exploiting the Markov structure of the unobserved types to invert the corresponding transition matrices and obtain consumers' stationary distribution over histories.

6 Identification

We show how both cross-sectional and panel variation in the data identify the model parameters. The identification strategy follows [Cohen and Einav \(2007\)](#) in identifying the joint distribution of unobserved demand and risk in any cross-section of the data, and then uses the panel dimension of the data to ask how consumers' unobserved types evolve over time.

We start with cross-sectional identification. [Cohen and Einav \(2007\)](#) show how price variation can non-parametrically identify the distribution of perfectly persistent two-dimensional unobserved types. In essence, indifference curves in two-dimensional type space between choosing two contracts – in our setting, between choosing to borrow and not – both shift and change their shape when prices change, allowing a researcher to trace out the marginal distribution of types by observing how both demand and realized risk respond to price changes.

We use two types of price variation for such cross-sectional identification. First, we build on the growing literature that studies the scheduled removal of negative information from consumer credit reports. Among other results, this literature confirms that credit supply shifts outward for a consumer who has negative information removed from a credit report ([Musto, 2004](#); [Bos and Nakamura, 2014](#); [Lieberman et al., 2019](#); [Gross et al., 2020](#); [Dobbie et al., 2020](#); [Jansen et al., 2022](#)). We use these supply shifts to identify both the credit demand and default risk of marginal borrowers, similar to how insurance premium changes in [Cohen and Einav \(2007\)](#) trace out the insurance demand and claim risk of marginal insurees. Second, conditional on any set of model parameters, the stochasticity of default also introduces price shocks for otherwise equivalent consumers: some consumers with the same unobserved type, same history, and same borrowing choice may default or not by chance in a given period, and defaulters then face higher prices than non-defaulters in subsequent periods. While [Cohen and Einav \(2007\)](#) show how the type distribution is in principle non-parametrically identified off of price variation, we use an approach similar to theirs in parameterizing a bivariate-normal distribution of hidden types to be identified by such variation.

Turning to the panel dimension, the second part of our identification strategy seeks to learn the persistence of consumer types over time. That is, for any given cross-sectional distribution of types, how quickly do consumers migrate across the type space? Are consumers' unobserved types perfectly persistent, or i.i.d. draws across time with zero persistence, or something in between?

To recover type persistence from the data, we use the panel dimension to ask how similar are consumer choices over time when facing similar histories and pricing. While all available panel

variation is in principle useful for this identification, some examples are particularly illustrative. For example, a consumer with credit history α might face the same reputational incentives and similar pricing as a consumer with credit history β , and history α might update into history β after one period for a given outcome in the current period. Comparing the choice probabilities of a consumer in history α , relative to the choice behavior of an equivalent consumer one period later in history β , reveals how persistent are types over time.¹²

Figure 6 illustrates model identification in the context of one of the key parameters to be estimated, the correlation ρ between demand and risk. Holding other parameter values constant, the figure shows how changes in ρ induce a rotation in the cloud of demand probabilities and default probabilities across all model histories, plotting each history as one point in this choice-probability space. As can be seen, higher values of ρ imply a higher correlation between borrowing probabilities and realized default probabilities, and, as shown by comparing the first and the middle panels of the figure, the model successfully recovers a ρ estimate that matches the correlation observed in the data.¹³

To illustrate identification in practice, Figure 5 shows convergence to a unique parameter vector estimate, across 50 different randomly chosen parameter starting values that are used as seeds for the outer-loop parameter search described in Section 5. The 50 different model runs are sorted by the value of the objective function, as shown in the lower-right panel of the figure, and each other panel in the figure then corresponds to a single model parameter. The red crosses in each panel show the randomly chosen starting values, and the blue circles in each panel show the converged value. While the converged value sometimes reflects a non-global minimum, this global search successfully identifies the same parameter vector for all 10 of the top-10 best-fitting parameter estimates across all 50 randomly chosen seeds.

7 Estimates

Our current estimates are preliminary and reflect a simplified version of our full model in which we target only moments related to unsecured borrowing. In ongoing work, we are using more richly detailed credit report data to allow us to more usefully target secured-loan outcomes as

¹²To further develop this example, consider the case where history β is the empty history, and history α is empty except for positive information (i.e., borrowing with no default) in the most distant time period. Because the long difference in the two histories is forgotten (i.e., removed from credit histories) by the next period, these consumers face identical reputational incentives. Changes in choice probabilities from history α to one period later in history β then reveal how the underlying types in these histories have changed.

¹³We note that even when there is no primitive demand-risk correlation ($\rho = 0$), these plots still exhibit positive correlation due to risk-based pricing and the persistence of types: consumers who selected into borrowing in high-default histories in the past on average had higher demand given the high prices they faced, and persistence in types then means that consumers currently in these histories tend to have high demand and high risk. However, as shown in this figure, ρ still governs how strong this positive correlation is.

well.¹⁴

This section describes the simplified model and its fit in subsection 7.1, presents model parameter estimates in subsection 7.2, and validates the model vis-à-vis untargeted moments in subsection 7.3.

7.1 Model Fit

To summarize our current approach, recall from Section 2.3 that our model features both secured and unsecured loan default; secured default occurs independently of a consumer’s choice of whether to borrow unsecured, and stochastically based on (only) the consumer’s risk type. In the simplified version of the model, we continue to allow secured default to occur but we do not target any empirical moments related to secured default, and we define borrower histories so that they include default if and only if there was an unsecured default.

Figure 7 illustrates the fit of such a model vis-à-vis our target moments related to unsecured borrowing. The left panel of the figure shows target moments, and the right panel shows corresponding model moments. Within each panel, a row corresponds to a credit history, and each column corresponds to potential outcomes for consumers who enter a given period with that history.

Starting with histories (the rows of Figure 7), recall that we represent a history as $[B_1, B_2 | D_1, D_2]$, where $B_1 \in \{0, 1\}$ records borrowing in the most recent period and $B_2 \in \{0, 1\}$ records borrowing in the preceding period, while D_1 and D_2 respectively record default over the same two periods, each taking a value in $\{0, 1, 2\}$ representing no default, mild default, or severe default respectively. The histories/rows in Figure 7 show all possible combinations of these components for unsecured borrowing. The row label for each history also shows the share of consumers in that history.

Next, the columns of Figure 7 show the probabilities of each possible current-period outcome for consumers entering that period with a given history: non-borrowing, borrowing with no default, borrowing with mild default, and borrowing with severe default. (The consumer’s probability of choosing to borrow is the sum of the last three of these probabilities.)

As can be seen comparing the two panels of the figure, the model is able to match these choice probabilities and default probabilities quite closely. For example, in the most commonly held history (a clean, thick file with a full history of borrowing and no history of default), the model finds that 70.9% of consumers choose to borrow and do not face a default; the corresponding probability in the data is 70.2%. Fit is also generally good for less common histories, though there are some exceptions. Most markedly the model tends to under-predict the share of consumers who choose non-borrowing in the rarely held histories 2, 3, and 4, which represent consumers

¹⁴An implication of this limitation is that we currently do not use price variation related to the scheduled removal of negative credit report information, as discussed in Section 6. This leaves us with the second of the two sources of price variation described in that section, which we do not believe are able to *non*-parametrically identify the model.

who have some credit history but did not borrow at all in the last period. Within these histories, the model fit is generally better as the share of consumers in the history becomes higher.

Next, Figure 8 examines how well the model fits the stationary distribution of consumers across histories, and how well the estimated equilibrium satisfies the zero-profit condition. The left panel shows the share of consumers in the data- and model-implied stationary distribution across histories. Consistent with the generally good fit to choice probabilities, the fit of the stationary distribution is quite close, though the model somewhat over-predicts the share of consumers with no credit history at all. The second panel of the figure shows that model-estimated prices are generally quite close to the default rate in each history, and there do not appear to be systematic over- or under-pricing errors across histories.

7.2 Parameter Estimates

Table 1 shows our estimates of model parameters specific to unsecured borrowing. We find moderately high correlation between demand and risk types, with $\rho = 0.65$. These types are also persistent: the autoregressive coefficients for demand and risk types are respectively 0.64 and 0.73. There is somewhat higher dispersion in unobserved risk than in unobserved demand.

To help interpret the magnitudes of these parameter estimates, we consider the implied probability of consumers' type changes. Table 2 shows one-period transition probabilities across joint demand ("D") and risk ("R") types, with higher type indices corresponding to higher demand and higher risk.¹⁵ Types are persistent, as shown by the substantial mass on the diagonal, but not fully persistent. Even the most persistent types – those in the tails of the distribution – have a higher chance of transitioning to a different type than remaining the same joint type in the next period.

Figure 9 illustrates the stationary distribution of hidden demand and risk types implied by these estimates. This distribution is an equilibrium object, so it reflects the influence of all estimated parameters, but it also particularly highlights the roles of the means and variances of unobserved demand and risk types $(\mu_1, \mu_2, \sigma_1, \sigma_2)$ and the correlation between demand and risk (ρ). The figure is an empirical analog to the calibrated examples in Figure 2 discussed above in Section 3, with similar color coding across efficient and inefficient exclusions and inclusions. As the plot shows, we estimate that the modal consumer borrows efficiently in the current US consumer credit market equilibrium, while a large share of excluded consumers (the yellow and blue dots) are *inefficiently* excluded (yellow). However, as shown by the location of the yellow dots, many inefficiently excluded consumers are close to indifferent between borrowing and not borrowing if priced at their marginal cost (i.e., they are located close to the 45-degree line), suggesting the welfare costs of these inefficient exclusions may be modest. We quantify these welfare

¹⁵For brevity, table 2 only shows selected types, so rows do not sum to 100%.

costs more formally in Section 8, where we also study welfare under counterfactual information structures.

Finally, these parameter estimates also interestingly show how much stochasticity there is in default outcomes. That is, how much of default could be anticipated by knowing consumers' types, as would be the case if most consumer types had default probabilities close to 0 or 1, and how much of default is ex ante uncertain even after knowing a consumer's type, as would be the case if most consumers had default probabilities on the interior of 0 and 1. In our current estimates, we find that the riskiest consumers have a 73% probability of becoming severely delinquent on their unsecured borrowing whereas lower risk types tend to have corresponding default probabilities below 2% over one period.

7.3 Model Validation

To validate the model against untargeted moments, we consider the implied price-elasticities of demand for each consumer type (θ_d, θ_r) and compare these to external estimates of consumer demand elasticities in the review by Karlan and Zinman (2019) of similar estimates in the literature. Table 3 shows that, at equilibrium prices, our estimated elasticities at the consumer-type level range from 2.3 (for low-demand types) to approximately zero (for the highest-demand types). By way of comparison, Karlan and Zinman (2019) report US unsecured credit demand elasticities that range from 0.8 to 1.3 (depending on the horizon over which quantity is measured), with 1.3 being the closest analog to our estimates. They show that other estimates from outside the US also fall within the range of elasticities we find. We find these results encouraging about the reasonableness of our estimates.

8 Counterfactual Information Structures

In this section, we study the welfare consequences of the current US credit reporting system by comparing welfare under the current equilibrium to two counterfactuals: full information about consumer unobserved types; and a no-history equilibrium with complete pooling of all consumers.¹⁶ We find that a large minority of US consumers are inefficiently excluded (or “priced out”) from borrowing under the current credit reporting system. However, the welfare costs of these inefficient exclusions are only moderate, because many of these consumers are close to indifferent between borrowing or not. As one illustration of these welfare costs, we estimate that total welfare falls by about 15% when moving from a full-information counterfactual to a no-information counterfactual, and we find that the current credit reporting system recovers about two-thirds of this welfare loss. We also find that the gains from credit reporting are heteroge-

¹⁶See also Glover and Corbae (2015); Bundorf et al. (2012); Chatterjee et al. (2020), who study similar counterfactuals.

neous: because adverse selection is not severe enough to induce extreme market unraveling in the absence of credit reporting (i.e., a situation where only the riskiest types borrow), many consumers with negative histories or with no credit history would in fact experience higher welfare under a shift to a no-information (i.e., fully pooling) equilibrium.

We implement the full-information and no-information counterfactuals as follows. Taking our estimates of the primitives of the consumer credit market that we report in Table 1, we reconstruct credit histories so that they either are a singleton (the no-information case) or they are indicators for consumers' current types (the full-information case). We then solve for zero-profit pricing at each history that accounts for adverse selection, in the sense of pricing to the average cost of consumers who hold a given history in the counterfactual equilibrium and who choose to participate at a given price. This means searching for prices that satisfy equation 4 for the counterfactual histories, accounting for consumers' choice behavior given those prices.¹⁷

Figure 10 illustrates these counterfactual equilibria in comparison to the equilibrium under current US credit reporting (which we refer to as the "baseline" equilibrium). The left panel of the figure shows the no-information case; the middle panel shows the baseline case; and the right panel shows the full-information case. Within each panel, four colored bars show the share of consumers who occupy each of the four groups defined by being efficiently or inefficiently included or excluded, as in the four quadrants of the earlier Figure 9. Starting with the no-information case, the figure shows that a majority ($\approx 55\%$) of consumers are inefficiently excluded from borrowing when credit scoring is made unavailable. Inefficient inclusions and efficient exclusions are both quite rare in this equilibrium.

Turning to the baseline case, in the middle panel of Figure 10, we see that credit reporting reduces the share of inefficiently excluded consumers, so that a slim majority of consumers is now efficiently included. However, the share of inefficient exclusions remains substantial, at over 40% of all consumers. Naturally, all of these inefficient exclusions are then efficiently *included* in the full-information case, as shown in the final panel of the figure.

While we find many consumers are inefficiently excluded under both the baseline equilibrium and the no-information counterfactual, we also find the welfare consequences of these exclusions are only moderate. Figure 11 shows total surplus in each of the three information environments; these total surplus estimates are also equivalent to consumer surplus in this zero-profit setting. We find that total surplus falls by roughly 15% in the no-information case relative to the full-information case, while credit reporting is able to restore about two-thirds of this welfare loss.

The earlier evidence in Figure 9 helps with understanding these two patterns across our counterfactuals. First, the welfare cost of moving from full information to no information is moderate, despite the rather high level of inefficient exclusion in the no-information case, because of how

¹⁷In case of multiple prices that satisfy that equation for a given history, we choose the *lowest* such price by designing our price search to ascend from zero.

close to indifferent many excluded consumers are from borrowing: their willingness to pay for credit is quite close to the social cost of making those loans. Moreover, even though credit reporting reduced the share of consumers who are inefficiently excluded only partially, credit reporting achieves most of the total surplus available under full information because it prevents the inefficient exclusion of consumers who are least indifferent – or among whom the surplus gains from inclusion are the highest.

We also find a nuanced picture of the distributional effects of credit reporting. Because the no-information equilibrium achieves roughly 85% of the total surplus available in the full-information case, and an even higher share of the total surplus in the baseline case, it would seem plausible that some groups of consumers in fact benefit from transitioning towards the no-information equilibrium. We therefore ask how total surplus changes from the baseline equilibrium to the no-information equilibrium for four groups of consumers: (1) those who borrowed in the last period and have no history of default, (2) those who borrowed in the last period but also have a history of default, (3) those who did not borrow in the last period and have no history of default, and (4) those who did not borrow in the last period but do have a history of default (e.g., from the prior period). We refer to groups (3) and (4) as those who played “out” in the prior period. (To be clear, these prior-period “out” consumers are potentially different from the consumers who play “out” in the current period.) .

Figure 12 shows the distributional effects of moving to different information environments for these four groups. The surplus losses from the transition to the no-information equilibrium are primarily concentrated among the “prime” consumers in the first group. In contrast, and precisely because we find that the credit market is not so adversely selected as to induce extreme unraveling in the no-information environment, we find that consumers in groups (2) and (3) (i.e., those with negative history or with no recent positive history), benefit from the no-information case. This is suggestive of potential for welfare gains among certain groups, including those with no recent history or no credit history at all, if the US were to transition to using less rich credit history information.¹⁸

Finally, we use the no-information counterfactual to help decompose the mechanisms underlying the equilibrium effects of credit reporting. In particular, note that credit reporting changes consumers’ borrowing decisions through two channels: a current-period price of borrowing; and a reputational incentive to avoid borrowing when having privately high default risk, in order to maintain a good credit record and access to lower future prices. To decompose credit reporting’s effects across these two channels, we adapt the four-quadrant framework used earlier in, for example, Figure 9, but now plotting willingness-to-pay (inclusive of reputational incentives) rather

¹⁸We explore this conjecture further in ongoing work with counterfactuals that consider intermediate cases between the no-information environment and the baseline environment. Interestingly, these results are in contrast with the narrative that the expansion of risk-based pricing has primarily benefited high-risk consumers (Edelberg, 2006).

than demand types on the y -axes. Such a plot is shown in the right panel of Figure 13. The other panels of Figure 13 then show how consumers are moved through this WTP-vs.-cost space when they counterfactually face only reputational incentives (but no heterogeneity in current-period pricing), in the middle panel; and when they face the no-information equilibrium, in the left panel. We find that reputational incentives account for only a modest re-allocation of consumers across this space and across quadrants. In this sense, much of the equilibrium effects of pricing arise from separating consumer types contemporaneously, rather than from introducing a risk of future credit market exclusion after default.

9 Conclusion

How does the availability of information about consumers drives market outcomes in the US consumer credit markets? To answer this question empirically, we develop a structural model of consumer credit market primitives, credit histories and banks' lending behavior to analyze how credit market equilibria change when the information structure changes.

We highlight how the persistence of consumers' risk and demand types λ as well the correlation between consumers' risk and demand ρ are key primitives determining much the scope to which information can shape credit market outcomes. Without demand-risk correlation there is no problem of adverse selection for which more information would be helpful. Without much persistence in consumers' risk and demand types, information about past behavior carries little value for inferring current risk.

Estimating our structural model on a large consumer credit panel, we find that types are persistent with λ estimates of 0.64 and 0.73 and that the correlation between demand and risk types is moderately high at around $\rho = 0.65$. This substantial adverse selection generates nontrivial welfare losses and excludes over 40% of all consumers inefficiently. We estimate that the current US credit reporting system recovers roughly two-thirds of the welfare that would be lost in a no-credit-reporting counterfactual, relative to a first-best full-information counterfactual.

Finally, we decompose the mechanisms underlying the equilibrium effects of credit reporting and demonstrate that the credit reporting effect operates primarily through the contemporaneous separation of consumer types rather than introducing strong reputational incentives in the form of risk of future credit market exclusion after default.

Bibliography

- AGARWAL, S., S. CHOMSISENGPHET, AND C. LIU (2010): "The Importance of Adverse Selection in the Credit Card Market: Evidence from Randomized Trials of Credit Card Solicitations," *Journal of Money, Credit and Banking*, 42.
- AGARWAL, S., S. CHOMSISENGPHET, N. MAHONEY, AND J. STROEBEL (2018): "Do Banks Pass through Credit Expansions to Consumers Who Want to Borrow?" *The Quarterly Journal of Economics*, 133.
- AGARWAL, S., J. GRIGSBY, A. HORTAÇSU, G. MATVOS, A. SERU, AND V. YAO (2020): "Searching for approval," Tech. rep., National Bureau of Economic Research.
- AGUIRRE, I., S. COWAN, AND J. VICKERS (2010): "Monopoly Price Discrimination and Demand Curvature," *American Economic Review*, 100.
- AGUIRREGABIRIA, V. (2021): "Chapter 4 - Dynamic Games in Empirical Industrial Organization," in *Handbook of Industrial Organization*, vol. IV.
- AGUIRREGABIRIA, V. AND P. MIRA (2007): "Sequential Estimation of Dynamic Discrete Games," *Econometrica*, 75.
- (2010): "Dynamic Discrete Choice Structural Models: A Survey," *Journal of Econometrics*, 156.
- ALLEN, J., R. CLARK, AND J.-F. HOUDE (2014): "The Effect of Mergers in Search Markets: Evidence from the Canadian Mortgage Industry," *American Economic Review*, 104.
- (2019): "Search Frictions and Market Power in Negotiated-Price Markets," *Journal of Political Economy*, 127.
- ALTONJI, J. G. AND L. M. SEGAL (1996): "Small-Sample Bias in GMM Estimation of Covariance Structures," *Journal of Business & Economic Statistics*, 14.
- AMBROSE, B. W., J. CONKLIN, AND J. YOSHIDA (2016): "Credit Rationing, Income Exaggeration, and Adverse Selection in the Mortgage Market," *The Journal of Finance*, 71.
- ARCIDIACONO, P. AND P. B. ELLICKSON (2011): "Practical Methods for Estimation of Dynamic Discrete Choice Models," *Annual Review of Economics*, 3.
- ARCIDIACONO, P. AND R. A. MILLER (2011): "Conditional Choice Probability Estimation of Dynamic Discrete Choice Models With Unobserved Heterogeneity," *Econometrica*, 79.
- ARGYLE, B. S., T. D. NADAULD, AND C. J. PALMER (2020): "Monthly payment targeting and the demand for maturity," *The Review of Financial Studies*, 33, 5416–5462.
- ATHREYA, K., X. S. TAM, AND E. R. YOUNG (2012): "A Quantitative Theory of Information and Unsecured Credit," *American Economic Journal: Macroeconomics*, 4.
- AUSUBEL, L. M. (1991): "The Failure of Competition in the Credit Card Market," *The American Economic Review*, 81.

- (1999): “Adverse Selection,” *Working Paper*.
- AZEVEDO, E. M. AND D. GOTTLIEB (2017): “Perfect Competition in Markets With Adverse Selection,” *Econometrica*, 85.
- BAJARI, P., C. DALTON, H. HONG, AND A. KHWAJA (2014): “Moral Hazard, Adverse Selection, and Health Expenditures: A Semiparametric Analysis,” *The RAND Journal of Economics*, 45.
- BALL, I. (2021): “Scoring Strategic Agents,” *Working Paper*.
- BAR-ISAAC, H., I. JEWITT, AND C. LEAVER (2021): “Adverse Selection, Efficiency and the Structure of Information,” *Economic Theory*, 72.
- BARTIK, A. AND S. NELSON (2020): “Deleting a Signal: Evidence from Pre-Employment Credit Checks,” *SSRN Electronic Journal*.
- BENETTON, M. (2021): “Leverage Regulation and Market Structure: A Structural Model of the U.K. Mortgage Market,” *The Journal of Finance*, 76.
- BERGEMANN, D., A. BONATTI, AND A. SMOLIN (2018): “The Design and Price of Information,” *American Economic Review*, 108.
- BERGEMANN, D., B. BROOKS, AND S. MORRIS (2015): “The Limits of Price Discrimination,” *American Economic Review*, 105.
- BERGEMANN, D. AND S. MORRIS (2019): “Information Design: A Unified Perspective,” *Journal of Economic Literature*, 57.
- BHASKAR, V. AND C. THOMAS (2017): “The Design of Credit Information Systems,” *Working Paper*.
- (2019): “Community Enforcement of Trust with Bounded Memory,” *The Review of Economic Studies*, 86.
- BHATTACHARYA, V., G. ILLANES, AND M. PADI (2019): “Fiduciary duty and the market for financial advice,” Tech. rep., National Bureau of Economic Research.
- BLATTNER, L. AND S. NELSON (2021): “How Costly Is Noise? Data and Disparities in Credit Scoring,” *Working Paper*.
- BONATTI, A. AND G. CISTERNAS (2020): “Consumer Scores and Price Discrimination,” *The Review of Economic Studies*, 87.
- BOND, P. AND A. KRISHNAMURTHY (2004): “Regulating exclusion from financial markets,” *The Review of Economic Studies*, 71, 681–707.
- BOOT, A. W. AND A. V. THAKOR (1994): “Moral hazard and secured lending in an infinitely repeated credit market game,” *International economic review*, 899–920.
- BOS, M., E. BREZA, AND A. LIBERMAN (2018): “The Labor Market Effects of Credit Market Information,” *The Review of Financial Studies*, 31.

- BOS, M. AND L. I. NAKAMURA (2014): "Should Defaults Be Forgotten? Evidence from Variation in Removal of Negative Consumer Credit Information," *Working Paper*.
- BROWN, M., A. FALK, AND E. FEHR (2004): "Relational contracts and the nature of market interactions," *Econometrica*, 72, 747–780.
- BROWN, M., T. JAPPELLI, AND M. PAGANO (2009): "Information sharing and credit: Firm-level evidence from transition countries," *Journal of Financial Intermediation*, 18, 151–172.
- BROWN, M. AND C. ZEHNDER (2007): "Credit reporting, relationship banking, and loan repayment," *Journal of Money, Credit and Banking*, 39, 1883–1918.
- BUCHAK, G., G. MATVOS, T. PISKORSKI, AND A. SERU (2020): "Why Is Intermediating Houses so Difficult? Evidence from iBuyers," *Working Paper*.
- BUNDORF, M. K., J. LEVIN, AND N. MAHONEY (2012): "Pricing and Welfare in Health Plan Choice," *American Economic Review*, 102.
- CASTELLANOS, S. G., D. J. HERNÁNDEZ, A. MAHAJAN, AND E. SEIRA (2018): "Financial Inclusion and Contract Terms: Experimental Evidence from Mexico," Working paper, National Bureau of Economic Research.
- CHATTERJEE, S., D. CORBAE, K. DEMPSEY, AND J.-V. RÍOS-RULL (2020): "A Quantitative Theory of the Credit Score," *Working Paper*.
- CHATTERJEE, S., D. CORBAE, M. NAKAJIMA, AND J.-V. RÍOS-RULL (2007): "A Quantitative Theory of Unsecured Consumer Credit with Risk of Default," *Econometrica*, 75.
- CHIAPPORI, P.-A. AND B. SALANIE (2000): "Testing for Asymmetric Information in Insurance Markets," *Journal of Political Economy*, 108.
- CLARK, R., J.-F. HOUDE, AND J. KASTL (2021): "Chapter 15 - The Industrial Organization of Financial Markets," in *Handbook of Industrial Organization*, ed. by K. Ho, A. Hortaçsu, and A. Lizzeri, Elsevier, vol. 5 of *Handbook of Industrial Organization, Volume 5*.
- COHEN, A. AND L. EINAV (2007): "Estimating Risk Preferences from Deductible Choice," *American Economic Review*, 97.
- CRAWFORD, G. S., N. PAVANINI, AND F. SCHIVARDI (2018): "Asymmetric Information and Imperfect Competition in Lending Markets," *American Economic Review*, 108.
- CROSS, R., T. DUTZIK, E. MIERZWINSKI, AND M. CASALE (2019): "Driving Into Debt," *US PIRG Education Fund*.
- CUESTA, J. I. AND A. SEPULVEDA (2021): "Price Regulation in Credit Markets: A Trade-Off between Consumer Protection and Credit Access," *Working Paper*.
- DEFUSCO, A. A., H. TANG, AND C. YANNELIS (2021): "Measuring the Welfare Cost of Asymmetric Information in Consumer Credit Markets," Tech. rep., National Bureau of Economic Research.

- DJANKOV, S., C. MCLIESH, AND A. SHLEIFER (2007): "Private credit in 129 countries," *Journal of financial Economics*, 84, 299–329.
- DOBBIE, W., P. GOLDSMITH-PINKHAM, N. MAHONEY, AND J. SONG (2020): "Bad Credit, No Problem? Credit and Labor Market Consequences of Bad Credit Reports," *The Journal of Finance*, 75.
- DOBBIE, W. AND P. M. SKIBA (2013): "Information Asymmetries in Consumer Credit Markets: Evidence from Payday Lending," *American Economic Journal: Applied Economics*, 5.
- DROZD, L. A. AND J. B. NOSAL (2008): "Competing for Customers: A Search Model of the Market for Unsecured Credit," *Working Paper*.
- DROZD, L. A. AND R. SERRANO-PADIAL (2017): "Modeling the Revolving Revolution: The Debt Collection Channel," *American Economic Review*, 107.
- DUBE, J.-P. AND S. MISRA (2022): "Personalized Pricing and Consumer Welfare," *Working Paper*.
- EDELBERG, W. (2006): "Risk-based pricing of interest rates for consumer loans," *Journal of monetary Economics*, 53, 2283–2298.
- EGAN, M., A. HORTAÇSU, AND G. MATVOS (2017): "Deposit Competition and Financial Fragility: Evidence from the US Banking Sector," *American Economic Review*, 107.
- EINAV, L., A. FINKELSTEIN, AND M. R. CULLEN (2010): "Estimating Welfare in Insurance Markets Using Variation in Prices," *The Quarterly Journal of Economics*, 125.
- EINAV, L., A. FINKELSTEIN, AND N. MAHONEY (2021): "Chapter 14 - The IO of Selection Markets," in *Handbook of Industrial Organization*, ed. by K. Ho, A. Hortaçsu, and A. Lizzeri, Elsevier, vol. 5 of *Handbook of Industrial Organization, Volume 5*.
- EINAV, L., A. FINKELSTEIN, S. P. RYAN, P. SCHRIMPF, AND M. R. CULLEN (2013a): "Selection on Moral Hazard in Health Insurance," *American Economic Review*, 103.
- EINAV, L., M. JENKINS, AND J. LEVIN (2012): "Contract Pricing in Consumer Credit Markets," *Econometrica*, 80.
- (2013b): "The Impact of Credit Scoring on Consumer Lending," *The RAND Journal of Economics*, 44.
- ELUL, R. AND P. GOTTARDI (2015): "Bankruptcy: Is It Enough to Forgive or Must We Also Forget?" *American Economic Journal: Microeconomics*, 7.
- EXLER, F., I. LIVSHITS, J. MACGEE, AND M. TERTILT (2020): "Consumer Credit With Over-Optimistic Borrowers," *Working Paper*.
- FINKELSTEIN, A. AND K. MCGARRY (2006): "Multiple Dimensions of Private Information: Evidence from the Long-Term Care Insurance Market," *The American Economic Review*, 96.
- FULFORD, S. L. AND É. NAGYPÁL (2020): "The Equilibrium Effect of Information in Consumer Credit Markets: Public Records and Credit Redistribution," *Working Paper*.

- GALENIANOS, M. AND A. GAVAZZA (2019a): “Regulatory Intervention in Consumer Financial Markets: The Case of Credit Cards,” Working paper.
- (2019b): “Regulatory Interventions in Consumer Financial Markets: The Case of Credit Cards,” *SSRN Electronic Journal*.
- GANONG, P. AND P. NOEL (2020a): “Liquidity versus Wealth in Household Debt Obligations: Evidence from Housing Policy in the Great Recession,” *American Economic Review*, 110.
- GANONG, P. AND P. J. NOEL (2020b): “Why do borrowers default on mortgages? A new method for causal attribution,” Tech. rep., National Bureau of Economic Research.
- GARCIA, D. AND M. TSUR (2021): “Information Design in Competitive Insurance Markets,” *Journal of Economic Theory*, 191.
- GERUSO, M., T. J. LAYTON, G. MCCORMACK, AND M. SHEPARD (2021): “The Two Margin Problem in Insurance Markets,” *Review of Economics and Statistics*.
- GLOVER, A. AND D. CORBAE (2015): “A Simple Dynamic Theory of Credit Scores Under Adverse Selection,” *Working Paper*.
- GROSS, D. B. AND N. S. SOULELES (2002): “An Empirical Analysis of Personal Bankruptcy and Delinquency,” *The Review of Financial Studies*, 15.
- GROSS, T., M. J. NOTOWIDIGDO, AND J. WANG (2020): “The Marginal Propensity to Consume over the Business Cycle,” *American Economic Journal: Macroeconomics*, 12.
- GUERRIERI, V. AND R. SHIMER (2018): “Markets with Multidimensional Private Information,” *American Economic Journal: Microeconomics*, 10.
- GUERRIERI, V., R. SHIMER, AND R. WRIGHT (2010): “Adverse Selection in Competitive Search Equilibrium,” *Econometrica*, 78.
- GUPTA, A. AND C. HANSMAN (2020): “Selection, Leverage, and Default in the Mortgage Market,” *Working Paper*.
- HAN, S. AND G. LI (2011): “Household Borrowing after Personal Bankruptcy,” *Journal of Money, Credit and Banking*, 43.
- HANDEL, B., I. HENDEL, AND M. D. WHINSTON (2015): “Equilibria in Health Exchanges: Adverse Selection versus Reclassification Risk,” *Econometrica*, 83.
- HANDEL, B. R. (2013): “Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts,” *American Economic Review*, 103.
- HEILBRON, J. (2022): “Credit Constraints and Bias in Hedonic Amenity Valuations,” *University of Chicago, PhD Dissertation*.


- HERTZBERG, A., A. LIBERMAN, AND D. PARAVISINI (2018): "Screening on Loan Terms: Evidence from Maturity Choice in Consumer Credit," *The Review of Financial Studies*, 31.
- HOTZ, V. J. AND R. A. MILLER (1993): "Conditional Choice Probabilities and the Estimation of Dynamic Models," *The Review of Economic Studies*, 60.
- HUNT, R. M. (2006): "Development and regulation of consumer credit reporting in the United States," *The Economics of consumer credit*, 301.
- (2007): "Collecting consumer debt in America," *Available at SSRN 993249*.
- INDARTE, S. (2021): "Moral Hazard versus Liquidity in Household Bankruptcy," *Working Paper*.
- JANSEN, M., F. NAGEL, C. YANNELIS, AND A. L. ZHANG (2022): "Data and Welfare in Credit Markets," *Working Paper*.
- JAPPELLI, T. AND M. PAGANO (2002): "Information sharing, lending and defaults: Cross-country evidence," *Journal of Banking & Finance*, 26, 2017–2045.
- JAPPELLI, T., M. PAGANO, ET AL. (2000): "Information sharing in credit markets: The European experience," *Unpublished Working Paper*, 35.
- JIANG, E. X. (2020): "Financing Competitors: Shadow Banks' Funding and Mortgage Market Competition," *Working Paper*.
- KAMBARA, M., P. GRIMM, AND K. P. BREVOORT (2015): "Data Point: Credit Invisibles," *The CFPB Office of Research*.
- KAMENICA, E. (2019): "Bayesian Persuasion and Information Design," *Annual Review of Economics*, 11.
- KAMENICA, E. AND M. GENTZKOW (2011): "Bayesian Persuasion," *American Economic Review*, 101.
- KARLAN, D. AND J. ZINMAN (2009): "Observing Unobservables: Identifying Information Asymmetries With a Consumer Credit Field Experiment," *Econometrica*, 77.
- (2019): "Long-Run Price Elasticities of Demand for Credit: Evidence from a Countrywide Field Experiment in Mexico," *The Review of Economic Studies*, 86.
- KARLAN, D. S. AND J. ZINMAN (2008): "Credit Elasticities in Less-Developed Economies: Implications for Microfinance," *American Economic Review*, 98.
- KEANE, M. P. AND K. I. WOLPIN (2009): "Empirical Applications of Discrete Choice Dynamic Programming Models," *Review of Economic Dynamics*, 12.
- KEHOE, T. J. AND D. K. LEVINE (1993): "Debt-constrained asset markets," *The Review of Economic Studies*, 60, 865–888.
- KHANDANI, A. E., A. J. KIM, AND A. W. LO (2010): "Consumer Credit-Risk Models via Machine-Learning Algorithms," *Journal of Banking & Finance*, 34.

- KOVASYUK, S. AND G. SPAGNOLO (2018): "Memory and Markets," SSRN Scholarly Paper, Social Science Research Network, Rochester, NY.
- KOVRIJNYKH, N., I. LIVSHITS, AND A. ZETLIN-JONES (2019): "Building Credit Histories," *Working Paper*.
- LEVIN, J. (2001): "Information and the Market for Lemons," *The RAND Journal of Economics*, 32.
- LIBERMAN, A., C. A. NELSON, L. OPAZO, AND S. ZIMMERMAN (2019): "The Equilibrium Effects of Information Deletion: Evidence from Consumer Credit Markets," *Working Paper*.
- LIVSHITS, I., J. MACGEE, AND M. TERTILT (2007): "Consumer Bankruptcy: A Fresh Start," *American Economic Review*, 97.
- LIVSHITS, I., J. MACGEE, AND M. TERTILT (2010): "Accounting for the Rise in Consumer Bankruptcies," *American Economic Journal: Macroeconomics*, 2.
- (2016): "The Democratization of Credit and the Rise in Consumer Bankruptcies," *The Review of Economic Studies*, 83.
- MAHONEY, N. AND E. G. WEYL (2017): "Imperfect Competition in Selection Markets," *The Review of Economics and Statistics*, 99.
- McFADDEN, D. (1989): "A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration," *Econometrica*, 57.
- MUSTO, D. K. (2004): "What Happens When Information Leaves a Market? Evidence from Postbankruptcy Consumers," *The Journal of Business*, 77.
- NELSON, S. T. (2022): "Private Information and Price Regulation in the US Credit Card Market," *Working Paper*.
- PAGANO, M. AND T. JAPPELLI (1993): "Information sharing in credit markets," *The journal of finance*, 48, 1693–1718.
- PERRIGNE, I. AND Q. VUONG (2011): "Nonparametric Identification of a Contract Model With Adverse Selection and Moral Hazard," *Econometrica*, 79.
- ROBLES-GARCIA, C. (2020): "Competition and Incentives in Mortgage Markets: The Role of Brokers," *Working Paper*.
- RUST, J. (1987): "Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher," *Econometrica: Journal of the Econometric Society*, 999–1033.
- (1994): "Chapter 51 Structural Estimation of Markov Decision Processes," *Handbook of Econometrics*, 4.
- SCHMALENSSEE, R. (1981): "Output and Welfare Implications of Monopolistic Third-Degree Price Discrimination," *The American Economic Review*.

- SCOTT, P. T. (2013): "Dynamic Discrete Choice Estimation of Agricultural Land Use," *Working Paper*.
- SHARMA, P. (2017): "Is More Information Always Better? A Case in Credit Markets," *Journal of Economic Behavior & Organization*, 134.
- STIGLITZ, J. E. AND A. WEISS (1981): "Credit Rationing in Markets with Imperfect Information," .
- TAUCHEN, G. (1986): "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions," *Economics Letters*, 20.
- TERRY, S. J. AND E. S. KNOTEK (2011): "Markov-Chain Approximations of Vector Autoregressions: Application of General Multivariate-Normal Integration Techniques," *Economics Letters*, 110.
- THOMAS, L. C. (2009): *Consumer credit models: pricing, profit and portfolios: pricing, profit and portfolios*, OUP Oxford.
- VARIAN, H. R. (1985): "Price Discrimination and Social Welfare," *American Economic Review*.
- VATTER, B. (2021): "Quality Disclosure and Regulation," *Working Paper*.
- VATTER, B. AND V. MARONE (2021): "Winners and Losers Under Counterfactual Health Risk Pooling," *Working Paper*.
- VEIGA, A. (2021): "Price Regulation vs Information Design," .
- YANG, K. H. (2022): "Selling Consumer Data for Profit: Optimal Market-Segmentation Design and Its Consequences," *American Economic Review*, 112.

10 Figures

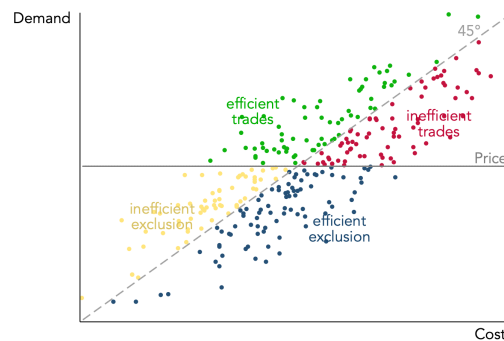
Figure 1: Model Setup: Sequencing of Choices and State Updating



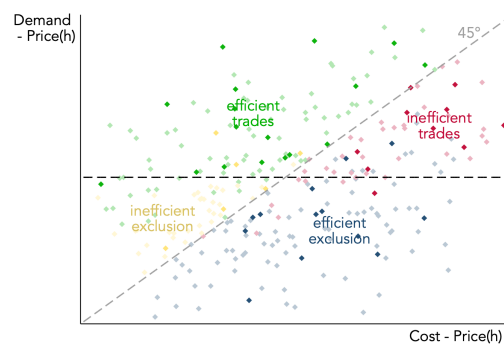
Notes: The figure shows sequences of choices and state updating in the model. A model period starts with the realisation of consumer types θ given consumer's previous type θ_{-1} according to $q(\theta|\theta_{-1})$. For each history h lenders offer loans at a interest rate $p(h)$ and consumers choose their borrowing $b(p(h)|\theta)$. Conditional on consumer's borrowing choice and default type, defaults are realized and histories are updated.

Figure 2: Mechanism Illustration: Persistence and Correlation

(a) Pooled Pricing




(b) Heterogeneous Pricing by History h




Notes: This figure illustrates the basic role persistence and demand-risk correlation plays in the design of consumer credit information. In this schematic, where we normalize demand (on the y-axis) and cost (on the x-axis), each consumer and his demand and cost is represented by a single dot. Trades (i.e providing consumers with a loan) are efficient if they lie above the 45° degree line. This splits consumers into four welfare groups: Efficient trades, inefficient trades, inefficient exclusions and efficient exclusions. Panel a) demonstrates the pooled pricing case in which all consumers face the same price. This homogeneous price is shown as a vertical dashed line. Consumers participate if they are above this price line. Panel b) demonstrates the heterogeneous pricing by history h case, in which the axis are further normalized by the price in each history. As such they show excess demand and excess cost above price. The dashed vertical line indicates at what level demand is exactly equal to the price for the given history.

Figure 3: Type Correlation vs. Persistence: Welfare Group Shares


(a) High Demand-Risk Correlation, High Persistence




(b) Low Demand-Risk Correlation, High Persistence



(c) High Demand-Risk Correlation, Low Persistence




(d) Low Demand-Risk Correlation, Low Persistence




Notes: This figure illustrates the share of consumers that fall into each of the four welfare groups for the four calibrated scenarios of high and low values of correlation ρ and high and low values of persistence λ across three different information environments. Within each panel, the first and second group of four bars show the consumer shares under the no information environment (A) and the baseline credit reporting environment (B) while the final (two) bars show the consumer shares under the full information environment (C). Blue and yellow refer to efficient and inefficient exclusions while green and red refer to efficient and inefficient inclusions.

Figure 4: Type Correlation vs. Persistence: Effects on Total Surplus


(a) High Demand-Risk Correlation, High Persistence




(b) Low Demand-Risk Correlation, High Persistence



(c) High Demand-Risk Correlation, Low Persistence

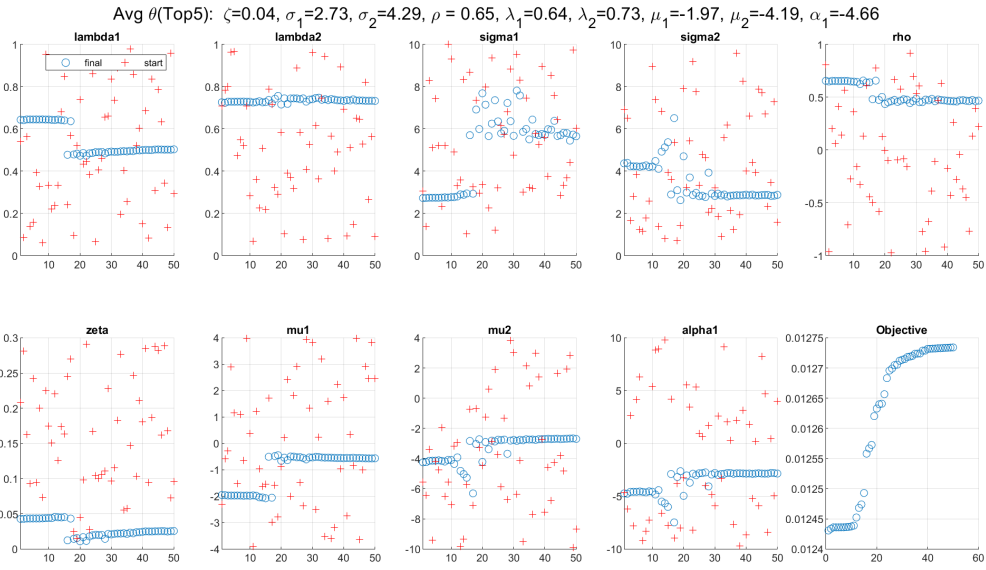


(d) Low Demand-Risk Correlation, Low Persistence



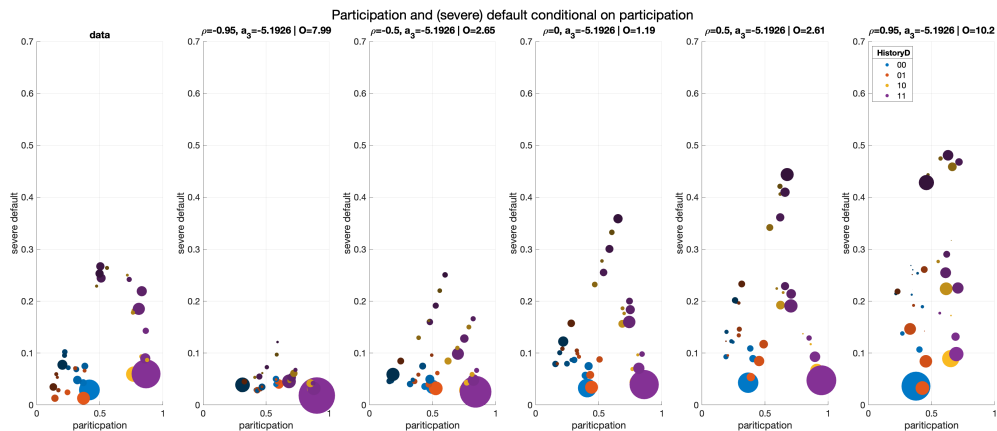
Notes: This figure illustrates the total surplus for the four calibrated scenarios of high and low values of correlation ρ and high and low values of persistence λ across three different information environments. Within each panel, the first and second bars show the total surplus under the no information environment (A) and the baseline credit reporting environment (B) while the third bars show the total surplus under the full information environment (C).

Figure 5: Outer Loop Convergence




Notes: This figure shows the parameter estimates to which the estimation algorithm converged for 50 different starting values. Starting values are marked by a red cross while final values are shown as blue circles. There is a subplot for each parameter while there is one final subplot on the bottom right that shows the obtained objective value. This global search demonstrates how we successfully identify the same parameter vector for all 10 of the top-10 best-fitting parameter estimates.

Figure 6: Illustration of Identification: ρ



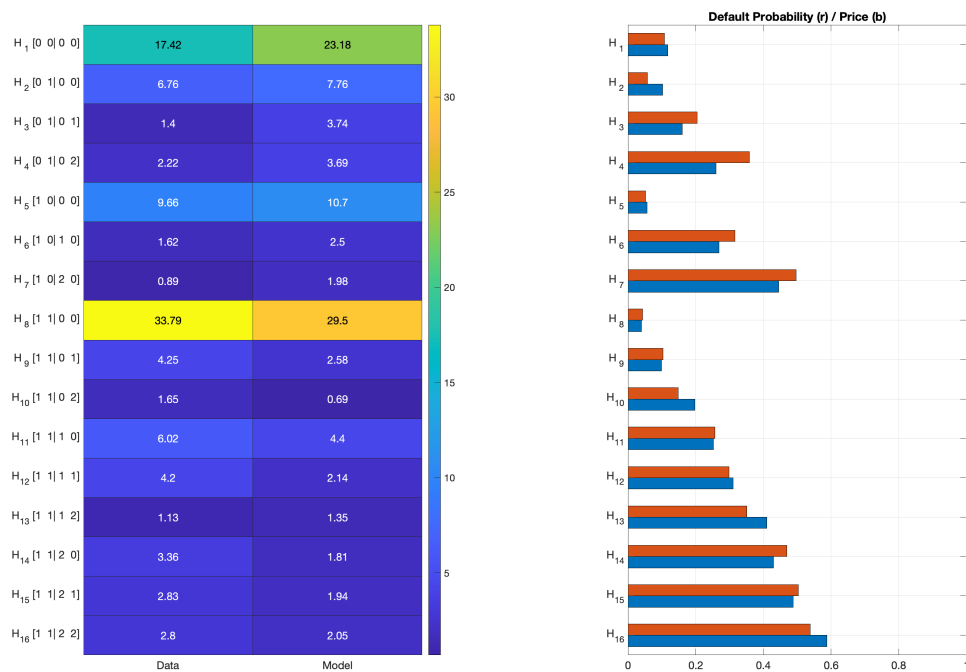
Notes: This figure shows how changes in the demand-risk correlation ρ induce a rotation in the cloud of demand probabilities and default probabilities across all model histories. Within each subplot, each consumer history is a single point. The share of consumers in each history is shown by the size of the point. All histories are grouped into four history groups based on their borrowing history $[B_1 \ B_2]$ as shown in the legend of the final subplot. While the first subplot shows the empirical cloud of demand probabilities and default probabilities as found in the data, the remaining subplots show the cloud for increasing values of ρ .

Figure 7: Target Moments and Model Fit: Transition Dynamics




Notes: This figure shows the history transition matrix as estimated from the data and as generated by the model under the estimated parameter vector. For each subplot, each cell shows the conditional choice probability with which a consumer with a given history chooses to borrow and realises default. Each row corresponds to a history of the form $[B_1 B_2 | D_1 D_2]$ and columns correspond to the borrowing and default choice $[BD]$. With just 9 parameters the model is able to match well the 64 targeted transition probabilities.

Figure 8: Target Moments and Model Fit: Stationary Distribution and Prices




Notes: This figure demonstrates the overall model fit by showing in the left panel the stationary distribution of consumers across histories while the right panel shows the default probabilities and prices across each history. Each history is of the form $[B_1 \ B_2 \ | \ D_1 \ D_2]$. In the left panel, the first column shows the empirical stationary distribution while the right column shows the model-generated stationary distribution for the estimated parameter vector. In the right panel, the red bars shows the default probabilities among consumers who choose to borrow for the given history while the blue bar shows the price lenders choose to charge in that history.

Figure 9: Estimated Distribution of Unobserved Demand and Risk Types




Notes: This figure shows the estimated distribution of unobserved demand and risk types. It is the direct analog of the schematic illustration in figure 2. Excess demand (Willingness-to-Pay minus price) is shown on y-axis and excess cost (marginal cost minus price) is shown on x-axis. For each history, a single marker represents a consumer type with his demand and risk type. The size of the marker gives a sense of the mass of consumers that are of the given type in the given history. Trades (i.e providing consumers with a loan) are efficient if they lie above the 45° degree line. This splits consumers into four welfare groups: Efficient trades (green), inefficient trades (red), inefficient exclusions (yellow) and efficient exclusions (blue).

Figure 10: Counterfactuals: Welfare Group Shares



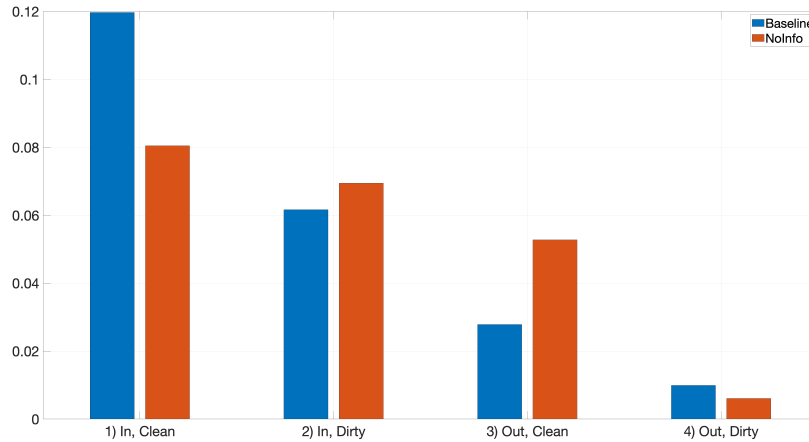
Notes: This figure shows the share of consumers that fall into each of the four welfare groups under counterfactual information environments for our estimated parameter vector. The first and second group of four bars show the consumer shares under the no information environment (A) and the baseline credit reporting environment (B) while the final (two) bars show the consumer shares under the full information environment (C). Blue and yellow refer to efficient and inefficient exclusions while green and red refer to efficient and inefficient inclusions.

Figure 11: Counterfactuals: Total Surplus



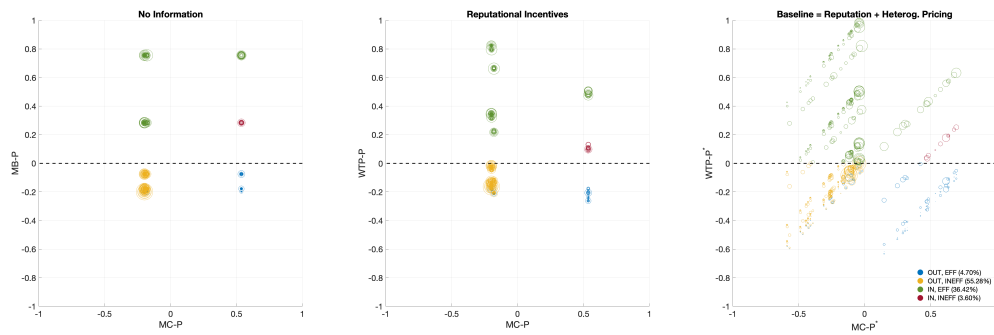
Notes: This figure shows the total surplus under counterfactual information environments for our estimated parameter vector. The first and second bars show the total surplus under the no information environment (A) and the baseline credit reporting environment (B) while the third bars show the total surplus under the full information environment (C).

Figure 12: Distributional Effects of Credit Reporting across History Groups



Notes: This figure shows the distributional effects of moving to counterfactual no information environment relative to the current credit reporting environment. Splitting consumers into four groups based on the credit history, the figure shows the total surplus for each consumer group under the current credit reporting environment (in blue) and the no information environment (in red). Given histories of the form $[B_1 \ B_2 \ | \ D_1 \ D_2]$, the first (1) In, Clean) and second (2) In, Dirty) set of bars refers to consumers with recent borrowing ($B_1 = 1$) and clean histories ($D_1 = 0, D_2 = 0$) or dirty histories ($D_1 \neq 0, D_2 \neq 0$). The third (3) Out, Clean) and fourth (4) Out, Dirty) set of bars refers to consumers with no recent borrowing ($B_1 = 0$) and clean histories ($D_1 = 0, D_2 = 0$) or dirty histories ($D_1 \neq 0, D_2 \neq 0$).

Figure 13: Reputational Incentives vs. Current Price Incentives



Notes: This figure shows the decomposition of the equilibrium effects of credit reporting. Across the three panels, the x-axis shows the excess costs (marginal costs minus price) for each consumer type-history combination where the price P is assumed to be homogeneous for the first two panels while the third panel uses the heterogeneous price P^* that is specific to each history. As before in the normalized willingness-to-pay vs cost space, Panel 1 shows consumers if they all information would be removed. Panel 2 introduces the reputational incentives from heterogeneous pricing while removing the contemporaneous effect from having to pay in the present the history-based price. Panel 3 shows the combined effect of reputational pricing and contemporaneous price heterogeneity under the current credit reporting environment.

11 Tables

Table 1: Model Parameter Estimates

	Estimate	Description
ρ	0.65	demand-risk correlation
λ_1	0.64	demand type persistence
λ_2	0.73	risk type persistence
σ_1	2.73	demand type variance
σ_2	4.29	risk type variance
μ_1	-1.97	demand type mean
μ_2	-4.19	risk type mean
ζ	0.04	taste shock variance
α_1	-4.66	severe-default relative risk

Notes: This table describes and shows the estimate parameter values in the model. Their definition can be found in section 2 while the corresponding estimation procedure is described in section 5.

Table 2: Type Transition Probabilities

	D1R1	D3R1	D5R1	D1R4	D3R4	D5R4	D1R7	D3R7	D5R7
D1R1	40.25	2.74	0.02	0.51	1.77	0.30	0.00	0.01	0.03
D3R1	12.60	14.51	1.65	0.01	0.71	2.57	0.00	0.00	0.06
D5R1	1.13	11.57	17.77	0.00	0.03	4.95	0.00	0.00	0.06
D1R4	5.78	0.02	0.00	11.15	2.27	0.03	0.37	1.93	0.67
D3R4	3.60	0.55	0.01	1.54	6.94	1.54	0.01	0.55	3.60
D5R4	0.67	1.93	0.37	0.03	2.27	11.15	0.00	0.02	5.78
D1R7	0.06	0.00	0.00	4.95	0.03	0.00	17.77	11.57	1.13
D3R7	0.06	0.00	0.00	2.57	0.71	0.01	1.65	14.51	12.60
D5R7	0.03	0.01	0.00	0.30	1.77	0.51	0.02	2.74	40.25

Notes: This table shows selected transition probabilities for the estimated Markov process of the unobserved consumer risk and demand types. Rows refer to the initial type and columns refer to the next period type. *D* refers to the demand type and *R* refers to the risk type. For instance, the demand and risk type 1 has a 40.25% probability of staying the same type but close to zero chance of becoming the highest risk type (*R7*) while staying at the lowest demand type (*D1*) but a 0.03% of becoming a highest demand, highest risk type *D5R7*. Repeated values in the table are a result of how the discretized type process preserves the symmetry of the normal distribution.

Table 3: Price Elasticities

	Risk Type 1	Risk Type 2	Risk Type 3	Risk Type 4	Risk Type 5	Risk Type 6	Risk Type 7
Demand Type 1	-1.08	-1.14	-1.24	-1.38	-1.59	-1.95	-2.34
Demand Type 2	-0.72	-0.79	-0.86	-0.97	-1.18	-1.62	-2.15
Demand Type 3	-0.16	-0.22	-0.28	-0.36	-0.52	-1.08	-1.82
Demand Type 4	-0.00	-0.01	-0.02	-0.06	-0.16	-0.46	-0.80
Demand Type 5	-0.00	-0.00	-0.00	-0.00	-0.00	-0.00	-0.08

Notes: This table summarizes the average price elasticities for each risk (columns) and demand type (rows). Price elasticities increase in magnitude with consumers' risk type and decrease with consumers' demand type. As discussed in section 7.3, the estimated price elasticities are close to typical estimates in the literature.

Table 4: Welfare Group Shares under different Information Structures

	Efficient Exclusion	Inefficient Exclusion	Efficient Trade	Inefficient Trade
A) No Information	4.70	55.28	36.42	3.60
B) Baseline	4.31	41.78	49.93	3.99
C) Full Information	8.30	0.00	91.70	0.00

Notes: This table shows the welfare group shares under the the baseline, credit reporting environment (B) and the two extreme counterfactual information environments, no information (A) and full information (C). Under full information all efficiencies can be eliminated and there are no inefficient exclusions or inclusions. For a full discussion of these welfare group shares see section 8

Online Appendix

A Equilibrium existence

An economy is a collection $E = [\Theta, \mathcal{D}, \mathcal{X}, \mathcal{V}, \mathcal{C}]$ where Θ is a type space, \mathcal{D} is a type dynamic, \mathcal{X} is a choice space and \mathcal{V} is a payoff structure and \mathcal{C} is a cost structure. Furthermore, an information environment is a collection $I = [H, X(H), \mathcal{G}]$ where H is a history space, $X(H)$ history-specific choice space and \mathcal{G} is a history updating function. The following equilibrium existence proof borrows some of its notation and structure from [Azevedo and Gottlieb \(2017\)](#), though we do not need to introduce a vanishing mass of behavioral high-cost types as in their proof; the unbounded support of consumer taste shocks ensures that all available contracts are traded even without the presence of behavioral types. The key innovation in our proof is generalizing the notion of an allocation from the static setup to a stationary equilibria in a dynamic context.

Definition 1: Given an E and an information environment I , the pair (p^*, α^*) is an equilibrium if

1. For each contract x , firms make no profits. Formally,

$$p^*(h, x) = \mathbb{E}_x[c|\alpha^*] = \mathbb{E}[c(\tilde{x}, \tilde{\theta})|\alpha^*, \tilde{x} = x]$$

2. Consumers select contracts optimally. Formally, for almost every $(\theta, h, x(h)) \in \Theta \times H \times X(H)$ with respect to α^* , we have

$$V(\theta, h, \omega, p^*) = \sup_{\omega' \in \Omega} \int v(\theta, h, \omega, p^*) + \beta E_{d,g}[V(\theta, h, \omega, p^*)] dF(\epsilon_{x,\theta}), \quad \epsilon_{x,\theta,i} \sim GEV$$

Step 0: Define tâtonnement correspondence

Let \mathcal{P} be the set of all possible price vectors in $[0, 1]^{H \times X(H)}$ and define the tâtonnement correspondence $T : \mathcal{P} \times \mathcal{A} \rightrightarrows \mathcal{P} \times \mathcal{A}$ in terms of two maps $T(p, \alpha) = \Phi(\alpha) \times \Psi(p)$ where

$$\begin{aligned} \Phi(\alpha) &= \{p \in \mathcal{P} : p(x) = \mathbb{E}_x[c|\alpha^*] \forall x \in X\} \\ \Psi(p) &= \arg \max_{\alpha \in \mathcal{A}} \int V(\theta, h, \omega, p^*) d\alpha \end{aligned}$$

The fixed points of T correspond to the equilibria of the economy. To see this, note that $p \in \Phi(\alpha)$ is equivalent to firms making 0 profits, and $\alpha \in \Psi(p)$ is equivalent to the consumers optimizing. Therefore, (p^*, α^*) is a fixed point of T if and only if (p^*, α^*) is an equilibrium. We will now prove the existence of a fixed point. The proof has three steps.

Step 1: $\Phi(\alpha)$ is nonempty, convex-valued, and has a closed graph.

Nonempty and convex-valued are trivial properties of Φ . To establish that it has a closed graph, consider a sequence $(\alpha^n, p^n)_{n \in \mathbb{N}}$ in the graph of Φ with limit (α, p) . We will show that $p(x)$ is a conditional expectation of cost given α . To see this, take an arbitrary set $S \subseteq X$. Let $\tilde{S} = \Theta \times S$. We have

$$\begin{aligned}
\int_{\tilde{S}} p(x) d\alpha &= \sum_{x \in S} p(x) \cdot \alpha(\Theta \times x) \\
&= \lim_{n \rightarrow \infty} \sum_{x \in S} p^n(x) \cdot \alpha^n(\Theta \times x) \\
&= \lim_{n \rightarrow \infty} \int_{\tilde{S}} p^n(x) d\alpha^n \\
&= \lim_{n \rightarrow \infty} \int_{\tilde{S}} c(x, \theta) d\alpha^n \\
&= \int_{\tilde{S}} c(x, \theta) d\alpha
\end{aligned}$$

Step 2: $\Psi(p)$ is nonempty, convex valued, and has a closed graph.

Clearly, Ψ is a much more complicated object as it maps a price vector p into all the type dynamics d which admit an allocation α that is optimal given the price vector. *Nonempty:* Ψ is nonempty because X is finite, and therefore $U(x, p(x), \theta)$ attains a maximum for every θ . *Convex-valued:* Ψ is trivially convex-valued because it is single-valued. *Closed graph:* Consider a sequence $(p^n, \alpha^n)_{n \in \mathbb{N}}$ in the graph of Ψ with limit (p, α) . For any $\alpha' \in \mathcal{A}$, we have

$$\int V(\theta, h, \omega, p^n) d\alpha' \leq \int V(\theta, h, \omega, p^n) d\alpha^n$$

Taking the limit, we have

$$\int V(\theta, h, \omega, p) d\alpha' \leq \int V(\theta, h, \omega, p) d\alpha$$

given that the LHS limit follows from the Dominated Convergence theorem while the RHS limit follows directly from

$$\int V(\theta, h, \omega, p^n) - V(\theta, h, \omega, p) d\alpha^n + \int V(\theta, h, \omega, p) d\alpha^n$$

The first integrand converges to 0 uniformly in x and θ because X is finite, and hence p^n converges uniformly to p , and because the continuous function U is uniformly continuous in the compact set where price vectors belong to the image of c . Therefore, the first integral converges to 0. The second integral converges to $\int V(\theta, h, \omega, p) d\alpha$ by the continuity of U and weak convergence of α^n to α .

Step 3: Existence of a fixed point.

The claims about Φ and Ψ imply that T is convex valued, nonempty, and has a closed graph. We have that the set $P \times D$ is compact, convex, and a subset of a locally convex topological vector space. Therefore, by the Kakutani–Glicksberg–Fan theorem, T has a fixed point. *Q.E.D.*

B Parametrization, Implementation and Computational Details

B.1 Consumer default

This subsection provides further details on the default process as described in section 2.3. For a given consumer with risk type θ_r who borrows both secured and unsecured, mild and severe delinquencies on both secured and unsecured borrowing are modeled as four independent Bernoulli draws

$$\begin{aligned}\tilde{d}_t^{m,u} &\sim \text{Bernoulli}(\omega^{m,u}(\theta_r)) \text{ with } \omega^{m,u}(\theta_r) = \frac{\exp(\theta_r)}{1 + \exp(\theta_r)} \\ \tilde{d}_t^{s,u} &\sim \text{Bernoulli}(\omega^{s,u}(\theta_r)) \text{ with } \omega^{s,u}(\theta_r) = \frac{\exp(\theta_r + \alpha_{1,r})}{1 + \exp(\theta_r + \alpha_{1,r})} \\ \tilde{d}_t^{m,s} &\sim \text{Bernoulli}(\omega^{m,s}(\theta_r)) \text{ with } \omega^{m,s}(\theta_r) = \frac{\exp(\theta_r + \alpha_{2,r})}{1 + \exp(\theta_r + \alpha_{2,r})} \\ \tilde{d}_t^{s,s} &\sim \text{Bernoulli}(\omega^{s,s}(\theta_r)) \text{ with } \omega^{s,s}(\theta_r) = \frac{\exp(\theta_r + \alpha_{3,r})}{1 + \exp(\theta_r + \alpha_{3,r})}\end{aligned}$$

while consumers that don't borrow unsecured only realize the two Bernoulli draws $\tilde{d}_t^{m,s}$ and $\tilde{d}_t^{s,s}$ corresponding to their default potential on their secured borrowing.

The consumer risk type θ_r has therefore a very precise interpretation: The risk type θ_r of a consumer is the logistic transform of his probability of having a mild delinquency if he borrows unsecured at any period in time. The additional parameters $\alpha_{1,r}$, $\alpha_{1,2}$ and $\alpha_{1,3}$ capture then the additional risk of becoming severely delinquent on his unsecured ($\alpha_{1,r}$) and mildly ($\alpha_{1,2}$) or severely ($\alpha_{1,3}$) delinquent on his secured borrowing relative to the risk of becoming mildly delinquent on unsecured borrowing.

We impose no parametric assumption on the individual binary default outcomes by taking the consumer's secured and unsecured default realization as independent binary random variables and allowing for separate Bernoulli probabilities for each of the four potential default outcomes. However, when estimating the model we impose some structure on the mapping from default types into the Bernoulli distributions of the default outcomes by assuming that the mapping from the risk type into the different Bernoulli probabilities is fixed across risk types. That is to say, we take $\alpha_{1,r} = \alpha_{1,s}$, $\alpha_{2,r} = \alpha_{2,s}$ and $\alpha_{3,r} = \alpha_{3,s}$ for all risk types r and s . Finally, we aggregate in the data and in the model the mild and severe delinquencies into a single default state for both secured and unsecured borrowing $k \in \{u, s\}$:

$$\tilde{d}_t^k = \begin{cases} 0 & \text{if } \text{o/w} \\ 1 & \text{if } \tilde{d}_t^{m,k} = 1 \\ 2 & \text{if } \tilde{d}_t^{s,k} = 1 \end{cases}$$

We therefore capture how consumers' default and delinquency behavior at any point in time provides valuable information to lenders about the riskiness of consumers.

B.2 Scorecard Pricing

Finally, we note that we parameterize the pricing of histories by using a scorecard that maps each element of a consumer's history into a price component. That is to say, each element in the history has a fixed contribution to the price in that history that is unaffected by the other elements in the history. Instead of solving for the price in each of the 36 histories, the introduction of the scorecard S reduces the dimensionality of the pricing problem to a vector of 10 entries:

$$S = [s_{B_1=0}, s_{B_1=1}, s_{B_2=0}, s_{B_2=1}, s_{D_1=0}, s_{D_1=1}, s_{D_1=2}, s_{D_2=0}, s_{D_2=1}, s_{D_2=2}] \quad (7)$$

The price $p(h)$ in history $h = [B_1, B_2 | D_1, D_2]$ is then equal to the sum of the four corresponding scorecard elements

$$p(h) = p([B_1, B_2 | D_1, D_2]) = s_{B_1} + s_{B_2} + s_{D_1} + s_{D_2} \quad (8)$$

It is straightforward to add additional interaction terms in the scorecard to increase the flexibility of the scorecard. A saturated scorecard would obviously be equivalent to fully flexible pricing. We experimented with increased scorecard flexibility but found that violations in the zero profit conditions tend to be small even without the use of any interaction.

This use of the scorecard is motivated on two levels: First, it is close to the actual pricing technology used in the credit card industry. Second, the use of a scorecard significantly reduces the problem of multiple equilibria by imposing restrictions on the prices across histories. By requiring that prices in histories with recent defaults are not higher than histories with no default, the scorecard eliminates equilibria in which the prices in histories with no default are so high that no consumer would want to borrow. The restriction imposed by the scorecard that the change in a single scorecard element affects prices in all histories with that scorecard element in the same way prevents similarly unraveling in individual histories.

B.3 Estimation Algorithm

1. Given a parameter vector $\gamma = (\lambda_d, \lambda_r, \sigma_d, \sigma_r, \rho, \mu_d, \mu_r, \zeta, \alpha_1, \alpha_2, \alpha_3)$, discretize the AR(1) process
2. Guess initial scorecard price vector s^i and initialize first choice-specific value function guess v^0

3. Inner loop iteration X

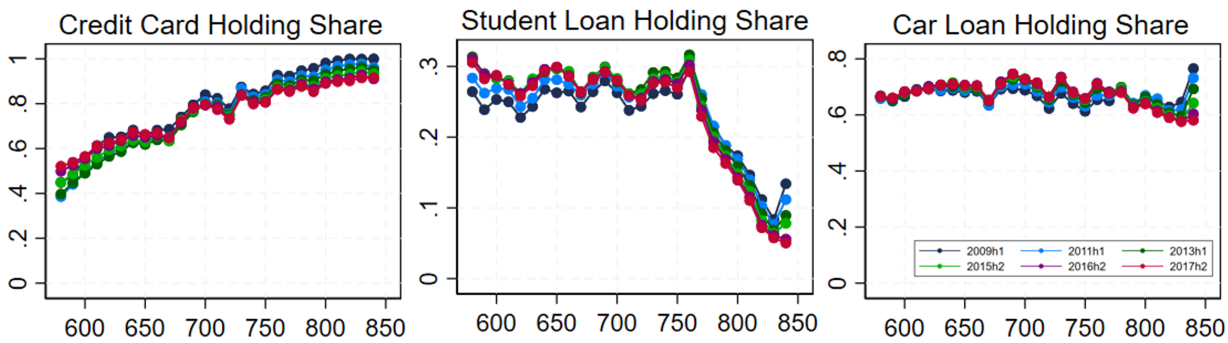
- 3.0 Given a guess for the scorecard price vector s^i , calculate implied price vector $p^i(h)$

- 3.1 Use initial guess for the choice-specific value function V^{i-1} and iterate on value function given the price vector $p^i(h)$
- 3.2 Construct history-type transition matrix based on the optimal borrowing choices as implied by the choice-specific value functions. Derive stationary distribution across histories
- 3.3 Find new optimal scorecard vector s^{i+1} that minimizes the distance across all histories between the average default rate of borrowers in that history (as captured in the stationary distribution $q(h, \theta)$) choosing to borrow at the given price.
- 3.4 If $d(s^{i+1}, s^i) > tol$ return back to 3.0.

4. Calculate model moments

C Appendix Figures and Tables

Appendix Figure 1: Secured and Unsecured Borrowing



Notes: This figure shows across the credit score distribution the share of consumers holding a given credit product. Consumers are grouped based on their credit score into buckets from 580-589, 590-599, to 840-849 and 850 and above. The first panel shows the share of consumers holding a credit card within each bucket. The second panel shows the share of consumers holding any student loan within each credit score bucket. The third panel shows the share of consumers holding any car loan within each credit score bucket. For each panel, these shares are shown at six points in time between the first half of 2009 and the second half of 2017 with a different line for each time period to demonstrate the stability of these holding shares over time.