Regulatory Intervention in Consumer Search Markets: The Case of Credit Cards*

Manolis Galenianos¶ Alessandro Gavazza§

August 2018

Abstract

Data on U.S. credit card markets display a large dispersion of interest rates at which consumers borrow. To understand this dispersion, we build a search model with two novel features: search effort/inattention and product differentiation. We calibrate the model to match statistics on the interest rate distribution that borrowers pay. The model fits these data well. Our analysis implies that low search effort accounts for almost all the dispersion in interest rates, whereas product differentiation is negligible. We use the calibrated model to study regulatory interventions in credit markets, such as caps on interest rates and higher compliance costs for lenders.

PRELIMINARY AND INCOMPLETE

*We thank Matthijs Wildenbeest for useful comments and discussions. William Matcham provided excellent research assistance. Alessandro Gavazza gratefully acknowledges support from the European Research Council (ERC-Consolidator grant award no. 771004).

¶Department of Economics, Royal Holloway, University of London. Egham Hill, Egham TW20 0EX, United Kingdom. Email: manolis.galenianos@gmail.com.

§Department of Economics, London School of Economics. Houghton Street, London WC2A 2AE, United Kingdom. Email: a.gavazza@lse.ac.uk.

1 Introduction

After the financial crisis, legislators and regulators in several countries have been intervening in markets for consumer financial products more aggressively than before. For example, in the United States the Dodd-Frank Act of 2010 established the Consumer Financial Protection Bureau, with the mission to supervise and regulate financial products for households, such as mortgages and credit cards. In the United Kingdom, the Financial Service Act of 2012 created a new regulatory framework for financial services, establishing the Financial Conduct Authority (FCA), with regulatory powers related to the marketing of financial services.

These regulations have taken many different forms, depending on the specific countries and on the specific products. In some cases, they have directly constrained the prices and the fees of some financial products, which have been either capped or banned, as many regulators viewed these fees as "predatory"—i.e., targeting unsophisticated and poorly-informed households—thereby increasing their debt levels. Specifically, the 2009 U.S. Credit Card Accountability Responsibility and Disclosure Act explicitly prohibited lenders from charging some fees on credit cards (Agarwal et al., 2015b). Similarly, in the U.K. the FCA has introduced regulatory caps for several financial products: in November 2014 it enacted a price structure for payday loans, capping the initial cost of a loan to a maximum of 0.8 percent per day; in November 2016, it restricted fees for individuals who want to access their pensions to a maximum of one percent. Furthermore, the FCA is currently evaluating limits on fees for other products, such as mutual fund fees (*The Financial Times*, Funds' lucrative entry fees under attack, May 26, 2016) and mortgage origination fees (*The Financial Times*, Mortgage lenders under FCA review for masking high fees, December 12, 2016).

In several other cases, regulatory agencies cracked down on lenders by increasing capital requirements and/or tightening enforcement, thereby increasing their operating costs.

The broad goal of this paper is to study these two main form of regulations—i.e., the regulation of fees and of their operating costs—in markets for consumer financial products, with a special focus on how these regulations affect consumers' incentives to acquire information about these products. More specifically, many of these price regulations are predicated on the assumptions that consumers may be poorly informed about some of these fees. Hence, we develop a modeling framework that explicitly considers these information frictions. Indeed, Sirri and Tufano (1998) and Hortaçsu and Syverson (2004) argue that information frictions play a prominent role in mutual fund markets, and Allen et al. (Forthcoming) and Woodward and Hall (2012) show their relevance in mortgage markets. Search theory allows exactly to incorporate these information frictions, also providing a flexible framework that has been successfully used for structural estimation.[1]

We tailor our model to the U.S. credit card market. Specifically, Stango and Zinman (2016) (henceforth SZ) report a large dispersion in the interest rates that consumers pay on their credit cards, and they document that consumer characteristics—most notably, their creditworthiness, as captured by their credit score—does not account for this large heterogeneity. Hence, to interpret this dispersion, we build a search model with two key features: search effort/inattention and product differentiation. We calibrate the model to match the statistics on the interest rate distribution that SZ report. The model fits these data reasonably well. Our analysis implies that low search effort accounts for almost all the dispersion in interest rates, whereas product differentiation is negligible.

We further use the calibrated model to understand the role of price caps and of higher operating costs on equilibrium outcomes. Specifically, Fershtman and Fishman (1994) and Armstrong et al. (2009) show that, in markets with search frictions *à la* Burdett and Judd (1983), price caps may have the unintended consequences of *increasing* the equilibrium prices paid by consumers. More precisely, they identify two opposing effects: 1) the direct effect of regulation is to reduce prices for uninformed consumers who, before the regulation, were paying high prices; and 2) the indirect effect is to reduce price dispersion, which reduces consumers' incentives to acquire information about prices, thereby increasing suppliers' market power and, thus, prices. Armstrong et al. (2009) further show that, if consumers are heterogenous in their costs of acquiring information, the introduction of a price cap has an ambiguous effect on the equilibrium price paid by consumers, thereby leading to the possibility that equilibrium prices may increase. Therefore, it is an empirical/quantitative question which of the two opposing effects dominates and, thus, whether or not price caps have the intended consequences. The calibrated model will allow us to determine which of the two opposing effect dominates and, thus, whether or not price caps are beneficial to consumers.

We further use the calibrated model to simulate alternative market structures through higher operating (i.e., entry) costs. Specifically, an interesting question in markets with search frictions is how the entry of new suppliers affects consumers' information acquisition, their search process, and, thus, welfare. Most notably, the insightful contribution of Janssen and Moraga-González (2004) shows that an increase in the number of entrants could reduce search, thereby leading to greater price dispersion and lower welfare. Hence, counter-

[1] Some recent papers that structurally estimate search models of consumer product markets include Hortaçsu and Syverson (2004); Hong and Shum (2006); Wildenbeest (2011); Allen et al. (Forthcoming); Gavazza (2016); Galenianos and Gavazza (2017).

factual simulations will help us to understand the empirical relevance of these considerations.

2 Related Literature

The paper contributes to several strands of the empirical literature. The first is the literature that studies credit card companies' market power. In an important contribution, Ausubel (1991) showed that interest rate on credit cards are substantially higher than lenders' funding costs and display limited intertemporal variability, citing search frictions as a potential departure from a competitive market. Calem and Mester (1995) present empirical evidence on consumers' limited search and switching behavior. Stango (2002) studies credit card pricing when consumers have switching costs. Grodzicki (2015) analyzes how credit card companies acquire new customers. We contribute to this literature by building a framework that allows us to quantify the effects of search frictions and consumer inertia on lenders' loan pricing and on consumers' cost of borrowing.

Second, a vast literature in household finance studies whether or not consumers behave optimally in credit markets: among others, Agarwal et al. (2008) and Agarwal et al. (2015a) analyze consumer mistakes in the credit card market. Ru and Schoar (2016) studies how credit card companies exploit consumers' mistakes. In this strand of literature, the most related paper is Woodward and Hall (2012), which studies consumers' shopping effort in the U.S. mortgage market. We contribute to this literature by developing and calibrating an equilibrium model in which consumers' shopping effort is endogenous, which allows us to analyze how it adjusts after regulatory interventions.

Third, many countries have recently enacted reforms and introduced new regulations in markets for consumer financial products (Campbell et al., 2011a,b). Several recent contributions provide descriptive analyses of the effects of these reforms. In the specific case of the U.S. credit card market, Agarwal et al. (2015b) and Nelson (2018) analyze how regulatory limits on credit card fees introduced by the 2010 CARD Act affect lender pricing and borrowing costs exploiting rich data. We complement these papers by analyzing some of these regulatory interventions in a quantitative model which focuses on key features—i.e., search frictions and consumer inattention—that account for pricing in the credit card market.

Finally, this paper is related to the literature on the structural estimation of consumer search models. Recent contributions include Hortaçsu and Syverson (2004), Hong and Shum (2006), Wildenbeest (2011), Allen et al. (Forthcoming), Galenianos and Gavazza (2017), and Salz (2016). We innovate on these previous empirical papers by incorporating in our model—

and, thus, by evaluating the quantitative importance of—additional features, such as endogenous search effort, that, according to some theoretical papers, could potentially offset the effects of the regulations that we study (Fershtman and Fishman, 1994; Armstrong et al., 2009; Janssen and Moraga-González, 2004).

3 Data

The available data dictate some of the modeling choices of this paper. For this reason, we describe the data before presenting the model. This description also introduces some of the identification issues that we discuss in more detail in Section 5.2.

3.1 Data Sources

Our quantitative analysis combines several sources of data. More specifically, we exploit some of the datasets that SZ use in their descriptive analysis of households' credit card terms, supplementing them with some aggregate statistics obtained from the Federal Reserve Board and the Survey of Consumer Finances. We now describe these datasets in more detail.

The first dataset is an account-level panel that samples individuals and reports the main terms of their credit-card accounts during (at most) 36 consecutive months between January 2006-December 2008, including their credit limit, the end-of-month balance, the revolving balance, the annual percentage rate (APR), and the cash advance APR. The dataset also reports limited demographic characteristics of the cardholders, such their household income and their FICO credit score.[2]

The second dataset reports the terms of all credit card offers that a sample of individuals receive in January 2007. This second dataset samples different individuals than those in the first dataset, but allows us to measure the dispersions in offers that individuals receive in a given month. As SZ emphasize, the dispersion in interest rates on all credit card offered to a given individual in a given month removes any effect of individual-specific factors on the cross-sectional distribution of interest rates on credit cards that individuals hold. We should point out that we do not have access to the individual survey data and, thus, we exploit data reported in tables of SZ.

We complement these datasets with some aggregate statistics: the fraction of individuals with credit card debt, computed from the 2007 Survey of Consumer Finances; the aggregate

[2]We are grateful to Victor Stango for sharing this dataset with us.

charge-off rate on credit card loans in the first quarter of 2007, reported by the Federal Reserve Board; and the risk-free rate, computed as the interest rate of the one-month Treasury bill on January 16th, 2007.[3]

3.2 Data Description

We use the first dataset on individuals' credit-card terms to sum up one of the main results of SZ's descriptive analysis: a large dispersion of the interest rate distribution persists, even after taking into account 1) different default risk across individuals, as measured by their FICO scores; 2) different revolving balances across borrowers; and 3) different card characteristics across borrowers, such as rewards. Specifically, the basic framework for this analysis is the following equation:

$$R_{ijt} = \gamma_X X_{it} + \gamma_Z Z_{ijt} + \epsilon_{ijt}, \tag{1}$$

where the dependent variable R_{ijt} is the APR that individual i pays on credit card j in month t; X_{it} are characteristics of individual i in month t, such as his default risk, measured by the FICO score; Z_{ijt} are characteristics of individual i's credit card j in period t, such as the credit limit, rewards, and the credit balance; ϵ_{ijt} are residuals.

The top panel of Table 1 reports the coefficient estimates of several specifications of equation (1); the bottom panel reports selected percentiles of the distribution of interest rates calculated as:

$$R'_{ijt} = \hat{\gamma}_X \overline{X}_{it} + \hat{\gamma}_Z \overline{Z}_{ijt} + \hat{\epsilon}_{ijt},$$

where $\hat{\gamma}_X$ and $\hat{\gamma}_Z$ are the coefficient estimates, $\overline{X}_{it}$ and $\overline{Z}_{ijt}$ are the sample averages of the covariates of each regression, and $\hat{\epsilon}_{ijt}$ are the estimates of the residuals.

Column (1) uses the raw data over the entire sample period, which exhibit a large dispersion of interest rates, i.e., the difference between the 90th and the 10th percentiles equals 18.5 percentage points. Column (2) restricts the data to January 2007 (this is the date of our other data sources), showing that the large dispersion of interest rates is almost identical to that in (1), for two reasons: a) there is limited aggregate variation in interest rates over time; and b) there is limited within-account variation in interest rates. Column (3) further restricts the data to cards without introductory "teaser" rates (i.e., low initial rates that reset to higher rates after an initial offer period); of course, interest rates increase relative to those displayed in column (2), but the increase is minimal and the difference between the

[3] We retrieved the last two values from FRED, Federal Reserve Bank of St. Louis, series https://fred.stlouisfed.org/series/CORCCACBS and https://fred.stlouisfed.org/series/DGS1MO, respectively.

Table 1: Dispersion of Interest Rates

	(1)	(2)	(3)	(4)	(5)	(6)
FICO SCORE				-0.024***	-0.022***	-0.021***
				(0.001)	(0.001)	(0.002)
REWARD CARD					0.260	0.135
					(0.178)	(0.273)
CREDIT LIMIT					-0.047***	-0.046*
					(0.013)	(0.024)
CREDIT BALANCE					0.065***	0.004
					(0.022)	(0.032)
R^2	0.749	0.752	0.789	0.127	0.130	0.092
OBSERVATIONS	382,567	9,167	8,575	4,375	4,299	2,578
PERCENTILES						
10TH	9.49	9.9	11.9	12.37	12.36	12.68
25TH	12.99	13.99	14.24	15.24	15.19	15.84
50TH	16.74	17.65	18.24	18.49	18.43	19.31
75TH	20.24	20.24	20.99	21.45	21.62	23.82
90TH	27.99	27.99	28.15	26.62	26.6	28.60

Notes: This table reports.

90th and the 10th percentiles slightly decreases to 16.25 percentage points. The specification of column (4) further controls for the credit risk of the individual through the FICO score (unfortunately, this is available for only 4,375 individuals out of 8,575); the corresponding distribution of interest rates still display a large dispersion. The specification of column (5) further controls for other card characteristics, such as the credit limit and an indicator variable which equals one if the card features some rewards (e.g., frequent flier miles or cash back) and zero otherwise, as well as the revolving balance. The specification of column (6) further restricts the sample to cards with a revolving balance (i.e., cards used for borrowing beyond the 25-day grace period), showing that the large dispersion of interest rates persists, i.e., the difference between the 90th and the 10th percentiles equals 15.92 percentage points.

Overall, Table 1 confirms that borrower and card characteristics account for a small fraction of the overall dispersion in interest rates that borrowers pay on their credit card, thereby suggesting that frictional dispersion is a pervasive feature of the credit card market. Hence, the model that we develop in Section 4 aims to bring about this dispersion through information frictions, and the calibration of Section 5 aims to quantitatively match the percentiles of specification (6) of Table 1.

Table 2: Empirical Targets

Panel A: Accepted Offers	
10th Percentile Accepted Offer Distribution	12.68
25th Percentile Accepted Offer Distribution	15.84
50th Percentile Accepted Offer Distribution	19.31
75th Percentile Accepted Offer Distribution	23.82
90th Percentile Accepted Offer Distribution	28.60
Panel B: Received Offers	
Fraction Receiving 2+ Offers (%)	75.00
Median Number of Offers Received, Conditional on 2+ Offers	3.00
Average Number of Offers Received, Conditional on 2+ Offers	4.00
10th Percentile Distribution of Differences in Offered Rates	0.00
30th Percentile Distribution of Differences in Offered Rates	2.25
50th Percentile Distribution of Differences in Offered Rates	4.34
70th Percentile Distribution of Differences in Offered Rates	7.25
90th Percentile Distribution of Differences in Offered Rates	9.25
Panel C: Aggregate Statistics	
Fraction with Credit Card Debt	36.70
Charge-Off Rate	4.01

Notes—This table provides the empirical targets of our calibrated model. Panel A reports statistics on the interest rates that borrowers pay on their credit cards. Panel B displays statistics on credit card offers that SZ report. Panel C reports auxiliary statistics.

Table 2 combines all empirical targets of our quantitative model. Panel A reproduces the percentiles of the distribution of interest rates derived in Table 1. Panel B reports statistics on credit card offers that SZ document. Specifically, Section 5.1 of SZ recounts that approximately 75 percent of individuals received more than one credit card offer during January 2007 and, among them, the median and the mean number of offers was three and four, respectively; for these individuals who received at least two offers, Table 4 of SZ reports key percentiles of the distribution of the difference between the highest and lowest offered interest rates at which balances incur interest charges after the expiration of any introductory "teaser" period (if any). Finally, Panel C reports auxiliary statistics on credit card markets, such the fraction of individuals with credit card debt in the 2007 Survey of Consumer Finances and the aggregate charge-off rate on credit card loans in the first quarter of 2007.

Table 2 provides a rich description of the credit card market, with some striking patterns. First, as we noted above, Panel A shows that the dispersion in the interest rates that borrowers pay on their credit card debt is very large, even after we control for observable borrower and card characteristics; this dispersion is informative of the extent of search frictions. Second, Panel B points out that many individuals receive multiple credit card offers at substantially different interest rates. Moreover, the within-individual dispersion in received offers in Panel B and cross-sectional dispersion in Panel A seem to suggest that many borrowers do not accept the offers with the lowest interest rates. Our model seeks to capture these features in two different ways: 1) consumers may not pay attention to all the offers that they receive; and 2) consumers may have idiosyncratic preferences for unobservable card features. In summary, these data seem to call for investigating the role of search frictions, search effort/inattention, and idiosyncratic preferences in the credit card market.

Despite all of their advantages, however, these data pose some challenges. First, they are mostly cross-sectional, and, therefore, we do not observe borrowers' and lenders' behavior over time. Specifically, we do not observe how frequently borrowers switch across credit cards. Hence, in the absence of more-detailed measurement on borrowers' switching behavior, we will seek to match the cross-sectional distribution through a static model. Moreover, while the theory can accommodate multidimensional heterogeneity of borrowers and/or lenders, our cross-sectional data make it difficult to identify such a model. Thus, we focus on a parsimonious framework with borrowers' heterogeneity in their willingness to pay for credit and lenders' heterogeneity in their funding costs, and we let other parameters be common across individuals. We discuss the implications of these data limitations for our results further in Section 7.

4 The Model

The economy is populated by measure 1 of borrowers and measure Λ of potential lenders. Borrowers are heterogeneous in their marginal valuation of a loan z, which is distributed according to a smooth distribution $M(\cdot)$ with connected support $[\underline{z}, \overline{z}]$. Potential lenders are heterogeneous in their marginal cost of providing a loan k, which is distributed according to a smooth distribution $\Gamma(\cdot)$ with connected support $[\underline{k}, \overline{k}]$. Potential lenders can pay cost χ to enter the market, where L and $G(\cdot)$ denote the measure of entrants and the distribution of their marginal costs, respectively. Every borrower wants to take a loan of size b to fund consumption and every lender has one loan of size b to give.[4]

Matching between borrowers and lenders is subject to frictions. Each lender sends one loan offer with an associated interest rate to a randomly-chosen borrower. Each borrower receives a random number of offers, examines every offer with some probability which depends on his search effort and decides which offer to accept.[5] Specifically, the number of offers received by a borrower follows a Poisson distribution with parameter L (the lender-borrower ratio) and the borrower examines each offer with independent probability $s \in [0,1]$, where s is his search effort. The borrower's effective number of offers, therefore, follows a Poisson distribution with parameter $s * L$. A borrower who exerts search effort s incurs cost $q(sL)$ which is strictly increasing, strictly convex and satisfies $q(0) = q'(0) = 0$.

Each loan offer is characterized by its cost to the borrower c, which has two components. The first component is the interest rate R that is chosen by the lender. The second component is an idiosyncratic utility draw e which is i.i.d. across borrowers and represents every other aspect of the loan that might affect the borrower's valuation. R is drawn from the distribution $F_R(\cdot)$ (to be determined in equilibrium) with support $[\underline{R}, \overline{R}]$ and e is drawn from a smooth distribution $F_e(\cdot)$ with zero mean and support in a connected subset of $(-\infty, \infty)$. The overall cost of a loan is $c = R+e$ and might be greater or lower than R depending on the idiosyncratic draw.

A type-z borrower's utility from taking a loan with interest rate R and idiosyncratic draw e is

$$b\big(z - R - e\big)$$

and his utility from not taking a loan is zero. Anticipating equilibrium behavior, a type-z borrower chooses the loan offer with the lowest overall cost among the offers that he examines, conditional on the cost being less than z. A loan offer with $R + e > z$ generates negative utility and is, thus, never accepted.

The ex ante value of a type-z borrower is equal to the expected value of his best loan $V_z(s)$ (which depends on search effort s) net of the cost of search effort, $q(sL)$:

$$V_z(s) - q(sL) \tag{2}$$

[4]A lender should be interpreted as a loan contract rather than a lending company (say, a credit card company). We do not model credit card companies explicitly.

[5]The random allocation of offers across borrowers leads to urn-ball matching which, for large numbers of borrowers and lenders, is described by a Poisson distribution. See Butters (1977).

We denote the optimal effort choice of a type-z borrower by $s(z)$.

A lender's revenues per-dollar-lent are equal to the interest rate R net of a cost $\rho(R)$ that represents the risk that the borrower does not repay. We assume that $\rho(0) = 0$, $\rho'(R) > 0$ and $\rho''(R) \geq 0$ to capture the observation that borrowers are more likely to default on higher interest-rate loans, possibly because of informational asymmetries (e.g. (Adams et al., 2009; Einav et al., 2012)) which, however, we will not model explicitly.

The lender's profit conditional on giving a loan is equal to revenues minus the lender's marginal cost, k. The expected profits of a type-k lender who offers interest rate R, $\pi_k(R)$, are given by the probability of making a loan, denoted by $P(R)$, times the loan's profits:

$$\pi_k(R) \quad = \quad b\Big(R(1-\rho(R)) - k\Big)P(R). \tag{3}$$

Notice that the borrower's idiosyncratic shock affects the lender's payoff only through the probability of making a loan, $P(R)$.

We denote the optimal interest rate choice of a type-k lender by $R(k)$ which, combined with lenders' entry decisions, determines the interest rate distribution $F_R(\cdot)$.

We are now ready to define the equilibrium.

Definition 1 *An equilibrium consists of borrowers' search effort* $\{s(z)\}$ *and lenders' entry and interest rate choices* $\{L, G(\cdot), F_R(\cdot)\}$ *such that borrowers maximize their ex ante value* (2)*, lenders maximize their expected profits* (3)*, there is free entry of lenders and the expected profits of all entrants exceed the entry cost* χ*.*

We proceed by determining the borrowers' and lenders' optimal choices separately and then proving the existence of equilibrium.

4.1 Borrowers' choices

We characterize borrowers' optimal search effort $s(z)$ taking as given the measure of lenders in the market L and some interest rate distribution $F_R(\cdot)$ (the type distribution of lenders $G(\cdot)$ does not directly affect borrowers' choices).

We begin by rewriting $V_z(s)$ in a more convenient way. Denote the value to a z-borrower of receiving n offers by $v_z(n)$, where $v_z(0) = 0$. The expected value of loan offers for a type-z

borrower who exerts search effort s is:

$$V_z(s) \;=\; \sum_{n=0}^{\infty} \frac{e^{-sL}(sL)^n}{n!} v_z(n) \tag{4}$$

Note that search effort only affects the arrival rate of offers and does not enter $v_z(n)$.

To determine $v_z(n)$ for $n \geq 1$, recall that the borrower chooses the loan offer with the lower overall cost c. Let $F_c(\cdot)$ denote the distribution of c. Since the overall loan cost c is the sum of two independent random variables (R and e) it is distributed according to

$$F_c(c) = \int_{-\infty}^{\infty} F_R(c-e)dF_e(e) = \int_{\underline{R}}^{\overline{R}} F_e(c-R)dF_R(R)$$

The distribution of the *lowest* cost out of $n \geq 1$ draws from $F_c(\cdot)$ is:

$$\bar{F}_c(c|n) \;=\; 1 - \big(1 - F_c(c)\big)^n$$

Therefore the value to a z-borrower of receiving $n \geq 1$ offers is:

$$v_z(n) \;=\; b\int_{-\infty}^{z} \big(z - c\big) d\bar{F}_c(c|n) \tag{5}$$

The optimal effort choice $s(z)$ solves:

$$V_z'(s) \;=\; q'(sL)L \tag{6}$$

The following proposition characterizes $s(z)$.

Proposition 2 *The optimal search effort of a type-z borrower, $s(z)$, is unique, continuous and strictly increasing and solves*

$$\sum_{n=0}^{\infty} \frac{e^{-sL}(sL)^n}{n!}\big(v_z(n+1) - v_z(n)\big) \;=\; q'(sL) \tag{7}$$

where $v_z(0) = 0$ and $v_z(n)$ is defined by equation (5) for $n \geq 1$.

Proof. We first show that $v_z(n)$ is strictly increasing and strictly concave in n. The cost distribution for a low n first order stochastically dominates that for a high n (proving that $v_z(n)$ is increasing in n) and the derivative of the cost distribution for a high n first order stochastically dominates that for a low n (proving concavity):

$$\begin{aligned}
\frac{d\bar{F}_c(c|n)}{dn} &= -\big(1-F_c(c)\big)^n \log\big(1-F_c(c)\big) > 0 \\
\frac{d^2\bar{F}_c(c|n)}{dn^2} &= -\big(1-F_c(c)\big)^n \Big(\log\big(1-F_c(c)\big)\Big)^2 < 0
\end{aligned}$$

Therefore $v_z(n+1) > v_z(n)$ and $v_z(n+2) - v_z(n+1) < v_z(n+1) - v_z(n)$ for all n.

Differentiating equation (4) with respect to s:

$$\begin{aligned}
V_z'(s) &= \sum_{n=1}^{\infty}\Big(-\frac{e^{-sL}(sL)^n}{n!}v_z(n)L + \frac{e^{-sL}(sL)^{n-1}}{(n-1)!}v_z(n)L\Big) \\
&= \Big(-\sum_{n=0}^{\infty}\frac{e^{-sL}(sL)^n}{n!}v_z(n) + \sum_{n=0}^{\infty}\frac{e^{-sL}(sL)^n}{n!}v_z(n+1)\Big)L \\
&= \sum_{n=0}^{\infty}\frac{e^{-sL}(sL)^n}{n!}\big(v_z(n+1) - v_z(n)\big)L > 0
\end{aligned}$$

As a result, the borrower's expected value of offers is strictly increasing in search effort.

Furthermore, the expected value of loan offers is strictly concave in search effort:

$$\begin{aligned}
V_z''(s) &= \sum_{n=1}^{\infty}\big(-\frac{e^{-sL}(sL)^n}{n!} + \frac{e^{-sL}(sL)^{n-1}}{(n-1)!}\big)\Big(v_z(n+1) - v_z(n)\Big)L^2 \\
&= \sum_{n=0}^{\infty}\frac{e^{-sL}(sL)^n}{n!}\Big(v_z(n+2) - v_z(n+1) - \big(v_z(n+1) - v_z(n)\big)\Big)L^2 < 0
\end{aligned}$$

Therefore, equation (6) has a unique solution $s(z)$ which yields the optimal search effort for a type-z borrower. Furthermore, the solution to that equation varies continuously with z.

Finally notice that:

$$\begin{aligned}
\frac{\partial v_z(n)}{\partial z} &= b\int_{-\infty}^{z} d\bar{F}_c(c|n) = b\Big(1 - \big(1-F_c(z)\big)^n\Big) > 0 \\
\Rightarrow \frac{\partial V_z(s)}{\partial z} &= \sum_{n=1}^{\infty}\frac{e^{-sL}(sL)^n}{n!}b\Big(1 - \big(1-F_c(z)\big)^n\Big) > 0
\end{aligned}$$

Therefore higher valuation borrowers put greater value to increasing arrival rates of lenders and therefore they exert more search effort. This completes the proof of proposition 2. ■

4.2 Lenders' choices

We characterize the optimal interest rate of a type-k lender ($R(k)$), aggregate the actions of lenders who enter the market to arrive at the interest rate distribution ($F_R(\cdot)$) and characterize the lenders' entry decisions (L, $G(\cdot)$), given borrowers' actions ($s(z)$).

To simplify notation, we denote the arrival rate of offers to a type-z borrower by $\alpha_z = s(z)L$.

We begin with a lemma that provides a partial characterization of entering lenders' types.

Lemma 3 *There is k^* such that lenders enter if only if $k \leq k^*$. k^* is lower than the interest rate which maximizes a loan's revenues.*

Proof. Equilibrium profits are strictly decreasing in k, as a low-cost lender can always offer the high-cost lender's interest rate and earn higher profits. Therefore, in equilibrium there is some cutoff k^* such that lender enter the market if and only if $k \leq k^*$.

Our assumption that $\rho(R)$ is increasing and convex imply that revenues conditional on giving a loan ($R(1-\rho(R))$) are maximized at some $\tilde{R}$. A potential lender with $k > \tilde{R}$ necessarily makes negative profits in the market and, therefore, $k^* \leq \tilde{R}$. ■

A loan with interest rate R is accepted if the borrower who receives it examines the offer, if the offer yields the lowest overall cost from every offer examined by the borrower taking into account the idiosyncratic component of all offers, and if the offer's overall cost is below the borrower's type z. The next lemma characterizes the probability that a loan with interest rate R is accepted, $P(R)$.

Lemma 4 *Given $F_R(\cdot)$, L and $s(z)$, the probability of making a loan when offering interest rate R is continuous and differentiable and it is given by:*

$$P(R) = \int_{\underline{z}}^{\overline{z}} s(z) \int_{-\infty}^{z-R} e^{-\alpha_z \int_{\underline{R}}^{\overline{R}} F_e(R+e-x)dF_R(x)} dF_e(e) dM(z) \tag{8}$$

Furthermore $P'(R) < 0$.

Proof. Denote the probability that a loan offer with total cost c is accepted by a type-z borrower by $P_c(c, z)$. If $z \geq c$ the offer is accepted if it is the lowest-cost offer received, which occurs with probability $\left(1 - F_c(c)\right)^n$ when the borrower examines n additional offers. If $c > z$

then the offer is not accepted. Therefore:

$$
\begin{aligned}
P_c(c,z) &= \sum_{n=0}^{\infty} \frac{e^{-\alpha_z}\alpha_z^n}{n!}\big(1-F_c(c)\big)^n \\
&= e^{-\alpha_z F_c(c)} \\
&= e^{-\alpha_z \int_{\underline{R}}^{\overline{R}} F_e(c-x)dF_R(x)}, \quad \text{if} \quad c \le z && (9) \\
P_c(c,z) &= 0, \qquad\qquad\qquad \text{if} \quad c > z && (10)
\end{aligned}
$$

Denote by $P_R(R,z)$ the probability that a type-z borrower accepts a loan offer with interest rate R. The offer is accepted if the loan's overall cost (including the idiosyncratic component) is less than z and if every other offer which is examined has higher cost. Integrating over the potential values of the idiosyncratic shock yields:

$$
\begin{aligned}
P_R(R,z) &= \int_{-\infty}^{\infty} P_c(R+e)dF_e(e) \\
&= \int_{-\infty}^{z-R} e^{-\alpha_z \int_{\underline{R}}^{\overline{R}} F_e(R+e-x)dF_R(x)} dF_e(e) && (11)
\end{aligned}
$$

A loan offer with interest rate R is accepted if the receiving borrower (of some type z) examines the offer (probability $s(z)$) and it is preferable to any other offer he examines (probability $P_R(R,z)$). Therefore, the probability a loan with interest rate R is accepted is:

$$
\begin{aligned}
P(R) &= \int_{\underline{z}}^{\overline{z}} s(z)P_R(R,z)dM(z) \\
&= \int_{\underline{z}}^{\overline{z}} s(z) \int_{-\infty}^{z-R} e^{-\alpha_z \int_{\underline{R}}^{\overline{R}} F_e(R+e-x)dF_R(x)} dF_e(e)dM(z) && (12)
\end{aligned}
$$

which yields equation (8).

Equation (12) shows that $P(R)$ is continuous and differentiable in R, since $F_e(\cdot)$ is assumed to be smooth. The probability of giving a loan is strictly decreasing the interest rate:

$$
\begin{aligned}
P'(R) = &-\int_{\underline{z}}^{\overline{z}} s(z)\Big[\int_{-\infty}^{z-R} e^{-\alpha_z \int_{\underline{R}}^{\overline{R}} F_e(R+e-x)dF_R(x)} \Big(\alpha_z \int_{\underline{R}}^{\overline{R}} F_e'(R+e-x)dF_R(x)\Big) dF_e(e) \\
&+ e^{-\alpha_z \int_{\underline{R}}^{\overline{R}} F_e(z-x)dF_R(x)} F_e'\big(z-R\big)\Big] dM(z) < 0.
\end{aligned}
$$

This completes the proof of lemma 4. ∎

We proceed to characterize the optimal interest rate schedule $R(k)$ and the distribution of interest rate offers $F_R(\cdot)$.

Proposition 5 *Given L, $G(\cdot)$ and $s(z)$, the lenders' optimal interest rate choice is characterized as follows:*

1. *The profit-maximizing interest rate of a type-k lender $R(k)$ is unique, continuous and increasing in k.*

2. *$R(k)$ solves the following functional equation:*

$$\int_{\underline{z}}^{\overline{z}} s_(z) \int_{-\infty}^{z-R(k)} e^{-\alpha_z \int_{\underline{k}}^{\overline{k}} F_e\left(R(k)+e-R(x)\right) dG(x)} dF_e(e) dM(z)$$
$$= \frac{R(k)(1-\rho(R(k)))-k}{1-\rho(R(k))-R\rho'(R(k))} \int_{\underline{z}}^{\overline{z}} s(z) \Big[\int_{-\infty}^{z-R(k)} e^{-\alpha_z \int_{\underline{k}}^{k^*} F_e\left(R(k)+e-R(x)\right) dG(x)} \Big(\alpha_z \int_{\underline{k}}^{k^*} F_e'\big(R(k)$$
$$+e-R(x)\big) dG(x) \Big) dF_e(e) + e^{-\alpha_z \int_{\underline{k}}^{k^*} F_e\left(z-R(x)\right) dG(x)} F_e'(z-R(k)) \Big] dM(z). \quad (13)$$

3. *The interest rate distribution is given by $F_R(x) = G\big(R^{-1}(x)\big)$.*

Proof. The optimal interest rate for a type-k lender solves:

$$\pi_k'(R) = b\Big(1-\rho(R)-R\rho'(R)\Big)P(R) + b\Big(R(1-\rho(R))-k\Big)P'(R) = 0.$$

Note that this expression is positive for $R = 0$ and negative at $\tilde{R}$ (the revenue-maximizing interest rate). Therefore an optimal choice $R(k) \in (0, \tilde{R})$ exists and is continuous in k. In the case of multiple roots, the lender chooses the solution that yields higher profits.

The cross-partial of profits with respect to lender type and interest rate is positive:

$$\frac{\partial \pi_k(R)}{\partial k} = -bP(R),$$
$$\frac{\partial^2 \pi_k(R)}{\partial k \partial R} = -bP'(R) > 0,$$

which implies that $R'(k) > 0$.

Since the optimal interest rate is strictly increasing in the lender's type, we have $F_R\big(R(k)\big) = G(k)$ for $k \in [\underline{k}, k^*]$ and therefore:

$$F_R(x) = G\big(R^{-1}(x)\big).$$

Using this feature, we can rewrite equation (8) as follows:

$$P\big(R(k)\big) = \int_{\underline{z}}^{\overline{z}} s(z) \int_{-\infty}^{z-R(k)} e^{-\alpha_z \int_{\underline{k}}^{k^*} F_e\left(R(k)+e-R(x)\right) dG(x)} dF_e(e) dM(z). \tag{14}$$

Equation (14) defines the probability that a type-k lender gives a loan when he makes the equilibrium choice $R(k)$. This expression does not directly depend on the interest rate distribution because it incorporates the result that the offered interest rate is strictly decreasing in lender type.

The profits of a type-k lender who follows the strategy of a type-$\hat{k}$ lender are:

$$\pi_k\big(R(\hat{k})\big) = b\big(R(\hat{k})(1-\rho(R(\hat{k}))) - k\big) \int_{\underline{z}}^{\overline{z}} s(z) \int_{-\infty}^{z-R(\hat{k})} e^{-\alpha_z \int_{\underline{k}}^{k^*} F_e\left(R(\hat{k})+e-R(x)\right) dG(x)} dF_e(e) dM(z).$$

Differentiating with respect to $\hat{k}$ we have:

$$\begin{aligned}\frac{\partial \pi_k(R(\hat{k}))}{\partial \hat{k}} &= bR'(\hat{k})(1-\rho(R(\hat{k})) - R\rho'(R(\hat{k}))) \int_{\underline{z}}^{\overline{z}} s(z) \int_{-\infty}^{z-R(\hat{k})} e^{-\alpha_z \int_{\underline{k}}^{k^*} F_e\left(R(\hat{k})+e-R(x)\right) dG(x)} dF_e(e) dM(z) \\ &\quad -b\big(R(\hat{k})(1-\rho(R(\hat{k}))) - k\big) \int_{\underline{z}}^{\overline{z}} s(z) \Big[\int_{-\infty}^{z-R(\hat{k})} e^{-\alpha_z \int_{\underline{k}}^{k^*} F_e\left(R(\hat{k})+e-R(x)\right) dG(x)} \Big(\alpha_z \int_{\underline{k}}^{k^*} F_e'\big(R(\hat{k}) \\ &\quad +e-R(x)\big) R'(\hat{k}) dG(x)\Big) dF_e(e) + R'(\hat{k}) e^{-\alpha_z \int_{\underline{k}}^{k^*} F_e\left(z-R(x)\right) dG(x)} F_e'\big(z-R(\hat{k})\big)\Big] dM(z).\end{aligned}$$

This derivative is equal to zero when $\hat{k} = k$. Therefore:

$$\begin{aligned}&(1-\rho(R(k)) - R\rho'(R(k))) \int_{\underline{z}}^{\overline{z}} s(z) \int_{-\infty}^{z-R(k)} e^{-\alpha_z \int_{\underline{k}}^{k^*} F_e\left(R(k)+e-R(x)\right) dG(x)} dF_e(e) dM(z) \\ &= \big(R(k)(1-\rho'(R(k))) - k\big) \int_{\underline{z}}^{\overline{z}} s(z) \Big[\int_{-\infty}^{z-R(k)} e^{-\alpha_z \int_{\underline{k}}^{k^*} F_e\left(R(k)+e-R(x)\right) dG(x)} \Big(\alpha_z \int_{\underline{k}}^{k^*} F_e'\big(R(k)+e \\ &-R(x)\big) dG(x)\Big) dF_e(e) + e^{-\alpha_z \int_{\underline{k}}^{k^*} F_e\left(z-R(x)\right) dG(x)} F_e'\big(z-R(k)\big)\Big] dM(z),\end{aligned}$$

which yields equation (13) that defines the interest rate schedule $R(k)$. This completes the proof of proposition 5. ■

The characterization of lenders' entry decisions is completed in the following proposition.

Proposition 6 *Given $s(z)$, lenders' entry is characterized as follows:*

1. *There is a unique cutoff cost k^* such that lenders enter if and only if $k \leq k^*$.*

2. *The measure of lenders in the market is given by* $L = \Lambda\Gamma(k^*)$ *and the cost distribution of entrants is* $G(k) = \frac{\Gamma(k)}{\Gamma(k^*)}$ *for* $k \leq k^*$*.*

3. *The cutoff is defined by the solution to:*

$$b\Big(R(k^*)(1-\rho(R(k^*))) - k^*\Big) \int_{\underline{z}}^{\overline{z}} s(z) \int_{-\infty}^{z-R(k^*)} e^{-s(z)\Lambda\Gamma(k^*)\int_{\underline{k}}^{k^*} F_e\left(R(k^*)+e-R(x)\right) d\frac{\Gamma(x)}{\Gamma(k^*)}} dF_e(e)dM(z) = \chi \tag{15}$$

Proof. Lemma 3 shows that a lender enters the market if and only if his cost is below some cutoff k^*. Consider a candidate cutoff for the market k^* and notice that the measure of lenders that enter the market is $L = \Lambda\Gamma(k^*)$. Denote the profits of the highest-cost lender by $\underline{\pi}_{k^*}$:

$$\underline{\pi}_{k^*}(R(k^*)) \;=\; b\Big(R(k^*)(1-\rho(R(k^*))) - k^*\Big)P(R(k^*))$$

where

$$P(R(k^*) \;=\; \int_{\underline{z}}^{\overline{z}} s(z) \int_{-\infty}^{z-R(k^*)} e^{-s(z)\Lambda\Gamma(k^*)\int_{\underline{k}}^{k^*} F_e\left(R(k^*)+e-R(x)\right) d\frac{\Gamma(x)}{\Gamma(k^*)}} dF_e(e)dM(z)$$

Note that the dependence of L and $G(\cdot)$ on k^* is made explicit.

The profits of the highest-cost lender are decreasing in his type:

$$\frac{d\underline{\pi}_{k^*}}{dk^*} \;=\; -bP\big(R(k^*)\big) + b\Big(R(k^*)\big(1-\rho(R(k^*))\big) - k^*\Big)\frac{\partial P(R(k^*))}{\partial L}\Lambda\Gamma'(k^*)$$

which is negative because an increase in k^* increases the measure of lenders in the market which reduces the probability of giving out a loan. Therefore, given $s(z)$, there is a unique k^* that characterizes lender entry.

The cutoff k^* is determined by equating the profits of the highest-cost seller with the entry cost χ, as shown in equation (15). ■

4.3 Equilibrium Existence and Offer Distribution

We can now prove the existence of equilibrium and provide some additional characterization results.

Based on the preceding results, the equilibrium can be described as a fixed point in the space of continuous functions $s(z)$ from $[\underline{z}, \overline{z}]$ to $[0, 1]$. From an arbitrary initial $s(z)$, propositions 5 and 6 characterize the lenders' entry and interest rate decisions $\{L, G(\cdot), F_R(\cdot)\}$; given lenders' actions, proposition 2 characterizes borrowers' search effort $s(z)$. The space is complete under the sup norm and therefore a fixed point (and an equilibrium) exist.

Proposition 7 *An equilibrium exists.*

In equilibrium, the distribution $H_R(R)$ of accepted offers equals:

$$H_R(R) \quad = \quad \frac{1}{1-\bar{P}} \int_{\underline{R}}^{R} P(R) dF_R(R), \tag{16}$$

where $\bar{P} = \int_{\underline{R}}^{\overline{R}} \big(1 - P(R)\big) dF_R(R)$ is the fraction of rejected offers, $P(R)$ is defined by equation (8) and $F_R(\cdot)$ is the distribution of offered interest rates. Intuitively, $H_R(R)$ weights each interest rate offer R by its probability $P(R)$ of being accepted and the distribution is rescaled to account for the aggregate rejection rate.

5 Quantitative Analysis

The model does not admit an analytic solution for all endogenous outcomes. Hence, we choose the parameters that best match moments of the data with the corresponding moments computed from the model's numerical solution. We then study the quantitative implications of the model evaluated at the calibrated parameters.

5.1 Parametric Assumptions

The calibration requires that we make several parametric assumptions. We borrow some parametric assumptions about the distributions of buyers' and sellers' heterogeneity from papers that structurally estimate search models of the labor market and our prior work on the retail market for illicit drugs (Galenianos and Gavazza, 2017). Specifically, given the similarity in modeling frameworks and empirical targets between this paper and those predecessors, we choose a lognormal distribution with parameters μ_z and σ_z for the distribution of buyers' preferences z. Moreover, we parametrize the distribution of sellers' costs k as the sum of the risk-free rate—we use the interest rate of the one-month Treasury bill at November 16th, 2007, which equals 3.78 percent—and an heterogeneous spread, drawn from a lognormal distribution with parameters μ_k and σ_k and upper-truncation point k^*.

We normalize the size of the loan to $b = 1$. We further assume that: 1) the product differentiation parameter e has a normal distribution with mean zero and standard deviation σ_e; 2) the cost of effort $q(s\Lambda)$ equals $\alpha_0 (s\Lambda)^{\alpha_1}$; and 3) the charge-off rate $\rho(R)$ equals βR.

Finally, we assume that the interest rates on accepted offers are measured with errors. Specifically, we assume that the reported accepted rates $\hat{R}$ and the "true" accepted rate R are related as: $\hat{R} = R\eta$, where η is a measurement error, drawn from a lognormal distribution with parameters (μ_η, σ_η), and with mean to equal 1—i.e., measurements are unbiased; hence, the parameters (μ_η, σ_η), satisfy $\mu_\eta = -0.5\sigma_\eta^2$. The assumption of measurement error on wages is quite common in the literature that structurally estimates search models of the labor market. In our application, it is plausible that surveyed borrowers may report the interest rates that they pay on their credit card debt with error. Measurement error could also account for some unobserved factor that our static model does not consider (i.e., adjustment of the interest rate after the offer is accepted), thereby allowing us to fit the distribution of accepted rates better. For example, Table 2 shows that this distribution displays a large dispersion, and the measurement η allows the model to capture this feature of the data.

5.2 Calibration

We choose the vector $\psi = \{\Lambda, \sigma_e, \mu_k, \sigma_k, \mu_z, \sigma_z, \sigma_\eta, \alpha_0, \alpha_1, \beta\}$ that minimizes the distance between the target moments m reported in Table 2 and the corresponding moments of the model.

Specifically, for any value of the vector ψ, we solve the model of Section 4 to find its equilibrium: the distribution $F_R(\cdot)$ of offered interest rates and borrowers' search effort $\lambda(s)$ that are consistent with each other. Once we solve for these optimal policy functions of borrowers and lenders, we can compute the equilibrium distributions of the number of received offers and of accepted rates. In practice, we simulate these distributions and the moments $\mathbf{m}(\psi)$ corresponding to those reported in Table 2 on received offers and on accepted offers, as well as the aggregate statistics on the fraction of credit card borrowers and on the charge-off rate.

We choose the parameter vector ψ that minimizes the criterion function

$$(\mathbf{m}(\psi) - m)' \Omega (\mathbf{m}(\psi) - m),$$

where $\mathbf{m}(\psi)$ is the vector of stacked moments simulated from the model evaluated at ψ and m is the vector of corresponding sample moments. Ω is a symmetric, positive-definite matrix;

we use the identity matrix.

5.2.1 Data Generating Process

Matching the moments reported in Table 2 requires that we account for the fact that the data generating process may be unusual, since we combine two separate datasets, collected for different purposes. Specifically, it seems plausible to us that the dataset on received offers reports *all* offers that borrowers receive (whose arrival rate is Λ), and not exclusively the offers that borrowers consider in equilibrium, which may be lower than the offered received because of borrowers' endogenous search effort/inattention s may be less than full effort $s = 1$. We derive in Appendix A the average number of offers and the distribution of the difference between the highest and lowest offers that borrowers receive under the assumption that the arrival rates if these offers is Λ.

However, lenders send these offers anticipating that borrowers will consider them according to their equilibrium λ_z. Hence, the moments on the empirical distribution of accepted offers reflect consumers' endogenous search effort λ_z.

5.2.2 Identification

The identification of the model is similar to that of other structural search models. Specifically, although the model is highly nonlinear, so that (almost) all parameters affect all outcomes, the identification of some parameters relies on some key moments in the data.

The moments on the number of offers that borrowers receive identify the offer rate parameter Λ. Moreover, the moments of the distribution of the difference between the highest and lowest offered interest rates identify the parameters of the distribution $G(k)$ of sellers' heterogeneity, and they contribute to the identification of the parameters of the distribution $M(z)$ of buyers' heterogeneity. More precisely, Appendix A shows how the distribution of the difference between the highest and lowest offered interest rates depends on the distribution $F_R(x)$ of offered rates and Proposition 5 shows that $F_R(x) = G\big(R^{-1}(x)\big)$, which would allow us to non-parametrically recover the distribution $G(k)$ of sellers' costs k; in practice, because we use the percentiles of the distribution of the difference between the highest and lowest offered interest rates reported in Panel B of Table 2 only, we specify a parametric lognormal distribution for $G(k)$. Moreover, because borrowers accept interest rates R only if they (together with the product differentiation e) are lower than their willingness to pay z, the distribution of offered rates—and the distribution of accepted rates, as well—is informative about the distribution of z.

Conditional on $G(k)$, the moments of the distribution of buyers' accepted interest rates identify the parameters of the search effort distribution α_0 and α_1, the parameter of the product differentiation distribution σ_e, and, as we argued above, they contribute to the identification of the parameters of the distribution $M(z)$ of buyers' preferences. Specifically, the difference between the distribution of offers and the distribution of accepted offers, along with the aggregate fraction of individuals with credit card debt, are informative about borrowers' search effort and product differentiation. In particular, the model with search effort but without product differentiation cannot fully account for the joint distribution of accepted offers and received offers, and the differentiation parameter e allows the empirical model to match this feature of the data more closely; thus, this feature of the model contributes to the identification of the parameter σ_e of the distribution of e. Moreover, as the comparative statics of Section 5.5 and Figures 5 and 6 will further clarify, high product differentiation and low search effort relax price competition and, thus, lenders offer credit cards high interest rates when the standard deviation σ_e of the product differentiation parameter e is large and/or when search effort is low. Hence, the model can account for the observed high level of offered interest rates and for observed difference between the offer distribution and the accepted offer distribution in two cases: 1) high value of the product differentiation; and 2) a low search effort. However, these two cases have opposite implications for the fraction of individuals with debt: the former increases this fraction, whereas the latter decreases it. Thus, aggregate fraction of individuals with credit card debt contributes to the identification of the parameters of the product differentiation distribution and of the search effort function.

Moreover, as in structural search models of the labor market, the data sometimes display events that should not occur according to the model, and these "zero-probability events" identify the parameters of the distribution of the measurement error η. In search models of the labor markets, these events include job-to-job transitions that feature a wage decrease, for example. Instead, our model without this measurement error cannot fully account for the observed heterogeneity in (i.e., the standard deviation of) accepted rates, and the measurement error η allows the empirical model to match this feature of the data; thus, this difference between the theory and the data identifies the parameter σ_η of the distribution of η.

Finally, the aggregate charge-off rate identifies the parameter β of the functional form of the charge-off rate $\rho(R)$. The fraction of borrowers further contributes to the identification of the parameters of the distribution of borrowers' willingness to pay z.

Table 3: Calibrated Parameters

Λ	2.650	σ_e	0.210
μ_k	0.843	σ_k	0.743
μ_μ	3.537	σ_μ	0.391
α_0	9.034	α_1	2.203
k^*	12.928	σ_η	0.278
β	0.002		

Notes—This table reports the calibrated parameters.

5.3 Calibrated Parameters and Model Fit

Table 3 reports the calibrated parameters.

The parameters μ_k, σ_k and k^* of the distribution of costs k imply that lenders' average costs equal 648 basis points and, thus, they display a small spread of 270 basis points over the risk-free rate. Moreover, the heterogeneity of lenders' costs in small—i.e., it equals 164 basis points. The parameters μ_z and σ_z of the distribution of z mean that borrowers' willingness to pay for credit is, on average, large and displays large heterogeneity. Figure 1 plots the distributions of lenders' costs (left panel) and borrowers' willingness to pay; the interval of values over which they overlap is small, implying that there are large gains from trade in this market.

The value of Λ indicates that sellers send, on average, approximately 2.7 credit card offers. However, the parameters α_0 and α_1 imply that buyers consider only a small fraction of these offers, as the cost of effort increases rapidly in λ: the cost of effort to evaluate an average number of offers per period equal to $\lambda = 1$ corresponds to $\alpha_0 = 21.103$, or 2,110 basis points, whereas the cost of effort to evaluate an average number of offers per period equal to $\lambda = 2$ corresponds to $\alpha_0 \lambda^{\alpha_1} = 94.62$, or 94,624 basis points. [These seem too large, obviously.]

The value of σ_e implies that the standard deviation of the product differentiation parameter is small, relative to the overall heterogeneity in lenders' costs and borrowers' preferences. Specifically, the variance of e accounts for approximately one percent of the overall variance of $c = R + e$. In the next section, we will perform some comparative statics that further illustrate how σ_e affects the equilibrium distributions of offered and accepted rates.

The parameter β implies, for example, that the charge-off rate βR on a credit card with an interest rate of 20 percent equals 4.4 percent and raising the interest rate to 25 percent

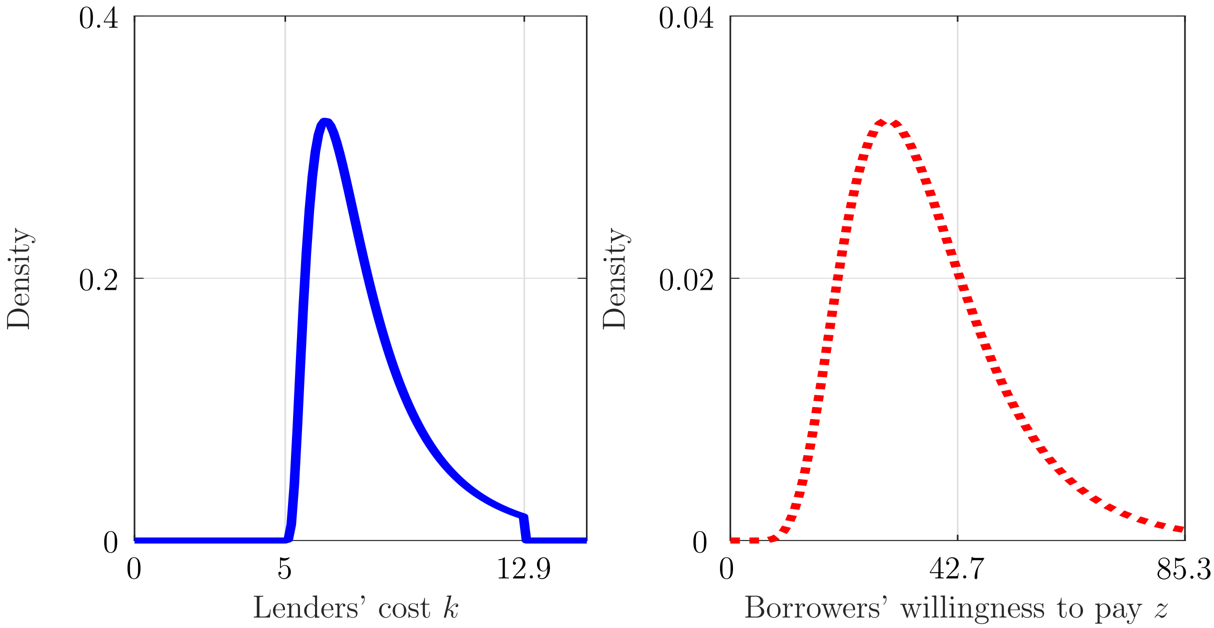

Figure 1: Distribution of lenders' costs (left panel) and borrowers' willingness to pay (right panel).

increases the charge-off rate by 1.1 percentage point to 5.5 percent.

Finally, the calibrate that σ_η equals 0.273, which means that the variance of the measurement error on the accepted rates equals 0.278. This value, along with those of the other estimates, implies that the model without measurement error η accounts for 55 percent of the dispersion of accepted interest rates observed in the data, and that the error η improves the fit, in particular by rationalizing the lowest- and the highest-R credit card loans.[6]

Table 4 presents a comparison between the empirical moments and the moments calculated from the model at the calibrated parameters. The model matches the percentiles of the distribution of accepted rates remarkably well, whereas it slightly overpredicts the percentiles of the distribution of the difference between the highest and lowest offered interest rates; finally, it matches quite closely the aggregate statistics on the fraction of credit card borrowers and the charge-off rate.

5.4 Model Implications

We study the implications of the model evaluated at the parameters reported in Table 3.

The left panel of Figure 2 displays lenders' optimal offered rate $R(k)$ as a function of their cost k, for values of the cost k from the risk free rate up to the cutoff value k^* that the free

[6] Of course, if we calibrate the model *without* measurement error, the estimated variance of accepted R increases.

Table 4: Model Fit

	Data	Model
10th Percentile Accepted Offer Distribution	12.68	12.87
25th Percentile Accepted Offer Distribution	15.84	15.53
50th Percentile Accepted Offer Distribution	19.31	19.09
75th Percentile Accepted Offer Distribution	23.82	23.65
90th Percentile Accepted Offer Distribution	28.60	28.54
Fraction Receiving 2+ Offers (%)	75.00	74.22
Median Number of Offers Received, Conditional on 2+ Offers	3.00	3.00
Average Number of Offers Received, Conditional on 2+ Offers	4.00	3.32
10th Percentile Distribution of Differences in Offered Rates	0.00	1.38
30th Percentile Distribution of Differences in Offered Rates	2.25	3.48
50th Percentile Distribution of Differences in Offered Rates	4.34	5.31
70th Percentile Distribution of Differences in Offered Rates	7.25	7.02
90th Percentile Distribution of Differences in Offered Rates	9.25	9.13
Fraction with Credit Card Debt	36.70	36.35
Charge-Off Rate	4.01	4.14

Notes—This table reports the values of the empirical moments and of the moments calculated at the calibrated parameters reported in Table 3.

entry condition 15 determines. Lenders' offered rates are strictly increasing in their costs k, as Proposition 5 states. The average markup, computed as $\frac{R(k)(1-\rho(R(k)))-k}{k}$ to take into account the charge-off rate $\rho(R(k))$, equals 197 percent and the median markup equals 212 percent, indicating that markups are skewed, as the lowest-cost lenders have markups in excess of 220 percent, whereas the highest-cost lenders have markups below 100 percent.

The right panel of Figure 2 displays borrowers' optimal search effort $\lambda(z)$ as a function of their willingness to pay z, for values of the willingness to pay z up to 200. Since the lowest-valuation borrowers (i.e., those whose valuations are below 15) have a willingness to pay that is below almost all offered interest rates and the product differentiation parameter e has a small variance, these borrowers do not exert any search effort. More generally, the search effort is low—i.e., on average, borrowers evaluate 0.53 offers—and only borrowers whose willingness to pay z is in the highest 15 percent of the distribution choose $\lambda(z)$ larger than 1.

Figure 3 displays the probability $P(R)$ that borrowers accept a credit card offer with an interest rate R. This probability is obviously decreasing, but perhaps its most striking feature is that, because of borrowers' low search effort, it is quite flat—i.e., borrowers' demand is quite

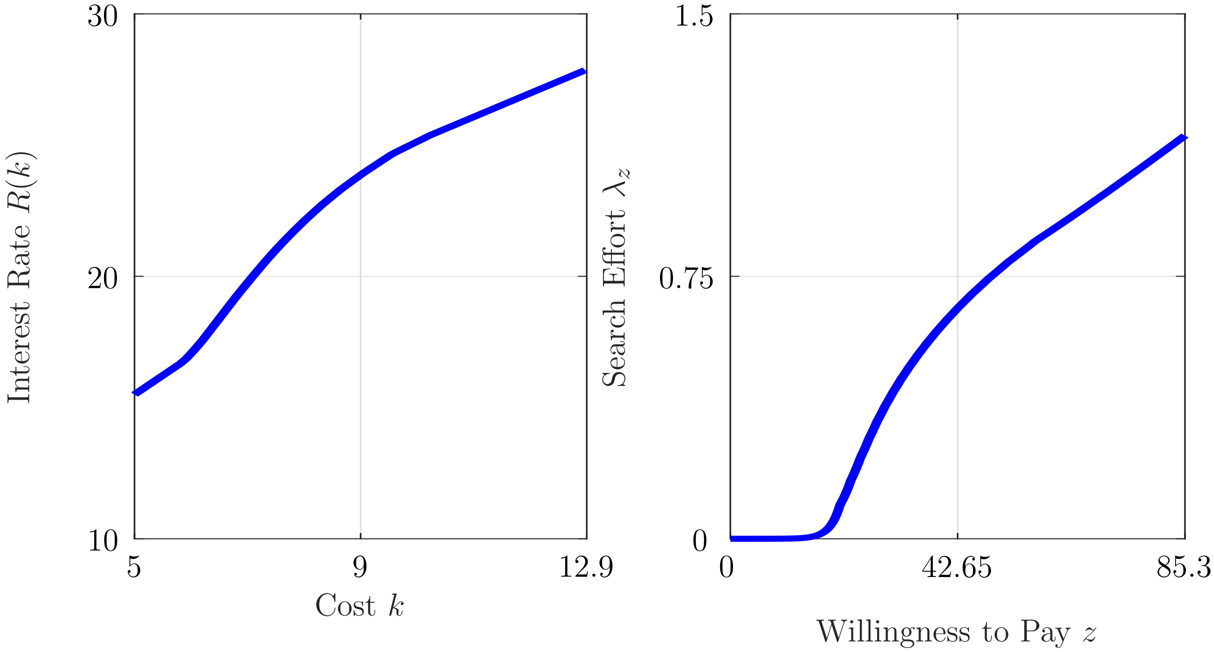

Figure 2: The left panel displays lenders' optimal interest rate $R(k)$ as a a function of their cost k, and the right panel displays borrowers' optimal search effort λ_z as a function of their willingness to pay z.

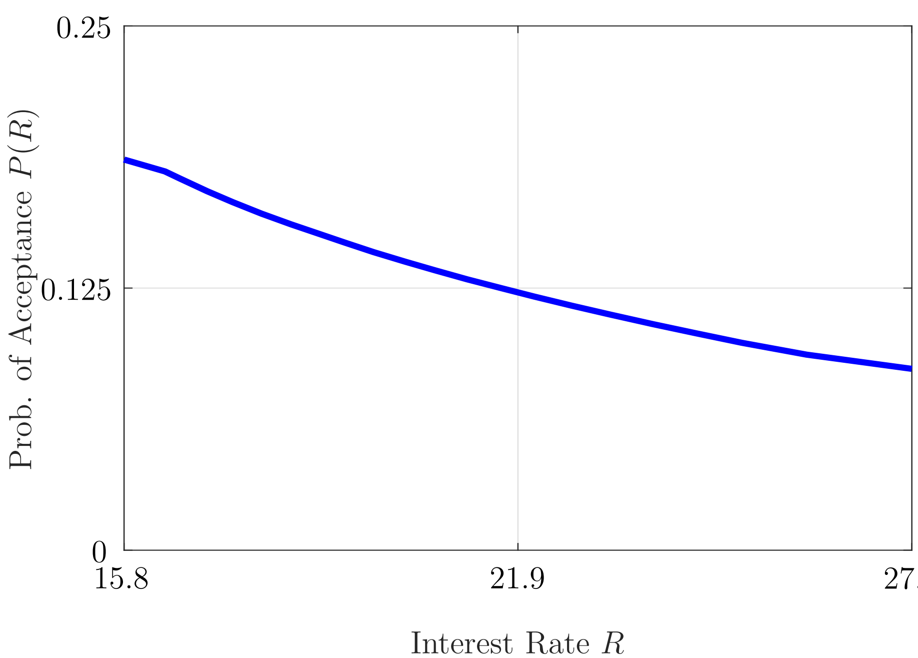

Figure 3: Probability $P(R)$ that borrowers accept an offer with interest R.

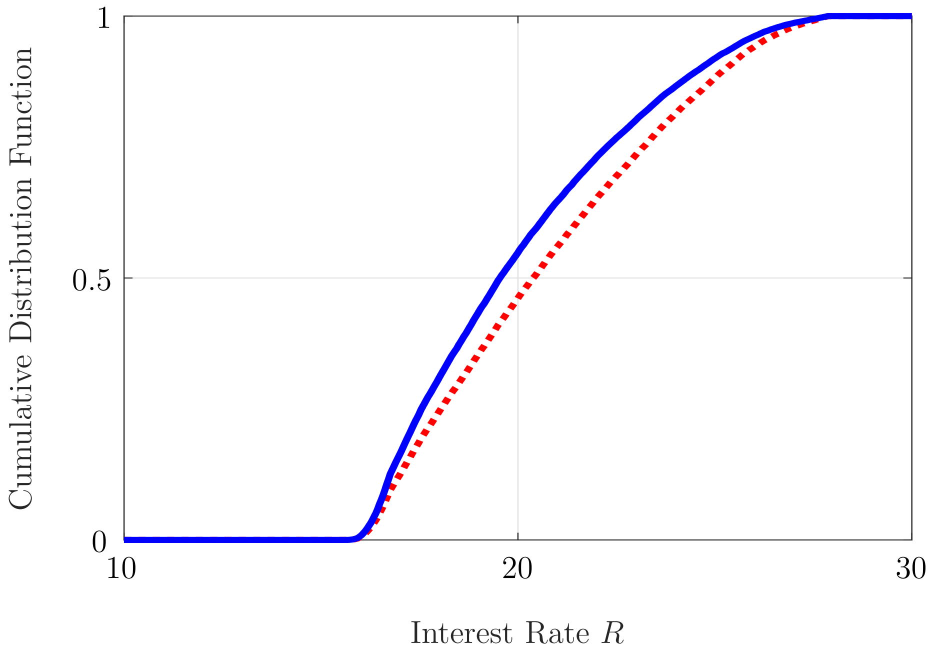

Figure 4: The solid line displays the cumulative distribution function $H_R(R)$ of accepted interest rates and the dashed line displays the cumulative distribution function $F_R(R)$ of offered rates.

inelastic: on average, it equals -1.26. The average probability $P(R)$ equals 0.13, which, scaled up by the aggregate mass of lenders $\Lambda \approx 2.7$, matches the aggregate fraction of individuals with credit card debt of 36 percent.

Figure 4 plots the distribution $F_R(R)$ of offered rates and the distribution $H_R(R)$ of accepted rates. Of course, the distribution of offered rates first-order stochastically dominates the distribution of accepted rates. However, the difference between the two distributions is small—e.g., the mean of the distribution of accepted rates equals 18.92, whereas that of the distribution of offered rates equals 19.72. Two reasons account for this small difference: 1) borrowers' search effort $\lambda(z)$ is low, as we recount above; and 2) borrowers do not always accept the offer with the lowest interest rate, because of the product-differentiation parameter e. However, this second factor is quantitatively negligible, as the standard deviation σ_e is small: the mean of the distribution of accepted rates would be almost identical if borrowers were to always choose the offer with the lowest interest rates.[7]

[7] Of course, this is not an equilibrium argument, as the endogenous distribution of offered rates $F_R(\cdot)$ depends on the product differentiation e.

5.5 Comparative Statics

We further illustrate the working of our model through three comparative statics. The first two vary the two set of parameters that are the main focus of our framework—i.e., the standard deviation σ_e of the product differentiation e and the parameter α_0 that scales the marginal cost of search effort. The third one varies the parameter β that affects the aggregate charge-off rate.

Product Differentiation. Figure 5 compares outcomes of the model at the calibrated parameters (solid line) to those of the model when we increase the standard deviation σ_e of the product differentiation e (dashed line) while holding all other parameters at their calibrated values. Since σ_e is calibrated to be small, we increase it fivefold.

The top panels show interesting outcomes. Most notably, the top left panel shows that the interest rate function $R(k)$ increases substantially when product differentiation is more important for borrowers. The reason is that, a larger σ_e means that product differentiation affects consumers' choice across lenders relatively more; thus, lenders compete less aggressively by offering higher interest rates.

The top right panel shows that a higher σ_e induces borrowers to decrease their search effort. This is the result of two opposite effects. Specifically, the increase in the product differentiation parameter induces borrowers to search more aggressively, since they are more likely to receive an offer with product feature e that they value highly. However, the increase in the product differentiation parameter induces lenders to increase their interest rates, thereby decreasing borrowers' expected surplus and, thus, their search effort. The increase in interest rates outweighs the increase in product differentiation and, thus, search effort decreases.

The bottom-left panel displays the probability $P(R)$ that borrowers accept an offer with interest rate R. Because all lenders offer higher interest rates when σ_e is higher, the acceptance probability of an offer with a given R increases relative that of the baseline case. Demand becomes more inelastic—i.e., the average elasticity equals -1.13 compared to -1.26 in the baseline case. However, the average acceptance probability across lenders as well as the fraction of individuals with credit card debt are lower than in the baseline case—i.e., 0.12 and 0.32 vs. 0.13 and 0.36, respectively—indicating that the increase in interest rates may outweigh borrowers' benefits from larger values of e and, thus, borrowers may be worse off when σ_e is higher.

The bottom-right panel displays the distribution of offered rates (thin lines) and of accepted rates (thick lines). Both distributions obtained in the model with a higher σ_e (dashed lines)

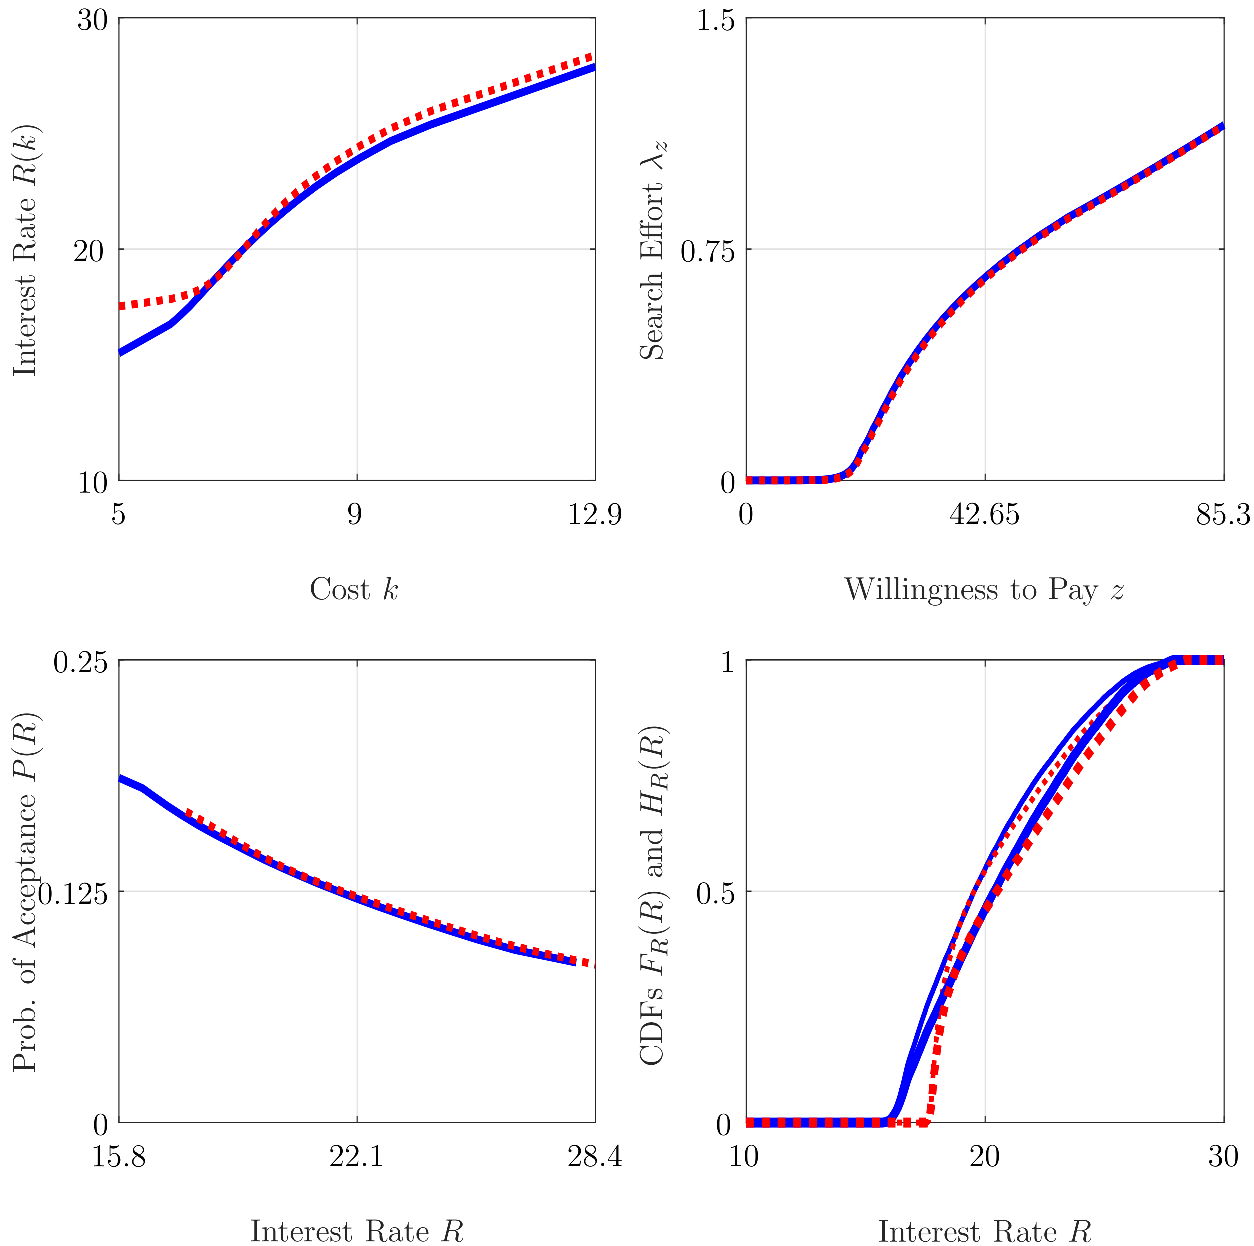

Figure 5: These panels display model outcomes model at the calibrated parameters (solid line) and in the case when $\sigma_e' = 5\sigma_e$ (dashed line). The top left panel displays lenders' optimal interest rate $R(k)$ as a a function of their cost k; the top right panel displays borrowers' optimal search effort λ_z as a function of their willingness to pay z; the bottom left panel displays the probability $P(R)$ that borrowers accept an offer with interest rate R; and the bottom right panel displays the distribution $F_R(R)$ of offered rates (thick lines) and the distribution $H_R(R)$ of accepted rates (thin lines).

first-order stochastically dominate the corresponding distributions obtained in the model at the calibrated σ_e (solid lines). In both cases, it is intuitive, as both offered rates and accepted rates are higher if product differentiation matters more for consumers' choices. The average offer rate increases from 19.78 in the baseline case at the calibrated parameters to 25.76 in the case with σ_e is higher, whereas the standard deviation of offers increases from 3.50 to 5.83, thereby reiterating that offered rates are higher and more dispersed if product differentiation matters more for borrowers. Similarly, the average accepted rate increases from 18.92 to 24.30, whereas the standard deviation of accepted rates increases from 3.42 to 5.45. Of course, accepted rates differ more substantially at higher percentiles than at lower percentiles: the 10th percentile increases by approximately four percentage points—i.e., from 14.73 to 18.80—whereas the 90th percentile increases by approximately nine percentage points—i.e., from 24.07 to 32.98.

Figure 5 also helps us understand why the calibrated model delivers a small value of σ_e : interest rates would be even higher than they currently are if σ_e was larger.

Cost of Effort. Figure 6 compares outcomes of the model at the calibrated parameters (solid line) to those of the model when we decrease the parameter α_0 of the cost of effort by 50 percent (dashed line) while holding all other parameters at their calibrated values.

The top-left panel shows that the interest rate function $R(k)$ is lower than that in the baseline case, as all lenders uniformly decrease their interest rates. The top-right panel explains why this happens: since the cost of effort is lower, on average borrowers increase their search effort—most notably high-valuation borrowers.

The bottom-left panel shows that the probability $P(R)$ that borrowers accept an offer with a given interest rate R increases relative that of the baseline case, since borrowers' search effort is higher. Demand becomes more elastic—i.e. the average elasticity equals -1.36—relative to the baseline case. Moreover, lenders decrease their rates and obviously borrowers accept offers with lower rates with a higher probability. These two forces increase the average acceptance probability across lenders as well as the fraction of individuals with credit card debt relative to their values in the baseline case—i.e., 0.17 and 0.46 vs. 0.13 and 0.36, respectively—indicating that borrowers are better off when α_0 is lower.

The bottom-right panel of Figure 6 displays the distribution of offered rates (thick lines) and of accepted rates (thin lines). Both distributions obtained in the model with a lower α_0 (dashed lines) are first-order stochastically dominated by the corresponding distributions obtained in the model at the calibrated α_0 (solid lines). This is because low-cost lenders lower their offered rates when the cost of effort is lower and, thus, borrowers' search more.

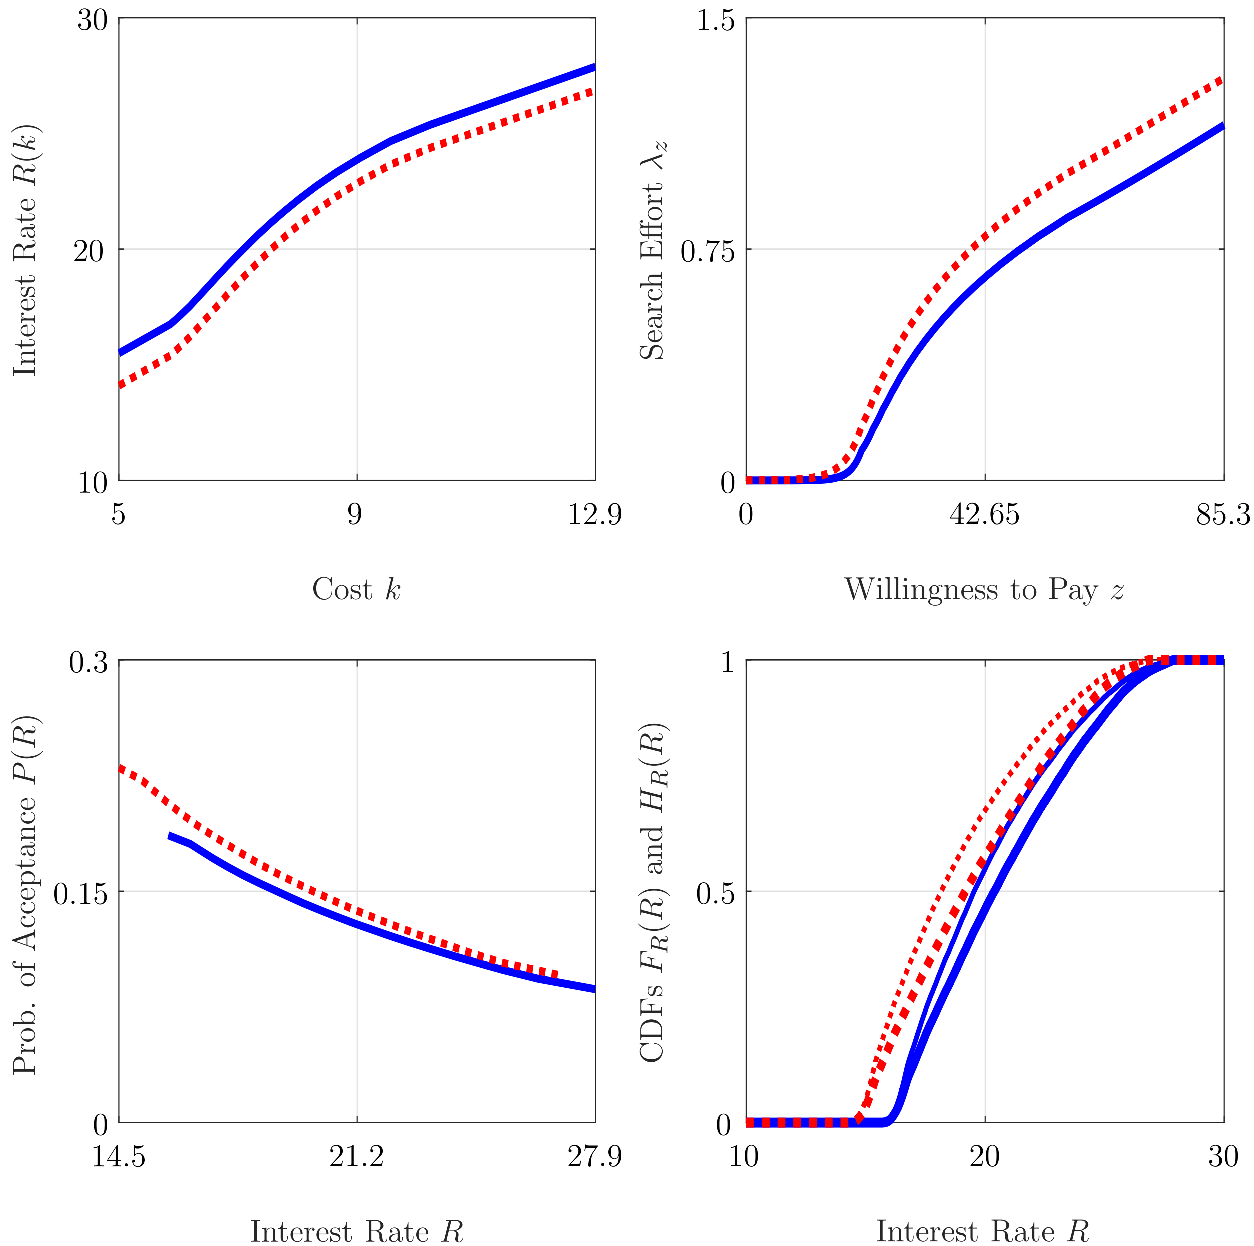

Figure 6: These panels display model outcomes at the calibrated parameters (solid line) and in the case when $\alpha_0' = 0.5\alpha_0$ (dashed line). The top left panel displays lenders' optimal interest rate $R(k)$ as a a function of their cost k; the top right panel displays borrowers' optimal search effort λ_z as a function of their willingness to pay z; the bottom left panel displays the probability $P(R)$ that borrowers accept an offer with interest rate R; and the bottom right panel displays the distribution $F_R(R)$ of offered rates (thick lines) and the distribution $H_R(R)$ of accepted rates (thin lines).

The average offered and accepted rates equal 17.37 and 16.18, respectively, and the standard deviation of offered and accepted rates equal 3.73 and 3.55, respectively, when the cost-of-effort parameter α_0 is half of the calibrated value. As the bottom plots shows, the lower cost of effort affects lower percentiles and higher percentiles in similar ways: for instance, the 10th percentile of the distribution of accepted rates equals 12.10 (versus 14.73 in the calibrated model) and the 90th percentile equals 21.69 (versus 24.07 in the calibrated model).

Charge-off Rate. Figure 7 compares outcomes of the model at the calibrated parameters (solid line) to those of the model when we increase the parameter β of the charge-off rate fourfold (dashed line), while holding all other parameters at their calibrated values.

The top-left panel shows that lenders charge lower interest rates, as the function $R(k)$ lies above that of the baseline case. The reason is that borrowers are more likely to default on higher-interest-rate loans and this curbs lenders' optimal rates; hence, if charge-off rate increases, as a higher parameter β captures, the restraint on lenders' rates increases and, thus, their optimal rates decrease. Notably, the decrease is larger for those lenders that charge higher interest rates, since these lenders are those that experience higher default rates and, thus, are harmed the most when the default rate increases.

The top-right panel shows how borrowers' search effort respond to these higher rates, displaying contrasting effects between low-valuation and high-valuation borrowers. Specifically, since interest rates are lower in this counterfactual case relative to the baseline case, the expected benefits of a credit card loan for low-valuation borrowers are higher; hence, they increase their search effort. High-valuation borrowers are more likely to receive multiple offers than low-valuation borrowers. Since the top-left panel shows that the dispersion of interest rates has decreased in the counterfactual case, the increase in the search effort of these high-valuation borrowers is lower than that of low-valuation borrowers because the gains of comparison shopping are lower when the dispersion of rates is lower. However, the top-right panel shows that the quantitative importance of these effects is minimal.

The bottom-left panel shows that the probability $P(R)$ that borrowers accept an offer with a given interest rate R decreases relative that of the baseline case, since lenders decrease their rates and, thus, borrowers are more likely to accept offers with lower rates. However, lenders decrease their rates and obviously borrowers accept offers with lower rates with a higher probability. These two forces push into opposite directions the changes in the average acceptance probability across lenders as well as the fraction of individuals with credit card debt relative to their values in the baseline case, so that the overall changes are negligible, though positive—i.e., 0.134 and 0.366 vs. 0.131 and 0.357, respectively. Demand becomes less

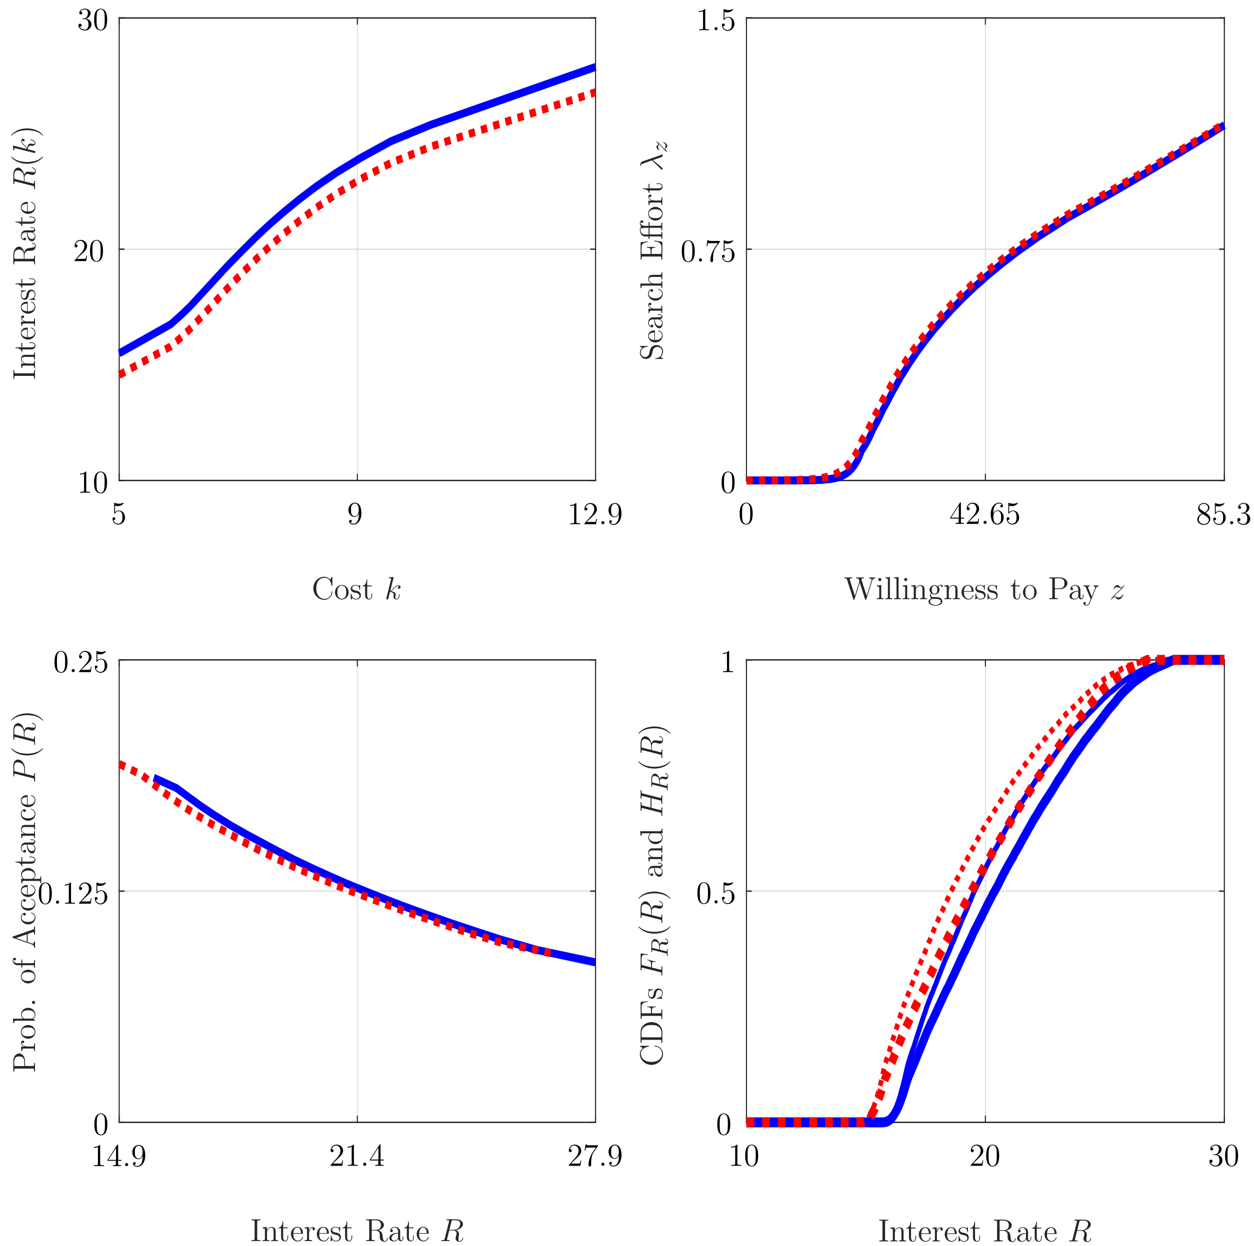

Figure 7: These panels display model outcomes at the calibrated parameters (solid line) and in the case when $\beta' = 4\beta$ (dashed line). The top left panel displays lenders' optimal interest rate $R(k)$ as a a function of their cost k; the top right panel displays borrowers' optimal search effort λ_z as a function of their willingness to pay z; the bottom left panel displays the probability $p(R)$ that borrowers accept an offer with interest rate R; and the bottom right panel displays the distribution $F_R(R)$ of offered rates (thick lines) and the distribution $H_R(R)$ of accepted rates (thin lines).

elastic relative to the baseline case: its average elasticity equals -1.20.

The bottom right-panel of Figure 6 displays the distribution of offered rates (thick lines) and of accepted rates (thin lines). Both distributions obtained in the model with a higher β (dashed lines) are first-order stochastically dominated by the corresponding distributions obtained in the model at the calibrated lower β (solid lines). The average offered and accepted rates equal 18.10 and 17.26, respectively, and the standard deviation of offered and accepted rates equal 3.40 and 3.32, respectively, when the parameter β of charge-off rate is larger than its calibrated value.

6 Policy Experiments

In this Section, we use our model to study two policy experiment, motivated by recent regulatory interventions: 1) A cap on the interest rate, captured by a maximum rate $\bar{R}$; 2) Higher compliance costs for lenders, captured by a higher fixed cost Φ. The goal of both experiments is to study how borrowers' search effort and lenders' offered rates respond and, thus, the equilibrium distribution of accepted rates.

6.1 Cap on Interest Rates

As we recount in the Introduction, several countries recently introduced price controls in markets for some consumer financial products, and are currently considering intervening in a larger number of these product markets. The goal of this section is to study the effects of a price cap on the equilibrium of our model.

Most notably, the literature highlights that these price caps may have unintended consequences: Fershtman and Fishman (1994) and Armstrong et al. (2009) show that, in markets with search frictions, price caps may *increase* the equilibrium prices paid by consumers. The reason is that price caps may reduce price dispersion and this reduction decreases consumers' search efforts, thereby increasing suppliers' market power and, thus, prices. This indirect effect may dominate the direct effect of price caps on those consumers who were paying prices higher than the cap before the regulation. Thus, the introduction of a price cap has a theoretically ambiguous effect on the equilibrium price paid by consumers. Therefore, it is an empirical/quantitative question which of the two opposing effects dominates and, thus, whether or not price caps benefit borrowers; our calibrated model—which features the margins of adjustment that the literature focuses on—is well suited to answer this question.

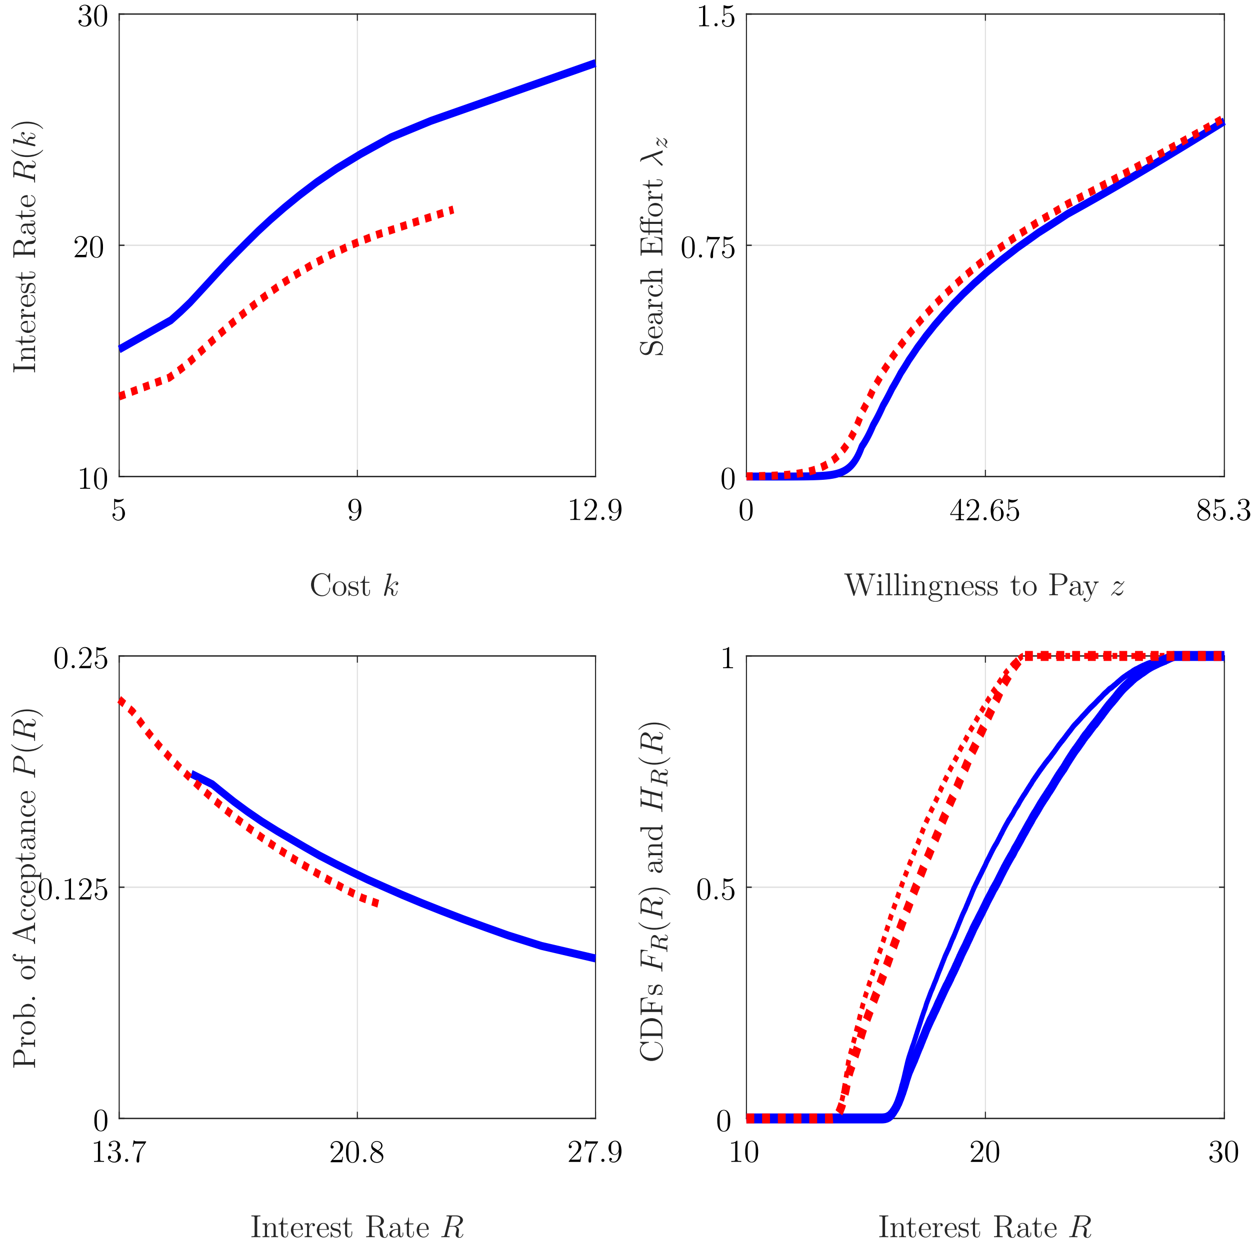

Figure 8: These panels display model outcomes at the calibrated parameters (solid line) and in the case when interest rates are capped at 22.5 percent (dashed line). The top left panel displays lenders' optimal interest rate $R(k)$ as a a function of their cost k; the top right panel displays borrowers' optimal search effort λ_z as a function of their willingness to pay z; the bottom left panel displays the probability $P(R)$ that borrowers accept an offer with interest rate R; and the bottom right panel displays the distribution $F_R(R)$ of offered rates (thick lines) and the distribution $H_R(R)$ of accepted rates (thin lines).

We study this counterfactual case in general equilibrium—i.e., we require that lenders' free entry condition (15) holds. Thus, some lenders may exit the market, in which case we decrease the aggregate arrival rate of offers to a new value Λ' proportionally. More precisely, the new arrival rate equals $\Lambda' = \Lambda \frac{G(k^{**})}{G(k^{*})}$, where k^{**} is the marginal cost of the marginal lender—i.e., the lender that satisfies the free entry condition (15)—in the counterfactual case, and k^{*} is the marginal cost of the marginal lender in the baseline case.

Figure 8 compares outcomes of the model at the calibrated parameters (solid line) to those of the model with the price cap of $\bar{R} = 22.5$ percent, while holding all other parameters at their calibrated values.

The top-left panel shows interesting outcomes. First, the highest-cost lenders exit the market, even tough the cap is above their marginal cost. Specifically, frictions are such that, even if these lenders were to decrease their interest rates substantially, their market share would not increase as much as to allow them to cover their fixed costs; hence they exit. Second, all surviving lenders charge lower interest rates, as the function $R(k)$ lies strictly below that of the baseline case. In particular, the lender with marginal cost k^{**} finds it worthwhile to drop its rate to satisfy the cap constraint, rather than exit; similarly, all other lenders with lower marginal costs charge slightly below their higher-cost competitors.

The top-right panel shows how borrowers' search effort respond to these higher rates, displaying the indirect and direct effects that Fershtman and Fishman (1994) and Armstrong et al. (2009) emphasize. Specifically, low-valuation borrowers increase their search effort to obtain a credit card loan, since the cap reduces the level of interest rates relative to the baseline case. However, the increase in search effort of high-valuation borrowers is lower than that of low-valuation borrowers, since the price cap reduces the dispersion of interest rates across lenders and, thus, reduces the benefits of comparison shopping.

The bottom-left panel displays the probability $P(R)$ that borrowers accept an offer with interest rate R. The acceptance probability of an offer with a given R decreases relative that of the baseline case because the price cap reduces lenders' rates and, thus, borrowers are less likely to accept an offer with a given R and more likely to accept offers with lower Rs. The average acceptance probability across lenders as well as the fraction of individuals with credit card debt increase relative to the baseline case, though the changes are small—i.e., 0.146 and 0.374 vs. 0.131 and 0.357, respectively. On average, demand is more elastic than in the baseline case: the average elasticity equals -1.33 vs. -1.26 in the baseline case. Overall, these changes indicate that the price cap increases consumer welfare.

The bottom right panel of Figure 8 displays the distribution of offered rates (thick lines) and of accepted rates (thin lines). Both distributions in the case of a interest rate ceiling (dashed lines) are first-order stochastically dominated by the corresponding distributions obtained in baseline case with no ceiling (solid lines). The average offered and accepted rates equal 16.36 and 15.68, respectively, and the standard deviation of offered and accepted rates equal 2.79 and 2.70, respectively.

6.2 Higher Compliance Costs

A second set of regulations that have been introduced since the Financial Crisis featured increased compliance costs on lenders. Through the lenses of our model, this can be interpreted as an increase in lenders' fixed cost Φ and, thus, our model is well-suited to understand the effect of this increase on the market equilibrium.

The increase in the fixed costs will share with our previous counterfactual about the introduction of a price cap the feature that highest-costs lenders will exit the market; thus, this counterfactual with larger fixed costs will allow us to understand how much the results displayed in Figure 8 obtain because of the exit of these highest-cost lenders. Moreover, Janssen and Moraga-González (2004) shows that a decrease in the number of active firms could increase search effort because it may decrease price dispersion, possibly leading to higher average prices.

To facilitate the comparison with our price cap experiment of Figure 8, we increase the fixed cost Φ so that the marginal lender has marginal cost equal to k^{**}—i.e., the marginal cost of the lender that satisfies the free entry (15) condition in the case of the price cap $\bar{R} = 22.5$. In practice, this implies that the new fixed cost Φ' is 16-percent larger relative to that of the baseline case. We further decrease the aggregate arrival rate of offers to a new value Λ' correspondingly—i.e., the new arrival rate equals $\Lambda' = \Lambda \frac{G(k^{**})}{G(k^{*})}$.

Figure 9 compares outcomes of the model at the calibrated parameters (solid line) to those of the model with a higher fixed cost Φ. Overall, the panels of Figure 9 show that the outcomes of this counterfactual case are quite similar to those of the baseline case. Specifically, the top-left panel shows that, while the highest-cost lenders exit the market, all other lenders charge rates that are almost identical to their optimal rates in the baseline case. As a result, borrowers search effort changes minimally, as the top-right panel documents. The bottom-left panel displays the probability $P(R)$ that borrowers accept an offer with interest rate R changes because borrowers do not accept offers with the highest interest rates, since borrowers that offered them in the baseline case are no longer in the market. The average acceptance

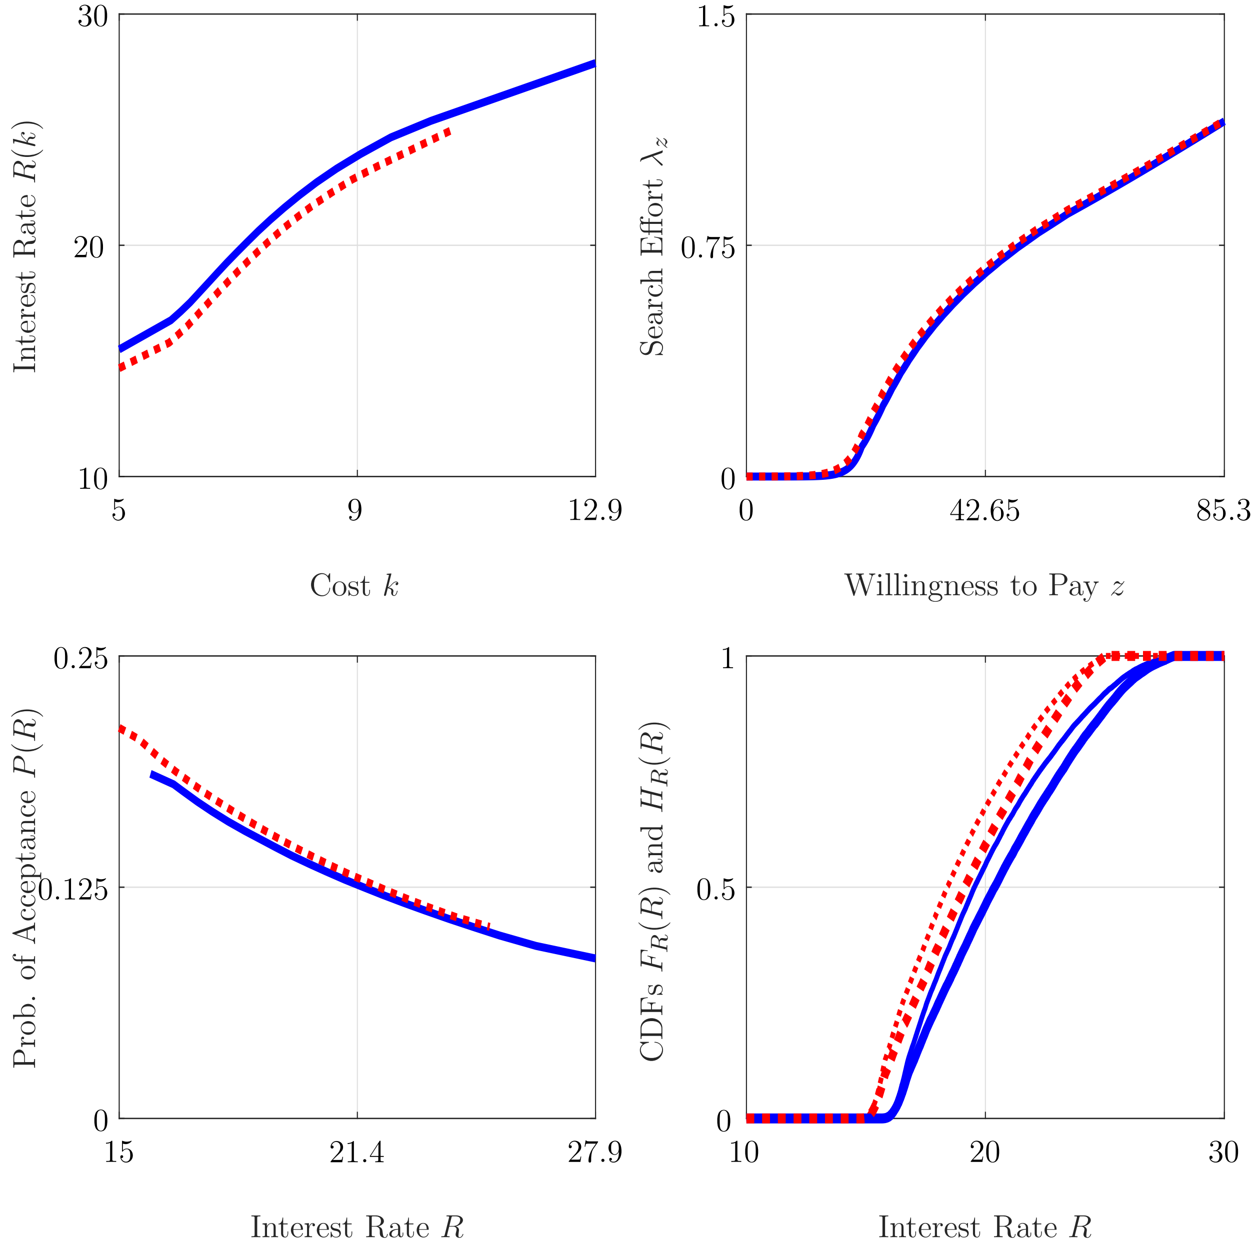

Figure 9: These panels display model outcomes at the calibrated parameters (solid line) and in the case when the fixed cost $\Phi' = 1.16\Phi$ (dashed line). The top left panel displays lenders' optimal interest rate $R(k)$ as a a function of their cost k; the top right panel displays borrowers' optimal search effort λ_z as a function of their willingness to pay z; the bottom left panel displays the probability $P(R)$ that borrowers accept an offer with interest rate R; and the bottom right panel displays the distribution $F_R(R)$ of offered rates (thick lines) and the distribution $H_R(R)$ of accepted rates (thin lines).

probability across lenders as well as the fraction of individuals with credit card debt increase relative to the baseline case, though the changes are small—i.e., 0.142 and 0.364 vs. 0.131 and 0.357, respectively. The average elasticity of demand equals -1.34 (vs. -1.26 in the baseline case). Overall, these changes indicate that higher compliance costs may increase consumer welfare, although this increase is lower than that obtained through price caps. The bottom tight panel of Figure 9 show that the distributions of offered rates (thick lines) and of accepted rates (thin lines) in a market with a higher fixed cost Φ (dashed line) first-order stochastically dominate the corresponding distributions obtained in the baseline case (solid lines), as the highest-rate lenders are no longer active. The average offered and accepted rates equal 18.39 and 17.61, respectively, and the standard deviation of offered and accepted rates equal 3.11 and 3.04 respectively. Hence, the magnitudes of the changes in these interest rate statistics relative to those of the baseline case are small, most notably when we compare them to those that we obtained under the experiment of the price cap displayed in Figure 8.

7 Conclusions

This paper develops a framework to study frictions in credit card markets. We focus on two features to explain the observed large dispersion in the interest rates that individuals pay on their credit cards: endogenous (low) search effort and product differentiation.

We calibrate the model using data on the U.S. credit card market. The model fits the data reasonably well. Our analysis implies that low search effort accounts for almost all the dispersion in interest rates, whereas product differentiation is negligible.

We should point out that these results obtain in a model with important limitations and, thus, we believe that it can be enhanced in several ways. As we recount in Section 3, our cross-sectional data impose some limitations on what our model can identify in the data, and richer data on borrowers and lenders would allow us to enrich our current framework. Specifically, multidimensional heterogeneity is difficult to identify with our data; thus, our model focuses on a single dimension of heterogeneity across borrowers—i.e., their willingness to pay for credit—and across lenders—i.e., their funding cost—and restricts other parameters to be homogeneous across individuals. Many structural search models share these features due to similar data constraints, and one objective of this paper is to adapt and to enrich these models to understand two key characteristics—i.e., product differentiation and consumer limited search effort/inattention—of credit card markets. Nonetheless, our theoretical framework delivers a large dispersion in credit card offers, and our quantitative analysis is successful in matching this large heterogeneity observed in the data.

For these main reasons, we view this paper as a first step in quantifying the role of search effort/inattention in search markets. The quantitative analysis clarifies the data requirements to calibrate/estimate such a model and how the parameters are identified, and the calibration delivers a sense of the magnitudes involved, allowing us to assess which forces dominate. Nonetheless, we hope that the future availability of richer data will allow us to incorporate additional features of credit card markets and other markets for consumer financial product.

References

Adams, W., L. Einav, and J. Levin (2009): "Liquidity Constraints and Imperfect Information in Subprime Lending," *American Economic Review*, 99, 49–84.

Agarwal, S., S. Chomsisengphet, C. Liu, and N. S. Souleles (2015a): "Do Consumers Choose the Right Credit Contracts?" *The Review of Corporate Finance Studies*, 4, 239–257.

Agarwal, S., S. Chomsisengphet, N. Mahoney, and J. Stroebel (2015b): "Regulating Consumer Financial Products: Evidence from Credit Cards," *The Quarterly Journal of Economics*, 130, 111–164.

Agarwal, S., J. C. Driscoll, X. Gabaix, and D. Laibson (2008): "Learning in the Credit Card Market," Working Paper 13822, National Bureau of Economic Research.

Allen, J., R. Clark, and J.-F. Houde (Forthcoming): "Search Frictions and Market Power in Negotiated Price Markets," *Journal of Political Economy*.

Armstrong, M., J. Vickers, and J. Zhou (2009): "Consumer Protection and the Incentive to Become Informed," *Journal of the European Economic Association*, 7, 399–410.

Ausubel, L. M. (1991): "The Failure of Competition in the Credit Card Market," *The American Economic Review*, 81, 50–81.

Burdett, K. and K. L. Judd (1983): "Equilibrium Price Dispersion," *Econometrica*, 51, 955–969.

Butters, G. R. (1977): "Equilibrium Distributions of Sales and Advertising Prices," *Review of Economic Studies*, 44, 465–491.

Calem, P. S. and L. J. Mester (1995): "Consumer Behavior and the Stickiness of Credit-Card Interest Rates," *The American Economic Review*, 85, 1327–1336.

Campbell, J. Y., H. E. Jackson, B. Madrian, and P. Tufano (2011a): "The Regulation of Consumer Financial Products: An Introductory Essay with Four Case Studies," in *Moving forward: the future of consumer credit and mortgage finance*, ed. by N. P. Retsinas and E. S. Belsky, Brookings Institution Press, chap. 7, 206–244.

Campbell, J. Y., H. E. Jackson, B. C. Madrian, and P. Tufano (2011b): "Consumer Financial Protection," *Journal of Economic Perspectives*, 25, 91–114.

Einav, L., M. Jenkins, and J. Levin (2012): "Contract Pricing in Consumer Credit Markets," *Econometrica*, 80, 1387–1432.

Fershtman, C. and A. Fishman (1994): "The 'Perverse' Effects of Wage and Price Controls in Search Markets," *European Economic Review*, 38, 1099–1112.

Galenianos, M. and A. Gavazza (2017): "A Structural Model of the Retail Market for Illicit Drugs," *American Economic Review*, 107, 858–896.

Gavazza, A. (2016): "An Empirical Equilibrium Model of a Decentralized Asset Market," *Econometrica*, 84, 1755–1798.

Grodzicki, D. (2015): "Competition and Customer Acquisition in the US Credit Card Market," Mimeo, Pennsylvania State University.

Hong, H. and M. Shum (2006): "Using Price Distributions to Estimate Search Costs," *The RAND Journal of Economics*, 37, 257–275.

Hortaçsu, A. and C. Syverson (2004): "Search Costs, Product Differentiation, and Welfare Effects of Entry: A Case Study of S&P 500 Index Funds," *The Quarterly Journal of Economics*, 119, 403–456.

Janssen, M. C. and J. L. Moraga-González (2004): "Strategic Pricing, Consumer Search and the Number of Firms," *The Review of Economic Studies*, 71, 1089–1118.

Nelson, S. (2018): "Private Information and Price Regulation in the US Credit Card Market," Mimeo, MIT.

Ru, H. and A. Schoar (2016): "Do Credit Card Companies Screen for Behavioral Biases?" Working Paper 22360, National Bureau of Economic Research.

Salz, T. (2016): "Intermediation and Competition in Search Markets: An Empirical Case Study," Mimeo, Columbia University.

Sirri, E. R. and P. Tufano (1998): "Costly Search and Mutual Fund Flows," *The Journal of Finance*, 53, 1589–1622.

Stango, V. (2002): "Pricing with Consumer Switching Costs: Evidence from the Credit Card Market," *The Journal of Industrial Economics*, 50, 475–492.

Stango, V. and J. Zinman (2016): "Borrowing High versus Borrowing Higher: Price Dispersion and Shopping Behavior in the US Credit Card Market," *Review of Financial Studies*, 29, 979–1006.

Wildenbeest, M. R. (2011): "An Empirical Model of Search with Vertically Differentiated Products," *The RAND Journal of Economics*, 42, 729–757.

Woodward, S. E. and R. E. Hall (2012): "Diagnosing Consumer Confusion and Sub-Optimal Shopping Effort: Theory and Mortgage-Market Evidence," *The American Economic Review*, 102, 3249–3276.

APPENDIX

A Auxiliary Results

A.1 Distribution of Offers in the Survey

We pointed out in Section 5.2 that it seems plausible that survey respondents report all offers that they receive, not only those that they would consider if they were not surveyed. We now derive the distribution of the number of offers and the distribution of the difference between the offers with the smallest and the largest interest rates under the assumption that respondents report all offers that they receive.

The expected number of offers for a borrower who receives $n \geq 2$ offers is:

$$\mathbf{E}[n|n \geq 2] \;=\; \frac{\Lambda\big(1 - e^{-\Lambda}\big)}{1 - e^{-\Lambda} - \Lambda e^{-\Lambda}}.$$

We obtain the unconditional probability that the difference between the highest and lowest offers is less than x by summing $\text{Prob}[R_H - R_L \leq x|n]$ over all possible numbers of offers, $n \geq 2$, whose arrival rate equals Λ:

$$D(x) \;=\; \frac{\Lambda e^{-\Lambda}}{1 - e^{-\Lambda} - \Lambda e^{-\Lambda}} \int_{\underline{R}}^{\overline{R}} F_R'(R_L)\Big(e^{\Lambda(F_R(R_L+x)-F_R(R_L))} - 1\Big)dR_L.$$

A.2 Distribution of Accepted Offers

We now provide an alternative way to derive the distribution of accepted interest rates by examining borrowers' choices.

Consider a type-z borrower. Let Q_z denote the probability that a type-z borrower does not get a loan. This occurs if he receives/examines no offer or if the cost of all offers is greater than his valuation. Therefore:

$$\begin{aligned}
Q_z &= \sum_{n=0}^{\infty} \frac{e^{-\lambda_z}\lambda_z^n}{n!}\big(1 - F_c(z)\big)^n \\
&= e^{-\lambda_z F_c(z)} \sum_{n=0}^{\infty} \frac{e^{-\lambda_z\left(1-F_c(z)\right)}\Big(\lambda_z\big(1 - F_c(z)\big)\Big)^n}{n!} \\
&= e^{-\lambda_z F_c(z)}
\end{aligned}$$

The share of borrowers without a loan is:

$$Q = \int_{\underline{z}}^{\overline{z}} e^{-\lambda_z F_c(z)} dM(z)$$

The lowest-cost offer when examining n offers is distributed according to $1 - \big(1 - F_c(c)\big)^n$. The cost distribution of accepted offers for type-z borrowers is:

$$\begin{aligned}
H_c(c,z) &= \frac{1}{1-Q_z}\Big(\sum_{n=0}^{\infty} \frac{e^{-\lambda_z}\lambda_z^n}{n!}\Big(1 - \big(1 - F_c(c)\big)^n\Big)\Big) \\
&= \frac{1}{1-Q_z}\Big(1 - e^{-\lambda_z F_c(c)} \sum_{n=0}^{\infty} \frac{e^{-\lambda_z\big(1-F_c(c)\big)}\Big(\lambda_z\big(1-F_c(c)\big)\Big)^n}{n!}\Big) \\
&= \frac{1 - e^{-\lambda_z F_c(c)}}{1-Q_z} \\
&= \frac{1}{1-Q_z}\Big(1 - e^{-\lambda_z \int_{\underline{R}}^{\overline{R}} F_e(c-x)dF_R(x)}\Big), \quad \text{if} \quad c \leq z \\
H_c(c,z) &= 1, \qquad\qquad c > z
\end{aligned}$$

The distribution of accepted interest rates for type-z borrowers is:

$$\begin{aligned}
H_R(R,z) &= \int_{-\infty}^{z-R} H_c(R+e,z)dF_e(e) \\
&= \frac{1}{1-Q_z}\int_{-\infty}^{z-R}\Big(1 - e^{-\lambda_z \int_{\underline{R}}^{\overline{R}} F_e(R+e-x)dF_R(x)}\Big)dF_e(e).
\end{aligned}$$

The distribution of accepted interest rates across all borrowers is:

$$H_R(R) = \frac{1}{1-Q}\int_{\underline{z}}^{\overline{z}} s_z \int_{-\infty}^{z-R}\Big(1 - e^{-\lambda_z \int_{\underline{R}}^{\overline{R}} F_e(R+e-x)dF_R(x)}\Big)dF_e(e)dM(z). \tag{A1}$$