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

Manolis Galenianos¶ Alessandro Gavazza§

March 2019

Abstract

We build a framework to understand the effects of regulatory interventions in credit markets, such as caps on interest rates and higher compliance costs for lenders. We focus on the credit card market, in which we observe a large dispersion of interest rates at which U.S. consumers borrow. Our framework includes two main features that may explain this dispersion: endogenous search effort/inattention and product differentiation. Our calibration suggests that low search effort accounts for almost all the dispersion in interest rates, whereas product differentiation is negligible. The calibrated model implies that price caps may curb lenders' market power, but may also reduce borrowers' search effort, with potentially ambiguous aggregate welfare effects.

PRELIMINARY AND INCOMPLETE

Comments Welcome

*We thank Victor Stango for sharing data with us, Scott Nelson, Matthijs Wildenbeest, and Glenn Ellison for useful comments and suggestions. 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 different forms, depending on the specific countries and on the specific products, with two broad ways being prevalent: 1) regulations 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. 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 financial press reports that 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). 2) Regulatory agencies cracked down on lenders by increasing capital requirements and/or tightening enforcement, thereby increasing their operating costs.

The objective of this paper is to study these two main form of regulations—i.e., the regulation of fees and of 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. Notably, we build on Stango and Zinman (2016) (henceforth SZ) and report a large dispersion in the interest rates that consumers pay on their credit cards, which does not seem to depend on observable borrower characteristics—including their creditworthiness, as captured by their credit score—and card characteristics (such as rewards). Hence, to interpret this dispersion, we develop 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. The literature highlights that these price caps may have unintended consequences, for two main reasons. First, if price caps prevent lenders to recover their costs, they may reduce the supply of credit, most notably to riskier borrowers who have higher default rates. Second, 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 borrowers' search efforts, thereby increasing lenders' market power and, thus, interest rates. 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.

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,

[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).

thereby leading to greater price dispersion and lower welfare. Hence, counterfactual 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 imperfect competition and frictions in credit card markets. 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 case of U.S. credit card markets, 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 administrative data. Similarly, in a contemporaneous contribution, Cuesta and Sepúlveda (2018) studies price regulation in the Chilean consumer loan market. 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 statistics obtained from credit card market reports of the Bureau of Consumer Financial Protection and from 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 pre-approved credit card offers that a sample of individuals received 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

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

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 additional statistics: the fraction of individuals with credit cards, computed from Campbell et al. (2016) and the 2007 Survey of Consumer Finances; the default rate on credit card loans in the first quarter of 2008, reported by Nelson (2018); the interest rate of the one-year Treasury bill on January 16th, 2007, which we use as the risk-free rate;[3] and Standard & Poor's U.S. Credit Card Quality Index at January 2007, which is a monthly performance index that aggregates information of securitized credit card receivables, most notably reporting a weighted average cost of funding (net of charge-offs) for credit card loans.

3.2 Data Description

We use the first dataset on individuals' credit-card terms to sum up and extend 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.

Based on regression equation (1), we calculate the centered interest rate residuals as:

$$R'_{ijt} = \hat{\gamma}_X \overline{X}_{it} + \hat{\gamma}_Z \overline{Z}_{ijt} + \hat{\epsilon}_{ijt}, \tag{2}$$

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.

We perform regression (1) and calculate interest rate residuals according to equation (2)

[3] We retrieved this value from FRED, Federal Reserve Bank of St. Louis, series `https://fred.stlouisfed.org/series/DGS1`.

separately for four different groups of cardholders based on their FICO score: 1) sub-prime borrowers, with FICO score strictly below 620; 2) near-prime borrowers, with FICO scores between 620 and 679; 3) prime borrowers, with FICO scores between 680 and 739; and 4) super-prime borrowers, with FICO scores above 740. These different groups constitute the main classification of borrowers used in the credit card industry (Bureau of Consumer Financial Protection, 2015), and performing separate regressions for each group allows us to capture in a flexible way the heterogeneity across these submarkets and, thus, to obtain a reasonably accurate measure of the dispersion in interest rates in each group of borrowers.

Table 1 reports coefficient estimates of several specifications of equation (1) and the main percentiles of the resulting distribution of interest rates based on equation (2). Column (1) uses the raw data over the entire sample period, which exhibit a large dispersion of interest rates: the difference between the 90th and the 10th percentiles equals 18 percentage points for subprime borrowers, it decreases for more-creditworthy borrowers, reaching a difference of 10 percentage points for super-prime borrowers. 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, e.g. the difference between the 90th and the 10th percentiles slightly decreases to 16 percentage points for subprime borrowers and 9 percentage points for super-prime borrowers. The specification of column (4) introduces the main individual characteristic that should affect pricing, i.e., the credit risk of the individual, measured by the FICO score; remarkably, the corresponding distribution of residual interest rates constructed as in equation (2) displays a dispersion that is almost identical to that computed from the raw data in column (3). 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); nevertheless, the large dispersion of residual interest rates persists almost unaffected.

Overall, Table 1 attests some striking features of credit card markets. First, while more creditworthy borrowers on average pay lower interest rates, the difference in interest rates

Table 1: Dispersion of Interest Rates, by Borrower Group

Subprime Borrowers	(1)	(2)	(3)	(4)	(5)	(6)
FICO Score				-0.018	-0.017	-0.018
				(0.006)	(0.006)	(0.006)
Reward Card					0.863	0.405
					(0.618)	(0.674)
Credit Limit					-0.170	-0.158
					(0.094)	(0.125)
Credit Balance					0.197	0.174
					(0.110)	(0.140)
R^2				0.010	0.015	0.013
Observations	27,024	903	877	877	871	766
10th Percentile	11.90	12.74	14.24	14.30	14.41	14.39
25th Percentile	16.15	16.95	17.24	17.55	17.46	17.58
50th Percentile	20.65	21.20	21.74	21.72	21.65	21.93
75th Percentile	27.49	27.99	27.99	27.52	27.45	27.80
90th Percentile	29.99	30.24	30.24	30.06	30.01	30.16
Near-prime Borrowers						
FICO Score				-0.046	-0.043	-0.052
				(0.011)	(0.011)	(0.013)
Reward Card					0.494	0.562
					(0.453)	(0.565)
Credit Limit					-0.211	-0.255
					(0.046)	(0.078)
Credit Balance					0.242	0.225
					(0.064)	(0.100)
R^2				0.019	0.044	0.043
Observations	27,059	944	900	900	885	661
10th Percentile	10.49	11.24	12.99	13.16	13.09	13.20
25th Percentile	14.90	14.99	15.94	16.01	16.06	16.55
50th Percentile	18.24	18.99	19.24	19.09	19.31	20.20
75th Percentile	23.15	23.24	23.30	24.07	23.87	25.72
90th Percentile	28.99	29.24	29.24	29.03	28.75	29.16
Prime Borrowers						
FICO Score				-0.053	-0.048	-0.051
				(0.011)	(0.011)	(0.015)
Reward Card					0.084	-0.257
					(0.374)	(0.520)
Credit Limit					-0.093	-0.070
					(0.026)	(0.049)
Credit Balance					0.080	0.019
					(0.038)	(0.058)
R^2				0.024	0.039	0.029
Observations	31,115	1,003	953	953	932	605
10th Percentile	9.90	9.90	11.99	11.79	11.85	11.56
25th Percentile	12.99	13.99	14.31	14.65	14.64	14.81
50th Percentile	16.74	17.31	18.24	17.74	17.70	17.93
75th Percentile	19.99	20.24	20.34	21.04	20.97	21.90
90th Percentile	25.99	27.24	28.15	27.32	27.25	28.68
Super-prime Borrowers						
FICO Score				-0.008	-0.008	-0.024
				(0.004)	(0.005)	(0.010)
Reward Card					0.442	0.346
					(0.220)	(0.471)
Credit Limit					0.010	0.028
					(0.014)	(0.031)
Credit Balance					0.058	-0.040
					(0.032)	(0.051)
R^2				0.002	0.007	0.012
Observations	56,880	1,741	1,645	1,645	1,611	546
10th Percentile	9.90	9.90	11.24	11.33	11.20	10.79
25th Percentile	12.99	13.24	14.15	14.05	13.97	13.82
50th Percentile	15.98	16.74	16.99	17.02	16.91	16.84
75th Percentile	18.24	18.24	18.24	18.48	18.52	19.54
90th Percentile	20.24	20.24	20.31	20.56	20.81	23.98

Notes: This table reports.

within groups is substantially larger than the difference across groups. Notably, the difference between the 90th and the 10th percentiles equals approximately 16 percentage points for subprime borrowers, near-prime borrowers, and prime borrowers, whereas it equals approximately 13 percentage points for super-prime borrowers. Second, observable card characteristics do not seem to have a noticeable effect on card pricing. A consequence of these two features is that a large dispersion of interest rates persists once we account for borrower and card characteristics. 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 combines all empirical targets of our quantitative model. Panel A reproduces the percentiles of the distributions 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 the lowest offered interest rates charged after the expiration of any introductory "teaser" period (if any).

Panel C reports auxiliary statistics on credit card markets. We compute the fraction of credit card revolvers in each group by combining the share of individuals with a credit card in 2007 reported in Campbell et al. (2016) with the probability of revolving conditional on having a credit card, which we compute directly as the ratio of observations in column (6) to observations in column (5) in Table 1. Interestingly, the share of individuals with a credit card is higher for borrowers with higher credit scores, whereas the the probability of revolving conditional on having a credit card is lower; hence, the fraction of credit card borrowers is non-monotonic as borrowers' credit scores increase.[4] Moreover, we obtain the charge-off rate for different borrowers based on their FICO scores from Nelson (2018), which we aggregate in our four groups using the distribution of FICO scores reported in Bureau of Consumer Financial Protection (2012). Finally, the average funding cost reported by Standard & Poor's Credit Card Quality Index is approximately 200 basis points above the risk-free rate.

Table 2 provides a rich description of the credit card market. First, as we noted above, Panel A shows that the dispersion in the interest rates that borrowers pay on their credit card

[4]The aggregate share of the population with a credit card and the aggregate share of revolvers, computed as the weighted averages of the corresponding group shares in our data, equal 0.761 and 0.457, respectively. These shares closely match the corresponding aggregate statistics in the 2007 Survey of Consumer Finances, which equal 0.729 and 0.439, respectively.

Table 2: Empirical Targets

Panel A: Accepted Offers	
10th Percentile Accepted Offer Distribution, Sub-Prime Borrowers	14.39
25th Percentile Accepted Offer Distribution, Sub-Prime Borrowers	17.58
50th Percentile Accepted Offer Distribution, Sub-Prime Borrowers	21.93
75th Percentile Accepted Offer Distribution, Sub-Prime Borrowers	27.80
90th Percentile Accepted Offer Distribution, Sub-Prime Borrowers	30.16
10th Percentile Accepted Offer Distribution, Near-Prime Borrowers	13.20
25th Percentile Accepted Offer Distribution, Near-Prime Borrowers	16.55
50th Percentile Accepted Offer Distribution, Near-Prime Borrowers	20.20
75th Percentile Accepted Offer Distribution, Near-Prime Borrowers	25.72
90th Percentile Accepted Offer Distribution, Near-Prime Borrowers	29.16
10th Percentile Accepted Offer Distribution, Prime Borrowers	11.56
25th Percentile Accepted Offer Distribution, Prime Borrowers	14.81
50th Percentile Accepted Offer Distribution, Prime Borrowers	17.93
75th Percentile Accepted Offer Distribution, Prime Borrowers	21.90
90th Percentile Accepted Offer Distribution, Prime Borrowers	28.68
10th Percentile Accepted Offer Distribution, Super-Prime Borrowers	10.79
25th Percentile Accepted Offer Distribution, Super-Prime Borrowers	13.82
50th Percentile Accepted Offer Distribution, Super-Prime Borrowers	16.84
75th Percentile Accepted Offer Distribution, Super-Prime Borrowers	19.54
90th Percentile Accepted Offer Distribution, Super-Prime Borrowers	23.98
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: Auxiliary Statistics	
Fraction with Credit Card Debt, Sub-Prime Borrowers	54.56
Fraction with Credit Card Debt, Near-Prime Borrowers	55.33
Fraction with Credit Card Debt, Prime Borrowers	54.00
Fraction with Credit Card Debt, Super-prime Borrowers	36.02
Charge-Off Rate, Sub-prime Borrowers	24.20
Charge-Off Rate, Near-prime Borrowers	7.62
Charge-Off Rate, Prime Borrowers	2.93
Charge-Off Rate, Super-prime Borrowers	0.83
Average Funding Cost	7.14

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.

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 may 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. Finally, the auxiliary statistics in Panel C give us precise information on the penetration of credit card loans across borrower groups, as well as on lenders' marginal costs of providing loans to these different groups of borrowers. In summary, all these data are well suited 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 cannot precisely assess 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 within group. We discuss the implications of these data limitations for our results further in Section 7.

4 The Model

The economy consists of J different markets, labeled by j, which are populated by borrowers and lenders. The different markets operate independently from each other and each agent (borrower or lender) is allocated to a single market.

Each market j has measure 1 of borrowers (a normalization) who have a market-specific default risk ρ_j and are heterogeneous in their marginal valuation of a loan, $\tilde{z}$.[5] The marginal valuation $\tilde{z}$ is distributed according to a market-specific smooth distribution $\tilde{M}_j(\cdot)$ with con-

[5] For the empirical work we consider four different borrower groups, as identified by the credit card industry: subprime, near-prime, prime and superprime.

nected support $[\underline{\tilde{z}}_j, \overline{\tilde{z}}_j]$. Every borrower wants to take a loan of size b_j.

Each market j has measure Λ_j of *potential* lenders who face entry cost χ_j to enter the market and are heterogeneous in their marginal cost of providing a loan, k. The marginal cost k is distributed according to a market-specific smooth distribution $\Gamma_j(\cdot)$ with connected support $[\underline{k}_j, \overline{k}_j]$. Every lender can give one loan of size b_j.[6] The measure of lenders who choose to enter market j and the distribution of their marginal costs are L_j and $G_j(\cdot)$, respectively.

Matching between borrowers and lenders in a market 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 which follows a Poisson distribution, examines every offer with some probability that depends on his search effort and decides which, if any, offer to accept.[7] The effective number of offers (i.e., offers received *and* examined) to a borrower who exerts effort s in a market with L_j lenders follows a Poisson distribution with parameter $\alpha_j(s, L_j)$. The function $\alpha_j(\cdot,\cdot)$ is strictly increasing and concave in both arguments and satisfies $\alpha_j(0, L_j) = 0$ and $\alpha_j(s, L_j) \leq L_j$ for all s. A borrower who exerts search effort s incurs cost $q_j(s, L_j)$ where $q(\cdot,\cdot)$ is strictly increasing and convex in both arguments and satisfies $q_j(0, L_j) = 0$ and $\lim_{s\to 0} \partial q_j(s, L_j)/\partial s = 0$.

Borrowers rank loan offers considering the sum of two components. The first component is the interest rate R, which is chosen by the lender and is drawn from the equilibrium distribution $F_{R_j}(\cdot)$. The second component is an idiosyncratic utility draw e, which is stochastic and represents every other aspect of the loan that might affect the borrower's valuation. The idiosyncratic utility component e is drawn from an exogenous distribution $F_{e_j}(\cdot)$ which is smooth, has zero mean and has support in a connected subset of $(-\infty, \infty)$. We call the sum $c = R + e$ as the cost of a loan, which might be higher or lower than R depending on the realization of the idiosyncratic draw.

The borrower only pays the cost if he does not default which occurs with probability $1 - \rho_j$, while if he defaults, which occurs with probability ρ_j, he incurs the utility cost of default $\tilde{\delta}_j$.[8] The utility of a type-$\tilde{z}$ borrower in market j who takes a loan with interest rate R and idiosyncratic draw e is $b_j\Big(\tilde{z} - (1-\rho_j)(R+e) - \rho_j\delta_j\Big)$, where $\delta_j \equiv \frac{\tilde{\delta}_j}{b_j}$. The borrower's utility from not taking a loan is zero.

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

[7] The random allocation of offers across borrowers in a context with finite numbers of borrowers and lender leads to urn-ball matching which, as the numbers of borrowers and lenders grows to infinite, is approximated by a Poisson distribution. See Butters (1977).

[8] We assume that defaulting occurs independently of any loan features, e.g. the interest rate R or the idiosyncratic shock e.

We define a borrower's preference net of expected default cost by $z = \frac{\tilde{z}-\rho_j\delta_j}{1-\rho_j}$ and note that it is distributed according to $M_j(z) = \tilde{M}_j\big((1-\rho_j)z + \rho_j\delta_j\big)$ with support $[\underline{z}_j, \overline{z}_j]$ where $\underline{z} = \frac{\underline{\tilde{z}}_j-\rho_j\delta_j}{(1-\rho_j)}$ and $\overline{z}_j = \frac{\overline{\tilde{z}}_j-\rho_j\delta_j}{(1-\rho_j)}$. We can, therefore, rewrite the utility of a type-z borrower from taking a loan with cost $R+e$ as:

$$b_j\,(1-\rho_j)\,(z-R-e)\,.$$

Anticipating equilibrium behavior, a type-z borrower chooses the loan offer with the lowest cost among the offers that he examines, conditional on the cost being less than z. A loan offer with a higher cost generates negative utility and, thus, the borrower will never accept it.

The ex ante value of a type-z borrower in market j equals the expected value of his best loan $V_{z_j}(s)$ (which depends on search effort s) net of the cost of search effort, $q(s, L_j)$:

$$V_{z_j}(s) - q(s, L_j). \tag{3}$$

We denote the optimal (utility-maximizing) effort choice of a type-z borrower in market j by $s_j(z)$.

The revenues per-dollar-lent for a lender in market j equal the interest rate R net of the market-specific default risk ρ_j: $R(1-\rho_j)$. The lender's profits conditional on giving a loan equal revenues minus the lender's marginal cost, k. The expected profits of a type-k lender from market j who offers interest rate R, $\pi_{k_j}(R)$, are given by the probability of making a loan, denoted by $P_j(R)$, times the loan's profits:

$$\pi_{k_j}(R) \;\;=\;\; b_j\,(R(1-\rho_j)-k)\,P_j(R). \tag{4}$$

Notice that the borrowers' idiosyncratic shock affects the lender's payoff only through the probability of making a loan.

We denote the optimal (profit-maximizing) interest rate choice of a type-k lender in market j by $R_j(k)$ which, combined with lenders' entry decisions, determines the interest rate distribution in market j, $F_{R_j}(\cdot)$.

We are now ready to define the equilibrium.

Definition 1 *An equilibrium consists of borrowers' search effort $s_j(\cdot)$ and lenders' entry and interest rate choices $\{L_j, G_j(\cdot), R_j(\cdot)\}$ such that in every market j borrowers maximize their ex ante value* (3)*, lenders maximize their expected profits* (4)*, the expected profits of all entrants* *exceed the entry cost χ_j and the expected profits of non-entrants would be strictly below χ_j if they entered.*

Since there is no interaction across markets, we henceforth drop the j subscript to ease notation. The reader should keep in mind, however, that all equilibrium outcomes are market-specific.

To proceed, we first determine the borrowers' and lenders' optimal choices separately and then prove the existence of equilibrium. Finally, we characterize the constrained efficient outcome.

4.1 Borrowers' Choices

We characterize borrowers' optimal search effort $s(\cdot)$ taking as given the measure L of lenders in the market and the interest rate distribution $F_R(\cdot)$, since the type distribution of lenders $G(\cdot)$ and interest rate choices $R(\cdot)$ affect borrowers' choices only through $F_R(\cdot)$.

We begin by rewriting $V_z(s)$ in a more convenient way. Denote the value of a z-borrower from receiving n offers by $v_{z,n}$, where $v_{z,0} = 0$. The expected value for a type-z borrower who exerts search effort s is:

$$V_z(s) \quad = \quad \sum_{n=0}^{\infty} \frac{e^{-\alpha(s,L)}\alpha(s,L)^n}{n!} v_{z,n}. \tag{5}$$

Notice that search effort affects the arrival rate of offers, but does not enter $v_{z,n}$ and, therefore, it is immediate from equation (5) that $V_z(s)$ is continuous and differentiable in s. As a result, the optimal effort choice $s(z)$ solves:

$$V_z'(s) \quad = \quad \frac{\partial q(s,L)}{\partial s}. \tag{6}$$

To determine $v_{z,n}$ for $n \geq 1$, recall that the borrower chooses the loan offer with the lower cost c. Let $F_c(\cdot)$ denote the distribution of c. Since the 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:

$$\bar{F}_{c,n}(c) \;\; = \;\; 1 - (1 - F_c(c))^n \,.$$

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

$$v_{z,n} \;\; = \;\; b(1-\rho) \int_{-\infty}^{z} (z - c) \, d\bar{F}_{c,n}(c). \tag{7}$$

The following proposition characterizes $s(\cdot)$.

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

$$\sum_{n=0}^{\infty} \frac{e^{-\alpha(s,L)} \alpha(s,L)^n}{n!} \left(v_{z,n+1} - v_{z,n}\right) \frac{\partial \alpha(s,L)}{\partial s} \;\; = \;\; \frac{\partial q(s,L)}{\partial s}, \tag{8}$$

where equation (7) define $v_{z,0} = 0$ and $v_{z,n}$ defined for $n \geq 1$.

4.2 Lenders' Choices

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

To ease notation, we denote the effective arrival rate of offers to a type-z borrower by $\alpha(z) \equiv \alpha\big(s(z), L\big)$.

A borrower accepts a loan offer with interest rate R if he examines this offer, if this offer yields the lowest cost from every offer that he examines (taking into account the idiosyncratic components of all offers), and if this offer yields net positive utility to the borrower. The next lemma characterizes the probability $P(R)$ that any borrower accepts a loan offer with interest rate R.

Lemma 3 *Given $F_R(\cdot)$, L and $s(\cdot)$, the probability $P(R)$ that borrowers accept a loan offer with interest rate R is continuous and differentiable in R and equals:*

$$P(R) \;\; = \;\; \int_{\underline{z}}^{\overline{z}} s(z) \int_{-\infty}^{z-R} e^{-\alpha(z) \int_{\underline{R}}^{\overline{R}} F_e(R+e-x) dF_R(x)} dF_e(e) dM(z). \tag{9}$$

Furthermore $P'(R) < 0$.

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

Proposition 4 *Given L, $G(\cdot)$ and $s(\cdot)$:*

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

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

$$\int_{\underline{z}}^{\overline{z}} s(z) \int_{-\infty}^{z-R(k)} e^{-\alpha(z)\int_{\underline{k}}^{\overline{k}} F_e(R(k)+e-R(x))dG(x)} dF_e(e)dM(z)$$
$$= \left(R(k) - \frac{k}{1-\rho}\right) \int_{\underline{z}}^{\overline{z}} s(z) \left(\int_{-\infty}^{z-R(k)} e^{-\alpha(z)\int_{\underline{k}}^{\overline{k}} F_e(R(k)+e-R(x))dG(x)} \left(\alpha(z) \int_{\underline{k}}^{\overline{k}} F_e'\left(R(k)\right.\right.\right.$$
$$\left.\left.\left. +e - R(x)\right) dG(x) \right) dF_e(e) + e^{-\alpha(z)\int_{\underline{k}}^{\overline{k}} F_e(z-R(x))dG(x)} F_e'(z - R(k)) \right) dM(z). \tag{10}$$

3. *The interest rate distribution equals $F_R(x) = G\left(R^{-1}(x)\right).$*

4. *The distribution of accepted offers equals*

$$H_R(R) = \frac{1 - \int_{\underline{z}}^{\overline{z}} e^{-\alpha(z)\int_{\underline{R}}^{R} F_e(z-x)dF_R(x)} dM(z)}{1 - \int_{\underline{z}}^{\overline{z}} e^{-\alpha(z)\int_{\underline{R}}^{\overline{R}} F_R(z-x)dF_R(x)} dM(z)} \tag{11}$$

The following proposition completes the characterization of lenders' entry decisions.

Proposition 5 *Given $s(\cdot)$, lenders' entry satisfies:*

1. *There is a unique cutoff cost $\hat{k}$ such that a lender enters if and only if $k \leq \hat{k}$.*

2. *The measure of lenders in the market equals $L = \Lambda\Gamma(\hat{k})$ and the cost distribution of entrants equals $G(k) = \frac{\Gamma(k)}{\Gamma(\hat{k})}$ for $k \leq \hat{k}$ and $G(k) = 1$ for $k > \hat{k}$.*

3. *The cutoff cost $\hat{k}$ solves:*

$$b\left(R(\hat{k})(1-\rho) - \hat{k}\right) \int_{\underline{z}}^{\overline{z}} s(z) \int_{-\infty}^{z-R(\hat{k})} e^{-\alpha\left(s(z),\Lambda\Gamma(\hat{k})\right)\int_{\underline{k}}^{\hat{k}} F_e\left(R(\hat{k})+e-R(x)\right)d\frac{\Gamma(x)}{\Gamma(\hat{k})}} dF_e(e)dM(z) = \chi. \tag{12}$$

4.3 Existence of the Equilibrium

Based on the preceding results, we can describe the equilibrium as a fixed point in the space of continuous functions $s(\cdot)$ from $[\underline{z}, \overline{z}]$ to $[0, 1]$. From an arbitrary initial $s(\cdot)$, Propositions 4 and 5 characterize the lenders' entry and interest rate decisions $\{L, G(\cdot), F_R(\cdot)\}$; given lenders' actions, Proposition 2 characterizes borrowers' search effort $s(\cdot)$. The space is complete under the sup norm and, therefore, a fixed point (and an equilibrium) exists.

Proposition 6 *An equilibrium exists.*

4.4 Constrained Efficiency

We now analyze the case of a social planner whose goal is to maximize aggregate welfare subject to frictions. The planner chooses lenders' entry decisions, as well as borrowers' search effort and trading decision rules (the interest rate simply redistributes surplus between borrowers and lenders and, thus, it does not matter for the planner's problem).

We denote the planner's optimal solution for the search effort of a type-z borrower by $s^*(z)$. The entry decision rule is, trivially, a cutoff rule and we denote the planner's optimal cutoff cost by $\hat{k}^*$. The surplus of a loan from a type-k lender with idiosyncratic cost e to a type-z borrower is $b(1-\rho)\big(z - e - \frac{k}{1-\rho}\big)$ where $z = \frac{\tilde{z}-\rho\delta}{1-\rho}$, as before, and the entry cost k is paid regardless of default. The planner's optimal trading decision rule is that the borrower trades with the lowest-cost lender as long as the surplus is positive.

The overall surplus for borrower search effort $s(\cdot)$ and lender entry cutoff $\hat{k}$ is:

$$\mathcal{W}\left(s(\cdot), \hat{k}\right) = \int_{\underline{z}}^{\overline{z}} \left(\mathcal{W}_z\left(s(z), \hat{k}\right) - q\left(s(z), L\right)\right) dM(z) - \chi L, \tag{13}$$

where $\mathcal{W}_z(s(z), \hat{k})$ is the expected surplus that a type-z borrower obtains from the offers that he receives; $L = \Lambda\Gamma(\hat{k})$ is the measure of lenders who enter the market; $q\big(s(z), L\big)$ is a z-borrower's cost of search effort; and χL are aggregate lenders' entry costs. Notice that there is no interaction among borrowers regarding their search effort and, thus, $\mathcal{W}_z$ only depends on the search effort of the type-z borrower and does not depend on the full search effort schedule.

The cost of a loan for the planner is $w \equiv \frac{k}{1-\rho} + e$. The planner's loan cost is distributed according to:

$$F_w(w) = \int_{\underline{k}}^{\hat{k}} F_e\left(w - \frac{k}{1-\rho}\right) dG(k)$$
$$= \frac{1}{\Gamma(\hat{k})} \int_{\underline{k}}^{\hat{k}} F_e\left(w - \frac{k}{1-\rho}\right) d\Gamma(k).$$

The distribution of the lowest w among n offers is:

$$\bar{F}_{w,n}(w) = 1 - \left(1 - \frac{1}{\Gamma(\hat{k})} \int_{\underline{k}}^{\hat{k}} F_e\left(w - \frac{k}{1-\rho}\right) d\Gamma(k)\right)^n.$$

The social value to a type-z borrower of receiving n offers is:

$$\mathcal{W}_{z,n}(\hat{k}) = b(1-\rho) \int_{-\infty}^{z} (z-w) d\bar{F}_{w,n}(w), \tag{14}$$

where $\mathcal{W}_{z,0}(\hat{k}) = 0$. Notice that these terms only depend on $\hat{k}$ and do not depend on search effort.

The surplus that a type-z borrower with search effort s generates from the offers that he receives when lenders' entry cutoff is $\hat{k}$ equals:

$$\mathcal{W}_z(s, \hat{k}) = \sum_{n=0}^{\infty} \frac{e^{-\alpha(s,L)}(\alpha(s,L))^n}{n!} \mathcal{W}_{z,n}(\hat{k}). \tag{15}$$

We now characterize the planner's optimal solution.

Proposition 7 *The constrained-efficient allocation is as follows:*

1. *The optimal search effort for a type-z borrower $s^*(z)$ satisfies:*

$$\sum_{n=0}^{\infty} \frac{e^{-\alpha(s^*(z),L^*)}(\alpha(s^*(z),L^*))^n}{n!} \left(\mathcal{W}_{z,n+1}(\hat{k}^*) - \mathcal{W}_{z,n}(\hat{k}^*)\right) \frac{\partial \alpha(s^*(z), L^*)}{\partial s} = \frac{\partial q\left(s^*(z), L^*\right)}{\partial s}, \tag{16}$$

where $\mathcal{W}_{z,n}(\hat{k}^*)$ is given by equation (14) and $L^* = \Lambda\Gamma(\hat{k}^*)$ is the optimal measure of lenders in the market. There is a unique solution $s^*(z)$ for each z.

2. *The optimal entry cost cutoff $\hat{k}^*$ for lenders satisfies:*

$$\int_{\underline{z}}^{\overline{z}} \left(\frac{\partial \mathcal{W}_z\left(s^*(z), \hat{k}^*\right)}{\partial \hat{k}} - \frac{\partial q(s^*(z), L^*)}{\partial L} \Lambda \Gamma'(\hat{k}^*) \right) dM(z) \;\; = \;\; \chi \Lambda \Gamma'(\hat{k}^*), \qquad (17)$$

where

$$\begin{aligned} \frac{\partial \mathcal{W}_z(s^*(z), \hat{k}^*)}{\partial \hat{k}} \;\; = \;\; & \sum_{n=0}^{\infty} \frac{e^{-\alpha(s^*(z),L^*)} \left(\alpha(s^*(z), L^*)\right)^n}{n!} \left(b \int_{-\infty}^{z} (z-w) d \left(\frac{\partial \bar{F}_{w,n}(w)}{\partial \hat{k}} \right) \right. \\ & \left. + \frac{\partial \alpha(s^*(z), L^*)}{\partial L} \Lambda \Gamma'(\hat{k}^*) \left(\mathcal{W}_{z,n+1}(\hat{k}^*) - \mathcal{W}_{z,n}(\hat{k}^*) \right) \right), \end{aligned} \qquad (18)$$

and

$$\begin{aligned} \frac{\partial \bar{F}_{w,n}(w)}{\partial \hat{k}} \;\; = \;\; & n \left(1 - \frac{1}{\Gamma(\hat{k})} \int_{\underline{k}}^{\hat{k}} F_e \left(w - \frac{k}{1-\rho} \right) d\Gamma(k) \right)^{n-1} \frac{\Gamma'(\hat{k})}{\Gamma(\hat{k})^2} \left(F_e \left(w - \frac{\hat{k}}{1-\rho} \right) \Gamma(\hat{k}) \right. \\ & \left. - \int_{\underline{k}}^{\hat{k}} F_e \left(w - \frac{k}{1-\rho} \right) d\Gamma(k) \right). \end{aligned}$$

The decentralized equilibrium of the economy features two potential sources of inefficiencies relative to the planner's allocation. First, for a given measure of lenders, some meetings in which it is efficient to trade (i.e., $z > \frac{k}{1-\rho} + e$) feature no trade because the interest rate of lenders is excessive (i.e., $R > z - e$) due to lenders' market power. Second, the measure of lenders is not optimal (i.e., $L \neq L^*$).

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 parametric assumptions for each of the four separate submarkets—i.e., sub-prime, near-prime, prime and super-prime borrowers.

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_j} and σ_{z_j} for the distribution of buyers' preferences z_j in submarket j. Moreover, we assume that sellers' costs k are common across submarkets and follow a right-truncated Pareto distribution with shape ξ, scale equal to the risk-free rate—we use the interest rate of the one-year Treasury bill at January 16th, 2007, which equals 5.06 percent—and an upper-truncation point $\hat{k}$. The assumption of a common cost distribution across markets means that we are implicitly assuming that the mass of potential lenders Λ_j varies across markets.

We normalize the size of the loan to $b = 1$. We further assume that: 1) the rate $\alpha(s, L_j)$ equals sL_j; 2) the cost of effort $q(s, L_j)$ equals $\beta_{0j}s^{\beta_1}$; 3) the charge-off rate in submarket j equals ρ_j; and 4) the product differentiation parameter e has a normal distribution with mean zero and standard deviation σ_{e_j}..

Finally, we calibrate two versions of the model: in the first one, we assume that the interest rates on accepted offers are measured without errors; in the second one, we assume that the interest rates on accepted offers are measured with errors. Specifically, in the second case, we assume that the reported accepted rates $\hat{R}_j$ and the "true" accepted rate R are related as: $\hat{R}_j = R_j\eta$, where η is a measurement error, drawn from a lognormal distribution with parameters (μ_η, σ_η), common across markets, 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 literature that structurally estimates search models of the labor market frequently assumes measurement error on wages. In our application, it is possible 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, such as adjustment of the interest rate after the offer is accepted as in Nelson (2018), thereby allowing us to fit the distributions of accepted rates better. For example, Table 2 shows that these distribution display a large dispersion, and the measurement η allows the model to more-precisely match this feature of the data quantitatively.

5.2 Calibration

We choose the vector $\psi = \left\{L_j, \mu_{z_j}, \sigma_{z_j}, \xi, \hat{k}, \rho_j, \sigma_{e_j}, \beta_{0j}, \beta_1, \sigma_\eta\right\}_{j \in J}$ that minimizes the distance between the target moments m reported in Table 2 and the corresponding moments of the model. In the calibration without measurement error, we impose $\sigma_\eta = 0$.

Specifically, for any value of the vector ψ, we solve the model of Section 4 to find its equilibrium: the distribution $F_{R_j}(k)$ of offered interest rates and borrowers' search effort $\alpha_(s(z))$ in each market j that are consistent with each other. Once we solve for these optimal policy functions of borrowers and lenders in each market j, we compute the equilibrium distributions of interest rates of received offers and of accepted offers. In practice, we simulate these distributions and compute 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 for each market j. Specifically, Panel A and Panel C in Table 2 report the distribution of accepted interest rates and the charge-off rate, respectively, for each group j, whereas Panel B reports moments of the distribution of the number and of the offered rates aggregated for the entire market. Hence, we use weights ω_j corresponding to the population share of each group j, obtained from Bureau of Consumer Financial Protection (2012), to aggregate the number of received offers and the distribution of their interest rates.

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; in practice, we use the identity matrix.

5.3 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 in group j receive, whose arrival rate is $\alpha(s = 1, L_j)$, 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 of these offers equal $\alpha\left(s=1, L_j\right)$.

However, lenders send these offers anticipating that borrowers will consider them according to their equilibrium $\alpha\left(s(z), L_j\right)$ Hence, the moments on the empirical distribution of accepted offers reflect consumers' endogenous search effort $\alpha\left(s(z), L_j\right)$.

5.4 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.

We obtain the default parameters ρ_j directly from the charge-off statistics, and we identify the parameter ξ of the distribution $G(k)$ of sellers' heterogeneity from the average funding cost reported by Standard and Poor's. The moments on the number of offers that borrowers receive identify the average offer rate $\sum_j \omega_j L_j$ and, thus, contribute to identify the group-specific offer rates L_j.

Thus, we know lenders' marginal costs and we observe how many offers they send out, along with the distribution of the difference between the highest and lowest offered interest rates R. Appendix A shows how the distribution of the difference between the highest and lowest offered interest rates depends on the distribution $F_{R_j}(x)$ of offered rates and, thus, we can recover the distribution $F_{R_j}(x)$. Moreover, the distribution of offer rates and the distribution of accepted rates in each group, along with its fraction of borrowers, identify the parameters of the distribution $M_j(z)$ of buyers' heterogeneity, since lenders offer interest rates R and borrowers accept them only if these rates (together with the product differentiation e) are lower than their willingness to pay z.

The comparison between lenders' costs k and interest rates R, adjusted for the default rate ρ_j, indicates that differences between them—i.e., markups—are large. Moreover, the comparison between the distribution of offered rates and the distribution of accepted offers indicate that many borrowers do not choose the lowest interest-rate offer. The model provides two different explanations for these data features: 1) product differentiation, captured by the parameter e; or 2) high search costs β_{0j}. However, two additional features of the data are at odds with the product differentiation explanation, whereas they are consistent with the search cost explanation: the moderate fraction of borrowers and the high dispersion of offered rates. Specifically, when the product differentiation parameter e in borrowers' utility is large:

a) Holding the distribution of offers F_{R_j} fixed, borrowers search aggressively and, thus, in equilibrium a large fraction of them ends up borrowing on a credit card. However, the data report that, on aggregate, only 40 percent of individuals borrow on credit cards. b) Low-cost lenders have little incentives to offer low interest rates. More precisely, product differentiation softens price competition, most notably for low-cost lenders, who increase their offered interest rates as σ_{e_j} increases. Thus, in equilibrium, the offered interest rate distribution F_{R_j} displays a lower dispersion when σ_{e_j} is larger. In summary, the model with large values of the product differentiation parameter e cannot simultaneously match the moderate fraction of borrowers and the high dispersion of offered rates that we observe in the data, whereas the model with high search costs β_{0j} can match these features quite precisely, as we document in the next section. The comparative statics of Section 5.7 and Figures 4 and 5 will further clarify the effects of product differentiation and of search cost on market outcomes.

Finally, lenders' free-entry condition (equation (12)) implies that we can recover lenders' fixed costs χ_j from the variable profits of the highest-cost lender in each market.

5.5 Calibrated Parameters and Model Fit

Table 3 reports the calibrated parameters of the model. Panel A refers to the version without measurement error ($\sigma_\eta = 0$) and Panel B to the version with measurement error ($\sigma_\eta > 0$). Overall, the parameters are almost identical across versions and, as we recount above and Table 4 will show in detail, the measurement error η allows the model to capture the dispersion of accepted offers more precisely.

The parameters μ_{z_j} and σ_{z_j} of the distributions of z in group j mean that borrowers' willingness to pay for credit is, on average, large and displays large heterogeneity within group, as well as across groups. Specifically, borrowers' average willingness to pay decreases as their creditworthiness increases. The standard deviation of the willingness to pay is non-monotonic in creditworthiness, with super-prime borrowers displaying a standard deviation more than four times larger than that of near-prime borrowers.

The parameters ξ and $\hat{k}$ of the distribution of costs k imply that lenders' average costs equal 640 basis points and, thus, they display a small spread of approximately 140 basis points over the risk-free rate. Moreover, the heterogeneity of lenders' costs in small—i.e., the standard deviation of costs equals 103 basis points. Thus, the model generates a large dispersion of offered rates even with a small dispersion of costs.

The values of L_j indicates that lenders send, on average, approximately 2.7 credit card

Table 3: Calibrated Parameters

Panel A: No Measurement Error			
μ_{z_1}	3.663	σ_{z_1}	0.125
μ_{z_2}	3.568	σ_{z_2}	0.091
μ_{z_3}	3.511	σ_{z_3}	0.120
μ_{z_4}	3.252	σ_{z_4}	0.239
ξ	3.268	$\hat{k}$	9.786
L_1	1.558	L_2	3.868
L_3	3.126	L_4	2.889
ρ_1	0.242	ρ_2	0.076
ρ_3	0.029	ρ_4	0.008
σ_{e_1}	0.100	σ_{e_2}	0.138
σ_{e_3}	0.150	σ_{e_4}	0.141
β_{01}	8.207	β_{02}	40.390
β_{03}	29.981	β_{04}	31.922
β_1	1.686	σ_η	0.000

Panel B: Measurement Error			
μ_{z_1}	3.640	σ_{z_1}	0.130
μ_{z_2}	3.553	σ_{z_2}	0.095
μ_{z_3}	3.497	σ_{z_3}	0.127
μ_{z_4}	3.247	σ_{z_4}	0.202
ξ	3.559	$\hat{k}$	9.711
L_1	1.577	L_2	3.888
L_3	3.143	L_4	2.898
ρ_1	0.242	ρ_2	0.076
ρ_3	0.029	ρ_4	0.008
σ_{e_1}	0.083	σ_{e_2}	0.141
σ_{e_3}	0.159	σ_{e_4}	0.156
β_{01}	8.233	β_{02}	41.265
β_{03}	30.637	β_{04}	31.921
β_1	1.691	σ_η	0.245

Notes—This table reports the calibrated parameters. Panel A refers to the version without measurement error ($\sigma_\eta = 0$) and Panel B to the version with measurement error ($\sigma_\eta > 0$)

offers, with considerable heterogeneity across groups—e.g., sub-prime borrowers receive less than half the offers that near-prime borrowers receive. The number of offers is non-monotonic in the creditworthiness of borrowers, thereby matching the patterns that Han et al. (2018) report. However, the parameters β_{0j} and β_1 imply that borrowers consider only a small fraction of these offers: the cost of effort to evaluate an average number of offers per period equal to $\alpha = 1$ corresponds to approximately 500 basis points, and it increases by approximately 1,000 basis points to evaluate an average number of offers per period equal to $\alpha = 2$.

The value of σ_{e_j} implies that the standard deviation of the product differentiation parameter is small, relative to the overall heterogeneity in borrowers' preferences. Notably, the variance of e accounts for approximately one percent of the overall variance of accepted $c = R + e$. However, this small contribution to the variance follows from the fact that the model implies that borrowers consider very few offers—e.g. often one only—and, thus, their choice would be almost identical even if product differentiation was more important for them. 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. Moreover, it is interesting to note that the value of σ_{e_j} increases with borrowers' creditworthiness, thereby indicating that non-price card attributes matter relatively more as borrowers' risk scores increase.

Finally, the calibrated σ_η equals 0.248, which means that the standard deviation of the measurement error on the accepted rates equals 0.251. This value is very small relative to the calibrated variances of accepted rates R in the version without measurement error, which equal 3.843 in the sub-prime market, 3.823 in the near-prime market, 3.401 in the prime market, and 1.984 in the super-prime market, respectively.

Table 4 presents a comparison between the empirical moments and the moments calculated from the model at the calibrated parameters reported in Panel A and in Panel B of Table 3, respectively—i.e., without ($\sigma_\eta = 0$) and with ($\sigma_\eta > 0$) measurement error. The model without measurement error matches the data reasonably well, though, as anticipated, it underpredicts the dispersion of accepted rates—i.e., it overpredicts the lower percentiles and it underpredicts the higher percentiles. It matches reasonably well the percentiles of the distribution of the difference between the highest and lowest offered interest rates, and almost perfectly the aggregate statistics on the fraction of credit card borrowers in each group, thereby reproducing the mild non-monotonicity of the fraction of borrowers as their creditworthiness increases observed in the data. The model with just a small measurement error on accepted offers matches the data almost perfectly.

Table 4: Model Fit

	Data	Model $\sigma_\eta = 0$	Model $\sigma_\eta > 0$
10th Percentile Accepted Rate, Sub-prime Borrowers	14.39	18.23	14.87
25th Percentile Accepted Rate, Sub-prime Borrowers	17.58	19.48	17.65
50th Percentile Accepted Rate, Sub-prime Borrowers	21.93	22.08	21.51
75th Percentile Accepted Rate, Sub-prime Borrowers	27.80	25.77	26.33
90th Percentile Accepted Rate, Sub-prime Borrowers	30.16	28.75	31.52
10th Percentile Accepted Rate, Near-prime Borrowers	13.20	16.82	13.78
25th Percentile Accepted Rate, Near-prime Borrowers	16.55	18.04	16.47
50th Percentile Accepted Rate, Near-prime Borrowers	20.20	20.60	20.18
75th Percentile Accepted Rate, Near-prime Borrowers	25.72	24.29	24.82
90th Percentile Accepted Rate, Near-prime Borrowers	29.16	27.23	29.80
10th Percentile Accepted Rate, Prime Borrowers	11.56	15.48	12.61
25th Percentile Accepted Rate, Prime Borrowers	14.81	16.55	15.00
50th Percentile Accepted Rate, Prime Borrowers	17.93	18.77	18.35
75th Percentile Accepted Rate, Prime Borrowers	21.90	22.08	22.50
90th Percentile Accepted Rate, Prime Borrowers	28.68	24.75	27.19
10th Percentile Accepted Rate, Super-prime Borrowers	10.79	14.37	11.70
25th Percentile Accepted Rate, Super-prime Borrowers	13.82	15.15	13.78
50th Percentile Accepted Rate, Super-prime Borrowers	16.84	16.70	16.54
75th Percentile Accepted Rate, Super-prime Borrowers	19.54	18.68	19.86
90th Percentile Accepted Rate, Super-prime Borrowers	23.98	20.24	23.50
Fraction Receiving 2+ Offers (%)	75.00	73.64	73.90
Median Number of Offers Received, Conditional on 2+ Offers	3.00	3.00	3.00
Average Number of Offers Received, Conditional on 2+ Offers	4.00	3.44	3.45
10th Percentile Distribution of Differences in Offered Rates	0.00	1.38	1.27
30th Percentile Distribution of Differences in Offered Rates	2.25	3.40	3.24
50th Percentile Distribution of Differences in Offered Rates	4.34	5.00	4.80
70th Percentile Distribution of Differences in Offered Rates	7.25	6.71	6.44
90th Percentile Distribution of Differences in Offered Rates	9.25	10.09	9.75
Fraction with Credit Card Debt, Sub-Prime Borrowers	54.56	55.44	54.85
Fraction with Credit Card Debt, Near-Prime Borrowers	55.33	56.17	55.26
Fraction with Credit Card Debt, Prime Borrowers	54.00	55.10	54.28
Fraction with Credit Card Debt, Super-prime Borrowers	36.02	36.06	35.93
Charge-Off Rate, Sub-prime Borrowers	24.20	24.20	24.20
Charge-Off Rate, Near-prime Borrowers	7.62	7.62	7.62
Charge-Off Rate, Prime Borrowers	2.93	2.93	2.93
Charge-Off Rate, Super-prime Borrowers	0.83	0.83	0.83
Average Funding Cost	7.14	6.16	6.11

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

5.6 Model Implications

We study the implications of the model evaluated at the parameters reported in Panel A of Table 3. Since these parameters are very similar to those of Panel B, the implications of the model evaluated at the latter parameters are very similar, as well.

Figure 1 displays lenders' and borrowers' equilibrium policies in each market j, shown in each row. The left column of Figure 1 display lenders' optimal offered rate $R(k)$ (solid line, left axis) to each group of borrowers as a function of their cost k, as well as the density of lenders' cost k (dotted line, right axis) for values of the cost k from the risk free rate up to the cutoff value $\hat{k}$ that the free entry condition (12) determines. Lenders' offered rates are strictly increasing in their costs k, as Proposition 4 states. The average markup, computed as $\frac{R(k)(1-\rho_j)-k}{k}$ to take into account the charge-off rate ρ_j, on average equals approximately 220 percent and it is decreasing in borrowers' creditworthiness—i.e., it equals 270 and 170 percent in the subprime and super-prime markets, respectively. Moreover, markups are quite heterogeneous across lenders, as the lowest-cost lenders have markups in excess of 300 percent, whereas the highest-cost lenders have markups below 100 percent.

The right panel of Figure 1 displays borrowers' optimal search effort $\alpha(z)$ (solid line, left axis) as a function of their willingness to pay z, as well as the density of borrowers' willingness to pay z (dotted line, left axis). Since the lowest-valuation borrowers 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 approximately 0.7 offers—and only borrowers whose willingness to pay z is in the highest 15 percent of the distribution choose $\alpha(z)$ larger than 1.

Figure 2 displays the probabilities $P(R)$ that borrowers accept a credit card offer with an interest rate R. These probabilities are obviously decreasing as R increase, but perhaps the most striking features of Figure 2 are that: 1) because of borrowers' low search effort, they are quite flat, which means that borrowers' demand is quite inelastic—i.e., the average elasticity equals approximately -1.45; and 2) for any R, the probability that subprime borrowers accept such an offer is at least twice as high as the probability that borrowers in any other risk group accept it, consistent with the evidence that Agarwal et al. (2017) report. The average acceptance probability $P(R)$ equals 0.17, which, scaled up by the mass of lenders $\sum_{j=1}^{J} L_j\omega_j \approx 2.7$, yields the aggregate fraction of individuals with credit card debt of 45.7 percent.

Figure 3 plots the distributions $F_R(R)$ of offered rates and the distributions $H_R(R)$ of

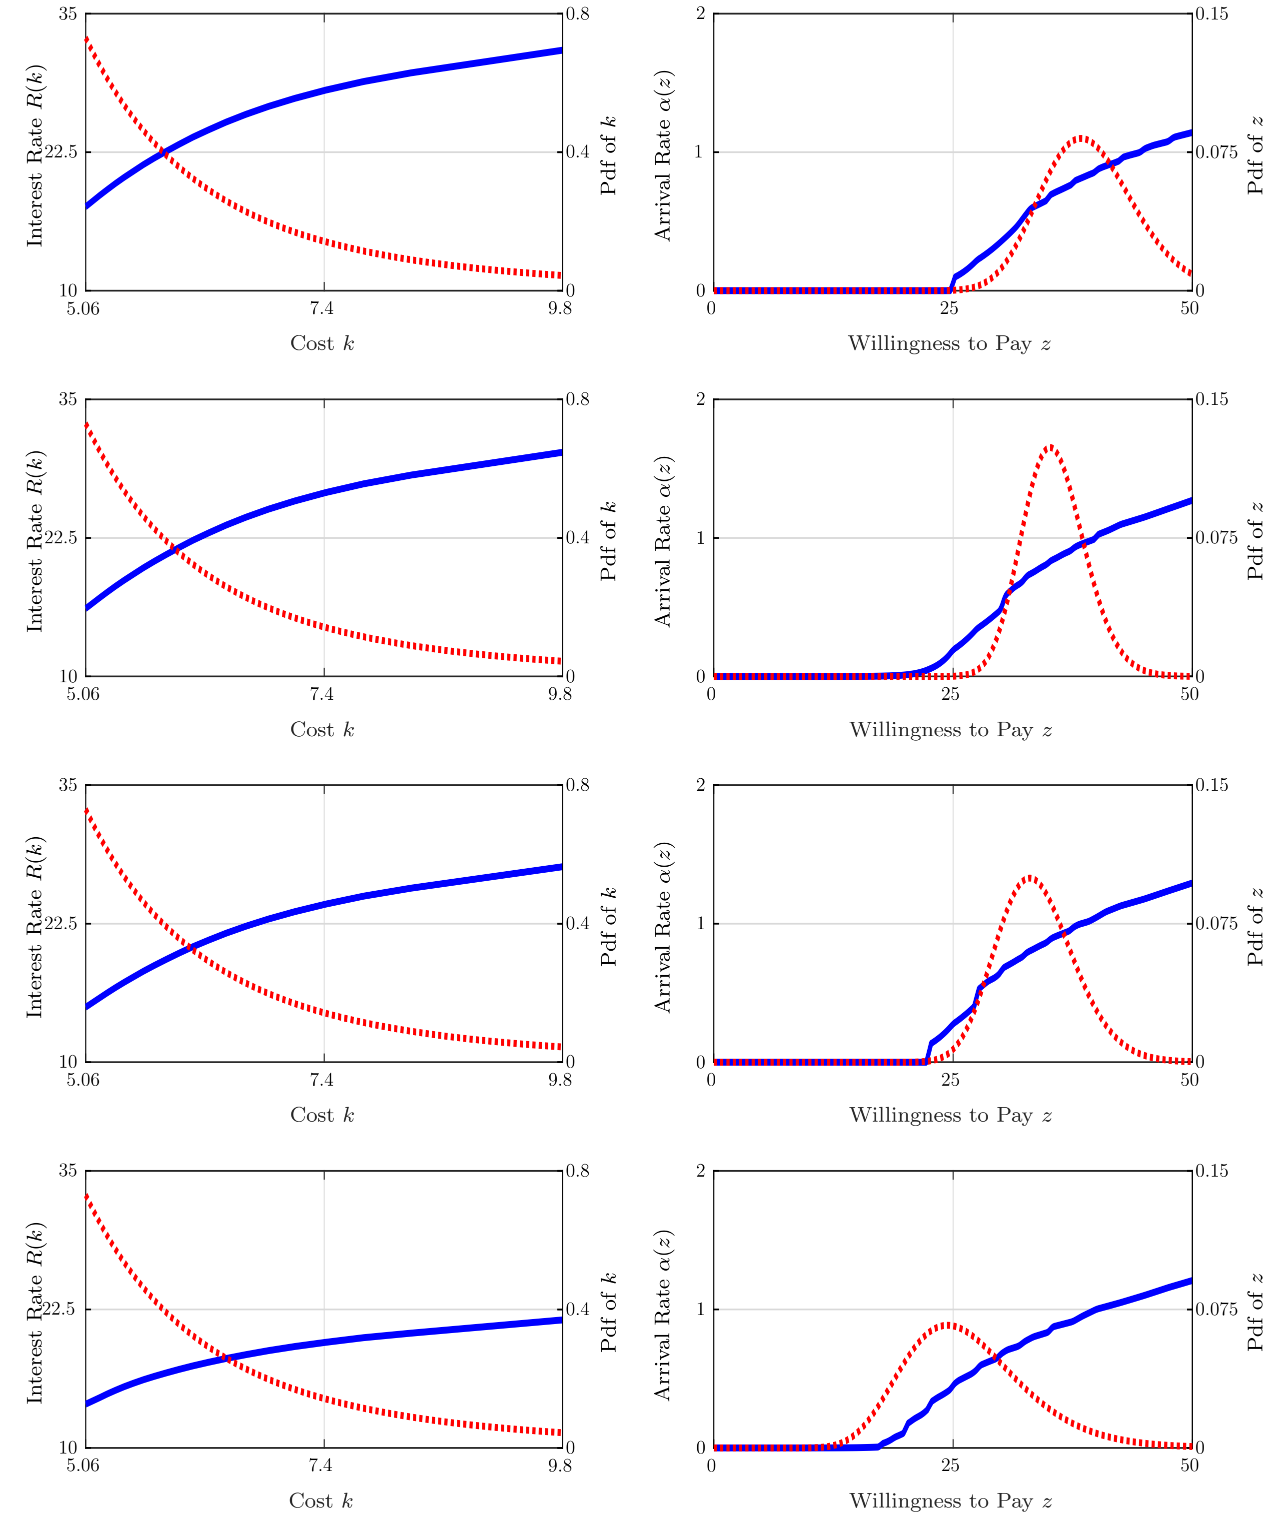

Figure 1: The left panels display lenders' optimal interest rate $R(k)$ (solid line, left axis) as a a function of their cost k, as well as the density of lenders' cost k (dotted line, right axis). The right panels display borrowers' optimal arrival rate $\alpha(z)$ (solid line, left axis) as a function of their willingness to pay z, as well as the density of borrowers' willingness to pay z (dotted line, left axis). The first row refers to the sub-prime market, the second row to the near-prime market, the third row to the prime market, and the fourth row to the super-prime market.

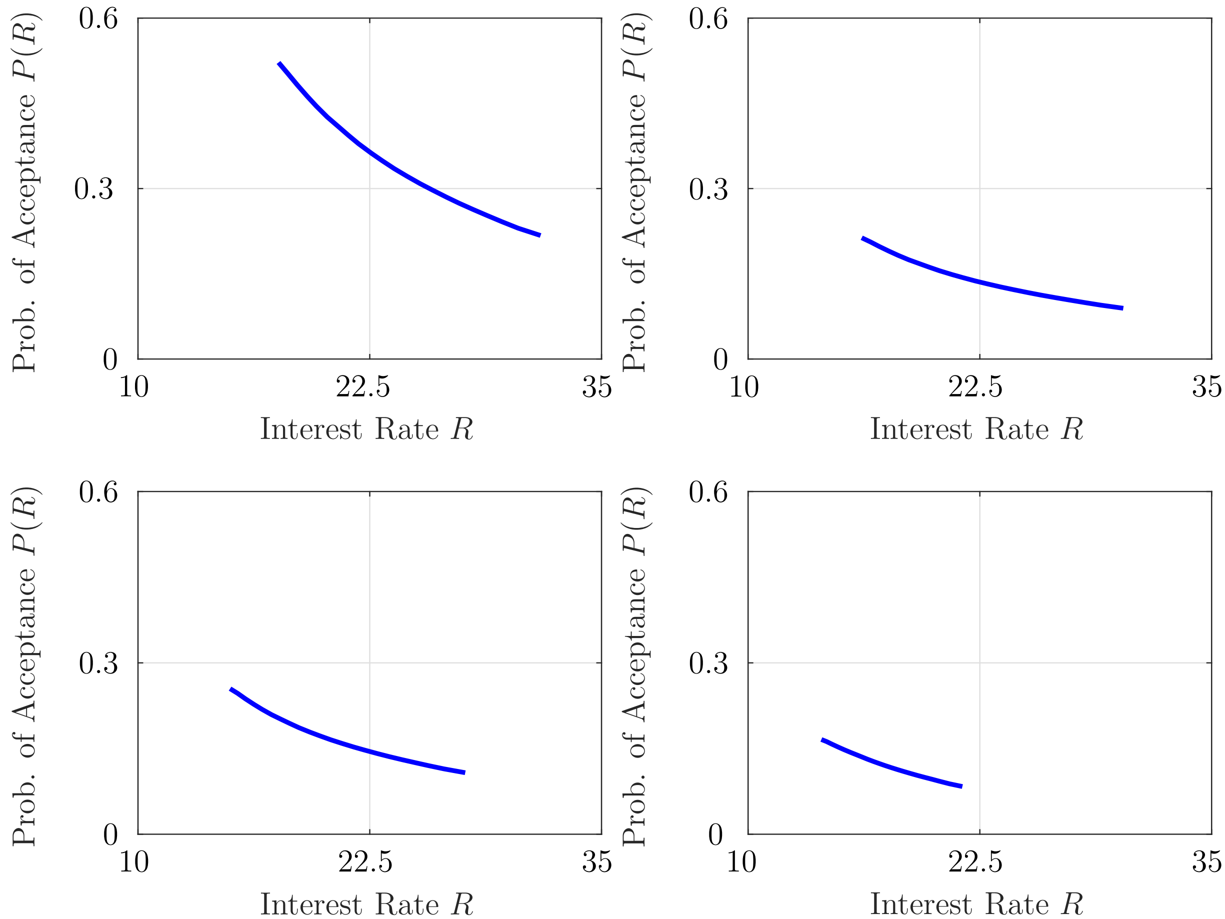

Figure 2: Probability $P(R)$ that borrowers accept an offer with interest R, for sub-prime borrowers (top-left panel), near-prime borrowers (top-right panel), prime borrowers (bottom-left panel), and super-prime borrowers (bottom-right panel).

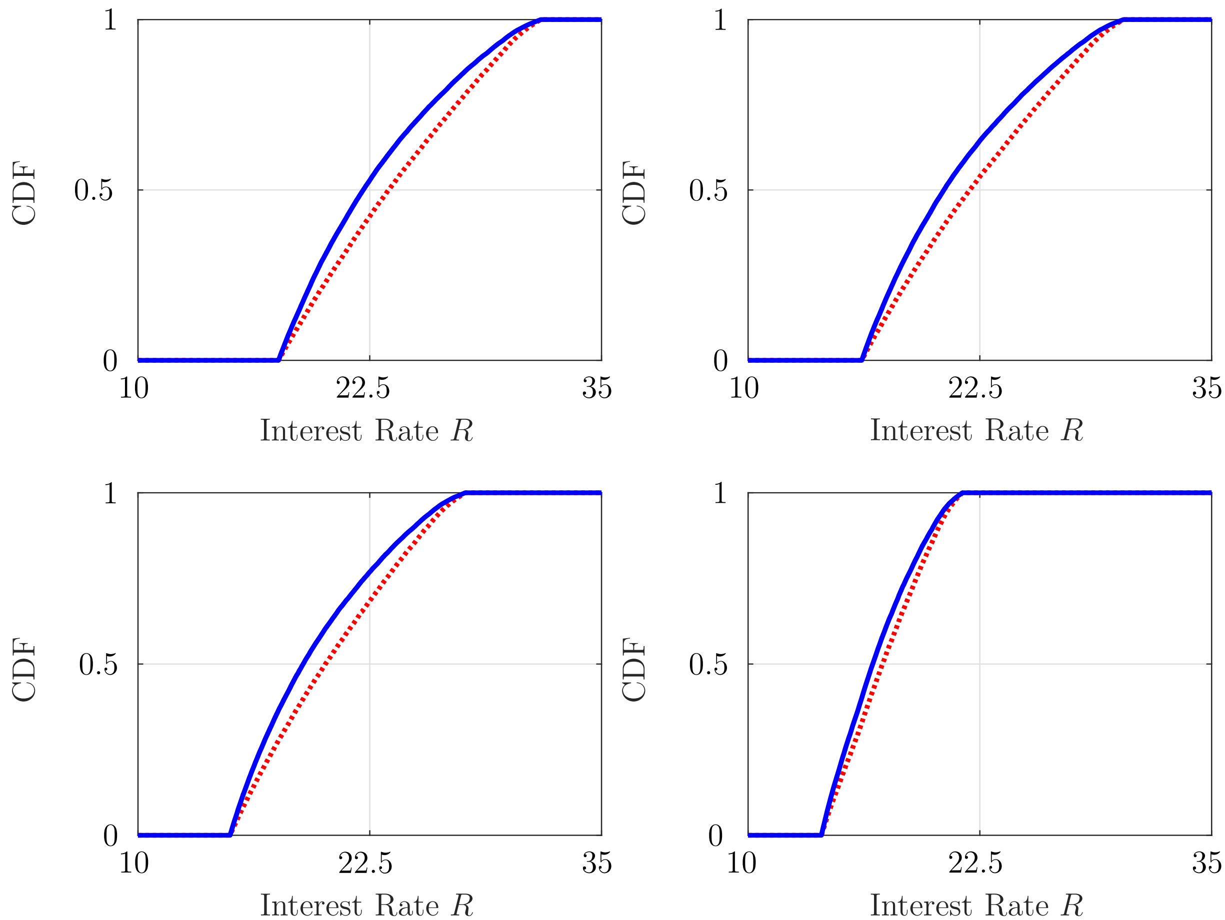

Figure 3: The solid line displays the cumulative distribution function $H_R(R)$ of accepted interest rates and the dashed line displays the cumulative distribution function $F_R(R)$ of offered rates.

accepted rates. Of course, the distributions of offered rates first-order stochastically dominate the distribution of accepted rates. However, the differences between the two distributions are small. Two reasons account for this small difference: 1) borrowers' low search effort implies that the rate $\alpha(z)$ at which they consider offers is low; 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, because the standard deviation σ_e is small and because, for e to have sizable effect, borrowers would need to consider more than one offer, which happens very infrequently due to their low search effort. Thus, 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.[9]

The top rows of Table 5 report summary statistics of market outcomes—i.e., prices and

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

Table 5: Market Outcomes and Welfare

	Sub-	Near-	Prime	Super-
Average Number of Offers per Borrower	1.56	3.87	3.13	2.89
Average Accepted Rate	22.88	21.34	19.60	16.98
Standard Deviation of Accepted Rates	3.83	3.83	3.43	2.03
Fraction of Borrowers	55.06	55.89	54.44	35.91
Consumer Surplus	4.33	4.50	4.68	2.61
Lender Profits	1.33	1.33	1.26	1.00
Welfare	5.65	5.83	5.95	3.61
Efficient Average Number of Offers per Borrower	1.76	4.38	3.54	3.27
Efficient Welfare	6.83	7.95	7.78	5.69

Notes—This table reports market outcomes and welfare in each market.

quantities—as well as consumer surplus, lenders' profits, and aggregate welfare in each market. The calibrated model implies that consumer surplus and aggregate welfare are highest in the subprime market and they decrease as borrowers' creditworthiness increases, which is due to the fact that average valuations are highest in the subprime market and they decline as risk scores increase. Lenders' profits are non-monotonic, peaking in the near-prime market.

The bottom two rows of Table 5 report the optimal number of offers and welfare in the constrained-efficient economy that we characterized in Section 4.4. As our theoretical analysis points out, the constrained-efficient allocation differs from the market allocation. Most notably, it requires a larger number of offers L_j in each market, on average by 13 percent; the resulting welfare gains would be large, ranging from 34 percent in the prime market to 56 percent in the super-prime market.

5.7 Comparative Statics

We further illustrate the working of our model through two comparative statics which vary the two parameters that are the main focus of our framework—i.e., the standard deviation σ_{e_j} of the product differentiation e and the parameter β_{0j} that scales the marginal cost of search effort.

We present the results of these comparative statics for near-prime borrowers—i.e., the group for which the model without measurement error matches the data most precisely, according to Table 4—but the outcomes for the other groups are similar.

Product Differentiation. Figure 4 compares outcomes of the model at the calibrated parameters (solid line) to those of the model when we increase the standard deviation σ_{e_j} of the product differentiation e (dashed line) while holding all other parameters at their calibrated values. Since σ_{e_j} is calibrated to be small, we increase it tenfold.

The top panels show interesting outcomes. Most notably, the top left panel shows that the interest rate function $R(k)$ increases for most lenders when product differentiation is more important for borrowers. The reason is that, a larger σ_{e_j} means that product differentiation affects consumers' choice across lenders relatively more; thus, lenders compete less aggressively on prices and offer higher interest rates. This is particularly important for low-cost lenders: when σ_{e_j} is low, these lenders maximize profits by charging lower prices, thereby gaining large market shares; when σ_{e_j} is higher, these lenders gain a lower market share and, thus, choose to increase their margins. As a result of these incentives, the interest rate function $R(k)$ is higher and flatter when σ_{e_j} is higher.

The top right panel shows that a higher σ_{e_j} has minimal effect on borrowers' search effort. This is the result of opposite effects. Specifically, holding the distribution of offered rates fixed, the increase in the product differentiation parameter induces borrowers to search more aggressively, since they are more likely to receive offers with product features e that they value more. However, the increase in the product differentiation parameter triggers two supply responses that reduce borrowers' search effort: 1) on average, lenders increase their interest rates, thereby decreasing borrowers' expected surplus and, thus, their search effort; and 2) the dispersion of offered interest rates decreases, which reduces borrowers' incentive to compare offers. As a result of these offsetting effects, search effort decreases, although minimally.

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_j} is higher, the acceptance probability of an offer with a given R increases relative that of the baseline case. Demand becomes more elastic—i.e., the average elasticity equals -1.433 compared to -1.397 in the baseline case—presumably due to higher interest rates. The average acceptance probability across lenders as well as the fraction of individuals with credit card debt are slightly lower than in the baseline case—i.e., 0.142 and 0.548 vs. 0.145 and 0.559, respectively—indicating that the increase in interest rates may outweigh borrowers' benefits from larger values of e.

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_j} (dashed lines) first-order stochastically dominate the corresponding distributions obtained in the model at the calibrated σ_{e_j} (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 22.369 in the baseline case at the calibrated parameters to

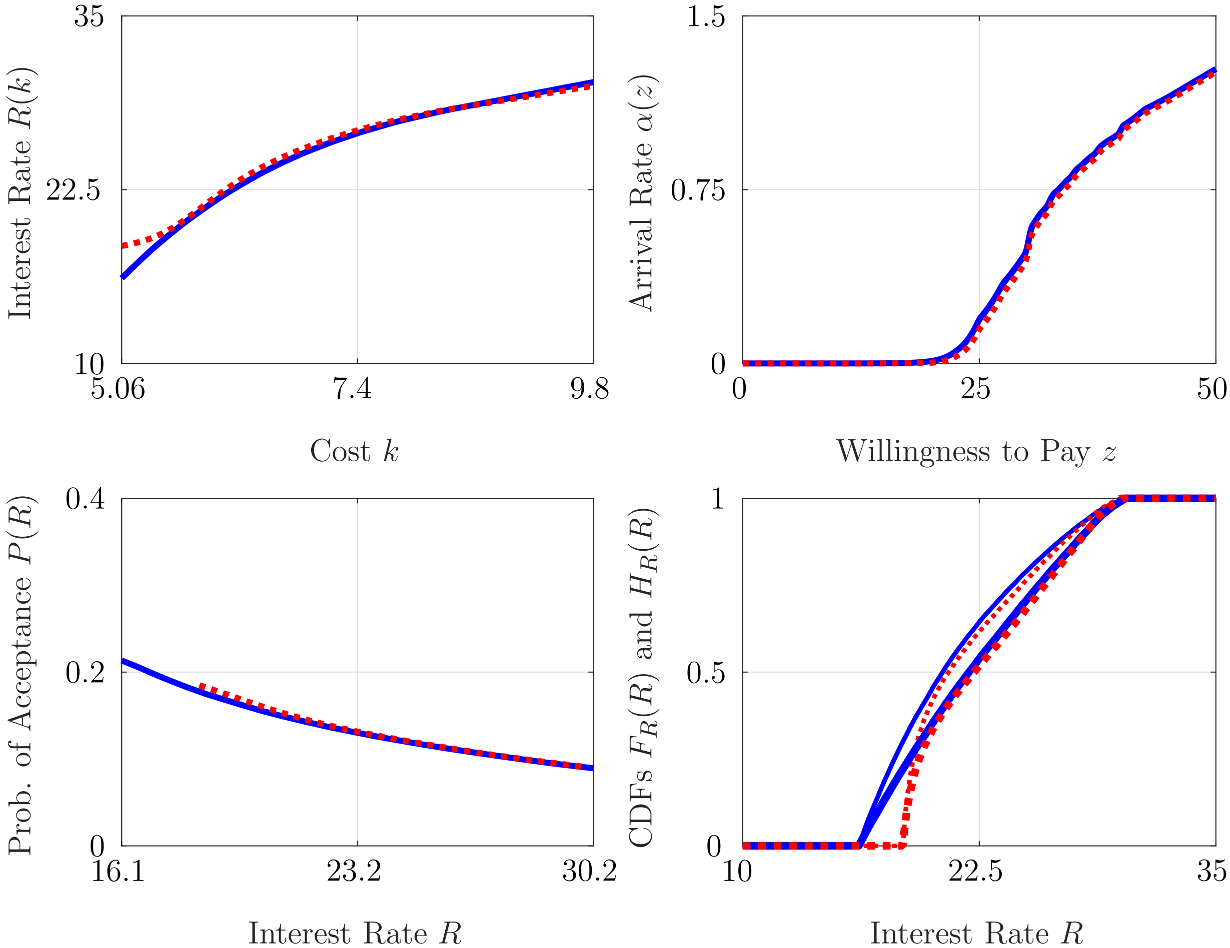

Figure 4: These panels display model outcomes model at the calibrated parameters (solid line) and in the case when $\sigma'_{e_j} = 10\sigma_{e_j}$ (dashed line) for near-prime borrowers. The top left panel displays lenders' optimal interest rate $R(k)$ as a a function of their cost k; the top right panel displays borrowers' optimal search effort $\alpha(z)$ 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).

22.845 in the case with σ_{e_j} is higher, whereas the standard deviation of offers decreases from 3.987 to 3.578, thereby reiterating that offered rates are higher and less dispersed if product differentiation matters more for borrowers. Similarly, the average accepted rate increases from 21.363 to 22.039, whereas the standard deviation of accepted rates decreases from 3.823 to 3.381. Of course, accepted rates exhibit larger differences at higher percentiles than at lower percentiles relative to those of the baseline case.

Figure 4 also helps us understand why the calibrated model calls for small values of σ_{e_j}: if they were larger, the level of interest rates would be higher and, most notably, their dispersion would be lower than those observed in the data. Therefore, the calibration finds that σ_{e_j} increases with borrowers' creditworthiness because the observed dispersion of accepted rates in the data decreases with borrowers' creditworthiness, as Table 2 reports.

Cost of Search Effort. Figure 5 compares outcomes of the model at the calibrated parameters (solid line) to those of the model when we decrease the parameters β_{0j} 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 decrease is larger for low-cost lenders than for high-cost lenders, since borrowers accept high-cost lenders' offers almost exclusively when borrowers consider one of these high offers only and, thus, high-cost lenders do not need to lower their rates as much as low-cost lenders. The top-right panel explains why lenders' offered rates are lower: 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 is higher than that of the baseline case for low values of R and lower for high values of R. The reason is that borrowers consider a larger number of offers and, thus, while the probability of accepting any offer increases, they are relatively less likely to accept high-interest rate offers. Demand becomes more elastic—i.e. the average elasticity equals -1.619—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., 0.202 and 0.779 vs. 0.145 and 0.559, respectively—indicating that borrowers are better off when β_{0j} is lower.

The bottom-right panel of Figure 5 displays the distribution of offered rates (thick lines) and of accepted rates (thin lines). Both distributions obtained in the model with a lower β_{0j} (dashed lines) are first-order stochastically dominated by the corresponding distributions

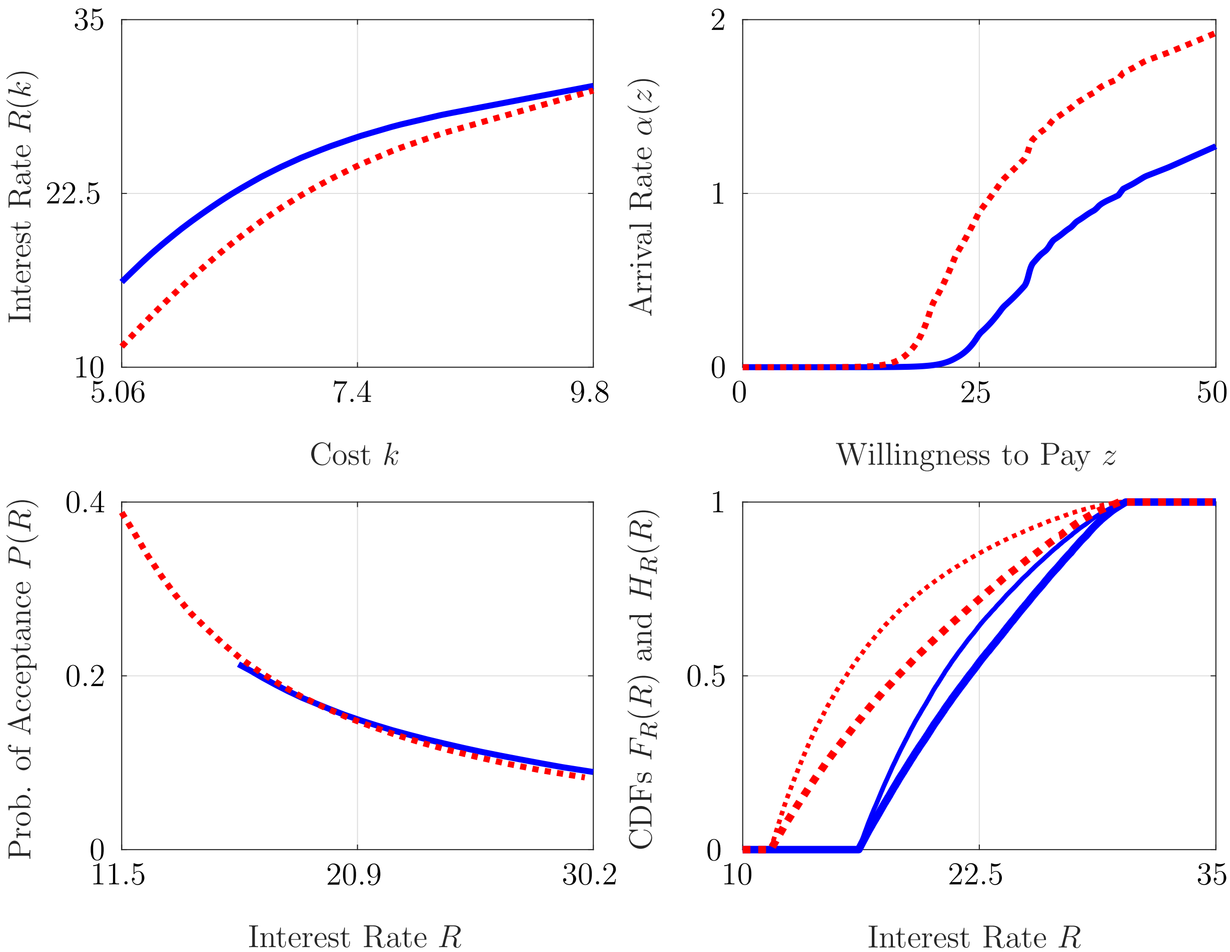

Figure 5: These panels display model outcomes at the calibrated parameters (solid line) and in the case when $\beta'_{0j} = 0.5\beta_{0j}$ (dashed line). The top left panel displays lenders' optimal interest rate $R(k)$ as a a function of their cost k; the top right panel displays borrowers' optimal rate $\alpha(z)$ 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).

obtained in the model at the calibrated β_{0j} (solid lines). This is because low-cost lenders lower their offered rates when the cost of effort is lower and borrowers' search more. The average offered and accepted rates equal 18.995 and 16.804, respectively, and the standard deviation of offered and accepted rates equal 5.195 and 4.617, respectively, when the cost-of-effort parameter β_{0j} is half of the calibrated value. As the bottom plots shows, the lower cost of effort affects lower percentiles relatively more than higher percentiles.

6 Policy Experiments

In this Section, we use our model to study two policy experiments, motivated by recent regulatory interventions: 1) A cap on the interest rate—i.e., a maximum rate R_{max}; 2) Higher compliance costs for lenders, captured by a higher fixed cost χ. The goal of both experiments is to study how borrowers' search effort and lenders' offered rates respond, thereby affecting market outcomes and welfare.

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 interest rate cap on the equilibrium of our model.

The theoretical literature points out that these caps may have unintended consequences, for two main reasons. First, if caps prevent lenders to recover their costs, they may reduce the supply of credit, most notably to riskier borrowers who have higher default rates. Second, Fershtman and Fishman (1994) and Armstrong et al. (2009) show that, in markets with search frictions, price caps may have the unintended consequences of *increasing* the equilibrium prices paid by consumers. More precisely, they identify two opposing effects: A) the direct effect of regulation is to reduce prices for uninformed consumers who, before the regulation, were paying high prices; and B) 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. Hence, the relative magnitude of these contrasting effects is an empirical/quantitative question. Our calibrated model allows us to determine which of these opposing effect dominates and, thus, whether or not price caps are beneficial to consumers.

To understand these issues, we set a cap of $R_{max} = 25$ percent. This cap does not bind in the super-prime market, but it does in all other markets: Table 2 shows that it corresponds to approximately the 65th, the 75th and the 95th percentiles of the distribution of accepted interest rates in the subprime, near-prime, and prime market, respectively. We study this counterfactual case in general equilibrium—i.e., we require that lenders' free entry condition (12) holds. Thus, some lenders may exit the market, in which case we decrease the aggregate arrival rate of offers to a new value L'_j proportionally. More precisely, the new arrival rate equals $L'_j = \Lambda_j G(\hat{k}')$, where $\hat{k}'$ is the marginal cost of the marginal lender—i.e., the lender that satisfies the free entry condition (12)—in the counterfactual case (the marginal cost of the marginal lender in the baseline case equals $\hat{k}$).

Figure 6 compares outcomes of the model at the calibrated parameters (solid line) to those of the model with $R_{max} = 25$ percent, while holding all other parameters at their calibrated values, for the near-prime market. 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, displaying the indirect and direct effects that Fershtman and Fishman (1994) and Armstrong et al. (2009) emphasize. Specifically, since some lenders exit the market, on average borrowers receive an 18-percent-lower number of offers than in the baseline case—i.e., $L'_j = 3.389$ when interest rates are capped versus $L_j = 3.873$ in the baseline case. Nevertheless, low-valuation borrowers slightly increase their search effort to more-than-offset the lower arrival rate of offers and, thus, the average number of offers $\alpha(z)$ that they consider is higher than in the baseline. This is because the cap reduces the level of interest rates relative to the baseline case, thereby increasing the expected payoff from a credit card loan for these lower-valuation borrowers. However, high-valuation borrowers respond differently than low-valuation borrowers, in that the average number of

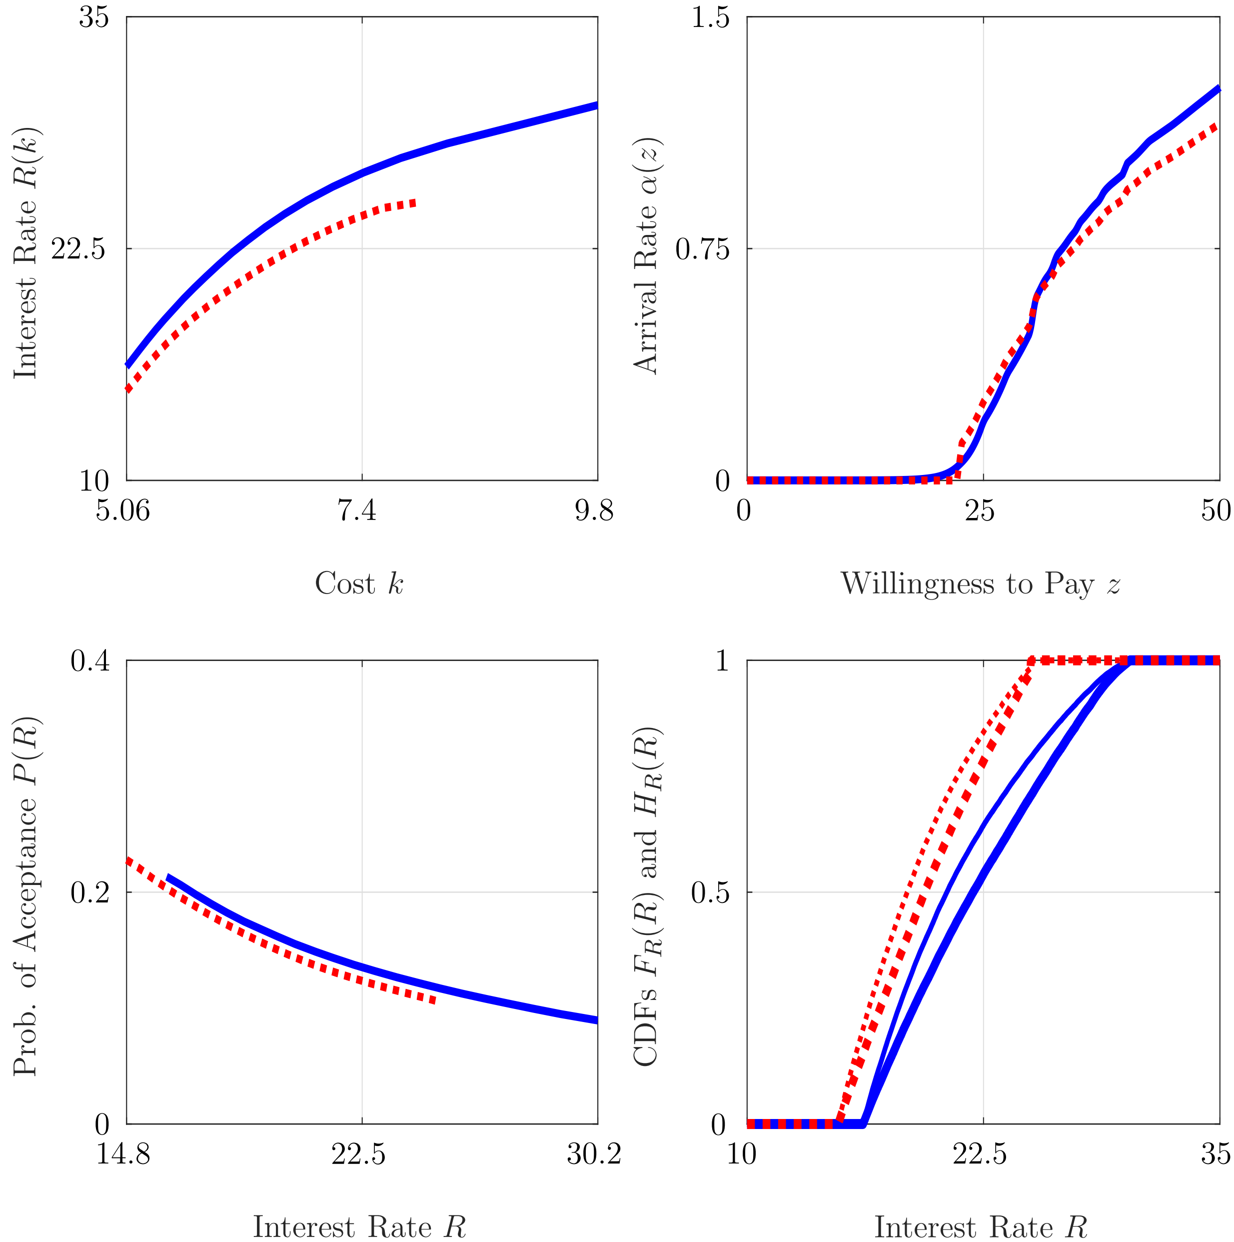

Figure 6: These panels display outcomes in near-prime market at the calibrated parameters (solid line) and in the case when interest rates are capped at 25 percent (dashed line). The top left panel displays lenders' optimal interest rate $R(k)$ as a a function of their cost k; the top right panel displays borrowers' optimal arrival rate $\alpha(z)$ 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).

offers $\alpha(z)$ that they consider is lower than in the baseline. This is because these borrowers already had positive gains-from-trade in the baseline case, but the cap reduces the dispersion of interest rates across lenders and, thus, reduces the benefits of considering multiple offers.

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: 1) on average, borrowers consider fewer offers; and 2) the 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 interest rates. The average acceptance probability across lenders increases relative to the baseline case—i.e., 0.158 vs. 0.145, respectively. On average, demand is more elastic than in the baseline case: the average elasticity equals -1.690 vs. -1.397 in the baseline case. Moreover, the fraction of individuals with credit card debt decreases to 0.536 from 0.559 in the baseline case. The reason is that the lower search effort of borrowers with a relatively higher valuation z leads them to consider and to accept fewer offers than in the baseline case; this decline more than offsets the increase due to borrowers with a relatively lower z, who search more aggressively and, thus, are more likely to accept an offer relatively to the baseline. Thus, one interesting outcome of this counterfactual analysis is that infra-marginal borrowers (i.e., those with a high valuation z) display a stronger response to the cap than marginal borrowers (i.e., those with a low valuation z).

The bottom right panel of Figure 6 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 19.434 and 18.874 respectively, and the standard deviation of offered and accepted rates equal 2.512 and 2.444, respectively. These values are considerably lower than those of the baseline, suggesting that the average surplus of those who borrow may be higher relative to the baseline.

Table 6 reports summary statistics of market outcomes, as well as consumer surplus, lenders' profits, and welfare for each group of borrowers when interest rates are capped. The cap induces a large redistribution of surplus from lenders to borrowers, but small aggregate welfare effects. Specifically, consumer surplus increases in all markets in which the cap is binding, with larger increases in markets in which lender pricing is more constrained: the increase in consumer surplus equals 14 percent in the submprime market, 11 percent in the near-rime market, and 2.8 percent in the prime market (it is zero in the super-prime market as the cap is not binding); weighting markets by the share of borrowers in each of them, the

Table 6: Market Outcomes and Welfare with a Price Cap

	Sub-	Near-	Prime	Super-
Average Number of Offers per Borrower	1.32	3.39	2.99	2.89
Average Accepted Rate	19.34	18.94	18.68	16.98
Standard Deviation of Accepted Rates	3.00	2.89	2.99	2.03
Fraction of Borrowers	52.68	53.80	54.09	35.91
Consumer Surplus	4.88	5.02	4.94	2.61
Lender Profits	0.54	0.75	1.04	1.00
Welfare	5.41	5.78	5.98	3.61

Notes—This table reports market outcomes and welfare in each market.

aggregate increase in consumer surplus equals 6.8 percent. Correspondingly, aggregate lender profits decline by 25 percent—i.e., they decline by 57 percent in the subprime market, by 49 percent in the near-prime market, and by 14 percent in the prime market. As a result, aggregate welfare is almost unchanged–i.e., it declines by 0.2 percent on aggregate, since it increases in the sub-prime market by 0.6 percent, whereas it declines by 1.5 percent in the near-prime market and it declines by 0.7 in the prime market.

6.2 Higher Compliance Costs

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

The increase in the fixed costs shares 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 6 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 6, we increase the fixed cost χ so that the marginal lender has marginal cost equal to $\hat{k}'$—i.e., the marginal cost of the lender that satisfies the free entry (12) condition in the case of the price cap $\bar{R} = 25$. In practice, this implies that the new fixed cost χ' is 15.1-percent larger than that of the baseline case in the near-prime market (17.4-percent and 6.3-percent larger in the sub-prime and prime market, respectively, whereas it does not change in the super-prime market as the -price cap was not binding in that market). We further decrease the aggregate arrival rate of offers to a new value L'_j correspondingly—i.e., the new arrival rate equals $L'_j = \Lambda_j G(\hat{k}')$.

Figure 7 compares outcomes of the model at the calibrated parameters (solid line) to those of the model with a higher fixed cost χ' for the near-prime market. It displays interesting patterns. Notably, the exit of high-cost lenders reduces interest rate dispersion (top-left panel), but it does not reduce the level of interest rates, as surviving lenders increased their rates due to lower competition. Hence, borrowers consider fewer offers than in the baseline case (top-right panel) for two reasons: 1) they receive fewer offers—i.e., $L'_j = 3.389$ when compliance costs are higher versus $L_j = 3.873$ in the baseline case;[10] and 2) they choose not to exert much effort because price dispersion is lower and, thus, the benefits of considering multiple offers are lower. The bottom-left panel shows that the probability $P(R)$ that borrowers accept an offer with a given interest rate R increases relative to the baseline case, because high and low offers are no longer available. However, the average acceptance probability across lenders decreases relative to the baseline case—i.e., 0.141 vs. 0.145. Similarly, the fraction of borrowers declines to 0.478 from 0.559 in the baseline.

The bottom-right panel of Figure 7 show that the distributions of offered rates (thick lines) and of accepted rates (thin lines) in a market with a higher fixed cost χ' (dashed line) intersect the corresponding distributions obtained in the baseline case (solid lines), as lenders no longer offer the lowest and the highest rates. The average offered and accepted rates are higher than those of the baseline (23.080 and 22.410 vs. 22.369 and 21.363, respectively), whereas the standard deviation of offered and accepted rates are lower (3.336 and 3.255 vs. 3.987 and 3.823, respectively).

Table 7 reports summary statistics of market outcomes, as well as consumer surplus, lenders' profits, and aggregate welfare for each group of borrowers when fixed costs are higher. The higher fixed costs reduces lender profits, as the price cap did, but it also decreases consumer surplus, with large negative welfare effects. Specifically, consumer surplus decreases in all markets in which the cap is binding: the decrease in consumer surplus equals 17 percent in the submprime market, 24 percent in the near-rime market, and three percent in the prime market (it is zero in the super-prime market as the fixed cost is the same as in the baseline case); weighting markets by the share of borrowers in each of them, the aggregate decrease in consumer surplus equals ten percent. Similarly, aggregate lender profits decline by 25

[10] L'_j in this counterfactual case is identical by construction to that of the counterfactual case of caps on interest rates.

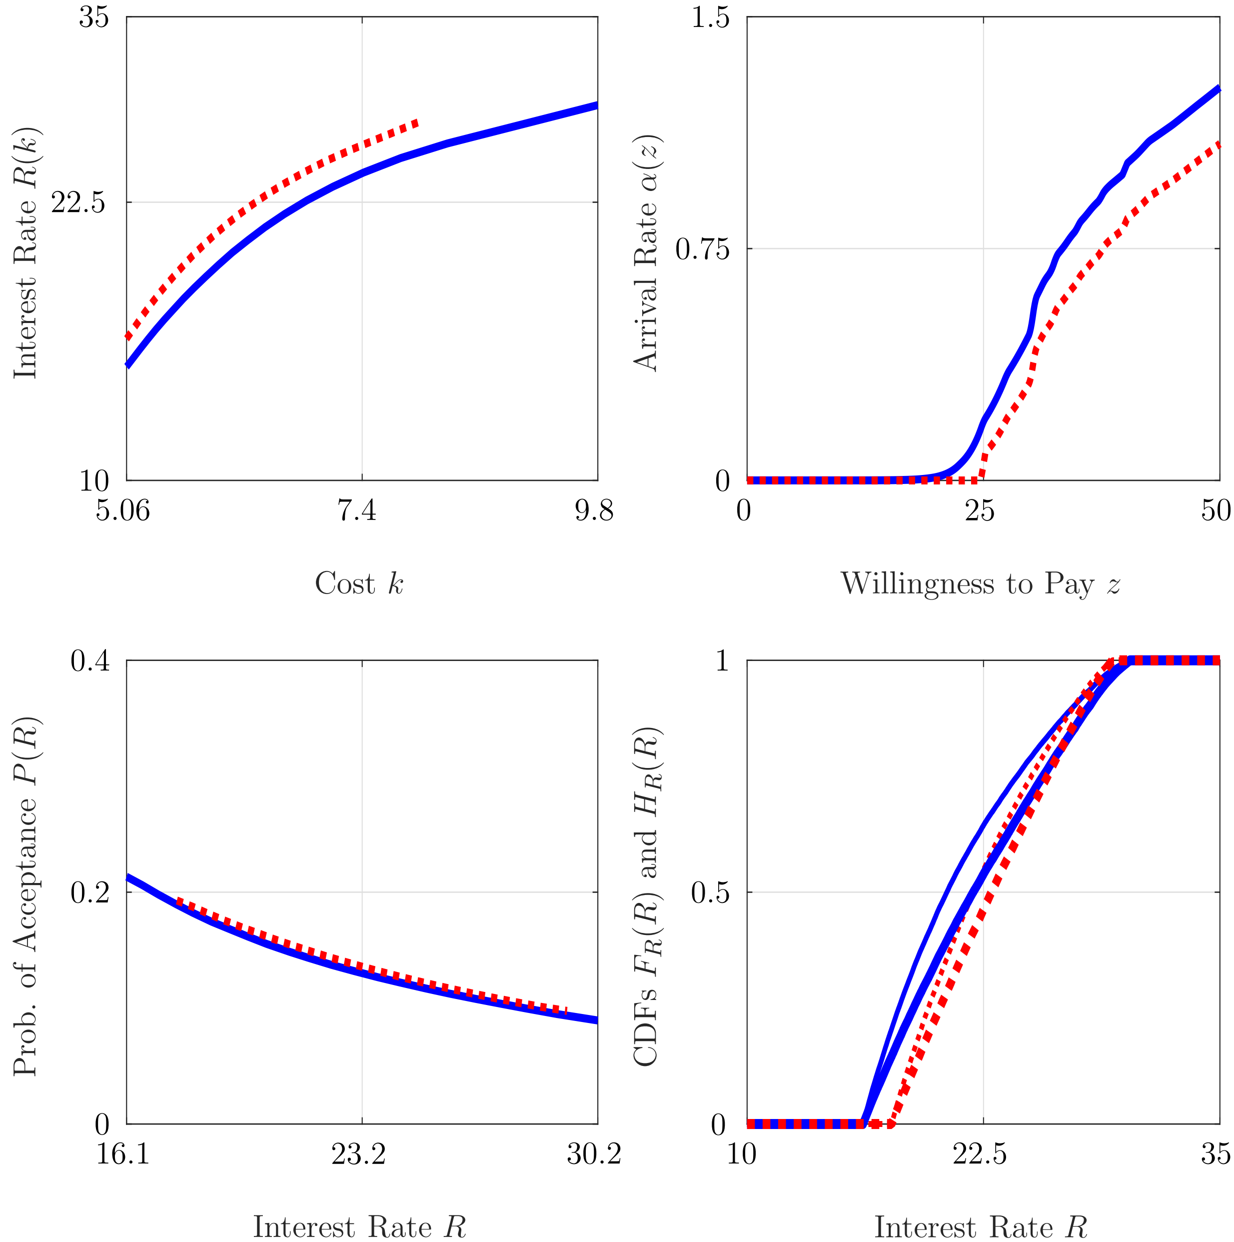

Figure 7: These panels display outcomes in near-prime market the at the calibrated parameters (solid line) and in the case when the fixed cost $\chi' = 1.151\chi$ (dashed line). The top left panel displays lenders' optimal interest rate $R(k)$ as a a function of their cost k; the top right panel displays borrowers' optimal search effort $\alpha(z)$ 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).

Table 7: Market Outcomes and Welfare with Higher Compliance Costs

	Sub-	Near-	Prime	Super-
Average Number of Offers per Borrower	1.32	3.39	2.99	2.89
Average Accepted Rate	23.03	22.43	19.39	16.98
Standard Deviation of Accepted Rates	3.05	3.25	3.20	2.03
Fraction of Borrowers	47.31	47.72	52.89	35.91
Consumer Surplus	3.54	3.42	4.57	2.61
Lender Profits	0.60	0.74	0.90	1.00
Welfare	4.14	4.16	5.47	3.61

Notes—This table reports market outcomes and welfare in each market.

percent—i.e., they decline by 53 percent in the subprime market, by 46 percent in the near-prime market, and by 18 percent in the prime market. As a result, aggregate welfare declines by 13 percent on aggregate—it declines by 23 percent in the sub-prime market, by 28 percent in the near-prime market and by eight percent in the prime market.

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

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.

——— (2017): "Do Banks Pass Through Credit Expansions to Consumers Who Want to Borrow?" *The Quarterly Journal of Economics*, 133, 129–190.

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.

Bureau of Consumer Financial Protection (2012): "Analysis of Differences between Consumer- and Creditor-Purchased Credit Scores," Tech. rep.

——— (2015): "The Consumer Credit Card Market," Tech. rep.

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, G., A. Haughwout, D. Lee, J. Scally, and W. van der Klauuw (2016): "Just Released: Recent Developments in Consumer Credit Card Borrowing," Liberty Street Economics (blog), Federal Reserve Bank of New York.

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.

Cuesta, J. I. and A. Sepúlveda (2018): "Price Regulation in Credit Markets: A Trade-off between Consumer Protection and Credit Access," Mimeo, University of Chicago.

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.

Han, S., B. J. Keys, and G. Li (2018): "Unsecured Credit Supply, Credit Cycles, and Regulation," *The Review of Financial Studies*, 31, 1184–1217.

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: Distribution of Offers

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:

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

Denote the probability distribution of the difference between the highest and the lowest interest rate of a borrower who receives $n \geq 2$ offers by $D(x)$. Denote the probability distribution of the difference between the highest and the lowest interest rate of a borrower who receives exactly n offers by $D_n(x)$ and note that

$$D(x) \quad = \quad \frac{1}{1 - e^{-L} - Le^{-L}} \sum_{n=2}^{\infty} \frac{e^{-L} L^n}{n!} D_n(x).$$

Consider a borrower who receives n offers. Denote the lowest offer by R_L and note that its distribution follows $\bar{F}_n(R_L) = 1 - (1 - F(R_L))^n$. Each of the other $n-1$ offers are distributed iid according to $\hat{F}(R|R_L) = \frac{F(R)-F(R_L)}{1-F(R_L)}$, for $R \geq R_L$. The highest among these $n-1$ offers is distributed according to $\hat{F}(R_H|R_L)^{n-1}$. As a result:

$$\begin{aligned}
D_n(x) \quad &= \quad \int_{\underline{R}}^{\overline{R}} \big(\frac{F(R_L + x) - F(R_L)}{1 - F(R_L)}\big)^{n-1} d\bar{F}_n(R_L) \\
&= \quad \int_{\underline{R}}^{\overline{R}} n\big(F(R_L + x) - F(R_L)\big)^{n-1} F'(R_L) dR_L.
\end{aligned}$$

Combining the above we have:

$$
\begin{aligned}
D(x) &= \frac{1}{1-e^{-L}-Le^{-L}} \sum_{n=2}^{\infty} \frac{e^{-L}L^n}{n!} \int_{\underline{R}}^{\overline{R}} n\big(F(R_L+x)-F(R_L)\big)^{n-1} F'(R_L)dR_L \\
&= \frac{1}{1-e^{-L}-Le^{-L}} \int_{\underline{R}}^{\overline{R}} \sum_{n=2}^{\infty} \frac{e^{-L}L^n}{(n-1)!} \big(F(R_L+x)-F(R_L)\big)^{n-1} F'(R_L)dR_L \\
&= \frac{Le^{-L}}{1-e^{-L}-Le^{-L}} \int_{\underline{R}}^{\overline{R}} \Big(e^{L(F_R(R_L+x)-F_R(R_L))}-1\Big) F_R'(R_L)dR_L.
\end{aligned}
$$

B Proofs

Proof of Proposition 2. We first show that the cost distribution for a low n first-order stochastically dominates that for a high n (thereby proving that $v_{z,n}$ is increasing in n) and that the derivative of the cost distribution for a high n first-order stochastically dominates that for a low n (thereby proving strictly decreasing differences):

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

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 (5) with respect to s (and noting that $v_{z,0} = 0$):

$$
\begin{aligned}
V_z'(s) &= \sum_{n=1}^{\infty} \Big(-\frac{e^{-\alpha(s,L)}\alpha(s,L)^n}{n!} v_{z,n} + \frac{e^{-\alpha(s,L)}\alpha(s,L)^{n-1}}{(n-1)!} v_{z,n}\Big) \frac{\partial \alpha(s,L)}{\partial s} \\
&= \Big(-\sum_{n=0}^{\infty} \frac{e^{-\alpha(s,L)}\alpha(s,L)^n}{n!} v_{z,n} + \sum_{n=0}^{\infty} \frac{e^{-\alpha(s,L)}\alpha(s,L)^n}{n!} v_{z,n+1}\Big) \frac{\partial \alpha(s,L)}{\partial s} \\
&= \sum_{n=0}^{\infty} \frac{e^{-\alpha(s,L)}\alpha(s,L)^n}{n!} \big(v_{z,n+1} - v_{z,n}\big) \frac{\partial \alpha(s,L)}{\partial s} > 0.
\end{aligned}
$$

As a result, the borrower's expected value of offers is strictly increasing in search effort and the optimal choice of effort is characterized by equation (8).

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

$$
\begin{aligned}
V_z''(s) &= \sum_{n=1}^{\infty}\Big(-\frac{e^{-\alpha(s,L)}\alpha(s,L)^n}{n!}+\frac{e^{-\alpha(s,L)}\alpha(s,L)^{n-1}}{(n-1)!}\Big)\Big(v_{z,n+1}-v_{z,n}\Big)\Big(\frac{\partial\alpha(s,L)}{\partial s}\Big)^2 \\
&\quad+\sum_{n=0}^{\infty}\frac{e^{-\alpha(s,L)}\alpha(s,L)^n}{n!}\Big(v_{z,n+1}-v_{z,n}\Big)\frac{\partial^2\alpha(s,L)}{\partial s^2} \\
&= \sum_{n=0}^{\infty}\frac{e^{-\alpha(s,L)}\alpha(s,L)^n}{n!}\Big(v_{z,n+2}-v_{z,n+1}-\big(v_{z,n+1}-v_{z,n}\big)\Big)\Big(\frac{\partial\alpha(s,L)}{\partial s}\Big)^2 \\
&\quad+\sum_{n=0}^{\infty}\frac{e^{-\alpha(s,L)}\alpha(s,L)^n}{n!}\Big(v_{z,n+1}-v_{z,n}\Big)\frac{\partial^2\alpha(s,L)}{\partial s^2}<0.
\end{aligned}
$$

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:

$$
\begin{aligned}
\frac{\partial v_{z,n}}{\partial z} &= b(1-\rho)\int_{-\infty}^{z} d\bar{F}_{c,n}(c) = b(1-\rho)\Big(1-\big(1-F_c(z)\big)^n\Big)>0, \\
\Rightarrow \frac{\partial V_z(s)}{\partial z} &= \sum_{n=1}^{\infty}\frac{e^{-\alpha(s,L)}\alpha(s,L)^n}{n!}b(1-\rho)\Big(1-\big(1-F_c(z)\big)^n\Big)>0.
\end{aligned}
$$

Thus, higher valuation borrowers exert more search effort, because they gain more from an increase in the arrival rates of lenders. This completes the proof of proposition 2. ■

Proof of Lemma 3. Denote the probability that a type-z borrower accepts a loan offer with total cost c by $P_c(c,z)$. If $c \le z$, the borrower accepts the offer if it is the lowest-cost offer received, which occurs with probability $\big(1-F_c(c)\big)^n$ when the borrower examines n additional offers. If $c > z$, then the borrower does not accept that offer. Therefore:

$$
\begin{aligned}
P_c(c,z) &= \sum_{n=0}^{\infty}\frac{e^{-\alpha(z)}\alpha_z^n}{n!}\big(1-F_c(c)\big)^n \\
&= e^{-\alpha(z)F_c(c)} \\
&= e^{-\alpha(z)\int_{\underline{R}}^{\overline{R}}F_e(c-x)dF_R(x)}, \quad \text{if } c\le z, \qquad \text{(B1)} \\
P_c(c,z) &= 0, \quad \text{if } c>z. \qquad \text{(B2)}
\end{aligned}
$$

Denote by $P_R(R,z)$ the probability that a type-z borrower accepts a loan offer with interest rate R. A borrower with valuation z accepts this offer if its cost (including the idiosyncratic component) is less than z and if all other offers which he examines have higher costs. Integrating over the potential values of the idiosyncratic shock yields:

$$P_R(R,z) = \int_{-\infty}^{\infty} P_c(R+e,z)dF_e(e)$$
$$= \int_{-\infty}^{z-R} e^{-\alpha(z)\int_{\underline{R}}^{\overline{R}} F_e(R+e-x)dF_R(x)} dF_e(e). \tag{B3}$$

A borrower with valuation z accepts a loan offer with interest rate R if he examines the offer (probability $s(z)$) and the offer is better to any other offer he examines (probability $P_R(R,z)$). Therefore, the probability that a randomly-drawn borrower accepts a loan with interest rate R equals:

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

which yields equation (9).

Since $F_e(\cdot)$ is smooth, $P(R)$ is continuous and differentiable in R. Differentiating $P(R)$ with respect to R yields:

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

Hence, the probability that borrowers accept a loan is strictly decreasing in the interest rate R. This completes the proof of lemma 3. ■

Proof of Proposition 4. The optimal interest rate for a type-k lender solves:

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

Note that $\pi_k(R) < 0$ for $R < \frac{k}{1-\rho}$, $\pi_k(\frac{k}{1-\rho}) = 0, \pi_k'(\frac{k}{1-\rho}) = b(1-\rho)P(\frac{k}{1-\rho}) > 0$ and $\lim_{R\to\infty}\pi_k(R) = 0$. Therefore, there exist some $\tilde{R} > \frac{k}{1-\rho}$ such that $\pi_k'(\tilde{R}) = 0$ and, thus, the optimal choice $R(k)$ exists. In the case of multiple roots, the lender chooses the solution that yields higher profits. Finally, since $\pi_k(R)$ is continuously differentiable in k, $R(k)$ is continuous and differentiable in k.

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

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

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

Since the optimal interest rate is strictly increasing in the lender's cost k, we have $F_R\big(R(k)\big) = G(k)$ for $k \in [\underline{k}, \overline{k}]$. Hence,

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

Using this feature, we can rewrite equation (9) as follows:

$$P\big(R(k)\big) \quad = \quad \int_{\underline{z}}^{\overline{z}} s(z) \int_{-\infty}^{z-R(k)} e^{-\alpha(z)\int_{\underline{k}}^{\overline{k}} F_e\left(R(k)+e-R(x)\right) dG(x)} dF_e(e) dM(z). \tag{B4}$$

Equation (B4) defines the probability that borrowers accept the loan of the cost-k lender when all lenders make their equilibrium choice. This probability does not directly depend on the interest rate distribution because it incorporates the result that the offered interest rate is strictly decreasing in a lender cost k.

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

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

Differentiating profits with respect to $\tilde{k}$, we obtain:

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

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

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

which yields equation (10) that defines the interest rate schedule $R(k)$.

We now calculate the distribution of accepted offers. Denote the probability that a type-z borrower gets a loan by Q_z and the probability that he gets a loan with interest rate less than R by $Q_z(R)$. Note that a type-z borrower gets a loan if he receives at least one offer with cost below z. Therefore:

$$
\begin{aligned}
Q_z &= 1 - e^{-\alpha(z)F_c(z)} \\
&= 1 - e^{-\alpha(z)\int_{\underline{R}}^{\overline{R}} F_e(z-x)dF_R(x)} \\
Q_z(R) &= 1 - e^{-\alpha(z)\int_{\underline{R}}^{R} F_e(z-x)dF_R(x)}
\end{aligned}
$$

Denote the probability that a borrower gets a loan by Q and the probability that he gets a loan with interest rate less than R by $Q(R)$:

$$
\begin{aligned}
Q &= \int_{\underline{z}}^{\overline{z}} Q_z dM(z) \\
&= 1 - \int_{\underline{z}}^{\overline{z}} e^{-\alpha(s)\int_{\underline{R}}^{\overline{R}} F_R(z-x)dF_R(x)} dM(z) \\
Q(R) &= 1 - \int_{\underline{z}}^{\overline{z}} e^{-\alpha(z)\int_{\underline{R}}^{R} F_e(z-x)dF_R(x)} dM(z)
\end{aligned}
$$

The distribution of accepted interest rates $H_R(R)$ gives the proportion of borrowers who get a loan with interest rate less than R among the borrowers who get a loan:

$$
\begin{aligned}
H_R(R) &= \frac{Q(R)}{1-Q} \\
&= \frac{1 - \int_{\underline{z}}^{\overline{z}} e^{-\alpha(z)\int_{\underline{R}}^{R} F_e(z-x)dF_R(x)} dM(z)}{1 - \int_{\underline{z}}^{\overline{z}} e^{-\alpha(s)\int_{\underline{R}}^{\overline{R}} F_R(z-x)dF_R(x)} dM(z)}
\end{aligned}
$$

This completes the proof of proposition 4. ■

Proof of Proposition 5. A lender's expected profits are strictly decreasing in his cost k, since a lender can always mimic the action of a higher-cost lender and make strictly higher profits.

Denote the highest-cost lender that enters the market by $\hat{k}$, where $\hat{k} \leq \overline{k}$, and note that the measure of lenders that enter the market is $L = \Lambda\Gamma(\hat{k})$. Denote the profits of the highest-cost lender by $\underline{\pi}_{\hat{k}}$:

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

where

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

This equation makes explicit the dependence of L and $G(\cdot)$ on $\hat{k}$.

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

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

which is negative because the first term equals zero by the envelope theorem, the second term reflects the cost increase and the third term reflects that an increase in $\hat{k}$ increases the measure of lenders in the market, which reduces the probability that borrowers accept a loan offer. Therefore, given $s(\cdot)$, there is a unique $\hat{k}$ that characterizes lenders' cutoff cost $\hat{k}$.

The cutoff $\hat{k}$ is determined by equating the profits of the highest-cost lender with the entry cost χ, as equation (12) shows. ■

Proof of Proposition 7. Differentiating equation (13) with respect to s and equating to zero for every z, we obtain:

$$\frac{\partial \mathcal{W}_z\big(s^*(z), \hat{k}^*\big)}{\partial s} \;=\; \frac{\partial q\big(s^*(z), L^*\big)}{\partial s}.$$

We use equation (15) to rearrange the above equation, obtaining equation (16). The solution is unique for reasons similar to the decentralized case.

Differentiating equation (13) with respect to $\hat{k}$ and equating to zero, we obtain equation

(17). Notice that:

$$
\begin{aligned}
\frac{\partial \mathcal{W}_z(s_z^*, \hat{k}^*)}{\partial \hat{k}} &= \sum_{n=0}^{\infty} \Big[\frac{e^{-\alpha(s^*(z), L^*)} \big(\alpha(s_z^*, L^*)\big)^n}{n!} \mathcal{W}'_{z,n}(\hat{k}^*) \\
&\quad + \mathcal{W}_{z,n}(\hat{k}^*) \Big(\frac{e^{-\alpha(s_z^*, L^*)} \big(\alpha(s_z^*, L^*)\big)^{n-1} \frac{\partial \alpha(s_z^*, L^*)}{\partial L} \Lambda \Gamma'(\hat{k}^*)}{(n-1)!} \\
&\quad - \frac{e^{-\alpha(s_z^*, L^*)} \frac{\partial \alpha(s_z^*, L^*)}{\partial L} \Lambda \Gamma'(\hat{k}^*) \big(\alpha(s_z^*, L^*)\big)^n}{n!} \Big) \Big] \\
&= \sum_{n=0}^{\infty} \frac{e^{-\alpha(s_z^*, L^*)} \big(\alpha(s_z^*, L^*)\big)^n}{n!} \Big[\mathcal{W}'_{z,n}(\hat{k}^*) + \frac{\partial \alpha(s_z^*, L^*)}{\partial L} \Lambda \Gamma'(\hat{k}^*) \Big(\mathcal{W}_{z,n+1}(\hat{k}^*) - \mathcal{W}_{z,n}(\hat{k}^*) \Big) \Big].
\end{aligned}
$$

Furthermore:

$$
\mathcal{W}'_{z,n}(\hat{k}^*) = b(1-\rho) \int_{-\infty}^{z} (z-w) d\Big(\frac{\partial \bar{F}_{w,n}(w)}{\partial \hat{k}} \Big).
$$

Combining the last two equations yields equation (18). ■