

REGULATORY INTERVENTIONS IN CONSUMER FINANCIAL MARKETS: THE CASE OF CREDIT CARDS

Manolis Galenianos
Royal Holloway, University of London,
UK

Alessandro Gavazza
London School of Economics, UK

Abstract

We build a framework to understand the effects of regulatory interventions in credit markets, such as caps on interest rates. We focus on the credit card market, in which we observe US consumers borrowing at high and very dispersed interest rates despite receiving many credit card offers. Our framework includes two main features to account for these patterns: the endogenous effort of examining offers and product differentiation. Our calibration suggests that most borrowers examine few of the offers they receive, and thereby forego cards with low interest rates and high non-price benefits. The calibrated model implies that interest-rate caps reduce credit supply and significantly curb lenders' market power, thereby increasing consumer surplus. Moderate caps may yield larger gains in consumer surplus than tighter ones. (JEL: D83, D14, G28)

Teaching Slides

A set of Teaching Slides to accompany this article are available online as [Supplementary Data](#).

1. Introduction

After the 2008 financial crisis, policymakers in several countries began intervening in markets for consumer financial products more aggressively than before, through both

The editor in charge of this paper was Guido Lorenzoni.

Acknowledgments: We thank Victor Stango for sharing data with us and the editor (Guido Lorenzoni), the referees, Mark Armstrong, Glenn Ellison, John Kennan, Scott Nelson, Nikita Roketskiy, and Matthijs Wildenbeest 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). This work was supported in part through the New York University IT High Performance Computing resources, services, and staff expertise.

E-mail: manolis.galenianos@gmail.com (Galenianos); a.gavazza@lse.ac.uk (Gavazza)

Journal of the European Economic Association 2022 20(5):1897–1932

<https://doi.org/10.1093/jeea/jvac016>

© The Author(s) 2022. Published by Oxford University Press on behalf of European Economic Association. This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

new legislation and the creation of new regulatory agencies, such as the Consumer Financial Protection Bureau (CFPB) in the United States and the Financial Conduct Authority (FCA) in the United Kingdom. Although the interventions have taken different forms, one broad, prevalent direction is a break from the recent past—that is, the imposition of direct constraints on some prices and fees for financial products—because policymakers viewed them as “predatory”, that is, as targeting unsophisticated and poorly informed households.¹

An important question is what effects these policies will have on the operation of markets for consumer financial products. Standard competitive theory predicts that binding price caps reduce market efficiency and lower market access, particularly for marginal borrowers. Policymakers’ primary motivation, however, is that markets for consumer financial products do not satisfy the conditions of perfectly competitive markets because of informational and other frictions.² The theoretical analysis of frictional markets provides ambiguous predictions on the effects of these policies on consumer surplus and on aggregate welfare. For example, although price caps have the direct effect of lowering some prices, they also reduce incentives to become informed, which may increase market power and lead to higher, rather than lower, average prices (Fershtman and Fishman 1994; Armstrong, Vickers, and Zhou 2009). Therefore, determining the overall effect of these policies is an empirical/quantitative question.

The objective of this paper is to quantitatively study the effects of these price interventions on the market for consumer financial products. We focus on the US credit card market and on individuals who use their credit cards to borrow, combining several sources of data from the period before the financial crisis. These data build on Stango and Zinman (2016) (henceforth *SZ*) and display two key patterns that appear, at first sight, somewhat contradictory. First, the interest rates borrowers pay on their credit cards are high in comparison to funding costs and are very dispersed, even after controlling for observable borrower characteristics—including their creditworthiness, as captured by their credit score—or card characteristics (e.g. rewards). Thus, this first pattern suggests that lenders enjoy a significant amount of market power. Second, the average consumer receives several preapproved credit card offers at very different interest rates every month. This second pattern suggests that lenders face considerable competition, because borrowers can choose low-interest-rate credit cards from the offers they receive. Interpreting these seemingly contradictory patterns and deriving the policy implications of our interpretation are the main contributions of our paper.

1. Specifically, the 2009 US Credit Card Accountability Responsibility and Disclosure Act explicitly prohibited lenders from charging some fees on credit cards (Agarwal et al. 2015). Similarly, in the United Kingdom 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% per day; in November 2016, it restricted fees for individuals who want to access their pensions to a maximum of 1%. Furthermore, the financial press reports that the FCA is currently evaluating limits on fees for other products, such as mutual fund fees (*Financial Times*, Funds’ lucrative entry fees under attack, May 26, 2016) and mortgage origination fees (*Financial Times*, Mortgage lenders under FCA review for masking high fees, December 12, 2016).

2. Sirri and Tufano (1998) and Hortaçsu and Syverson (2004) show that information frictions play a prominent role in mutual fund markets, and Allen, Clark, and Houde (2019) and Woodward and Hall (2012) in mortgage markets.

We develop a rich but tractable modeling framework that accounts for these patterns. Our framework includes two features to help interpret high and dispersed interest rates: information frictions, such as the costs of examining and evaluating different offers and product differentiation.³ In our model, lenders, who are heterogeneous in their funding costs, choose whether to enter a particular market characterized by borrower creditworthiness (i.e. subprime, near-prime, prime, and super-prime), choose what interest rate to offer, and send credit card offers to borrowers. Borrowers, who are heterogeneous in their willingness to pay for a loan, choose how much effort to exert in examining the offers they receive and decide which offer, if any, to accept. The acceptance decision depends on the offer's interest rate and on an idiosyncratic match-specific attribute, interpreted as product differentiation. Hence, a borrower might reject a low-interest-rate credit card because he does not examine it, or because he does not like the idiosyncratic attribute.

We calibrate the model to match the statistics on the distribution of interest-rate offers, on the distributions of accepted offers, and on the fractions of borrowers in each market, as well as the lenders' average funding costs and charge-off rates. The model fits the data well. The calibration implies that whereas product differentiation affects borrowers' choice and lenders' pricing, our model requires that borrowers have a high cost of examining offers in order to match all empirical patterns in the data. The reason is as follows: Given a distribution of offered rates with high dispersion, both high examination costs and high product differentiation are potentially consistent with the high level of and the high dispersion in accepted interest rates. However, high examination costs are required to account for the high dispersion of offered rates and the moderate fraction of borrowers in the data, whereas the high product differentiation leads to a lower dispersion of offered rates and a larger fraction of borrowers in the population than those observed in the data.

We use our calibrated model to perform counterfactual experiments. We consider the effect of the introduction of three interest-rate caps, which are common across markets: 27.5, 25, and 22.5 percentage points (pps). These caps have relatively small effects on outcomes in the markets for creditworthy borrowers (prime and super-prime), because they are rarely binding in our data. In markets for riskier borrowers—that is, subprime and near-prime—the caps are binding for many borrowers; for example, the 25-pp cap is binding for 35% and 25% of subprime and near-prime borrowers in our data, respectively, and thus the effects are more pronounced: The number of offers and the average accepted interest rate decline substantially. However, the number of people getting a loan changes by a small amount in all cases and increases in most markets, because borrowers respond to the more favorable interest-rate distribution by examining a larger share of the offers they receive. The overall effect of any cap is a large redistribution of surplus from lenders to borrowers—that is, consumer surplus

3. Carlin (2009) argues that producers of retail financial products strategically make their prices more complex for consumers, thereby increasing consumers' costs of evaluating different products. Similarly, Ellison and Wolitzky (2012) develop a search model of obfuscation, and Ellison and Ellison (2009) provide empirical evidence on obfuscation among online retailers.

increases substantially and lender profits decline steeply—and a smaller change in aggregate welfare. Moreover, we find that caps can increase both consumer surplus and welfare—aggregate welfare increases relative to the baseline when the cap equals 25 or 27.5 pps—and a moderate cap may be preferable to a tighter one—consumer surplus and aggregate welfare are higher with a 25-pp cap than with a 22.5-pp cap.

We view the large positive effect of an interest-rate cap on the consumer surplus and welfare to be interesting results: In a perfectly competitive market with complete information, such caps will reduce supply precisely toward marginal borrowers and will, therefore, negatively affect their surplus. Our results are reminiscent of monopolistic markets, though our model explicitly accounts for the large number of credit card offers borrowers receive. Most notably, the presence of informational frictions—that is, the costs of examining and evaluating credit card offers—rationalizes the appearance of a competitive market with the reality of high lender market power.

The paper proceeds as follows: Section 2 reviews the literature and highlights our contributions. Section 3 describes the data. Section 4 presents the theoretical model. Section 5 presents our calibration of the model and illustrates its main quantitative implications. In Section 6, we conduct our counterfactual analyses. Section 7 concludes. The online appendices report further results and collect all proofs.

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.⁴ In an important contribution, Ausubel (1991) showed that interest rates on credit cards are substantially higher than lenders' funding costs and display limited intertemporal variability, and cites 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. Galenianos, Law, and Nosal (2021) study the interaction between imperfect competition, lender market power, and default in a quantitative dynamic model of consumption and savings. We contribute to this literature by building a framework that allows us to quantify the effects of product differentiation and choice frictions on lenders' loan pricing and consumers' cost of borrowing.

Second, a vast literature in household finance studies whether consumers behave optimally in credit markets: Among others, Agarwal et al. (2008) analyze consumer mistakes in the credit card market, and Ru and Schoar (2016) study how credit card companies exploit consumers' mistakes. In this strand of the literature, the most closely related paper is Woodward and Hall (2012), who study consumers' shopping effort in

4. A strand of the theoretical monetary literature examines how credit might improve market outcomes in the presence of search and trading frictions. See Diamond (1990) and Shi (1996) for early contributions.

the US mortgage market. We contribute to this literature by developing and calibrating an equilibrium model of a differentiated product market with endogenous consumers' shopping effort, 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 US credit card markets, Agarwal et al. (2015) and Nelson (2020) analyze how regulatory limits on credit card pricing introduced by the 2009 CARD Act affect borrowing costs by exploiting rich administrative data. Similarly, in a contemporaneous contribution, Cuesta and Sepúlveda (2019) study price regulation in the Chilean consumer loan market. We complement these papers by analyzing some of these regulatory interventions in a quantitative model that features product differentiation and borrowers' cost of examining offers, and we evaluate their importance for market power and 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), Gavazza (2016), Galenianos and Gavazza (2017), Allen, Clark, and Houde (2019), and Salz (2022). Our theoretical framework expands on these empirical papers by building on the models of Butters (1977) and Burdett and Judd (1983) and by combining product differentiation, search frictions, and consumers' endogenous shopping effort. Fershtman and Fishman (1994), Armstrong, Vickers, and Zhou (2009), and Janssen and Moraga-González (2004) show that consumers' shopping effort could potentially offset the effects of the regulations we focus on; thus, our framework that incorporates it seems well suited for a quantitative analysis of these policy interventions.⁵

3. Data

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

3.1. Data Sources

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

5. Knittel and Stango (2003) show that price ceilings served as a focal point for tacit collusion in the US credit card market during the 1980s. However, they also show that price dispersion was a lot more limited in that period than in the period of our data.

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 and 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 cardholders, most notably their FICO credit score.⁶ The second dataset samples different individuals from those in the first dataset and reports information on all preapproved credit card offers these individuals received in January 2007; most notably the number of offers and their interest rates.⁷ As *SZ* emphasize, there is a large dispersion in interest rates offered to a given individual in a given month. We should point out that we do not have access to individual survey data, and thus, we exploit data reported in *SZ*'s tables.

We complement these datasets with some additional statistics: the shares of subprime, near-prime, prime, and super-prime borrowers in the US population, calculated using the distribution of FICO scores reported by the Consumer Financial Protection Bureau (2012);⁸ the fraction of individuals with credit cards, computed from Campbell et al. (2016) and the 2007 Survey of Consumer Finances; the charge-off rate on credit card loans in the first quarter of 2007, reported by the Board of Governors of the Federal Reserve System; the interest rate of the one-year Treasury bill in January 16, 2007, which we use as the risk-free rate;⁹ and the Standard & Poor's US Credit Card Quality Index in January 2007, which is a monthly performance index that aggregates information on securitized credit card receivables, most notably reporting an average cost of funding (i.e. excluding expected 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 risks across individuals, as measured by their FICO scores; (2) different card characteristics across borrowers, such as rewards; and (3) different revolving balances across borrowers.

Specifically, the basic framework for this analysis is the following equation:

$$R_{ijt} = \gamma_X X_{it} + \gamma_Z Z_{ijt} + \varepsilon_{ijt}, \quad (1)$$

6. We are grateful to Victor Stango for sharing this dataset.

7. In the US market, lenders often send personalized preapproved credit card offers in the mail that commit to terms with borrowers. During our sample period, these preapproved offers were the main customer acquisition channel of credit card companies.

8. These shares equal 0.215, 0.140, 0.166, and 0.479, respectively.

9. We retrieved the values of the charge-off rate and the interest rate of the one-year Treasury bill from the FRED, Federal Reserve Bank of St. Louis, series <https://fred.stlouisfed.org/series/CORCCT100S> and <https://fred.stlouisfed.org/series/DGS1>, respectively.

where the dependent variable R_{ijt} is the APR individual i pays on credit card j in month t ; X_{it} are characteristics of individual i in month t —namely, his default risk, measured by the FICO score;¹⁰ Z_{ijt} are characteristics of individual i 's credit card j in period t —namely, the credit limit, rewards, and the credit balance; and ε_{ijt} are residuals.

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

$$R'_{ijt} = \hat{\gamma}_X \bar{X}_{it} + \hat{\gamma}_Z \bar{Z}_{ijt} + \hat{\varepsilon}_{ijt}, \quad (2)$$

where $\hat{\gamma}_X$ and $\hat{\gamma}_Z$ are the coefficient estimates, \bar{X}_{it} and \bar{Z}_{ijt} are the sample averages of the covariates of each regression, and $\hat{\varepsilon}_{ijt}$ are the estimates of the residuals. Hence, equation (2) removes the variation in R_{ijt} due to the variation in X_{it} and Z_{ijt} , while keeping that due to ε_{ijt} .

We perform regression (1) and calculate interest-rate residuals according to equation (2) separately for four groups of cardholders based on their FICO score: (1) subprime 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 groups constitute the main classification of borrowers used in the credit card industry (Consumer Financial Protection Bureau 2015). Hence, performing separate regressions for each group allows us to capture in a flexible way the heterogeneity across them, and thus to obtain a reasonably accurate measure of the dispersion in interest rates within each group of borrowers.

Table 1 reports coefficient estimates of several specifications of (1) and the main percentiles of the resulting distribution of interest rates based on (2). Odd-numbered columns report unweighted regressions and the resulting percentiles of residual interest rates, whereas even-numbered columns report regressions weighted by balance and the resulting percentiles.

Specifically, 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 pps for subprime borrowers, and it decreases for more-creditworthy borrowers, reaching a difference of 10 pps for super-prime borrowers.¹¹ However, many people use their credit cards as a means of payment and repay their balance in full at the end of each month. For such transactors, the interest rate is arguably not a salient feature of their credit card, since they never actually pay interest charges. To address this issue, column (2) weighs each observation used in column (1) by its revolving balance in the corresponding month; hence, the distribution does

10. The dataset reports household income brackets for approximately 50% of the individuals in the sample. In order to have larger sample sizes, we report results obtained without including income among the individual characteristics. However, we estimated equation (1) including income as an individual characteristic, and obtained results very similar to those reported in Table 1.

11. This difference exhibits a correlation with the standard deviation of interest rates within each market equal to 0.956.

TABLE 1. Dispersion of interest rates by borrower group.

Subprime borrowers	(1)	(2)	(3)	(4)	(5)	(6)
FICO score					-0.018 (0.006)	-0.022 (0.008)
Reward card					0.405 (0.674)	0.069 (0.535)
Credit limit					-0.158 (0.125)	-0.124 (0.116)
Credit balance					0.174 (0.140)	0.219 (0.108)
R^2					0.013	0.037
Observations	27,024	27,024	877	877	766	766
10th percentile	11.90	9.90	14.24	12.99	14.39	13.22
25th percentile	16.15	15.90	17.24	16.24	17.58	16.43
50th percentile	20.65	21.40	21.74	22.90	21.93	22.05
75th percentile	27.49	28.24	27.99	29.23	27.80	27.75
90th percentile	29.99	29.99	30.24	30.24	30.16	30.27
Near-prime borrowers						
FICO score					-0.052 (0.013)	-0.076 (0.014)
Reward card					0.562 (0.565)	-0.253 (0.504)
Credit limit					-0.255 (0.078)	-0.173 (0.061)
Credit balance					0.225 (0.100)	0.053 (0.072)
R^2					0.043	0.090
Observations	27,059	27,059	900	900	661	661
10th percentile	10.49	9.90	12.99	12.25	13.20	13.73
25th percentile	14.90	14.24	15.94	15.81	16.55	16.99
50th percentile	18.24	18.24	19.24	19.24	20.20	20.96
75th percentile	23.15	24.24	23.30	25.40	25.72	25.67
90th percentile	28.99	29.74	29.24	29.99	29.16	29.81
Prime borrowers						
FICO score					-0.052 (0.015)	-0.054 (0.015)
Reward card					-0.240 (0.520)	-0.614 (0.503)
Credit limit					-0.065 (0.049)	-0.100 (0.045)
Credit balance					0.013 (0.059)	0.078 (0.047)
R^2					0.029	0.033
Observations	31,115	31,115	953	953	604	604
10th percentile	9.90	9.90	11.99	11.24	11.55	11.63
25th percentile	12.99	12.99	14.31	14.24	14.81	14.73
50th percentile	16.74	16.99	18.24	18.24	17.90	18.00
75th percentile	19.99	20.34	20.34	21.24	21.90	21.84
90th percentile	25.99	28.99	28.15	28.99	28.65	28.88

TABLE 1. Continued

Subprime borrowers	(1)	(2)	(3)	(4)	(5)	(6)
Super-prime borrowers						
FICO score					-0.024 (0.010)	-0.039 (0.011)
Reward card					0.346 (0.471)	-0.226 (0.498)
Credit limit					0.028 (0.031)	0.030 (0.029)
Credit balance					-0.040 (0.051)	0.008 (0.040)
R^2					0.012	0.028
Observations	56,880	56,880	1,645	1,645	546	546
10th percentile	9.90	7.99	11.24	10.09	10.79	10.53
25th percentile	12.99	11.74	14.15	13.49	13.82	13.07
50th percentile	15.98	15.24	16.99	17.15	16.84	16.63
75th percentile	18.24	18.74	18.24	19.99	19.54	19.76
90th percentile	20.24	24.24	20.31	24.24	23.98	24.67

Note: This table reports OLS coefficient estimates of equation (1) and the corresponding percentiles of the distribution of centered interest rates as in equation (2).

not include interest rates of accounts whose balances are paid in full. For all borrower groups, the percentiles of the distributions reported in column (1) and column (2) are very similar.

Column (3) further restricts the data to January 2007 (i.e. before the financial crisis started) and excludes introductory “teaser” rates (i.e. low initial rates that reset to higher rates after an initial offer period). The restriction to January 2007 data has almost no effect on dispersion, because interest rates display limited aggregate and within-account variations over time. Removing teaser rates has a more meaningful effect on the distribution, and thus interest rates increase relative to those displayed in column (1), but the increase is small; for example, the difference between the 90th and the 10th percentiles slightly decreases to 16 pps for subprime borrowers and 9 pps for super-prime borrowers. Column (4) weighs each observation used in column (3) by the average revolving balance calculated over all available months in the panel period of January 2006–December 2008. Again, for all borrower groups, the distributions reported in columns (3) and (4) display very similar levels and overall dispersions of interest rates.

The specification of column (5) includes the main individual characteristic that should affect pricing; that is, the credit risk of the individual, measured by the FICO score; and controls for card characteristics, such as the credit limit, an indicator variable that equals one if the card features some rewards (e.g. frequent flier miles or cash back) and zero otherwise, and the revolving balance. Within all groups, higher-risk individuals face higher interest rates. Averaging across all groups, a 10-point increase

in the FICO score corresponds approximately to a 30-basis-point decrease in interest rates, which is almost identical to the magnitude Nelson (2020) estimates.¹²

The regressions reported in column (6) further weigh each observation used in the regressions of column (5) by the average revolving balance calculated over all available months in the panel period of January 2006–December 2008. The coefficient estimates are similar between column (5) and column (6); most notably, that of the FICO score. The percentiles reported in the bottom part of column (6) further weigh these residual interest rates by their average revolving balance over the sample period—that is, both the coefficient estimates and the distribution of residuals weigh each observation by the amount of the average revolving balance. For all borrower groups, the percentiles of the distributions reported in column (5) and in column (6) are strikingly similar, with both displaying a large dispersion of residual interest rates.¹³

Overall, Table 1 attests to some remarkable features of credit card markets. First, although more creditworthy borrowers on average pay lower interest rates, the difference in interest rates within groups is substantially larger than the difference across groups. Notably, the difference between the 90th and the 10th percentiles equals approximately 17 pps for subprime borrowers, near-prime borrowers, and prime borrowers, whereas it equals approximately 14 pps for super-prime borrowers. Hence, because the average outstanding revolving balance of borrowers equals approximately \$4,000 in our sample, moving from the 90th percentile to the 10th percentile of interest rates would reduce borrowers' annual payment by approximately \$500–\$600. Second, observable credit card characteristics do not seem to have a major 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.

Table 2 combines all empirical targets of our quantitative model. Panel A reproduces the percentiles of the distributions of interest rates derived in column (6) of Table 1. Panel B reports the statistics on credit card offers that SZ document. Specifically, Section 5.1 of SZ states that approximately 75% of individuals received two or more credit card offers during January 2007; among them, the median and mean number of offers was three and four, respectively. For these individuals who received two or more 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

12. Moreover, Nelson (2020) shows that (1) the interest rate on a credit card changes in response to a change in the credit cardholder's FICO score over time; and (2) the magnitude of the change in response to a change in the FICO score over time is almost identical to the cross-sectional difference between individuals with different FICO scores at credit card origination. These two observations imply that the long-term nature of the credit card contract does not affect the magnitude of the correlation between the FICO score and the interest rate in our data.

13. The R^2 of the regressions of Table 1 are lower than those reported in Table of 3 of SZ. The difference is due to the fact that we perform our regressions separately within each of the four groups of cardholders based on their FICO score.

TABLE 2. Empirical targets.

<i>Panel A: Accepted offers</i>	
10th Percentile accepted offer distribution, sub-prime borrowers	13.22
25th Percentile accepted offer distribution, sub-prime borrowers	16.43
50th percentile accepted offer distribution, sub-prime borrowers	22.05
75th percentile accepted offer distribution, sub-prime borrowers	27.75
90th percentile accepted offer distribution, sub-prime borrowers	30.27
10th percentile accepted offer distribution, near-prime borrowers	13.73
25th percentile accepted offer distribution, near-prime borrowers	16.99
50th percentile accepted offer distribution, near-prime borrowers	20.96
75th percentile accepted offer distribution, near-prime borrowers	25.67
90th percentile accepted offer distribution, near-prime borrowers	29.81
10th percentile accepted offer distribution, prime borrowers	11.63
25th percentile accepted offer distribution, prime borrowers	14.73
50th percentile accepted offer distribution, prime borrowers	18.00
75th percentile accepted offer distribution, prime borrowers	21.84
90th percentile accepted offer distribution, prime borrowers	28.88
10th percentile accepted offer distribution, super-prime borrowers	10.53
25th percentile accepted offer distribution, super-prime borrowers	13.07
50th percentile accepted offer distribution, super-prime borrowers	16.63
75th percentile accepted offer distribution, super-prime borrowers	19.76
90th percentile accepted offer distribution, super-prime borrowers	24.67
<i>Panel B: Received offers</i>	
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
<i>Panel C: Auxiliary statistics</i>	
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	4.01
Average funding cost	7.02

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

a credit card in 2007 reported by Campbell et al. (2016) with the probability of revolving conditional on having a credit card, which we compute directly in our data. Interestingly, the share of individuals with a credit card is lower for borrowers with lower credit scores, whereas the probability of revolving conditional on having a credit card is higher, exceeding 80% and 90% for near-prime and subprime borrowers, respectively. Hence, the fraction of credit card borrowers is non-monotonic in borrowers' credit scores; on average, 46% of individuals borrow on their credit

card.¹⁴ Finally, the aggregate charge-off rate equals approximately 4 pps, and the average funding cost reported by the Standard & Poor's Credit Card Quality Index is approximately 2 pps above the risk-free rate.

3.3. *Implications for Modeling*

Table 2 provides an interesting description of the credit card market and informs the model we develop in Section 4. We focus on two key data patterns. First, Panel A shows that the dispersion in the interest rates similar borrowers pay on their credit card debt is very large, even after we control for observable borrower and card characteristics. Second, Panel B points out that many individuals receive several credit card offers at substantially different interest rates. Hence, public information on individuals' repayment probabilities, as measured by FICO scores, does not account for the dispersion observed in Panel A and, even more so, in Panel B. Moreover, if all individuals chose the credit card with the lowest interest rate among all their offers, then the level and the dispersion of the accepted interest rate distributions would be considerably lower than those reported in Table 2.

These striking patterns motivate some of our modeling assumptions. We focus on borrowers (rather than transactors), and we allow two possible explanations for why these borrowers do not accept credit card offers with the lowest interest rates: (1) They may not examine all of the offers they receive; and (2) they may have idiosyncratic preferences for card attributes that our data do not report. The calibration of Section 5 aims to quantitatively match the dispersion of interest rates as well as the sizable number of offers individuals receive, and it allows us to assess the contributions of the two theoretical explanations for the empirical patterns. Nevertheless, our calibration admits other (unmodeled) factors that may feed into the large variance of the cross-sectional distribution of accepted rates, such as adjustment of the interest rate after the offer is accepted, as in Nelson (2020).

Despite all of their advantages, however, we acknowledge that our data have some limitations. First, they are mostly cross-sectional, and therefore, we do not have precise information on borrowers' and lenders' behavior over time. Specifically, we cannot precisely assess how frequently borrowers switch across credit cards. Hence, in the absence of a more-detailed measurement of borrowers' switching behavior, we seek to match the cross-sectional distribution through a static model.¹⁵ Moreover, although the theory can accommodate a large multidimensional heterogeneity of borrowers, our cross-sectional data make identifying such a model difficult. Hence, our framework allows for flexible heterogeneity across markets (i.e. subprime, near-prime, prime, and

14. 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 0.76 and 0.46, respectively. These shares closely match the corresponding aggregate statistics in the 2007 Survey of Consumer Finances: 0.73 and 0.44, respectively.

15. Search frictions/costs deliver a major part of the substance, if not the form, of switching costs in that they restrain competition and prevent the realization of every trade that yields a positive surplus.

super-prime), as well as within-market heterogeneity in borrowers' willingness to pay for credit and in borrowers' valuation of non-price card attributes, whereas some other parameters are common across borrowers within markets. Most notably, we do not observe individual default, and thus, we abstract from within-market heterogeneity in repayment risk as well as from asymmetric information between lenders and borrowers (though we note again that differences in observable FICO scores across individuals do not account for the dispersion in interest rates in Panel A and in Panel B). We further discuss the implications of these data limitations for our results 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 independent of each other, and each agent (borrower or lender) participates in a single market. Our calibration of Section 5 will consider four markets that correspond to the general classifications of creditworthiness used in the credit card industry: subprime, near-prime, prime, and super-prime.

Each market j has measure 1 of borrowers (a normalization) who have market-specific default risk ρ_j , want to take a loan of market-specific size b_j , and are heterogeneous in their marginal valuation of a loan \tilde{z} . We abstract from within-market heterogeneity in repayment probability for reasons we describe in Section 3.3. Furthermore, we abstract from the intensive-margin decision of how much to borrow for the sake of tractability, following Allen, Clark, and Houde (2019), Crawford, Pavanini, and Schivardi (2018), and Nelson (2020), among others. We allow for unobserved heterogeneity in borrowers' marginal valuation \tilde{z} , which is distributed according to a market-specific discrete distribution $\tilde{M}_j(\cdot)$ with an N_j -point support $\tilde{Z} = \{\tilde{z}_1, \dots, \tilde{z}_{N_j}\}$, where $\tilde{z}_1 \leq \dots \leq \tilde{z}_{N_j}$.¹⁶ We define $s_{\tilde{z}}$ to be the share of type- \tilde{z} borrowers, where $\tilde{z} \in \tilde{Z}$.

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, \tilde{k} . The marginal cost \tilde{k} follows a market-specific smooth distribution $\tilde{\Gamma}_j(\cdot)$ with connected support $[\tilde{k}_j, \bar{\tilde{k}}_j]$. The measure of lenders who choose to enter market j and the distribution of their marginal costs are L_j and $\tilde{G}_j(\cdot)$, respectively. Every entering lender can give one loan of size b_j .¹⁷

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 random borrower. Each borrower chooses his examination effort and then receives a random number of offers

16. We assume a discrete distribution of borrower types to facilitate some technical derivations. The interaction between borrowers and lenders does not hinge on that assumption.

17. A lender should be interpreted as a loan contract rather than a lending (or credit card) company. We do not model lending companies explicitly.

that follows a Poisson distribution, examines every offer with a probability that depends on his effort, and decides which, if any, to accept.¹⁸ The realized effective number of offers (i.e. offers received *and* examined) to a borrower who exerts examination effort e in a market with L_j lenders follows a Poisson distribution with parameter $\alpha \equiv e * L_j$, that is, α is a borrower's effective arrival rate of offers. A borrower who exerts examination effort e incurs cost $q_j(e, L_j)$, where $q_j(\cdot, \cdot)$ is strictly increasing and convex in effort e and satisfies $q_j(0, L_j) = 0$, $\lim_{e \rightarrow 0} \partial q_j(e, L_j) / \partial e = 0$ and $\lim_{e \rightarrow 1} q_j(e, L_j) = \infty$.

Borrowers consider two components in ranking loan offers. The first is the net interest rate R , which is chosen by the lender and is drawn from the equilibrium offer distribution $F_{R_j}(\cdot)$. The second component is an idiosyncratic (i.e. borrower-specific) attribute a , which is stochastic and represents every other aspect of the loan that might affect the borrower's valuation. The attribute draw captures the importance of horizontal product differentiation in this market, whose value may vary across borrowers. Idiosyncratic attribute a is drawn from an exogenous distribution $F_{a_j}(\cdot)$ that is smooth, has zero mean, and has support in a connected set $[\underline{a}, \bar{a}] \subset (-\infty, \infty)$. In this section and in the calibration of Section 5, we assume attribute draw a is independent across lenders. Online Appendix D considers the case in which some lenders' offers might consistently draw better values of a . We call the sum $c = R + a$ the (net) cost of a loan to the borrower, which might be higher or lower than R depending on attribute a .

If the borrower does not default, which occurs with probability $1 - \rho_j$, then he pays the cost of the loan; if he defaults, which occurs with probability ρ_j , then he incurs the utility cost of default $\tilde{\delta}_j$.¹⁹ The expected utility of a type- \tilde{z} borrower in market j who takes a loan with interest rate R , and attribute a is $b_j(\tilde{z} - (1 - \rho_j)(1 + R + a) - \rho_j \delta_j)$, where $\delta_j \equiv \tilde{\delta}_j / b_j$.²⁰ The borrower's utility from not taking a loan is zero.

We define a borrower's preference for a loan net of expected default cost and principal repayment as $z = (\tilde{z} - \rho_j \delta_j) / (1 - \rho_j) - 1$, and note that it is distributed according to $M_j(z) = \tilde{M}_j((1 - \rho_j)(z + 1) + \rho_j \delta_j)$ with support $Z = \{z_1, \dots, z_{N_j}\}$. We can therefore rewrite the utility of a type- z borrower from taking a loan with cost

18. The random allocation of offers across borrowers in an environment with a finite number of borrowers and lenders leads to urn-ball matching, which, as the numbers of borrowers and lenders grow large, is approximated by a Poisson distribution. See Butters (1977) for an early application of urn-ball matching to a similar setting.

19. We assume defaulting occurs independent of any loan features; that is, the interest rate R or the attribute draw a .

20. Our assumption implies that borrowers do not enjoy the benefits of the idiosyncratic product differentiation draw a in the event of default. This specification captures the fact that many credit card benefits expire upon default. However, a simple redefinition of the variable a encompasses the case in which borrowers enjoy the idiosyncratic attribute in the event of default, without changing any theoretical or quantitative results.

$R + a$ as

$$b_j(1 - \rho_j)(z - R - a). \tag{3}$$

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 offer $V_{z,j}(e)$ (which depends on effort e) net of the cost of effort, $q_j(e, L_j)$:

$$V_{z,j}(e) - q_j(e, L_j). \tag{4}$$

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

The expected profits per dollar lent for a type- \tilde{k} lender in market j equal the difference between the expected revenues, given by the gross interest rate $(1 + R)$ times the market-specific repayment probability $(1 - \rho_j)$, and the costs, given by the lender's gross cost of funds $(1 + \tilde{k})$: $(1 + R)(1 - \rho_j) - (1 + \tilde{k})$. We define the lender's expected marginal cost inclusive of non-repayment risk as $k = \tilde{k} + \rho_j$, which means that per-dollar expected profits equal $R(1 - \rho_j) - k$. We note that for potential lenders, k is distributed according to the smooth distribution $\Gamma_j(k) = \tilde{\Gamma}_j(k - \rho_j)$ with support $[\underline{k}_j, \bar{k}_j]$, where $\underline{k}_j = \tilde{\underline{k}}_j + \rho_j$ and $\bar{k}_j = \tilde{\bar{k}}_j + \rho_j$; for entrants, k is distributed according to $G_j(k) = \tilde{G}_j(k - \rho_j)$.²¹

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 expected profits

$$\pi_{k,j}(R) = b_j(R(1 - \rho_j) - k)P_j(R). \tag{5}$$

Notice that idiosyncratic attribute a 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' effort $e_j(\cdot)$ 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 (4), lenders maximize their expected profits (5), 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.

21. It will prove convenient for some derivations to maintain the assumption $z_{N_j} - a > \bar{k}_j / (1 - \rho_j)$ (recall that $a < 0$). Nothing important hinges on this assumption.

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

4.1. Borrowers' Choices

We characterize borrowers' optimal effort $e(\cdot)$ of examining offers, taking as given the measure L of lenders in the market and the interest rate offer distribution $F_R(\cdot)$. We should point out that the type distribution of lenders $G(\cdot)$ and interest-rate choices $R(\cdot)$ affect borrowers' choices only through $F_R(\cdot)$.

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

$$V_z(e) = \sum_{n=0}^{\infty} \frac{e^{-eL} (eL)^n}{n!} v_{z,n}. \quad (6)$$

Notice that effort e affects the arrival rate of offers but does not enter $v_{z,n}$; therefore, it is immediate from equation (6) that $V_z(e)$ is continuous and differentiable in e . As a result, the optimal effort choice $e(z)$ solves

$$V'_z(e) = \frac{\partial q(e, L)}{\partial e}. \quad (7)$$

To determine $v_{z,n}$ for $n \geq 1$, recall that the borrower chooses the loan offer with the lowest cost c , if $c \leq z$. Let $F_c(\cdot)$ denote the distribution of c . Because the loan cost c is the sum of two independent random variables (R and a), it is distributed according to

$$F_c(c) = \int_R^{\bar{R}} F_a(c - R) dF_R(R).$$

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 examining $n \geq 1$ offers is

$$v_{z,n} = b(1 - \rho) \int_{-\infty}^z (z - c) d\bar{F}_{c,n}(c). \quad (8)$$

The following proposition characterizes borrowers' optimal effort $e(\cdot)$ of examining offers, the resulting distribution of accepted rates, and the fraction of borrowers who get a loan, conditional on lenders' actions.

PROPOSITION 1. Given $F_R(\cdot)$ and L :

1. The optimal effort of a type- z borrower, $e(z)$, is unique and strictly increasing in z and solves

$$\sum_{n=0}^{\infty} \frac{e^{-eL}(eL)^n}{n!} (v_{z,n+1} - v_{z,n})L = \frac{\partial q(e, L)}{\partial e}, \tag{9}$$

where $v_{z,0} = 0$, and equation (8) defines $v_{z,n}$ for $n \geq 1$.

2. The distribution of accepted offers equals

$$H_R(R) = \frac{1 - \sum_{z \in Z} s_z e^{-e(z)L} \int_R^R F_a(z-x) dF_R(x)}{1 - \sum_{z \in Z} s_z e^{-e(z)L} \int_R^{\bar{R}} F_a(z-x) dF_R(x)}. \tag{10}$$

3. The fraction of borrowers who get a loan is

$$Q = 1 - \sum_{z \in Z} s_z e^{-e(z)L} \int_R^{\bar{R}} F_a(z-x) dF_R(x). \tag{11}$$

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 offer distribution $F_R(\cdot)$, and finally characterize lenders' entry decisions L and $G(\cdot)$ given borrowers' effort $e(\cdot)$. To ease notation, we denote the effective arrival rate of offers to a type- z borrower by $\alpha(z) \equiv e(z) * L$.

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 he examines (taking into account their attributes a), and if this offer yields net positive utility to the borrower. The next lemma characterizes the probability $P(R)$ that a random borrower accepts a loan offer with interest rate R .

LEMMA 1. Given $F_R(\cdot)$, L , and $e(\cdot)$, the probability $P(R)$ that borrowers accept a loan offer with interest rate R is continuous and differentiable in R and equals

$$P(R) = \sum_{z \in Z} s_z e(z) \int_{-\infty}^{z-R} e^{-\alpha(z) \int_R^{\bar{R}} F_a(R+a-x) dF_R(x)} dF_a(a). \tag{12}$$

Furthermore, $P'(R) < 0$.

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

PROPOSITION 2. Given L , $G(\cdot)$, and $e(\cdot)$:

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

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

$$\begin{aligned} & \sum_{z \in Z} s_z e(z) \int_{-\infty}^{z-R(k)} e^{-\alpha(z) \int_k^{\bar{k}} F_a(R(k)+a-R(x)) dG(x)} dF_a(a) \\ &= \left(R(k) - \frac{k}{1-\rho} \right) \sum_{z \in Z} s_z e(z) \left(\int_{-\infty}^{z-R(k)} e^{-\alpha(z) \int_k^{\bar{k}} F_a(R(k)+a-R(x)) dG(x)} \right. \\ & \quad \times \left. \left(\alpha(z) \int_k^{\bar{k}} F'_a(R(k) + a - R(x)) dG(x) \right) dF_a(a) \right. \\ & \quad \left. + e^{-\alpha(z) \int_k^{\bar{k}} F_a(z-R(x)) dG(x)} F'_a(z - R(k)) \right). \end{aligned} \tag{13}$$

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

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

PROPOSITION 3. Given borrowers' effort $e(\cdot)$, lenders' entry satisfies the following:

1. A cutoff cost \hat{k} exists 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 $G(k) = \Gamma(k) / \Gamma(\hat{k})$ for $k \leq \hat{k}$ and $G(k) = 1$ for $k > \hat{k}$.
3. The cutoff cost \hat{k} solves

$$\begin{aligned} & b \left(R(\hat{k})(1-\rho) - \hat{k} \right) \sum_{z \in Z} s_z e(z) \\ & \quad \times \int_{-\infty}^{z-R(\hat{k})} e^{-e(z) \Lambda \Gamma(\hat{k}) \int_k^{\hat{k}} F_a(R(\hat{k})+a-R(x)) d \frac{\Gamma(x)}{\Gamma(\hat{k})}} dF_a(a) = \chi. \end{aligned} \tag{14}$$

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 markets: subprime, near-prime, prime, and super-prime borrowers.

We borrow some parametric assumptions about the distributions of borrowers' and lenders' 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 discretized lognormal distribution with parameters μ_{z_j} and σ_{z_j} , and $N_j = 20$ support points for the distribution $M_j(z)$ of buyers' preferences z in market j . Moreover, we assume that the distribution $\tilde{G}_j(\tilde{k})$ of sellers' costs \tilde{k} is common across markets and follows 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 in January 16, 2007, which equals 5.06%—and an upper-truncation point \hat{k} . The assumption of a common cost distribution across markets means that the mass of potential lenders Λ_j varies across markets.

We normalize the loan size to $b_j = 1$. We further assume the following: (1) the effort cost of examining offers $q_j(e, L_j)$ equals $\beta_{0j}e^{\beta_1}$; (2) the charge-off rate in market j equals ρ_j ; and (3) attribute a is unobserved in our data, and it follows a normal distribution with mean zero and standard deviation σ_{a_j} , and is uncorrelated with R (Online Appendix D reports the results of the calibration in the case in which attribute a is correlated with the costs k and, thus, with the interest rate R as well). Hence, the effort costs of examining offers, the charge-off rates, and the standard deviations of the product attributes vary across markets.

Finally, we assume that the reported accepted rates \hat{R}_j and the “true” accepted rate R_j are related as $\hat{R}_j = R_j\eta$, where η is a random variable, identically and independently distributed across observations, drawn from a lognormal distribution with parameters (μ_η, σ_η) , common across markets, and with mean to equal 1. Hence, reported rates are unbiased and 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 that wages are measured with error. In our application, surveyed borrowers may report the interest rates they pay on their credit card debt incorrectly. Moreover, the random variable η may also account for some additional factor our model does not consider, such as adjustment of the interest rate after the offer is accepted, as in Nelson (2020). Table 2 shows that the distributions of accepted rates display a large dispersion, and these random η allow the model to more precisely match this feature of the data quantitatively.

5.2. Calibration

We choose the vector $\psi = \{L_j, \mu_{z_j}, \sigma_{z_j}, \xi, \hat{k}, \rho_j, \sigma_{a_j}, \beta_{0j}, \beta_1, \sigma_\eta\}_{j \in J}$ that minimizes the distance between the target moments m reported in Table 2 and the corresponding moments of the model. We calibrate two versions of the model: In the first, we impose $\sigma_\eta = 0$; and in the second, σ_η can take any positive value.

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' effective arrival rate $\alpha_j(z)$ in each market j that are consistent with each other. Once we solve for these policy functions of borrowers and lenders in each market j , we compute the equilibrium distributions of the interest rates of received offers and of accepted offers. In practice, we simulate these distributions and compute the moments $\mathbf{m}(\psi)$ that correspond to those reported in Table 2 on received offers and on accepted offers, as well as the aggregate fraction of credit card borrowers in each market j . Panels A and 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 that correspond to the population share of each group j (see footnote 8) to aggregate the number of received offers and their interest-rate distribution.

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 because we combine two separate datasets, collected for different purposes. Specifically, the dataset on received offers reports *all* offers that borrowers in group j receive, whose arrival rate is L_j , and not exclusively the offers that borrowers examine in equilibrium, which may be lower than the offers received because borrowers' endogenous examination effort e may be less than full effort $e = 1$. We derive in Online Appendix A the average number of offers and the distribution of the difference between the highest and the lowest offers borrowers receive under the assumption that the arrival rates of these offers equal L_j .

However, lenders send these offers anticipating that borrowers examine only a subset of them according to their equilibrium effort. Hence, the moments of the empirical distribution of accepted offers reflect borrowers' endogenous examination effort.

5.4. Sources of 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 more heavily on certain moments in the data.

The moments on the number of offers borrowers receive identify the average offer rate $\sum_j \omega_j L_j$ (where ω_j are the known shares of borrowers in each group), and thus contribute to the identification of group-specific offer rates L_j .²² Similarly, the aggregate charge-off rate is informative about the group-specific default rates ρ_j . Moreover, we identify the parameter ξ of the distribution $G(k)$ of sellers' heterogeneity from the average funding cost reported by Standard and Poor's.

Furthermore, we observe the distribution of the difference between the highest and lowest offered interest rates R that borrowers receive. We show in Online Appendix A that this distribution depends in a precise way on the offer distribution $F_{R_j}(\cdot)$, which allows us to recover $F_{R_j}(\cdot)$.

With this knowledge, we still have to recover three sets of parameters that determine (1) the distribution of borrowers' preferences; (2) borrowers' examination effort; and (3) the extent of product differentiation (variation in attributes). Propositions 1 and 2 show that these three sets of parameters shape three mappings between observable outcomes: (A) the mapping between the distribution $G(k)$ of costs k and the distribution $F_{R_j}(R)$ of offered rates $R(k)$; (B) the mapping between the offer distribution $F_{R_j}(R)$ and the distribution of accepted rates $H_{R_j}(R)$; and (C) the mapping between the offer distribution $F_{R_j}(R)$ and the fraction of borrowers who get a loan Q_j . Hence, these three outcomes jointly identify the remaining three sets of parameters.

Intuitively, given lenders' costs $G(k)$, the dispersion of offers (i.e. mapping A) increases in the dispersion of borrowers' preferences and decreases in the standard deviation σ_a of the product attribute a , because lenders (most notably low-cost lenders) charge similar rates, anticipating that consumers' choices depend relatively less on interest rates when a displays larger values. Furthermore, given the offer distribution F_{R_j} , the dispersion of accepted interest rates (i.e. mapping B) increases in examination costs—because borrowers examine fewer offers when costs are high—and in the standard deviation σ_a of the product differentiation, because larger values of a imply that interest rates affect consumers' choices relatively less than smaller values. Similarly, given the offer distribution F_{R_j} , equation (11) shows that the fraction of borrowers (i.e. mapping C) increases as examination effort $e(z)$ increases. Our discussion of the calibrated parameters in Section 5.5 and the comparative statics in Section 5.7 will further clarify how examination costs and product attribute a differentially affect market outcomes.

22. We should point out that the main implications of the model do not particularly rely on the specific values of L_j , but rather on borrowers' effective arrival rates $\alpha(z) = e(z)L_j$, which are identified from the distributions of costs, of offered rates, and of accepted rates, as we explain shortly. Hence, different values of L_j would imply different equilibrium values of $e(z)$ (and thus different costs β_{0j}), keeping $\alpha(z)$ unchanged.

TABLE 3. Calibrated parameters.

<i>Panel A: No measurement error</i>				<i>Panel B: Measurement error</i>			
μ_{z_1}	3.644	σ_{z_1}	0.143	μ_{z_1}	3.575	σ_{z_1}	0.123
μ_{z_2}	3.563	σ_{z_2}	0.082	μ_{z_2}	3.532	σ_{z_2}	0.108
μ_{z_3}	3.525	σ_{z_3}	0.157	μ_{z_3}	3.444	σ_{z_3}	0.127
μ_{z_4}	3.242	σ_{z_4}	0.342	μ_{z_4}	3.224	σ_{z_4}	0.191
ξ	3.626	\hat{k}	11.048	ξ	4.489	\hat{k}	9.661
L_1	1.440	L_2	3.683	L_1	1.552	L_2	3.947
L_3	3.116	L_4	3.156	L_3	3.228	L_4	2.995
ρ_1	0.015	ρ_2	0.007	ρ_1	0.040	ρ_2	0.030
ρ_3	0.004	ρ_4	0.003	ρ_3	0.020	ρ_4	0.010
σ_{a_1}	0.158	σ_{a_2}	0.143	σ_{a_1}	0.077	σ_{a_2}	0.118
σ_{a_3}	0.155	σ_{a_4}	0.101	σ_{a_3}	0.144	σ_{a_4}	0.125
β_{01}	9.069	β_{02}	34.502	β_{01}	8.629	β_{02}	42.408
β_{03}	28.051	β_{04}	30.075	β_{03}	28.938	β_{04}	32.663
β_1	1.555	σ_η	0.000	β_1	1.739	σ_η	0.284

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 ($\sigma_\eta > 0$).

Finally, lenders' free-entry condition (equation (14)) 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 thus we now discuss only those in Panel A of Table 3. As we report above and Table 4 shows 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 both within each group and across groups. Specifically, borrowers' average willingness to pay decreases as their creditworthiness increases. The standard deviation of willingness to pay, equal to $\sqrt{e^{2\mu_{z_j} + \sigma_{z_j}^2} (e^{\sigma_{z_j}^2} - 1)}$, is non-monotonic in creditworthiness, with super-prime borrowers displaying a standard deviation approximately ten times larger than near-prime borrowers.

The parameters ξ and \hat{k} of the distribution of costs \hat{k} imply that the average costs of all entrants (not weighted by market shares) equal 648 basis points (the average funding cost used in the calibration weighs lenders by their market shares, and equals 616 basis points at the calibrated parameters). Thus, average costs display a spread

TABLE 4. Model fit.

	Data	Model $\sigma_\eta = 0$	Model $\sigma_\eta > 0$
10th percentile accepted rate, subprime borrowers	13.22	17.53	14.22
25th percentile accepted rate, subprime borrowers	16.43	18.76	17.20
50th percentile accepted rate, subprime borrowers	22.05	21.40	21.28
75th percentile accepted rate, subprime borrowers	27.75	25.12	26.46
90th percentile accepted rate, subprime borrowers	30.27	28.19	31.96
10th percentile accepted rate, near-prime borrowers	13.73	17.26	13.54
25th percentile accepted rate, near-prime borrowers	16.99	18.52	16.36
50th percentile accepted rate, near-prime borrowers	20.96	21.22	20.38
75th percentile accepted rate, near-prime borrowers	25.67	24.95	25.32
90th percentile accepted rate, near-prime borrowers	29.81	27.89	30.67
10th percentile accepted rate, prime borrowers	11.63	15.34	12.20
25th percentile accepted rate, prime borrowers	14.73	16.43	14.77
50th percentile accepted rate, prime borrowers	18.00	18.78	18.30
75th percentile accepted rate, prime borrowers	21.84	22.03	22.80
90th percentile accepted rate, prime borrowers	28.88	24.76	27.90
10th percentile accepted rate, super-prime borrowers	10.53	13.75	11.14
25th percentile accepted rate, super-prime borrowers	13.07	14.60	13.39
50th percentile accepted rate, super-prime borrowers	16.63	16.41	16.43
75th percentile accepted rate, super-prime borrowers	19.76	18.87	20.12
90th percentile accepted rate, super-prime borrowers	24.67	20.96	24.28
Fraction receiving 2+ offers (%)	75.00	74.43	74.70
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	3.49	3.49
10th percentile distribution of differences in offered rates	0.00	1.60	1.21
30th percentile distribution of differences in offered rates	2.25	4.03	2.98
50th percentile distribution of differences in offered rates	4.34	5.89	4.37
70th percentile distribution of differences in offered rates	7.25	7.74	5.85
90th percentile distribution of differences in offered rates	9.25	10.19	8.61
Fraction with credit card debt, subprime borrowers	54.56	55.60	54.86
Fraction with credit card debt, near-prime borrowers	55.33	55.78	55.25
Fraction with credit card debt, prime borrowers	54.00	54.74	54.20
Fraction with credit card debt, super-prime borrowers	36.02	35.70	36.00
Charge-off rate	4.01	0.72	2.29
Average funding cost	7.02	6.16	5.95
Criterion function		139.63	20.36

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

of approximately 140 basis points over the risk-free rate. Moreover, the heterogeneity of lenders' costs is small; that is, the standard deviation of costs equals 121 basis points. Thus, the model generates a large dispersion of offered rates even with a small dispersion of costs.

The values of L_j indicate that lenders send, on average, approximately 2.8 credit card offers, with considerable heterogeneity across groups—subprime borrowers receive less than half the offers that near-prime, prime, and super-prime borrowers

receive. The number of offers is non-monotonic in the creditworthiness of borrowers, and thereby matches the patterns Han, Keys, and Li (2018) report. However, the parameters β_{0j} and β_1 imply that borrowers examine only a small fraction of these offers: The cost of effort to examine an average number of offers equal to $\alpha = 1$ is approximately 470 basis points, which corresponds to approximately \$190 given that borrowers' average outstanding revolving balance equals approximately \$4,000 in our sample. It increases by approximately 900–1,000 basis points across borrower groups—approximately \$350–\$400—to examine an average number of offers equal to $\alpha = 2$.

The value of σ_{a_j} implies that the standard deviation of the product attribute a is not large, relative to the overall heterogeneity in borrowers' preferences.

To understand why the calibration results in a sizable role for examination costs and a more modest one for product differentiation, we turn to the parts of the model that most contribute to their identification; namely, the mappings from lenders' costs to lenders' offers and from lenders' offers to borrowers' outcomes. Key features of the data are: (1) the level of offered interest rates is high relative to lenders' costs; (2) the distribution of offered rates is very dispersed; (3) the accepted rate distribution is similar to the offered rate distribution; (4) borrowers receive several credit card offers; and (5) the share of households with credit card debt is moderate. The central question is why borrowers do not end up with low interest rates, given that they receive many offers from a highly dispersed offer distribution.

Examination costs and product differentiation are two features that could rationalize the fact that individuals do not borrow at low interest rates. If examination costs are high, then borrowers examine few of the offers they receive—often only one—which results in high and dispersed accepted rates. If examination costs are low, then borrowers examine several offers, which tends to reduce the level and dispersion of accepted rates. High accepted rates could still occur with low examination costs if borrowers choose an offer because the value of its attribute a is large, which occurs more often when product differentiation is an important feature of the market; that is, when σ_a is large. However, the combination of low examination cost and high product differentiation leads to the following: (1) On the supply side, lenders set similar interest rates, because they matter less for borrowers' choices when the variance of the attribute a is larger; (2) holding the distribution of offered rates constant, on the demand side, many borrowers take out loans because borrowing is attractive due to large values of a . However, the data show that the dispersion of offered rates is very high, and that the aggregate fraction of individuals with credit card debt is moderate (approximately 45%).

In summary, the model calls for high examination costs β_{0j} in order to match the moderate fraction of borrowers, as well as the high dispersion of offered and of accepted rates we observe in the data. The comparative statics of Section 5.7 further illustrate these issues.

Finally, the calibrated σ_η equals 0.284, which means that the standard deviation of the measurement error on the accepted rates equals 0.289. This value is small

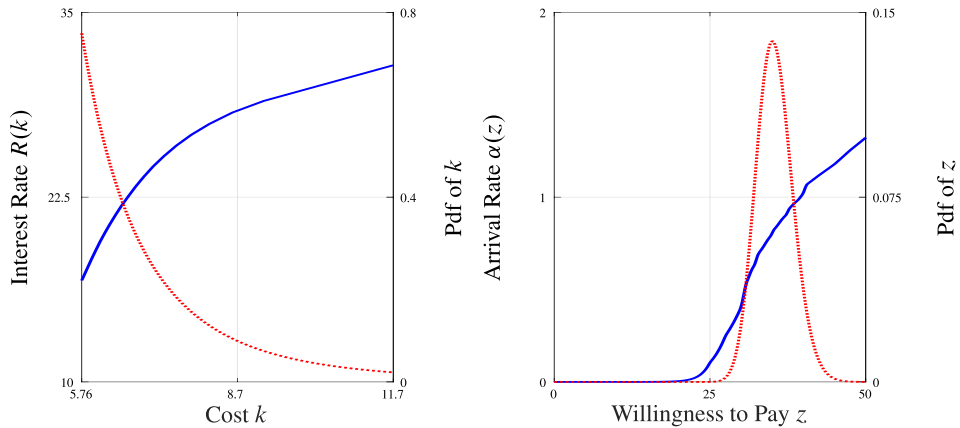


FIGURE 1. The left panel displays lenders’ optimal interest rate $R(k)$ (solid line, left axis) and the density of their cost k (dotted line, right axis) in the near-prime market. The right panel displays near-prime borrowers’ optimal arrival rate $\alpha(z)$ (solid line, left axis) and the density of their willingness to pay z (dotted line, right axis).

relative to the calibrated standard deviations of accepted rates R in the version without measurement error, which equal 3.89 in the subprime market, 3.88 in the near-prime market, 3.39 in the prime market, and 2.56 in the super-prime market.

Table 4 presents a comparison between the empirical moments and the moments calculated from the model at the calibrated parameters reported in Panels A and B of Table 3, respectively. The model without measurement error matches the data well; however, as anticipated, it underpredicts the dispersion of accepted rates—that is, it overpredicts the lower percentiles and underpredicts the higher percentiles. It matches reasonably well the percentiles of the distribution of the difference between the highest and the lowest offered interest rates, and almost perfectly matches 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, as observed in the data. The model with a small measurement error on accepted offers matches the data almost perfectly. Perhaps the most notable difference between the model and the data is the fact that the model underpredicts the aggregate charge-off rate.

5.6. Model Implications

We study the implications of the model evaluated at the parameters reported in Panel A of Table 3. Because 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 the near-prime market (Online Appendix Figure C.1 displays them in the other markets). The left panel displays lenders’ optimal offered rate $R(k)$ (solid line, left axis) 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 (14) determines. Lenders' offered rates are strictly increasing in their costs k , as Proposition 2 states. Markups, computed as $(R(k)(1 - \rho_j) - k)/k$, are non-monotonic in borrowers' creditworthiness: They equal 179%, 219%, 189%, and 153% in the subprime, near-prime, prime, and super-prime markets, respectively.

The right panel of Figure 1 displays near-prime borrowers' effective arrival rate of offers $\alpha(z)$ (solid line, left axis), which is the outcome of their optimal examination effort e , as a function of their willingness to pay z , as well as the density of z (dotted line, right axis). Because the lowest-valuation near-prime borrowers have a willingness to pay that is below almost all offered interest rates, and the product attribute a has a small variance, these borrowers do not exert any effort to examine offers. More generally, the examination effort is low—on average, borrowers examine approximately 0.7 offers—and only borrowers with the highest willingness to pay spend enough effort to attain $\alpha(z)$ larger than 1.

Moreover, borrowers' demand functions—that is, the probabilities $P_j(R)$ that borrowers accept a credit card offer with an interest rate R , displayed in Online Appendix Figure C.2—show some noteworthy features: (1) For any R , the probability that the 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. (2018) report; and (2) because of borrowers' low examination effort, borrowers' demand is quite inelastic—the average elasticity equals approximately -1.50 —which is similar to the elasticity Nelson (2020) estimates. These external comparisons seem to suggest that our calibration yields reasonable parameters. Overall, the average acceptance probability $P_j(R)$ equals 0.16, which, scaled up by the mass of lenders $\sum_{j=1}^J \omega_j L_j \approx 2.8$, yields an aggregate fraction of individuals with credit card debt of 46.4%.

Moreover, the difference between the distributions $F_{R_j}(R)$ of offered rates and the distributions $H_{R_j}(R)$ of accepted rates (displayed in Online Appendix Figure C.3) are small, for two reasons: (1) Borrowers' low examination 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 differentiation attribute a . However, this second factor is quantitatively smaller than the first one, because the standard deviation σ_a is small, and because, for a to have sizable effect, borrowers would need to consider more than one offer, which happens very infrequently due to their high costs of examining them. 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 rate.²³

Finally, Table 5 reports summary statistics of market outcomes—prices and quantities—as well as consumer surplus, lenders' profits, and aggregate welfare in

23. Of course, this is not a full equilibrium argument, as the endogenous distribution of offered rates $F_R(\cdot)$ depends on the product attribute a .

TABLE 5. Market outcomes and welfare.

	Subprime	Near-prime	Prime	Super-prime
Average number of offers per borrower	1.44	3.68	3.12	3.16
Average accepted rate	22.24	22.20	19.16	16.47
Standard deviation of accepted rates	3.90	3.89	3.39	2.56
Fraction of borrowers	55.17	54.80	54.92	37.03
Consumer surplus	5.48	4.11	5.19	3.56
Lender profits	1.80	1.91	1.73	1.20
Welfare	7.28	6.02	6.92	4.77

Note: This table reports market outcomes and welfare in each market.

each market. Interestingly, consumer surplus, lenders' profits, and aggregate welfare are all lowest in the super-prime market because the estimated gains from trade are lowest.

5.7. Comparative Statics

We further illustrate the working of our model through two comparative statics that vary the parameters that are the main focus of our framework: (1) the parameter β_{0j} that affects the effort cost of examining offers; and (2) the standard deviation σ_{a_j} of the product attribute a .

We focus on the results of these comparative statics for near-prime borrowers—that is, 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.

Cost of Examining Offers. We decrease the parameters β_{0j} of the cost of effort by 30% while holding all other parameters at their calibrated values.

Figure 2 compares outcomes of the model with lower examination costs (dotted line) with those of the baseline model at the calibrated parameters (solid line). The 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, because 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 right panel explains why lenders' offered rates are lower: Because the cost of effort is lower, borrowers increase their search effort.

Based on these changes in lenders' rates and borrowers' efforts, 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 their probability of accepting any offer increases—but they are relatively less likely to accept high-interest-rate offers. Demand becomes more elastic relative to that of the baseline case. Moreover, because lenders decrease their rates and borrowers accept offers with lower rates

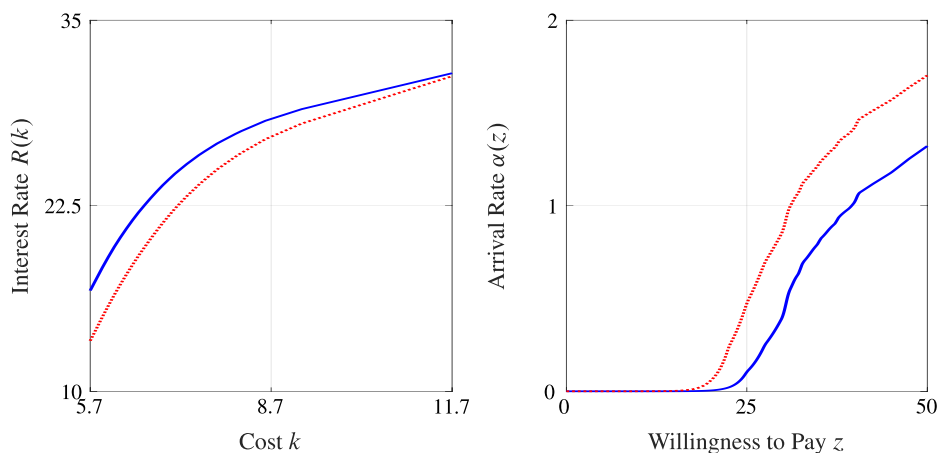


FIGURE 2. These panels display model outcomes at the calibrated parameters (solid line) and in the case when $\beta'_{0j} = 0.7\beta_{0j}$ (dotted line) for the near-prime market. The left panel displays lenders' optimal interest rate $R(k)$ as a function of their cost k ; the right panel displays borrowers' effective arrival rate $\alpha(z)$ as a function of their willingness to pay z .

with a higher probability, the fraction of individuals with credit card debt increases considerably relative to its value in the baseline case—from 54.7% to 70.2%.

The distribution of offered rates and of accepted rates obtained in the model with a lower β_{0j} are first-order stochastically dominated by the corresponding distributions in the baseline model. The reason is that low-cost lenders decrease their offered rates, because borrowers compare more offers if their effort to examine them is less costly. The lower cost of effort affects lower percentiles relatively more than higher percentiles.

These comparative statics help us to understand why the calibrated model calls for relatively large effort costs: If they were smaller, then the level of offered and of accepted interest rates would be lower, and the fraction of borrowers would be higher than those observed in the data.

Product Differentiation. We compare outcomes of the baseline model at the calibrated parameters with those of the model with lower search costs and those of the model when we further increase the standard deviation σ_{a_j} of the product attribute a , while holding all other parameters at their calibrated values. Because σ_{a_j} is calibrated to be small, we increase it by a factor of 30, which makes the value of the interquartile range of a similar to that of R observed in the data.

Figure 3 shows interesting outcomes. Most notably, the left panel shows that the interest-rate function $R(k)$ flattens when product differentiation is more important for borrowers. The reason is that a larger σ_{a_j} means that interest rates affect consumers' choice across lenders relatively less, and thus all lenders charge similar rates.

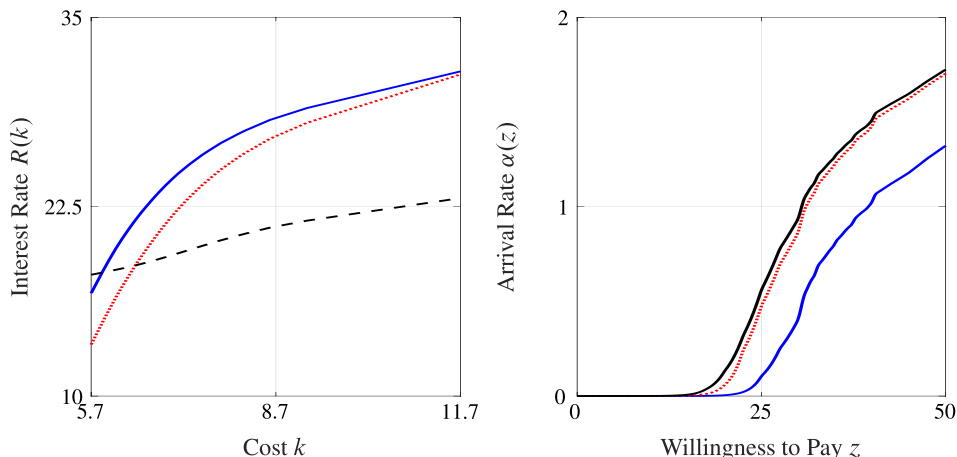


FIGURE 3. These panels display model outcomes at the calibrated parameters (solid line), in the case when $\beta'_{0j} = 0.7\beta_{0j}$ (dotted line) and in the case when $\beta'_{0j} = 0.7\beta_{0j}$ and $\sigma'_{aj} = 30\sigma_{aj}$ (dashed line) for the near-prime market. The left panel displays lenders' optimal interest rate $R(k)$ as a function of their cost k ; the right panel displays borrowers' effective arrival rate $\alpha(z)$ as a function of their willingness to pay z .

The comparison between the dashed and the dotted lines in the right panel shows that a higher σ_{aj} has a small effect on borrowers' search effort. This small change in effort is the result of opposite effects. Holding the distribution of offered rates fixed, the increase in the product-differentiation parameter induces borrowers to search more aggressively, because they are more likely to receive offers with product features a that they value more. However, the dispersion of offered interest rates decreases, which decreases borrowers' incentives to search. As a result of these offsetting effects, borrowers' effort to examine offers changes minimally.

Because lenders offer similar interest rates when σ_{aj} is higher (as the left panel shows), the probability $P(R)$ that the borrowers accept an offer with interest rate R increases relative to the case with identical costs of effort but a lower σ_{aj} . Holding the distribution of offered rates fixed, this increase, cumulated over the range of R , would lead to a non-trivial increase in the fraction of individuals who borrow on credit cards, as we recount in Section 5.5.

The changes displayed in Figure 3 have large effects on the distribution of offered rates and of accepted rates. Both distributions obtained in the model with a higher product differentiation and lower search costs intersect the corresponding distributions obtained in the baseline model, as well as those obtained with lower search costs only. This crossing is intuitive, because offered rates—and thus accepted rates as well—are less dispersed if the product attribute a matters more for consumers' choices. The average offer rate and the average accepted rate decrease relative to the baseline.

However, the most striking effects are on the standard deviation of offered and accepted rates, which decrease approximately by a factor of three from their baseline values.

These comparative statics further help us to understand why the calibrated model disfavors low costs of effort and large values of σ_{a_j} : If costs were low and σ_{a_j} were large, then the dispersion of interest rates would be significantly lower than those observed in the data.

6. Policy Experiments: Caps on Interest Rates

As we recount in the Introduction section, several countries recently introduced price controls in markets for some consumer financial products and are considering intervening in other markets as well. The goal of this section is to study the effects of interest-rate caps on the equilibrium of our model—specifically, how borrowers' examination effort and lenders' offered rates respond, thereby affecting market outcomes and welfare.

The theoretical literature points out that price caps may have undesirable consequences, for two main reasons. First, caps reduce profit margins and thus may reduce the supply of credit, most notably to riskier borrowers who have higher default rates. Second, in frictional markets, price caps reduce price dispersion and thus reduce consumers' incentives to acquire information about prices, thereby increasing suppliers' market power and possibly posted prices. This indirect effect may outweigh the direct effect of reducing the highest market prices, leading to consumers paying higher average prices (Fershtman and Fishman 1994; Armstrong, Vickers, and Zhou 2009). Hence, the relative magnitude of these contrasting effects is an empirical/quantitative question. Our calibrated model allows us to determine which of these opposing effects dominates, and thus whether price caps are beneficial to consumers and increase welfare.

To understand these issues, we study three cases, each with a different cap R_{\max} common across markets: 27.5, 25, and 22.5 pps. These caps are binding in most markets: For example, Table 2 shows that the 25-pp cap corresponds approximately to the 65th, 75th, and 85th percentiles of the distributions of accepted interest rates in the subprime, near-prime, and prime market, respectively. These three cases shed light on the rich equilibrium interactions between lenders and borrowers.

Under the policy, borrowers and lenders optimize as in the baseline model, with the additional constraint that lenders' interest rates satisfy the cap; that is, $R_j(k) \leq R_{\max}$. We study these counterfactual cases in general equilibrium; that is, we require that lenders' free-entry condition (14) holds. Thus, some lenders may exit (or enter) the market, in which case we change the aggregate arrival rate of offers to a new value L'_j proportionally. Formally, the new arrival rate equals $L'_j = \Lambda_j G(\hat{k}')$, where \hat{k}' is the marginal cost of the marginal lender—that is, the lender that satisfies the free-entry condition (14)—in the counterfactual case (the marginal cost of the marginal lender in the baseline case equals \hat{k}).

Figure 4 illustrates the main economic forces of the model by comparing outcomes of the baseline case at the calibrated parameters (solid line) with those of the model

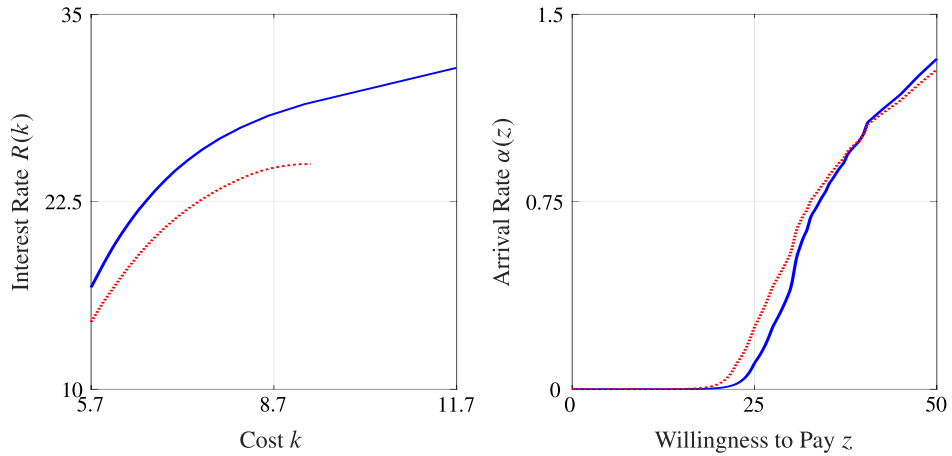


FIGURE 4. These panels display outcomes in the near-prime market at the calibrated parameters (solid line) and in the case when interest rates are capped at 25 pps (dotted line). The left panel displays lenders’ optimal interest rate $R(k)$ as a function of their cost k ; the right panel displays borrowers’ optimal arrival rate $\alpha(z)$ as a function of their willingness to pay z .

with the intermediate cap of $R_{\max} = 25$ pps, while holding all other parameters at their calibrated values for the near-prime market.²⁴ The left panel shows that the highest-cost lenders exit the market, even though 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 enough to allow them to cover their fixed costs; hence, they exit. All surviving lenders charge lower interest rates than in the baseline case. The lender with marginal cost \hat{k}' drops its rate to satisfy the constraint rather than exit. In turn, all other lenders with lower marginal costs charge slightly below their higher-cost competitors.

The right panel shows how borrowers’ effective arrival rate of offers adjusts, reflecting the indirect and direct effects that Fershtman and Fishman (1994) and Armstrong, Vickers, and Zhou (2009) emphasize. Specifically, because some lenders exit the market, on average, borrowers receive fewer offers than in the baseline case. Nevertheless, low-valuation borrowers increase their effort to more than offset the lower arrival rate of offers, and thus, the average effective number of offers $\alpha(z)$ they examine is higher than in the baseline. The reason is that 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

24. In particular, we keep the standard deviations σ_{a_j} of the product attribute a constant. We should point out that (1) reducing these standard deviations does not affect our main counterfactual results, as these standard deviations are very small relative to the standard deviation of R_j ; and (2) in the counterfactuals of Online Appendix D, which considers the case in which R and a are correlated, the standard deviations of the product attribute endogenously adjust when R is capped.

TABLE 6. Market outcomes and welfare with a price cap.

	Subprime	Near-prime	Prime	Super-prime
<i>Panel A: Cap = 27.5 pps</i>				
Average number of offers per borrower	0.97	0.98	1.00	1.00
Average accepted rate	0.92	0.92	1.00	1.00
Standard deviation of accepted rates	0.88	0.92	1.00	1.00
Fraction of borrowers	1.03	1.05	1.00	1.00
Consumer surplus	1.14	1.23	1.00	1.00
Lender profits	0.73	0.75	1.00	1.00
Welfare	1.04	1.08	1.00	1.00
<i>Panel B: Cap = 25 pps</i>				
Average number of offers per borrower	0.90	0.91	0.99	1.00
Average accepted rate	0.86	0.85	0.97	1.00
Standard deviation of accepted rates	0.76	0.80	0.94	1.00
Fraction of borrowers	1.01	1.04	1.01	1.00
Consumer surplus	1.19	1.34	1.06	1.00
Lender profits	0.49	0.52	0.90	1.00
Welfare	1.02	1.08	1.02	1.00
<i>Panel C: Cap = 22.5 pps</i>				
Average number of offers per borrower	0.80	0.82	0.94	1.00
Average accepted rate	0.80	0.78	0.89	1.02
Standard deviation of accepted rates	0.62	0.67	0.81	1.02
Fraction of borrowers	0.94	0.99	1.02	0.98
Consumer surplus	1.15	1.35	1.16	0.97
Lender profits	0.30	0.32	0.63	1.02
Welfare	0.94	1.02	1.03	0.98

Note: This table reports market outcomes and welfare in each market, as ratios of those of the baseline case.

of offers $\alpha(z)$ they consider is lower than in the baseline: 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 examining multiple offers.

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, as ratios of those of the baseline case for the three different cases: Panel A for the 27.5-pp cap, Panel B for the 25-pp cap, and Panel C for the 22.5-pp cap. In markets in which the cap is not binding (e.g. the 27.5-pp cap in the super-prime market), all ratios equal 1 by construction. The table shows several patterns, and we emphasize the most interesting in our opinion.

First, the average and the standard deviation of offered and accepted rates are lower than those of the baseline in almost all cases, suggesting that caps increase the surplus of those who retain access to credit in those cases. The case of the 22.5-pp cap in the super-prime market represents a notable exception, and we describe this case in more detail below.

Second, the number of offers sent out by lenders decreases in almost all cases, again with the exception of the case of the 22.5-pp cap in the super-prime market. Nevertheless, access to credit, as measured by the fraction of borrowers, might increase or decrease modestly relative to the baseline case in response to a binding interest-rate cap. The reason is that borrowers exert higher examination effort, on average, in order to take advantage of the lower offered interest rates, and this additional effort mitigates or, in some cases, fully outweighs the reduction in the number of offers received. Interestingly, the effort choice of marginal borrowers (i.e. those with a low valuation z) display a stronger response to the cap than infra-marginal borrowers (i.e. those with a high valuation z). Overall, we find that interest rate caps of 25 or 27.5 pps weakly increase access to credit for all borrowers, whereas the tighter 22.5-pp cap reduces access to credit for subprime, near-prime, and super-prime borrowers and increases it for prime borrowers.

Third, most caps induce a large redistribution of surplus from lenders to borrowers, again with the exception of the 22.5-pp cap in the super-prime market. Consumer surplus increases and aggregate lender profits decline dramatically in almost all markets with a binding cap relative to the baseline. As a result of the increase in consumer surplus and the decrease in lender profits, aggregate welfare displays smaller changes than its components, and notably increases when the cap equals 27.5 or 25 pps relative to the baseline.

Fourth, and perhaps most surprising, for subprime borrowers, consumer surplus is non-monotonic in the cap and highest with a 25-pp cap. The reason is that, as explained above, a tighter cap reduces both the number of offers per borrower and the average accepted rate; the first effect reduces surplus, while the second one increases it, which result in a non-monotonic composite effect. In turn, aggregate welfare may also be non-monotonic in the cap, since it is slightly higher (by 0.02%) with the 25-pp cap than with the 27.5-pp cap in the near-prime market. Hence, this analysis suggests that a moderate cap can increase both consumer surplus and welfare, and may be preferable to a lower one.

Fifth, the super-prime market with the 22.5-pp cap showcases an example in which the undesirable consequences emphasized by Fershtman and Fishman (1994) and Armstrong, Vickers, and Zhou (2009) dominate. Specifically, the cap is just binding and the resulting reduction in interest-rate dispersion lowers borrowers' search efforts, which increases lenders' market power. Hence, average offered and accepted rates increase, the fraction of borrowers and consumer surplus decrease, and lender profits increase. In this case, there is also a minor increase (by 0.2%) in the entry of the higher-cost lenders than in the baseline case.

Online Appendix D shows that these results are quite similar in the case in which attribute a is correlated with interest rate R , since the data nevertheless seem to reject that the unobserved attribute, whether correlated with R or not, has a large variance.

The results reported in Table 6 are broadly consistent with the empirical findings of Agarwal et al. (2015), who report an increase in consumer surplus and a decrease in lender profits after the 2009 Credit Card Act banned overlimit fees on credit cards. More generally, the aggregate welfare reported in Table 6 assigns equal weights to consumer

surplus and lender profits. Of course, any larger weight assigned to consumer surplus relative to that assigned to lender profits increases the assessments of the benefits of interest-rate caps.

Finally, Online Appendix E further considers counterfactual cases, in which we increase lenders' fixed costs χ_j . The increase in the fixed costs and the case of interest-rate caps share the feature whereby the highest-costs lenders exit the market; hence, the counterfactuals of Online Appendix E allow us to understand how much the results of Table 6 obtain because of the exit of these highest-cost lenders. However, Online Appendix E shows that, in contrast to price caps, higher operating costs unambiguously harm borrowers and decrease aggregate welfare. The reason is that the exit of lenders reduces competition relative to the baseline case, and thus, all surviving lenders increase their interest rates.

7. Conclusion

This paper develops a framework that captures the observed the large number of credit card offers individuals receive and the high level and large dispersion of the interest rates individuals pay on their credit cards. We focus on two main reasons: the endogenous (low) effort of examining offers and product differentiation. We calibrate the model using data on the US credit card market, which fits them well. Our analysis implies that the low effort of examining offers mostly accounts for the observed patterns in the data, whereas product differentiation plays a smaller role. We further use the calibrated model to perform policy experiments. Most notably, we find that interest-rate caps could generate quite large gains in consumer surplus, because they decrease lenders' market power.

We should point out that these results obtain in a model with some limitations, and thus future research could enhance it 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 further enrich our current framework. Specifically, extensive multidimensional heterogeneity is difficult to identify with our data. Many structural search models share this limitation due to similar data constraints, and one contribution of this paper is to adapt and enrich these models to incorporate two key features—consumer limited examination effort and product differentiation—that may rationalize the large number of credit card offers and the large dispersion of interest rates we observe in the data.

For these main reasons, we view this paper as a first step in quantifying the role of effort in examining and evaluating offers 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 retail financial markets.

References

- Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel (2015). "Regulating Consumer Financial Products: Evidence from Credit Cards." *The Quarterly Journal of Economics*, 130, 111–164.
- Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel (2018). "Do Banks Pass Through Credit Expansions to Consumers Who Want to Borrow?" *The Quarterly Journal of Economics*, 133, 129–190.
- Agarwal, Sumit, John C. Driscoll, Xavier Gabaix, and David Laibson (2008). "Learning in the Credit Card Market." NBER Working Paper No. 13822.
- Allen, Jason, Robert Clark, and Jean-François Houde (2019). "Search Frictions and Market Power in Negotiated Price Markets." *Journal of Political Economy*, 127, 1550–1598.
- Armstrong, Mark, John Vickers, and Jidong Zhou (2009). "Consumer Protection and the Incentive to Become Informed." *Journal of the European Economic Association*, 7, 399–410.
- Ausubel, Lawrence M. (1991). "The Failure of Competition in the Credit Card Market." *American Economic Review*, 81(1), 50–81.
- Burdett, Kenneth and Kenneth L. Judd (1983). "Equilibrium Price Dispersion." *Econometrica*, 51, 955–969.
- Butters, Gerard R. (1977). "Equilibrium Distributions of Sales and Advertising Prices." *Review of Economic Studies*, 44, 465–491.
- Calem, Paul S. and Loretta J. Mester (1995). "Consumer Behavior and the Stickiness of Credit-Card Interest Rates." *American Economic Review*, 85(5), 1327–1336.
- Campbell, Graham, Andrew Haughwout, Donghoon Lee, Joelle Scally, and Wilbert van der Klauuw (2016). "Just Released: Recent Developments in Consumer Credit Card Borrowing." Liberty Street Economics (Blog), Federal Reserve Bank of New York.
- Campbell, John Y., Howell E. Jackson, Brigitte C. Madrian, and Peter Tufano (2011a). "Consumer Financial Protection." *Journal of Economic Perspectives*, 25(1), 91–114.
- Campbell, John Y., Howell Edmunds Jackson, Brigitte Madrian, and Peter Tufano (2011b). "The Regulation of Consumer Financial Products: An Introductory Essay with Four Case Studies." In *Moving Forward: The Future of Consumer Credit and Mortgage Finance*, edited by Nicolas Paul Retsinas and Eric S. Belsky. Brookings Institution Press, Chap. 7, pp. 206–244.
- Carlin, Bruce I (2009). "Strategic Price Complexity in Retail Financial Markets." *Journal of Financial Economics*, 91, 278–287.
- Consumer Financial Protection Bureau (2012). "Analysis of Differences between Consumer- and Creditor-Purchased Credit Scores." Technical report, Consumer Financial Protection Bureau.
- Consumer Financial Protection Bureau (2015). "The Consumer Credit Card Market." Technical report, Consumer Financial Protection Bureau.
- Crawford, Gregory S., Nicola Pavanini, and Fabiano Schivardi (2018). "Asymmetric Information and Imperfect Competition in Lending Markets." *American Economic Review*, 108(7), 1659–1701.
- Cuesta, José Ignacio and Alberto Sepúlveda (2019). "Price Regulation in Credit Markets: A Trade-off between Consumer Protection and Credit Access." Working paper, Stanford University.
- Diamond, Peter (1990). "Pairwise Credit in Search Equilibrium." *The Quarterly Journal of Economics*, 105, 285–319.
- Ellison, Glenn and Sara Fisher Ellison (2009). "Search, Obfuscation, and Price Elasticities on the Internet." *Econometrica*, 77, 427–452.
- Ellison, Glenn and Alexander Wolitzky (2012). "A Search Cost Model of Obfuscation." *The RAND Journal of Economics*, 43, 417–441.
- Fershtman, Chaim and Arthur Fishman (1994). "The 'Perverse' Effects of Wage and Price Controls in Search Markets." *European Economic Review*, 38, 1099–1112.
- Galenianos, Manolis and Alessandro Gavazza (2017). "A Structural Model of the Retail Market for Illicit Drugs." *American Economic Review*, 107(3), 858–896.
- Galenianos, Manolis, Tzuo Hann Law, and Jaromir Nosal (2021). "Market Power in Credit Markets." Working paper, Royal Holloway, University of London.
- Gavazza, Alessandro (2016). "An Empirical Equilibrium Model of a Decentralized Asset Market." *Econometrica*, 84, 1755–1798.

- Grodzicki, Daniel (2015). "Competition and Customer Acquisition in the US Credit Card Market." Working paper, Pennsylvania State University.
- Han, Song, Benjamin J. Keys, and Geng Li (2018). "Unsecured Credit Supply, Credit Cycles, and Regulation." *The Review of Financial Studies*, 31, 1184–1217.
- Hong, Han and Matthew Shum (2006). "Using Price Distributions to Estimate Search Costs." *The RAND Journal of Economics*, 37, 257–275.
- Hortaçsu, Ali and Chad 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, Maarten C. W. and José Luis Moraga-González (2004). "Strategic Pricing, Consumer Search and the Number of Firms." *The Review of Economic Studies*, 71, 1089–1118.
- Knittel, Christopher R. and Victor Stango (2003). "Price Ceilings as Focal Points for Tacit Collusion: Evidence from Credit Cards." *American Economic Review*, 93(5), 1703–1729.
- Nelson, Scott (2020). "Private Information and Price Regulation in the US Credit Card Market." Working paper, University of Chicago.
- Ru, Hong and Antoinette Schoar (2016). "Do Credit Card Companies Screen for Behavioral Biases?" NBER Working Paper No. 22360.
- Salz, Tobias (2022). "Intermediation and Competition in Search Markets: An Empirical Case Study." *Journal of Political Economy*, 130, 310–345.
- Shi, Shouyong (1996). "Credit and Money in a Search Model with Divisible Commodities." *The Review of Economic Studies*, 63, 627–652.
- Sirri, Erik R. and Peter Tufano (1998). "Costly Search and Mutual Fund Flows." *The Journal of Finance*, 53, 1589–1622.
- Stango, Victor (2002). "Pricing with Consumer Switching Costs: Evidence from the Credit Card Market." *The Journal of Industrial Economics*, 50, 475–492.
- Stango, Victor and Jonathan 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, Matthijs R. (2011). "An Empirical Model of Search with Vertically Differentiated Products." *The RAND Journal of Economics*, 42, 729–757.
- Woodward, Susan E. and Robert E. Hall (2012). "Diagnosing Consumer Confusion and Sub-Optimal Shopping Effort: Theory and Mortgage-Market Evidence." *American Economic Review*, 102(7), 3249–3276.

Supplementary Data

Supplementary data are available at [JEEA](#) online.