

PRIVATE INFORMATION AND PRICE REGULATION IN THE US CREDIT CARD MARKET

SCOTT T. NELSON

Booth School of Business, University of Chicago

The 2009 CARD Act limited credit card lenders' ability to raise borrowers' interest rates on the basis of new information. Pricing became less responsive to public and private signals of borrowers' risk and demand characteristics, and price dispersion fell by one-third. I estimate the efficiency and distributional effects of this shift toward more pooled pricing. Prices fell for high-risk and price-inelastic consumers, but prices rose elsewhere in the market and newly exceeded willingness to pay for over 30% of the safest subprime borrowers. On net, average traded prices fell and consumer surplus rose at all credit scores. Higher consumer surplus was partly driven by a fall in lender profits, and partly by the Act's insurance value to borrowers who could retain favorable pricing after adverse changes to their default risk. The relatively high level of pre-CARD-Act markups was crucial for realizing these surplus gains.

KEYWORDS: Adverse selection, imperfect competition, price discrimination.

1. INTRODUCTION

LENDERS TYPICALLY LEARN new information about their borrowers over time (Petersen and Rajan (1995)). What are the consequences of restricting how lenders use such information, and what does this reveal about the role of such information in credit markets?

I study these questions in the context of the US credit card market and the Credit Card Accountability Responsibility and Disclosure (CARD) Act of 2009. The CARD Act restricted credit card lenders' ability to discretionarily raise borrowers' interest rates over time and also restricted fees that could otherwise substitute for such interest rate increases, broadly limiting the responsiveness of borrower-specific pricing to new information.

Understanding the effects of the CARD Act price restrictions is important both because of these restrictions' economic interest and because of the credit card market's central role in the supply of US consumer credit. Among the estimated 85 million US households with credit cards, roughly 60% use credit cards for at least occasional borrowing, accessing over \$3 trillion in open credit lines. Reliance on credit card borrowing is

Scott T. Nelson: Scott.Nelson@chicagobooth.edu

This paper is a revised version of the first chapter of my PhD dissertation, and I am indebted to my dissertation committee for their generous help at all stages of this project: Antoinette Schoar, Jonathan Parker, and especially my committee chair Jim Poterba. For thoughtful discussions and suggestions, I thank Nikhil Agarwal, Alex Bartik, Vivek Bhattacharya, Ron Borzekowski, Ken Brevoort, Brian Bucks, Amy Finkelstein, Daniel Green, Daniel Grodzicki, Jacob Hartwig, Luu Nguyen, Richard Schmalensee, David Silberman, Lauren Taylor, Daniel Waldinger, and Mike Whinston, as well as conference discussants Johannes Stroebel and Simcha Barkai. Sarah Griebel, Jacob Hartwig, and Marco Loseto provided invaluable research assistance with revisions of this paper. Financial support from the University of Chicago Booth School of Business is gratefully acknowledged. The views expressed herein are those of the author and do not necessarily reflect those of the Consumer Financial Protection Bureau (CFPB) or the United States. I am also grateful to Shaista Ahmed, Michelle Kambara, Joe Remy, Stefano Sciolli, and Susan Singer for their support at the CFPB, to the Lynde and Harry Bradley Foundation for financial support at MIT, and to the National Science Foundation Graduate Research Fellowship for support under grant number 1122374.

© 2025 The Author. *Econometrica* published by John Wiley & Sons Ltd on behalf of The Econometric Society. Scott T. Nelson is the corresponding author on this paper. This is an open access article under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

especially pronounced for subprime consumers, for whom the share of accounts used for at least occasional borrowing exceeds 95%.¹

In this paper, I quantify the distributional and efficiency consequences of the CARD Act price restrictions. I analyze two channels through which informational restrictions on pricing can influence credit market outcomes. First, if lenders learn over time about borrower demand, these price restrictions may limit lenders' ability to adjust markups personalized to each consumer. Second, such restrictions may also limit lenders' ability to adjust prices for risk, which may exacerbate the effect of information asymmetries and induce either partial or complete market unraveling. Both channels matter in the short run and also dynamically in this setting: consumers face changes to their demand and risk over time, and consumers may value low pricing more in some states than in others. The interplay of these forces may cause interest rates to fall for some consumers and rise for others; total welfare may also either rise or fall.

I study these effects using two large administrative data sets. The first contains monthly account-level data from the near-universe of US credit card accounts. These data have detailed price measures, including both interest rates paid and fees incurred, as well as measures of outstanding consumer debt, new borrowing, and repayment. The second data set is a large, randomly sampled panel of US consumer credit reports. These data reveal patterns not observable in the account data, including which consumers are not credit card holders at a given time.

I first present new facts about how credit card pricing changed with the implementation of the CARD Act. I show that the class of interest rate increases restricted by the Act affected over 50% of borrowing accounts annually prior to the CARD Act, but this rate of incidence dropped to nearly zero once the Act took effect. The elimination of these interest rate increases had immediate effects on the price distribution: as lenders became unable to discretionarily raise some borrowers' interest rates, price dispersion (as measured by the interquartile range of interest rates) on new cohorts of mature accounts dropped immediately by one-third. The *bottom* of the price distribution was also compressed, albeit not immediately: within credit score, the bottom quartile of interest rates rose over time relative to the mean by over 100 basis points for most prime borrowers and by over 200 basis points for subprime borrowers. The credit score segments that saw the greatest increase in the left tail of the price distribution also experienced the greatest rates of consumer exit. This is consistent with (partial) market unraveling as the market shifted toward greater pooling (Akerlof (1970)).

I also describe the dynamic pricing of risk- and demand-relevant information after the CARD Act. Prior to the Act, interest rates were strongly responsive to changes in risk after origination, whereas after the Act, such "emergent" risk became nearly 75% cheaper for a borrower, relative to risk observable at origination. Lenders appear to face higher adverse retention of risky borrowers as a result. I find that the Act also restricted lenders from adjusting interest rates in response to private information about borrower demand, and that lenders' excess returns from privately high-demand borrowers then fell. These two results underscore the importance of both demand- and risk-relevant information in studying the Act's effects.

With these descriptive results in mind, I develop and estimate a structural model of the credit card market as a tool for studying the CARD Act price restrictions. The structural model features consumers who face changes in their risk and demand over time,

¹See Bricker et al. (2017) and this paper's Appendix Table 3.

differentiated lenders who acquire private information about borrowers, and flexible correlation between borrower demand and risk. I estimate the model on the pre-CARD-Act equilibrium observed in the market. I then impose the CARD Act price restrictions in the model and analyze their effects for different types of consumers and for total welfare when the market reequilibrates. Consequently, this exercise quantifies one precise sense of the Act's restrictions' effects: the *ceteris paribus* effects holding consumer characteristics and other features of the pre-CARD-Act environment constant, rather than the effects of these restrictions in conjunction with other contemporaneous changes, such as the Great Recession and coincident regulation that both accompanied the Act.

In estimating the model, I estimate several key parameters related to the workings of the US credit card market that to my knowledge are not available in previous academic work. I use a novel source of quasi-experimental price variation—portfolio-wide repricing of existing accounts—to estimate borrowers' sensitivities to price. Other demand estimates indicate that consumers' setup (or switching) costs for opening new credit card accounts are relatively high, contributing to persistence in lending relationships even when consumers face rising prices over time, and that markups in the pre-CARD-Act credit card market were substantial, for example, exceeding 40 percent for the median-risk subprime consumer.

Imposing the CARD Act price restrictions in the model reveals several interrelated effects. On net, the restrictions cause average traded prices to fall throughout the market and especially on subprime accounts, consistent with the results in Agarwal, Chomsisengphet, Mahoney, and Stroebel (2015b). At the same time, some consumers who previously could access the cheapest credit within their credit score segment tend to face higher prices and exit from borrowing. This type of partial unraveling is especially pronounced among subprime consumers, and I estimate that prices newly exceed willingness to pay for over 30% of borrowers. Nonetheless, given the effect of lower prices for consumers with the strongest demand for credit, consumer surplus rises throughout the market.

This rise in consumer surplus is larger than the fall in lender profits. Besides this transfer from lenders, another source of consumer surplus gain is the insurance value of the CARD Act restrictions for consumers whose credit scores deteriorate over time. This insurance is most relevant for superprime borrowers.² However, the Act's insurance value also affects the interpretation of surplus gains among subprime borrowers. The subprime borrowers who benefit most are those whose credit score has recently fallen below prime, since these restrictions allow them to retain favorable pricing from loans originated at prime scores. In contrast, subprime borrowers looking to open a new credit card—for example, a young borrower or a long-time subprime consumer—feel the effects of market unraveling more severely.

This paper makes several contributions relative to existing literature. In a seminal paper, Agarwal et al. (2015b) also study how the CARD Act affected credit card pricing, finding through a difference-in-differences strategy that the Act reduced the average, fee-inclusive price of credit card borrowing. They also estimate the effects of several nonprice provisions of the Act not studied here, such as the Act's nudges for consumers to repay balances more quickly. I complement their analysis by examining which consumers benefited from CARD-Act-induced price decreases, and which consumers may have instead exited the market as they were pooled with their peers; I also translate these price changes

²I refer to FICO scores below 660 as subprime, FICO from 660 to 720 as prime, FICO scores of 720 and above as superprime. If there appears no risk of confusion, I also at times use prime as an antonym for subprime.

and exit patterns into estimates of consumer and total surplus gains. Furthermore, I highlight the importance of reducing market power from private information rents, and of the insurance value in the Act's restrictions, as countervailing forces that made it possible for the Act to increase surplus despite the Act making it more difficult for lenders to price risk. Finally, I show these surplus gains would be mitigated or reversed were it not for the relatively high markups in pre-CARD-Act pricing.

Other research on the CARD Act includes [Keys and Wang \(2019\)](#), who study the Act's nudges for borrowers to pay more than their minimum required payment each month, [Jambulapati and Stavins \(2014\)](#), who describe patterns of account closures and credit line changes coinciding with the Act and the Great Recession, [Debbaut, Ghent, and Kudryak \(2016\)](#), who focus on the Act's particular restrictions to protect young borrowers, and [Han, Keys, and Li \(2017\)](#), who compare direct-mail offers for credit cards with those for other financial products before and after the CARD Act to conclude, consistent with my results on partial market unraveling among subprime accounts, that the Act partially curtailed supply among subprime credit cards. There is also a growing body of theoretical work focused on the CARD Act price restrictions in particular: [Hong, Hunt, and Serfes \(2023\)](#) and [Pinheiro et al. \(2016\)](#) examine pricing and welfare effects of repricing restrictions in a perfectly competitive setting, whereas lender market power plays a central role in my study.

This paper also joins a long literature examining the competitiveness of, and sources of market power in, the credit card industry. After [Ausubel \(1991\)](#) showed limited pass-through of lenders' cost of funds to borrowers, many have explored whether and why the industry may be imperfectly competitive, including for reasons of search costs ([Berlin and Mester \(2004\)](#), [Galenianos and Gavazza \(2022\)](#)), consumer irrationality ([Brito and Hartley \(1995\)](#)),³ adverse selection ([Stavins \(1996\)](#)), and lender concentration ([Herkenhoff and Raveendranathan \(2024\)](#)). My work integrates many of these potential sources of market power in a single model—including switching costs across firms, adverse selection, and lender private information—and provides an estimation framework that helps identify their relative importance. My results on the particular importance of switching costs across firms join a growing recent literature on the role of switching costs in selection markets, including [Handel \(2013\)](#) and [Illanes \(2017\)](#).

I also provide new evidence on consumer demand for credit card borrowing and how consumers respond to changes in their credit terms. To date, much of the research on this front has focused on how spending or borrowing responds to changes in credit limits ([Gross and Souleles \(2002\)](#), [Agarwal, Chomsisengphet, Mahoney, and Stroebel \(2018\)](#), [Gross, Notowidigdo, and Wang \(2020\)](#), [Fulford \(2015\)](#)). Research on how borrowers respond to interest rates and fees has been more limited.⁴ To help fill this gap, I estimate

³Research on behavioral consumers in the credit card market has remained quite active, including work by [Heidhues and Koszegi \(2010\)](#), [Meier and Sprenger \(2010\)](#), [Heidhues and Kőszegi \(2015\)](#), [Ru and Schoar \(2016\)](#), and [Kuchler and Pagel \(2021\)](#). Related work focuses on how consumers learn over time how to avoid apparent mistakes with credit cards ([Agarwal, Driscoll, Gabaix, and Laibson \(2008, 2009\)](#)) and how the probability of mistakes also falls as consumers face higher stakes, for example, higher balances borrowed ([Agarwal, Chomsisengphet, Liu, and Souleles \(2015a\)](#)). However, for some contrasting evidence on this point, see [Gathergood, Mahoney, Stewart, and Weber \(2019\)](#).

⁴The available evidence does find a nontrivial elasticity of borrowing with respect to interest rates, although this evidence tends to use price variation generated either by (1) the prescheduled expiration of promotional interest rates ([Gross and Souleles \(2002\)](#)), which may predominantly affect a particularly price-sensitive subset of borrowers who serially shop for promotional rates, or (2) within-account interest rate changes over time ([Grodzicki, Alexandrov, Bedre-Defolie, and Koulayev \(2023\)](#)), which, as I detail in Section 3.2, can arise endogenously as lenders respond to shifts in individual borrowers' risk or demand.

borrower price elasticities across a range of borrower risk types, and I also estimate primitives of a rich demand model—including switching costs, liquidity costs, and disutility from price—that predict how price elasticities change nonlocally as pricing changes. Estimates of these primitives help not only for understanding the CARD Act price restrictions, but also for future research in consumer credit markets.⁵

This paper is organized as follows. In Section 2, I provide background on the credit card market, the CARD Act and the two data sets that I analyze. In Section 3, I report descriptive analyses of how lenders used CARD-Act-restricted repricing prior to the Act and how the market responded to the implementation of the Act. I develop and estimate my model of the credit card market in Section 4. Section 5 presents results using the model to study how the CARD Act's pricing restrictions affect prices, borrowing and welfare in equilibrium. Section 6 concludes.

2. BACKGROUND AND DATA

The US credit card market is an important source of credit for many households (Bricker et al. (2017)). Before the CARD Act, the price of this credit depended largely on two types of information: consumer credit bureaus and credit scores; and cardholder behavior after origination. The former is public information in that it is available to all lenders. Much of the latter is private to a lender and not observable to competitors, including a consumer's shopping and borrowing behavior. Lenders' pricing of individual accounts based on this private information is also typically not observable to competitors.

The CARD Act substantially limited lenders' ability to price this private information, through restrictions both on interest rate increases and on behavior-contingent fees, such as over-limit fees or late fees. Further institutional details are discussed in the Supplemental Appendix's section A.2 (Nelson (2025)).

I use two main data sets in my analysis. Both are anonymized, administrative data sets furnished by industry and maintained by the Consumer Financial Protection Bureau (CFPB). One data set is the CFPB's Credit Card Database (CCDB), a near-universe of deidentified credit card account data in a monthly panel from 2008 to present (Consumer Financial Protection Bureau (2013b)). The data include all open credit card accounts held by 17 to 19 large and midsize credit card issuers (lenders), which together cover roughly 90% of outstanding general-purpose US credit card balances. For each account in each month, the data show totals of all aggregate quantities that would appear on a monthly account statement, including total purchases in dollars, amount borrowed and repaid, interest charges and fees, payment due dates, and delinquencies. These data represent a modest superset of the credit card data used in Agarwal et al. (2015b) and Agarwal et al. (2018), including 9 to 10 additional midsize issuers that cover an additional 17% to 23% of outstanding balances.

The second data set is the CFPB's Consumer Credit Panel (CCP), a large, randomly sampled, anonymized panel of consumer credit reports drawn from one of the three

⁵Beyond the credit card market, this paper contributes to a growing literature using tools from industrial organization to study consumer financial markets (Allen, Clark, and Houde (2014, 2019), Benetton (2021), Bhattacharya, Illanes, and Padi (2025), Buchak, Matvos, Piskorski, and Seru (2022), Chatterjee, Corbae, Dempsey, and Rios-Rull (2023), Clark, Houde, and Kastl (2021), Cuesta and Sepulveda (2021), Dempsey and Faria-e-Castro (2021), Egan, Hortaçsu, and Matvos (2017), Robles-Garcia (2020), Galenianos and Gavazza (2022), Jiang (2023)) and the welfare implications of personalized pricing more broadly (Grunewald, Lanning, Low, and Salz (2023), Dubé and Misra (2023), Buchholz, Doval, Kastl, Matějka, and Salz (2024), Rhodes and Zhou (2024)).

nationwide consumer credit reporting agencies (Consumer Financial Protection Bureau (2013a)). The CCP makes it possible to study borrower entry and exit from credit-card holding. Neither accounts nor account-holders can be linked between the CCDB and CCP. Appendix A.3 provides additional detail on both data sets.

3. DESCRIPTIVE EVIDENCE

This section presents descriptive evidence on credit card market changes after the CARD Act.

3.1. Price Changes and Price Dispersion

In Figure 1, panel (a), I show the prevalence over time of the type of discretionary interest rate increases restricted by the Act (these restrictions are discussed in detail in Appendix A.2). Forty-eight to 54% of borrowing accounts experienced such a discretionary interest rate increase at least once a year before the CARD Act. The prevalence of interest rate increases then dropped sharply to nearly zero after the Act.

In panel (b) of the figure, I document that this decrease in the prevalence of interest rate increases coincided with an immediate compression in the dispersion of interest rates across accounts. In this figure, I show the interquartile range (IQR) of interest rates after controlling for origination FICO score, with one data point presented for each quarterly origination cohort. For cohorts reaching maturity before the Act's repricing restrictions went into effect, these IQRs are consistently about 7.5 percentage points; for cohorts reaching maturity after these restrictions took effect, these IQRs fell sharply to around 5 percentage points.

In Appendix Table 1, I present further evidence on which percentiles of the price distribution compressed with the implementation of the Act. As discussed in Appendix A.3.3, I show increases of several hundred basis points for the lower percentiles of the price distribution and substantial consumer exit in the same market segments that saw increases in the low-cost tail of the price distribution. While these patterns are only suggestive, they are consistent with partial market unraveling as the market transitioned to more pooled pricing.

I further discuss pricing before and after the Act in Appendix Sections A.2 and A.3.3.

3.2. Pricing of Public Information: Risk and FICO Scores


This subsection studies the pricing of public information about default risk, as captured by FICO scores, before and after the CARD Act. I compare two types of risk information treated differently under the Act: risk observable at the time of origination, which I term origination risk; and risk that becomes observable after origination, which I term emergent risk. The Act restricted lenders' ability to adjust pricing in response to emergent risk but not origination risk, so comparing the relative pricing of these two types of risk before and after the Act provides one indication of how risk-pricing changed with the Act.

To study the pricing of origination risk, I estimate a linear relationship between interest rates (APRs) $r_{i,0}$ and FICO scores at origination, $\text{FICO}_{i,0}$,

$$r_{i,0} = a + b\text{FICO}_{i,0} + e_{i,0}. \quad (3.1)$$

I estimate the pricing of emergent risk using a similar model, where I estimate the relationship between interest rates and *change* in FICO since origination,

$$r_{i,t} = \alpha_{\tau_{i,t}} + \alpha_{\text{FICO}_{i,0}, \tau_{i,0}} + \beta(\text{FICO}_{i,t} - \text{FICO}_{i,0}) + \epsilon_{i,t}. \quad (3.2)$$



(a) Interest Rate Repricing



(b) Interquartile Ranges in Credit Card Interest Rates by Vintage

FIGURE 1.—Repricing probabilities and price dispersion. *Notes:* Panel (a) shows the incidence of interest rate increases on current borrowers over 1-month, 6-month, and 12-month horizons, excluding most interest rate increases permitted by the CARD Act (i.e., increases coinciding with the expiration of a promotional rate, with changes in an index rate, or with delinquencies of 60 days or more). Dotted lines extrapolate from the most recent available datapoint when these horizons overlap with the implementation of the CARD Act’s repricing restrictions. Panel (b) shows the interquartile range (IQR) of annual percentage rates on borrowing accounts by origination cohort, after partialling out origination credit score and origination month. The date shown for each cohort is at an age of 18 months, by which point introductory promotional rates have typically expired. Credit score controls are 20-point bins, and the sample is restricted to include only accounts in the same credit score bin at the date observed as at origination. The vertical black lines show CARD Act’s repricing restrictions’ implementation date in February 2010.

This regression also includes fixed effects α for origination credit score $FICO_{i,0}$ interacted with origination cohorts $\tau_{i,0}$, which are included to absorb variation in interest rates $r_{i,0}$ from the time of origination, as well as fixed effects for account age $\tau_{i,t}$, which absorb average changes in interest rates over the life of an account due to, for example, promotional

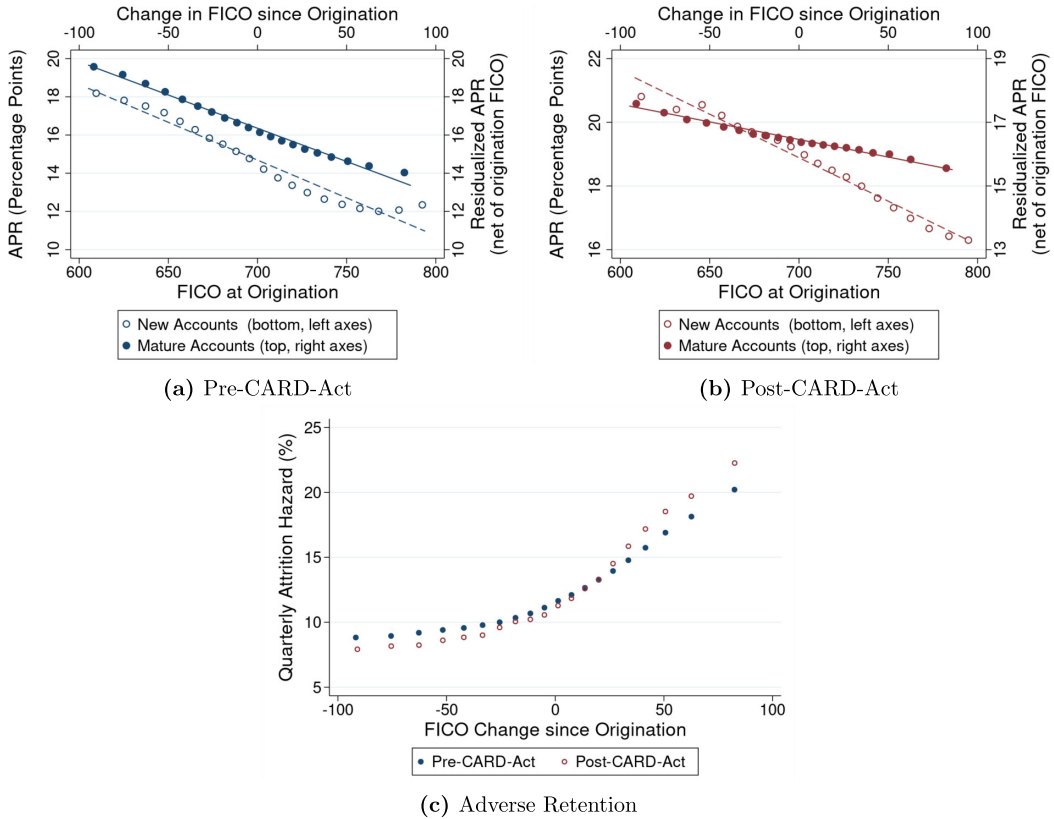


FIGURE 2.—Pricing of origination risk and emergent risk. *Notes:* Panel (a) shows two gradients of risk in the pre-CARD-Act sample (2008Q3 to 2009Q2) on two pairs of axes. On the left and bottom axes, I plot the average annual percentage rate (APR) on newly originated accounts across quantiles of the credit score distribution, together with a line of best fit. On the right and top axes, the figure plots the average current APR on mature accounts across quantiles of those accounts’ change in credit score since origination, after partialling out origination credit score, together with a line of best fit. See equations (3.1) and (3.2) in the text. Panel (b) presents the same price-risk gradients as in panel (a) but in post-CARD-Act data (2011Q3 to 2014Q2). The two y-axes have the same axis scale, but the axis ranges are shifted to facilitate comparison of the two gradients. Panel (c) presents quarterly attrition rates from borrowing (including both account closure and debt repayment) across quantiles of borrowing accounts’ changes in FICO score since origination, separately in pre-CARD-Act data and post-CARD-Act data. See equation (3.3) in the text.

rates expiring over time. Appendix A.4 further discusses the role of these fixed effects and shows robustness to alternatives, including a first-differences specification. Given these fixed effects, the estimated coefficient β then shows the correlation between changes in FICO score since origination and changes in (average) interest rate since origination.

In Figure 2, panel (a), I show the estimates of the coefficients b and β from specifications (3.1) and (3.2) in pre-CARD-Act data, with accompanying binned scatterplots. The price gradient of origination risk (b) is plotted against the bottom-left axes; the price gradient of emergent risk (β) is plotted with the same scaling against the upper-right axes. The two gradients are nearly indistinguishable: for both origination risk and emergent risk, borrowers on average face a difference in price of about 30 basis points in annualized interest for every 10 FICO-point difference in risk. Thus, in the figure, I show how

the credit card market set a consistent price of risk, on average, in the pre-CARD-Act data, regardless of whether the risk was evident at origination or emergent later.

In Figure 2, panel (b), I repeat the same analysis in post-CARD-Act data. Here, there is a divergence between the two gradients: whereas origination risk is priced at 26 basis points annualized per 10 points of FICO score difference, lenders are only able to price risk that emerges after origination at less than a third of that rate, at 7 basis points per 10 FICO points.

The gap between these gradients weakens incentives for newly risky borrowers to attrite from borrowing and incentivizes newly safe borrowers to attrite. I look for evidence of such adverse retention by estimating the relationship between borrower retention and changes in FICO score since origination, using a specification similar to equation (3.2),

$$A_{i,t} = \alpha_{\tau_{i,t}} + \alpha_{\text{FICO}_{i,0}, \tau_{i,0}} + \beta(\text{FICO}_{i,t} - \text{FICO}_{i,0}) + \eta_{i,t}, \quad (3.3)$$

where $A_{i,t}$ is an indicator for attrition from borrowing, and, as in equation (3.2), the fixed effects α control for age $\tau_{i,t}$ since origination and FICO score at origination. The coefficient β therefore captures how quarterly linear-probability hazards from borrowing to nonborrowing change as a function of FICO score differences since origination.

I estimate this attrition model separately in the pre-CARD-Act and post-CARD-Act data and show corresponding binned scatterplots in Figure 2, panel (c). The gap between the two plotted relationships shows the difference between attrition hazards at each credit score. Borrowers who become safer (riskier) over time become more (less) likely to attrite from borrowing after the Act relative to before; I estimate that for every 100 basis points by which emergent risk is priced below origination risk, the quarterly hazard of attrition from borrowing falls by 0.7 percentage points. Appendix A.4 includes further discussion of these results, alternative specifications, and evidence on heterogeneity across accounts.

3.3. Borrower Behavior and Private Information Rents

I now consider the pricing of lenders' private information. I focus on two account behaviors that are privately observed because they are typically not reported to credit bureaus: delinquencies of less than 30 days, and transactions exceeding an account's credit limit. I show that in the pre-CARD-Act period, lenders increased prices on accounts exhibiting these behaviors and then earned higher than usual returns on these accounts, suggesting that these behaviors revealed low price sensitivities, enabling higher markups and generating information rents for lenders.

In Table I, I illustrate these patterns. Columns group accounts into four categories: (1) accounts with "baseline" behavior, which I define as accounts with no delinquency, no over-limit transactions, and no credit score decline of 30 or more points in the prior quarter (results are not sensitive to using other similar thresholds); (2) nondelinquent accounts with an over-limit transaction; (3) accounts that are less than 30 days delinquent and not over-limit; and (4) all other accounts, the vast majority of which are 30 or more days delinquent. Across rows, I then show realized returns, revenues, default losses, and price changes following the observed behavior.

In column (1) of the table, I show that 12-month returns on prime, borrowing accounts average 9.80% annualized in pre-CARD-Act data, and that a typical account sees little to no change in its fee-inclusive price in a month where it exhibits baseline behavior: the 25th, 50th, and 75th percentiles of price changes are -0.5 , 0 , and 0.6 percentage points annualized. In contrast, in columns (2) and (3), I show that following an over-limit transaction or a less-than-30-day-late delinquency, typical price changes can be 10 or more

TABLE I
REPRICING AND PRIVATE INFORMATION RENTS.

Account Behavior:	Baseline	Over-Limit	<30 Days Late	Other
<i>Panel A: Prime Accounts</i>				
Share of Accounts	0.85	0.02	0.07	0.06
12-Month Expected Performance				
Returns (p.p. ann.)	9.80	10.47	10.67	-2.58
Average Monthly Balance (\$)	5141	6133	4989	6024
Revenue Yield	0.13	0.20	0.21	0.14
Charge-off Rate	0.04	0.09	0.10	0.16
Next-month Fee-Inclusive Price Change				
10th Percentile	-1.7	-0.6	2.3	.
25th Percentile	-0.5	1	6.8	.
50th Percentile	0	6.9	15.5	.
75th Percentile	0.6	20.1	35.1	.
90th Percentile	2.1	51.7	88	.
<i>Panel B: Subprime Accounts</i>				
Share of Accounts	0.51	0.10	0.09	0.31
12-Month Expected Performance				
Returns (p.p. ann.)	8.45	8.56	-1.12	-47.66
Average Monthly Balance (\$)	3039	2343	3547	3800
Revenue Yield	0.21	0.30	0.25	0.15
Charge-off Rate	0.13	0.21	0.27	0.63
Next-month Fee-Inclusive Price Change				
10th Percentile	-2.2	-1.3	4.1	.
25th Percentile	-0.9	0.5	10.8	.
50th Percentile	0	12.3	26.2	.
75th Percentile	0.9	49.4	73.1	.
90th Percentile	2.3	95	139.7	.

Note: The table summarizes price changes, revenues, costs, and lender returns for revolving credit card accounts that exhibit various behaviors in the pre-CARD-Act period. The first column shows accounts with “baseline” behavior, defined as no delinquency, no over-limit transactions, and no credit score change of more than 30 points relative to the prior month. The second and third columns show accounts with over-limit transactions and delinquencies of less than 30 days. The fourth column shows all accounts not included in the first three. The revenue yield is revenue as a share of balances, the charge-off rate is default losses as a share of balances, and lender returns are the difference of these. The fee-inclusive price is defined in Appendix Section A.2. Price changes are not shown in the “other” column due to the lack of a clear benchmark period for some of the “other” account behaviors (e.g., a multimonth delinquency spell).

percentage points annualized. This revenue increase is not transitory: over the next 12 months, the revenue yield (i.e., monthly revenue as a share of monthly balance) is over 50% higher on these accounts; returns also rise.

These higher returns are in spite of these accounts also presenting higher risk: as I show in the table, expected charge-offs also increase substantially after observing these behaviors. Hence, these behaviors may be understood as revealing both higher expected costs and, after price increases, higher expected revenue. Subprime accounts exhibit similar patterns as prime accounts, though delinquencies of less than 30 days predict lower rather than higher returns for subprime accounts, despite 4 percentage-point increase in revenue yield over the next 12 months.

4. A MODEL OF THE CREDIT CARD MARKET

In this section, I develop and estimate a model of the credit card market. Differentiated credit card lenders compete à la Bertrand to lend to consumers of various public and

hidden types. Consumers' types change over time, determine credit demand and risk, and generate adverse selection as lenders compete for new borrowers. Consistent with evidence in the prior section, lenders learn new information about each consumer's risk and demand over time, and in the pre-CARD-Act regime, respond to this information by changing pricing.

4.1. Model Exposition

4.1.1. The Consumer's Problem

The consumer solves a dynamic discrete choice problem (see, e.g., Rust (1994), Aguirregabiria and Mira (2010)). Each period, a consumer either chooses to hold no credit card at all, denoted $l' = 0$, or chooses a credit card lender $l' \in L$. If she chooses a credit card, she also chooses whether or not to borrow on it, $b' \in \{\text{debt, no debt}\}$. Denote $a = (l', b')$.⁶

The consumer's state $s = (\theta, l, b)$ in each period consists of her type θ (described more below), her borrowing status b at the end of the prior period, and the lender l from which she held a credit card at the end of the prior period. The consumer's problem is dynamic both because she faces adjustment costs in changing l or b , as described below, and because her type θ changes over time. These model features are motivated by the persistence of borrower behavior and the dynamics in borrower characteristics (e.g., risk) documented in Section 3.

Flow utility from choice a is $u(a, s) + \epsilon(a)$. In Table II, I summarize both the model primitives and the prices that determine $u(a, s)$. First, card quality $\xi_{\theta l'}$ determines type- θ 's utility of choosing a credit card from lender l' ; this captures a consumer's taste for that card's brand, card features such as cash-back rewards, and unobserved product quality more generally (Berry (1994), Berry, Levinsohn, and Pakes (1995)). Borrowing on a card yields additional utility $\delta_{\theta l'}$, reflecting a consumer's demand for credit. Utility from the outside good is normalized to zero.

Consumers face two adjustment costs. As I show in Table II, if a consumer borrowed on a card from lender l in the preceding period but does not borrow on any card this period, she incurs a liquidity cost $\lambda_{\theta l}$ to pay off her past-period debt. Second, if a consumer switches to holding a credit card from a different lender, such that $l' \neq l$, she incurs an account setup cost $\kappa_{\theta l'}$.

In Table II, I also show the prices $p_{tl}(\theta)$ consumers pay to borrow. In the first period of holding a credit card from lender l' (i.e., whenever $l' \neq l$), the price of borrowing is $p_{0l'}(\theta)$. These $t = 0$ "teaser" prices reflect the prevalence of new-account promotional rates, and as will be discussed in Section 4.1.3, the more limited information a lender may have about borrowers on newly opened accounts. In all subsequent periods that the account remains open, the price of borrowing is $p_{1l'}(\theta)$. Hence, the time subscript t on the pricing function can be understood as an indicator for whether $l = l'$. A price coefficient γ_{θ} determines disutility from these prices.⁷

⁶These two modeling decisions—that consumers single-home over lenders and choose extensive rather than intensive-margin borrowing—are primarily made for sake of tractability; similar modeling choices appear in Crawford, Pavanini, and Schivardi (2018), Einav, Finkelstein, and Schrimpf (2010), and Einav, Finkelstein, and Cullen (2010). In Appendix A.8, I discuss evidence for why these modeling choices do not depart much from realism in the credit card market.

⁷I use the fee-inclusive price defined in Section 3.1 when I calculate these prices empirically, so these prices are one-dimensional and are the appropriate marginal prices to use when modeling the extensive margin of borrowing. These prices exclude any annual fees that credit card lenders in practice may charge for all cardholders regardless of borrowing choice. The model treats the card-quality parameter $\xi_{\theta l'}$ as the benefit of card-holding (e.g., airline points) net of any annual fee charged for just holding the card.

TABLE II
CONSUMERS' ONE-PERIOD PAYOFFS BY STATE.

	State s From Prior Period:	
	Borrower on Card From Lender l	Nonborrower With Lender l or Noncardholder ^{††}
<i>Choice this Period:</i>		
Same lender l :		
Borrower	$\delta_{\theta l} + \xi_{\theta l} - \gamma_{\theta} p_{1l}(\theta)$	$\delta_{\theta l} + \xi_{\theta l} - \gamma_{\theta} p_{1l}(\theta)$
Nonborrower	$\xi_{\theta l} - \lambda_{\theta l}$	$\xi_{\theta l}$
New lender l' :		
Borrower	$\delta_{\theta l'} + \xi_{\theta l'} - \kappa_{\theta l'} - \gamma_{\theta} p_{0l'}(\theta)$	$\delta_{\theta l'} + \xi_{\theta l'} - \kappa_{\theta l'} - \gamma_{\theta} p_{0l'}(\theta)$
Nonborrower	$\xi_{\theta l'} - \kappa_{\theta l'} - \lambda_{\theta l}$	$\xi_{\theta l'} - \kappa_{\theta l'}$
No card from any lender	$-\lambda_{\theta l}$	0

Note: The table shows one-period deterministic utilities $u(a, s)$ depending on the consumer's state s at the start of the period (by column) and the consumer's choice a (by row). The parameters shown include the flow utility from borrowing, $\delta_{\theta l}$, the flow utility from holding a credit card without borrowing, $\xi_{\theta l}$, disutility from price (marginal utilities of income), γ_{θ} , and two adjustment costs, including setup costs for opening new accounts, $\kappa_{\theta l'}$, and liquidity costs for paying off existing balances, $\lambda_{\theta l}$. The subscripts l and l' can refer to any lender in the set of lenders L with $l \neq l'$. Subscripts θ refer to consumer types. (††) Noncardholders at the start of the period cannot make a "same lender" choice in the first two rows of the table.

Consumers face two stochastic outcomes at the end of the period. First, if a consumer chose to borrow on her card, she defaults with probability $\chi(\theta)$. When default occurs, a consumer has her credit card account closed and she holds the outside good ($l = 0, b = \text{no debt}$) at the start of the next period. Default occurs after all flow utilities are realized in the period. Second, each consumer draws a new type $\theta' \sim F(\cdot|\theta)$ for the next period. Type transitions occur independently of default, consumer choices, and taste shocks.

Let $V(s)$ be the expected continuation value when starting a period in state s . The consumer's problem each period can then be written as

$$\max_a u(a, s) + \epsilon(a) + \beta \mathbb{E}_{\theta} [V(\tilde{s})|a, s]. \tag{4.1}$$

Here, the deterministic utility $u(a, s)$ is as described above, and the expectation \mathbb{E}_{θ} over the next-period state \tilde{s} reflects the default risk $\chi(\theta)$ and type updating under $\theta' \sim F(\cdot|\theta)$. Suppose the taste shocks $\epsilon(a)$ that capture other unobservable determinants of consumer choice are i.i.d. Gumbel-distributed. Then V can be written recursively as

$$V(s) = \log \left(\sum_a \exp(u(a, s) + \beta \mathbb{E}_{\theta} [V(\tilde{s})|a, s]) \right). \tag{4.2}$$

As an example of how $u(a, s)$ can take a particular form based on Table II, consider a consumer who starts the period with a credit card from lender l on which she borrowed last period. Her payoff from borrowing again from lender l this period is then $\delta_{\theta l} + \xi_{\theta l} - \gamma_{\theta} p_{1l}(\theta) + \beta \mathbb{E}_{\theta} [V(\tilde{s})|a, s] + \epsilon_a$ where the terms $\delta_{\theta l} + \xi_{\theta l} - \gamma_{\theta} p_{1l}(\theta)$ are taken from the upper-left-most cell of Table II. Writing choice probabilities in the usual logit form, this means her probability of borrowing from l is

$$\frac{\exp(\delta_{\theta l} + \xi_{\theta l} - \gamma_{\theta} p_{1l}(\theta) + \beta \mathbb{E}_{\theta} [V(\tilde{s})|a, s])}{\sum_{\hat{a}} \exp(u(\hat{a}, s) + \beta \mathbb{E}_{\theta} [V(\tilde{s})|\hat{a}, s])}. \tag{4.3}$$

Further examples using Table II are collected in Appendix A.5.

Expressions (4.1) through (4.3) capture several intertemporal trade-offs faced by the consumer. Because adjustment costs make consumer-lender relationships and borrowing behavior somewhat persistent, a consumer chooses a lender not just based on current-period pricing and preferences, but also based on lenders’ anticipated future pricing (e.g., changes from teaser prices $p_{0l'}(\theta)$ to mature-account prices $p_{1l'}(\theta)$). Also, because default is costly—in the sense that it necessitates paying an account setup cost $s_{0l'}$ to reenter the credit card market later—consumers may on the margin choose to avoid borrowing when they are a high-risk type, in order to avoid default.

The role of default may be clarified by decomposing the expectation \mathbb{E}_θ . Let $T_{\theta\theta'}$ denote the Markov transition matrix across types implied by $F(\cdot|\theta)$. Then for a consumer who chooses to borrow from l' (i.e., $a = (l', \text{debt})$), expected continuation values can be written as

$$\mathbb{E}_\theta[V(\tilde{s})|a, s] = \underbrace{(1 - \chi(\theta))T_{\theta\theta'}(\theta)V(\theta', l', \text{debt})}_{\text{no default}} + \underbrace{\chi(\theta)T_{\theta\theta'}(\theta)V(\theta', 0, \text{no debt})}_{\text{default}}, \quad (4.4)$$

where I have expanded the components of $s = (\theta, l, b)$ on the right-hand side for clarity, and where the θ argument in $T_{\theta\theta'}(\theta)$ selects the relevant row of the matrix $T_{\theta\theta'}$. In contrast, for consumers who do not choose to borrow (i.e., $a = (l', \text{no debt})$), the expectation \mathbb{E}_θ does not depend directly on default rates and takes the form

$$\mathbb{E}_\theta[V(\tilde{s})|a, s] = T_{\theta\theta'}(\theta)V(\theta', l', \text{no debt}). \quad (4.5)$$

To summarize, the primitives of the consumer’s problem are the parameters in $u(a, s)$, that is, $\{\delta_{\theta l}, \xi_{\theta l}, \kappa_{\theta l}, \lambda_{\theta l}, \gamma_{\theta}\}_{(\theta, l) \in \Theta \times L}$, as described in Table II, along with transition probabilities $T_{\theta\theta'}$ across types and the mapping $\chi(\cdot)$ from types into default probabilities.

4.1.2. Consumer Types

To reflect the two broad types of information discussed in Section 2, I allow types θ to have two dimensions, one private component $\psi \in \Psi$ and one “public” component $x \in X$. The latter is public in the sense that it is observable to all firms in the market; it is best thought of as a credit score, which is designed to be a composite of all available public information. The private type is known to the consumer, but not to the lender at the time a new account is opened.

Two assumptions will prove useful in recovering private information types ψ from the data. The first is that borrower default rates depend only on types, and in particular do not depend on prices. I refer to this as price-invariance of default:

$$\chi = \chi(\theta) \quad \forall l, p_{0l}, p_{1l}. \quad (4.6)$$

Several pieces of evidence support this being a reasonable assumption in the credit card market. First, there is direct evidence that price changes have little to no effect on default rates;⁸ second, the effect of a change in credit card pricing on a typical consumer’s overall debt service is arguably negligible, suggesting that changes in credit card pricing

⁸Using the same price variation highlighted below in Section 4.2.2, Appendix Section A.7 shows that the effect of a 100 basis point increase in interest rates on default rates is statistically indistinguishable from zero, and I can reject resultant increases in default rates of more than 0.04 percentage points. This null result is supported by similar findings in the randomized controlled trial of Castellanos, Hernández, Mahajan, and Seira (2025).

are unlikely to affect solvency;⁹ third, related research in consumer finance suggests the channel through which credit card price changes could affect default rates is limited, as default is often driven by short-run liquidity rather than the long-run value of a loan contract for a borrower (Bhutta, Dokko, and Shan (2017), Guiso, Sapienza, and Zingales (2013), Ganong and Noel (2020), Indarte (2023)). This assumption also follows on other research that has used structural models of selection markets without moral hazard, for example, Cohen and Einav (2007) and Einav, Finkelstein, and Schrimpf (2010).

Given this assumption, it is without loss of generality to order private types ψ by the default rates they induce. Essentially, private types become an index of residual default risk, similar to how residuals have been used in control function approaches to address unobserved heterogeneity (Imbens and Newey (2009), Train (2009), Adams, Einav, and Levin (2009), Agarwal (2015), Crawford, Pavanini, and Schivardi (2018)). I order private types ψ at each public type x such that default is increasing in ψ ,

$$\psi' > \psi \implies \chi(x, \psi') > \chi(x, \psi) \quad \forall x. \quad (4.7)$$

A second assumption is nonadvantageous selection. I explain the reason why this is a “selection” assumption more below, but formally the assumption is that higher-risk private types also face higher pricing in equilibrium:

$$\psi' > \psi \implies p_{1l}(x, \psi') > p_{1l}(x, \psi) \quad \forall x, l. \quad (4.8)$$

As I describe more in Section 4.1.3, lenders learn about private types ψ by the time they are setting mature-account prices, so these prices p_{1l} can be conditioned on ψ . This conditioning then makes clear how this definition of nonadvantageous selection has a tight relationship with traditional definitions of adverse selection described in terms of the correlation between residual risk and residual demand (Chiappori and Salanie (2000), Einav, Finkelstein, and Cullen (2010), Mahoney and Weyl (2014), Crawford, Pavanini, and Schivardi (2018)). Adverse selectedness in these definitions depends on whether, within observables (here, public types x), residual risk is positively correlated with residual demand; advantageous selection in contrast is a negative correlation. My requirement of nonadvantageous selection is slightly weaker than requiring adverse selection: I require only that, across private types ψ , demand be either positively correlated with, or sufficiently weakly negatively correlated with risk, such that lenders never set a lower price for a higher-risk ψ . Note however that this assumption embeds some restrictions on the competitive environment, namely that one lender’s relative quality advantage (as expressed in differences across l in demand parameters such as the flow utility from borrowing, δ_{0l}) does not change so drastically with ψ such that lenders in fact face lower demand as private risk rises. That is, residual demand curves are nonadvantageously selected in the pre-CARD-Act equilibrium. This assumption on residual demand curves is appealing because these are the demand curves that existing RCT evidence confirms are adversely selected (Ausubel (1999), Agarwal, Chomsisengphet, and Liu (2010)).

It is worth noting that the assumption of one-dimensional private types ψ and one-dimensional prices $p_{1l}(x, \psi)$ is crucial for expression (4.8). To the extent that credit card

⁹In CCP data, the median consumer incurs less than a \$2 change in their monthly minimum payment summed across all credit card accounts in response to a 100 basis point change in their credit card interest rate. Likewise, for the median consumer the minimum monthly payments due on a credit card are only 17% of total minimum payments due across all other loans including mortgages, auto loans, student loans, and other liabilities.

pricing is instead a two-part tariff, expression (4.8) could fail to hold if lenders offer low-rate, but high-annual fee cards to some high-demand and high-risk consumers. Evidence in Appendix Table 2 helps address this concern, showing that annual fees typically did not play this role in the pre-CARD-Act period.

4.1.3. *The Lender’s Problem*

Lenders are risk-neutral and solve two related problems: setting promotional or “teaser” prices $p_{0l}(\theta)$ for new accounts, and setting prices $p_{1l}(\theta)$ on mature accounts.

One important assumption about lenders is informational: I suppose lenders observe only a consumer’s public type x at the time of account origination, so there is asymmetric information about private types ψ . Accordingly, even though I write $p_{0l}(\theta)$ for new account prices, these prices are constrained to be the *same* for all types $\theta = (x, \psi)$ with the same public dimension x ; new account pricing does not vary by the private dimension ψ . Meanwhile, I assume private types ψ are learned through a relationship with a consumer, and I suppose full types $\theta = (x, \psi)$ are observed after one period of account holding so that ongoing account relationships do not feature asymmetric information between a lender and consumer.¹⁰ Thus mature-account prices $p_{1l}(\theta)$ can be conditioned on both the public and private components of θ .

Lenders post prices at the start of the period and lack commitment power over future periods’ pricing strategies. It therefore will be convenient to denote the expected *future* market price vector as $\bar{p} = \{\bar{p}_{0l}, \bar{p}_{1l}\}_{l \in L}$. As I discuss more in Section 4.1.4, I restrict my attention to a stationary equilibrium where optimal current prices are equal to these expected future prices.

I suppose credit card lenders’ costs also come in two types. First, on new accounts, lenders incur acquisition costs $c_{0l}(\theta)$ related to underwriting, marketing, and account setup. Second, on mature accounts, lenders incur costs $c_{1l}(\theta)$, which include costs of day-to-day account management plus the expected costs of default net of recoveries. The expected present value of profits on a mature account, held by a consumer in state $s = (\theta, l, b)$, can now be defined recursively as

$$\begin{aligned} \Pi_{1l}(p_l, s|\bar{p}) &= \sigma(l, \text{debt}, p_l|s, \bar{p}) \underbrace{(p_{1l}(\theta) - c_{1l}(\theta))}_{\text{flow profit}} \\ &\quad + \sigma(l, \text{debt}, p_l|s, \bar{p}) \underbrace{\beta(1 - \chi(\theta))T_{\theta\theta'}(\theta)\Pi_{1l}(p_l, \theta', l, \text{debt}|\bar{p})}_{\text{expected continuation profit | borrow}} \\ &\quad + \sigma(l, \text{no debt}, p_l|s, \bar{p}) \underbrace{\beta T_{\theta\theta'}(\theta)\Pi_{1l}(p_l, \theta', l, \text{no debt}|\bar{p})}_{\text{expected continuation profit | not borrow}}. \end{aligned} \tag{4.9}$$

The σ notation denotes choice probabilities, as in expression (4.3) (see also further examples in Appendix A.5). The $t = 1$ subscript on Π_{1l} refers to mature accounts, p_l is the lender’s pricing strategy, \bar{p} is the market price vector described above, and s is the consumer’s state.

¹⁰It would be interesting to relax this assumption and to study how lenders learn about borrower private types over time or from contract choice at take-up (Han, Keys, and Li (2017)). However, such learning dynamics and related borrower incentives to engage in signal jamming are beyond the scope of this paper.

The economic forces in this profit function are straightforward. Borrowing probabilities $\sigma(l, \text{debt}, \cdot)$ play the role of quantity and are decreasing in the price of borrowing. Lenders trade off downward-sloping demand against the incentive to raise margins. Higher prices also, through the downward slope in σ , lead to lost continuation values $\Pi_{1l}(\cdot, \text{debt}|\cdot)$ on borrowing accounts, though this loss is partly offset by continuation values $\Pi_{1l}(\cdot, \text{no debt}|\cdot)$ among a subset of marginal consumers who substitute to nonborrowing rather than account closure.

As noted above, lenders lack commitment power in the pre-CARD-Act regime. Therefore, a one-shot change in the current period's prices only affects consumers who pay those prices this period; in particular, a change in $p_{1l}(\theta)$ is only relevant for a type- θ consumer holding a mature account with lender l . This implies lenders optimize current-period, mature-account pricing type-by-type, solving $|\Theta|$ independent one-dimensional problems. Let μ_{0lb} denote the stationary distribution of types θ who start the period with an account from lender l , after having chosen borrowing status b in the past period. Then each θ -specific problem can be written,

$$\max_{p_{1l}(\theta)} \sum_b \mu_{0lb} \Pi_{1l}(p_l, \theta, l, b|\bar{p}), \tag{4.10}$$

where the sum is over consumers' prior-period borrowing choices, $b \in \{\text{debt}, \text{no debt}\}$.

Using the closed form for logit own-price and cross-price demand derivatives, the profit-maximizing price $p_{1l}^*(\theta)$ for each type θ therefore satisfies the first-order condition,

$$\begin{aligned} \sum_b \mu_{0lb} \sigma_{1bd\theta}^* &= \sum_b \gamma_\theta \mu_{0lb} \sigma_{1bd\theta}^* (1 - \sigma_{1bd\theta}^*) \\ &\quad \times (p_{1l}^*(\theta) - c_{1l}(\theta) + \beta(1 - \chi(\theta))T_{\theta\theta'}(\theta)\Pi_{1l}(\bar{p}_l, \theta', l, \text{debt}|\bar{p})) \\ &\quad - \sum_b \gamma_\theta \mu_{0lb} \sigma_{1bd\theta}^* \sigma_{1bn\theta}^* \beta T_{\theta\theta'}(\theta)\Pi_{1l}(\bar{p}_l, \theta', l, \text{no debt}|\bar{p}), \end{aligned} \tag{4.11}$$

where I use the shorthand, $\sigma_{1bd\theta}^* \equiv \sigma(l, \text{debt}, p_{1l}^*(\theta)|\theta, l, b, \bar{p})$ and $\sigma_{1bn\theta}^* \equiv \sigma(l, \text{no debt}, p_{1l}^*(\theta)|\theta, l, b, \bar{p})$.

Lenders' pricing problem for *new* accounts is similar. However, because private types ψ are unobserved at origination, new account pricing may also be shaped by adverse selection—a model feature that I confirm quantitatively using my estimates in Section 4.3. Specifically, lenders price new accounts based on an expected composition of private types who select into account opening at a given price:

$$\begin{aligned} &\Pi_{0l}(p_l, x|\bar{p}) \\ &= \sum_{\tilde{s} \in S(x)} [\mu_{\tilde{s}} \sigma(l, \text{debt}, p_l|\tilde{s}, \bar{p}) \underbrace{(p_{0l}(\theta(\tilde{s})) - c_{0l}(\theta(\tilde{s})))}_{\text{flow profit}} \\ &\quad + \underbrace{\mu_{\tilde{s}} \sigma(l, \text{debt}, p_l|\tilde{s}, \bar{p}) \beta(1 - \chi(\theta(\tilde{s})))T_{\theta\theta'}(\theta(\tilde{s}))\Pi_{1l}(\bar{p}_l, \theta', l, \text{debt}|\bar{p})}_{\text{expected continuation profit | borrow}} \\ &\quad + \underbrace{\mu_{\tilde{s}} \sigma(l, \text{no debt}, p_l|\tilde{s}, \bar{p}) \beta T_{\theta\theta'}(\theta(\tilde{s}))\Pi_{1l}(\bar{p}_l, \theta', l, \text{no debt}|\bar{p})}_{\text{expected continuation profit | not borrow}}]. \end{aligned} \tag{4.12}$$

Relative to mature-account profits Π_{1l} in the earlier equation (4.10), the key differences here for new account profits Π_{0l} are the sum over all potential consumer states \tilde{s} with public type x (i.e., a sum over all competing lenders, all borrowing statuses, and all private types, which I denote $S(x)$), and the dependence on the stationary distribution of consumers over those types, with type masses denoted $\mu_{\tilde{s}}$ (similar to in expression (4.11)).

Expression (4.12) makes clear the lender's asymmetric information problem on new accounts: optimal pricing depends on the expected type composition across hidden types ψ conditional on the fact that consumers are switching. Consistent with institutional details in the US credit card market, I suppose lenders cannot price discriminate based on which lender a consumer is switching from (even though lenders may have an incentive to do so, as the distribution of borrower types varies across lenders in equilibrium, à la Villas-Boas (1999)). Furthermore, even though competing lenders' pricing on mature accounts fully incorporates consumers' private types, there is no information "leakage" to other lenders offering new accounts to these consumers; as discussed in Appendix Section A.2, the pricing on a particular credit card account typically is not observed by a lender's competitors. Hence, there remains asymmetric information at the time of new account origination, even if the consumer opening the account held a mature credit card account with some lender in the prior period.

It should be noted that even though lenders face no adjustment costs to prices in this model, the model is still consistent with the phenomenon of "sticky prices" (Ausubel (1991), Grodzicki (2012)) in the credit card market, a term that has been used to refer to limited pass-through of cost shocks to borrowers. The potential for limited pass-through in this model is generated by imperfect competition and adverse selection, even in the absence of lender adjustment costs.

4.1.4. Equilibrium

A stationary Markov perfect Bayesian equilibrium of the model consists of:

1. Lender pricing strategies p_{1l} and p_{0l} that (1) maximize (4.10) and (4.12), and (2) are equal to expected future pricing $(\bar{p}_{0l}, \bar{p}_{1l})$, for all lenders l .
2. Choice probabilities $\sigma(l, b, p_l | s, \bar{p})$ that satisfy the consumer's problem (4.1) for all lenders l and all consumer states s .
3. Stationary distributions $\mu_{\theta lb}$ of consumer types θ over past-period lenders l and past-period borrowing status b that are generated by the equilibrium borrower choice probabilities, by default rates $\chi(\theta)$, and by consumer type transitions $T_{\theta\theta'}$.

In the zero-probability event of off-path play, I suppose that (i) lenders believe, consistent with the information structure in the lender's problem, that the stationary distribution $\mu_{\theta lb}$ is unchanged, and (ii) both consumers and lenders expect future play to remain on-path.

4.2. Model Estimation

4.2.1. Demand Estimation: Borrower Private Types

The first step in demand estimation is recovering a type θ for each borrower in each time period, as well as the probabilities that consumers experience type changes over time.

Recall types θ are a duple of public and private types, $\theta = (x, \psi)$. I allow each borrower's public type to be a binned version of her FICO score, as FICO scores are designed to be a one-dimensional composite of all publicly available information predicting default. I use 20-point FICO score bins ranging from 580 to 780 for a total of 11 distinct public types x .

To recover private types ψ , my approach builds on other literatures that seek to identify unobservable ex ante types from ex post outcomes, for example, the public economics literature on annuities markets that estimates ex ante frailty using ex post mortality (Finkelstein and Poterba (2004), Einav, Finkelstein, and Schrimpf (2010)). Here, I use a similar outcome, loan default, to recover ex ante borrower types. Because borrower types change over time, and also because default is only observed at most once for each account, this exercise is more complex than simply estimating individual-level residual default risk after controlling for FICO. Rather, I recover these private types from the observed pricing that each borrower faces in each period.

Here, I make use of the two assumptions in expressions (4.6) and (4.8) above. These two, together with the fact that default rates χ are, by construction, increasing in private types ψ , imply that default rates and equilibrium prices p_{1l} are increasing with respect to each other,

$$\hat{\chi}_{lx}(p_{1l}(x, \psi)) \nearrow p_{1l}(x, \psi), \quad (4.13)$$

where $\hat{\chi}_{lx}$ is the default rate as an indirect function of prices charged to each type in equilibrium, among borrowers with FICO score x for lender l . Using the inverse of χ implied by equation (4.13), private types can then be recovered by inverting default rates observed at each price. Suppressing dependence on x , this inversion is

$$\psi = \chi^{-1}(\hat{\chi}_l(p_{1l}(\psi))). \quad (4.14)$$

Note that equilibrium price schedules p_{1l} are lender-specific, as are the indirect functions $\hat{\chi}_l$ relating these prices to realized default rates. However, the inverse χ^{-1} maps default rates, which are *common* for all borrowers of a given type, back to types. So, in estimation, $\hat{\chi}_l$ is estimated separately by lender and by FICO score x , while χ^{-1} is estimated across all lenders—that is, for the market as a whole—within each FICO score.

To do this inversion in practice, I first use isotonic regression to estimate $\hat{\chi}_l$ for each lender l and FICO score group x . The default measure I use is delinquencies of 90+ days within the following 2 years, as this is the outcome FICO scores themselves are specified to predict. In a few cases where the fitted isotonic functions for a particular lender map onto a strict subset of the population distribution of default rates at a given FICO score, I use linear interpolation or extrapolation to extend the estimated function. This procedure results in $\hat{\chi}_l$ being a consistent estimate of actual default rates at each price level, as I prove in the Supplemental Appendix.

To define the inverse $\chi_x^{-1}(\cdot)$, I use the fact that private types ψ are an index of default risk (see equation (4.7)) across all lenders, and I therefore specify $\chi_x^{-1}(\cdot)$ to return quantiles of the population distribution of estimated default rates, for a desired number of quantiles. In my baseline estimation, I take 5 such quantiles (i.e., quintiles). This yields 5 private types for each of the 11 public types, for a total of 55 consumer types θ . Appendix A.1 shows consistency for continuously distributed private types as the number of quantiles used grows large. I then also coarsen each lender's pricing function $p_{1l}(x, \psi)$ to equal that lender's average price observed for each type.

This process is illustrated for two actual lenders in the data in Figure 3. As can be seen, a borrower of a given type shares a common default rate regardless of her current lender, while the price faced by each borrower is different depending on the lender she chooses; indeed lenders' price schedules can cross each other, as I show in the figure, when they set higher prices relative to default risk in some market segments than in others.

The raw data also show that the fit of the isotonic regressions is quite good, that is, true pricing functions do appear to be (nearly) monotone in default rates. Across all lenders

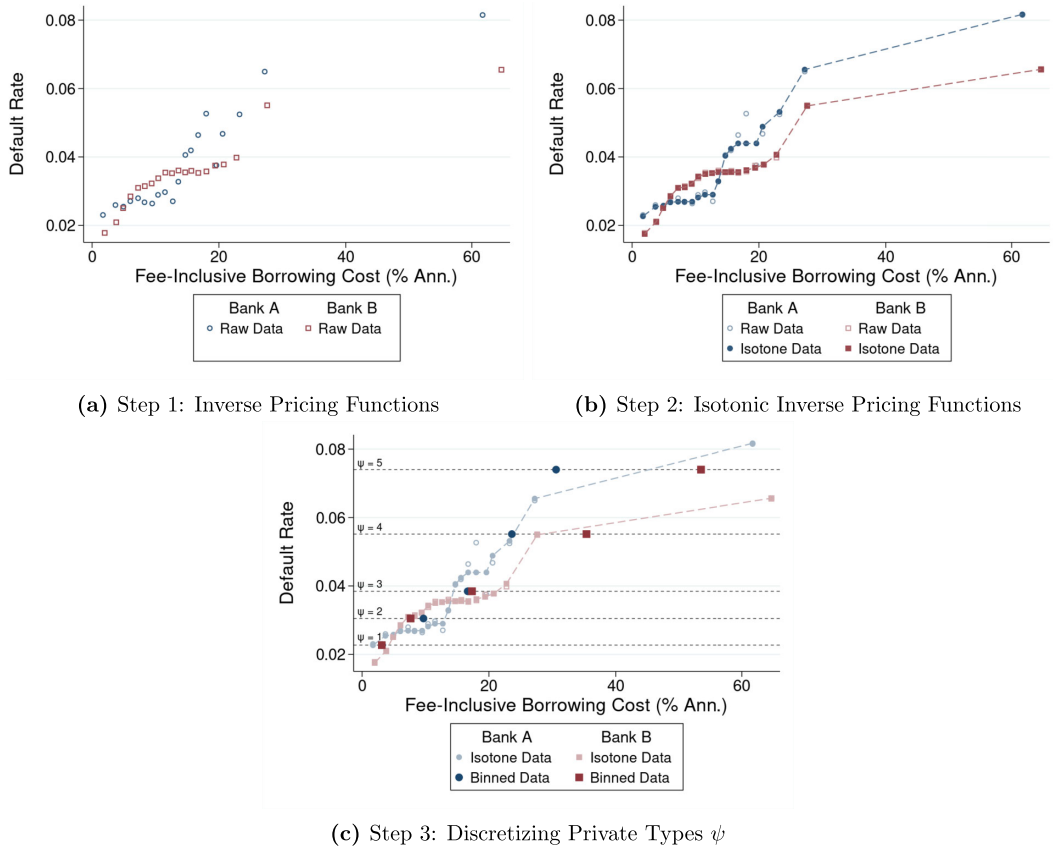


FIGURE 3.—Recovering private information types. *Notes:* The figure illustrates the process of recovering private information types from observed equilibrium pricing in pre-CARD-Act data, as described in equations (4.13) and (4.14) in the text. Panel (a) shows an example of raw data on default rates at quantiles of price levels on two different lenders, labeled Bank A and Bank B, for borrowers with public type corresponding to credit scores 760–779. Default is defined as delinquencies of 90+ days at any time over the subsequent 2 years. Panel (b) shows isotonic regression estimates of the relationship between default and equilibrium pricing, together with the raw data from panel (a). Panel (c) then shows how borrowers at different quantiles of the population distribution of default rates within this credit score range are grouped into discrete private information types ψ that share a common default rate, but face different prices depending on their choice of lender.

and credit scores, the R-squared of these isotonic fits in explaining the actually observed default rates is 97.9%. The McFadden pseudo R-squared (Menard (2000)) for predicting individual-level default outcomes is 20.96%, which is high relative to what is achievable when default is stochastic and default probabilities are between 0 and 1. For example, if default probabilities were uniformly distributed on $[0, 1]$, a McFadden pseudo R-squared using the *true* probabilities to explain simulated binary outcomes would be just 27.9%.

In Table III, I summarize the estimated private types and their characteristics. For brevity, I average over public types by pooling credit score tiers of subprime, prime, and superprime consumers (see footnote 2 above), while columns of the table show the five private types. There is considerable heterogeneity across private types in default risk, pricing, borrowing behavior, and profitability. First, note how the riskiest private types ($\psi = 5$) have about twice the default risk of the safest private types ($\psi = 1$), holding credit scores (public types) constant. The riskiest types also face substantially higher prices. Despite

TABLE III
ESTIMATED PRIVATE TYPE.

Private Type (ψ):	1	2	3	4	5
Annual Default Rate (p.p. ann.)					
Subprime	14.27	23.15	28.56	32.62	38.19
Prime	6.89	8.01	8.54	9.78	11.35
Superprime	1.62	2.20	2.66	3.78	4.57
Borrowing Cost (p.p. ann.)					
Subprime	15.92	25.83	43.07	78.64	143.52
Prime	10.40	20.18	29.31	38.19	56.23
Superprime	5.71	10.20	16.15	23.38	42.41
Revolving Share					
Subprime	0.24	0.24	0.26	0.28	0.48
Prime	0.33	0.33	0.37	0.41	0.66
Superprime	0.28	0.31	0.30	0.31	0.56
Cardholder Revolving Share					
Subprime	0.85	0.82	0.87	0.81	0.76
Prime	0.76	0.74	0.77	0.73	0.74
Superprime	0.50	0.54	0.51	0.49	0.65
Markup over Marginal Cost (%)					
Subprime	22.54	33.40	40.83	44.51	43.39
Prime	13.98	13.67	16.84	19.64	12.54
Superprime	8.28	7.71	7.28	9.37	7.60
EPDV Profit per Account					
Subprime	365.8	364.0	367.3	362.5	247.9
Prime	320.9	319.7	323.2	316.7	246.7
Superprime	244.3	245.6	246.2	248.7	219.8

Note: The table shows characteristics of estimated borrower private types. The default risk of these types is increasing in the type index ψ by construction. Borrowing cost refers to the fee-inclusive price of borrowing. The revolving share is the share of consumers that hold a credit card and borrow on it, while the cardholder revolving share is the probability that a consumer holding a credit card borrows from the same lender in the next quarter. The EPDV (expected present discounted value) of profits per account is reported in price units, rather than being dollarized as in Figure 7 (see Section 5.3.2 for discussion of these two units for profits). All statistics are weighted by market share. Subprime refers to accounts with FICO scores below 660; prime refers to accounts with FICO scores of 660 or above but below 720; superprime refers to accounts with FICO scores of 720 and above.

these higher prices, the revolving share—or the probability a consumer is borrowing on a credit card—also rises monotonically with private type. Anticipating later estimates of lenders' marginal costs and profitability, remarkably the highest-price consumers are not the most profitable, both in terms of within-period markups and in terms of lifetime profitability.

The consumer types estimated in this process also can be used to estimate how types change over time. In particular, the transition matrix $T_{\theta\theta}$ can be estimated nonparametrically off of type-to-type transition rates for borrowers who are observed in two successive periods. This takes advantage of the independence of type transitions from choices: type transitions do not depend on borrower choices and borrowers do not choose entry or exit from the market in anticipation of type transitions, as these transitions are not yet realized at the time choices are made.

The estimated transition matrix is illustrated as a contour plot in Figure 4. Here, the integer-labeled type indices correspond to the 11 different FICO score groups described earlier, while the subticks within each integer index correspond to the 5 discrete private types ψ within each FICO group. As can be seen, types are strongly but not perfectly




FIGURE 4.—Transition rates among public and private types. *Notes:* The figure displays a contour plot of period-to-period transition probabilities among consumer types. These probabilities are estimated quarterly among borrowers observed for two subsequent quarters, using the duple of public and private types recovered through the process illustrated in Figure 3. The integer values of the index correspond to the public dimension of types, in order of increasing credit score; for example, the range [0,1) corresponds to the 580–599 FICO score group, and the range [1,2) corresponds to the 600–619 FICO score group. Within integers, the subticks correspond to the five private information types recovered at each FICO score level, in order of increasing risk.

persistent, in both public and private dimensions. The rippling pattern evident in the plot shows how private types are predictive of future changes in public types, as borrowers of highly risky *private* types are more likely to be downgraded to a lower FICO score next period than other borrowers are.

Finally, after verifying that the estimated transition matrix $T_{\theta\theta'}$ is ergodic, this matrix can be used to recover the probability distribution over types μ_θ , which satisfies $\mu_\theta = T_{\theta\theta'}\mu_\theta$.

4.2.2. Demand Estimation: Elasticities

The next step in demand estimation is to estimate borrowers' price sensitivity. I first describe the variation I use for identification and then describe the estimation procedure in detail.

The pricing variation I use is, to my knowledge, novel: occasional repricing campaigns where lenders increase interest rates on entire credit card portfolios—for example, a sub-prime portfolio, or an airline card portfolio—at once. Former industry participants have confirmed in conversation that these repricing campaigns occur for a variety of reasons, sometimes in response to demand or risk perceived by an individual portfolio manager, or potentially also in response to lender-level cost shocks.

To use the most clearly exogenous pricing variation available, I focus on a lender's repricing campaign where investor relations materials confirm that the lender in question was seeking to shrink its credit card portfolio in anticipation of acquiring another lender. This variation therefore appears to come from a cost shock—in particular, a change in the lender's opportunity cost of capital in anticipation of the acquisition—rather than a demand shock, making it useful pricing variation for identifying demand. The particular merger or acquisition was not consummated until several quarters after the repricing event in question, and I can find no evidence that other dimensions of the lender's product quality, such as its product branding, changed in the months around this event.

This event also has the advantage of affecting nearly all of the given lender's accounts. All other repricing events that I can identify in my study period, with one exception,¹¹ affect no more than 10% of a given lender's mature revolving accounts; for an average lender in an average month, the share of mature revolving accounts treated by that lender-month's *modal* nonzero price change is just 2%. Identifying which criteria were used to select accounts into these other repricing events is difficult, so they are less useful sources of price variation.

To illustrate the particular repricing campaign I use, Figure 5's left panel shows, in red, all nine deciles of the APR distribution for this lender's credit card portfolio over time; the lender in question is labeled as "Bank A." All deciles of Bank A's APR distribution shift upward by exactly 100 basis points in a month labeled as event time 0, after a preceding period with minimal price change. This campaign occurred more than a year before the implementation of the CARD Act and occurred at a time when, as depicted by the figure's dashed blue line, other lenders' pricing was on average unchanged.

In the right panel of Figure 5, I then show Bank A's retention of its existing credit card borrowers around this price change, relative to competitors' retention of their existing credit card borrowers. The retention rate for Bank A falls relative to other lenders immediately after the repricing campaign. This pattern appears clearly despite strong seasonal effects on borrowing that occur during this time period, as retention rates peak annually in or around the month labeled as event time 0. Appendix A.6 provides further detail on how these retention rates are calculated, and formally tests (and fails to reject) equality of pretrends before the event.

I now describe how price coefficients γ are estimated using this variation. Recall from Section 4.1.3 that $\sigma(l, \text{debt}, p_l | s, \bar{p})$ denotes, for $s = (\theta, l, \text{debt})$, a type- θ consumer's probability of continuing to borrow from lender l after borrowing from l in the past period. Here, I shorten this to $\sigma_{ld\theta}$ for ease of notation. Differentiating $\sigma_{ld\theta}$ with respect to price $p_{1l}(\theta)$ to derive a demand elasticity yields

$$\text{dlog}(\sigma_{ld\theta}) = -\gamma_{\theta} p_{1l}(\theta)(1 - \sigma_{ld\theta})\text{dlog}(p_{1l}(\theta)). \quad (4.15)$$

I then use difference-in-differences in logs as empirical analogs of infinitesimal changes in logs,

$$\log\sigma_{ld\theta t} = \alpha_{\theta l} + \alpha_t + \beta_l t - \gamma_{\theta} \log P_{\theta l t} + \epsilon_{\theta l t}. \quad (4.16)$$

Here, the fixed effects α implement difference-in-differences, the β_l term absorbs lender-specific linear trends, and the term $P_{\theta l t}$ is a price term scaled as in equation (4.15), with scalars taken from the period immediately prior to the repricing event denoted here by $t = 0$:

$$\log P_{\theta l t} \equiv (1 - \sigma_{ld\theta 0}) p_{1l0}(\theta) \log(p_{1l t}(\theta)). \quad (4.17)$$

The base-period values $p_{1l0}(\theta)$ and $(1 - \sigma_{ld\theta 0})$ in equation (4.17) are chosen because they correspond to demand elasticities at the time of the repricing.

I estimate price coefficients γ using two-stage least squares (2SLS) with instruments that isolate the repricing variation in Figure 5. Specifically, I instrument for the price term $P_{\theta l t}$ with a dummy instrument $Z_{l t}$ equal to unity for Bank A following its repricing campaign, interacted with consumer type. Note that these instrumental variables address

¹¹The one exception is a small lender with a single repricing event affecting 37% of mature revolving accounts.

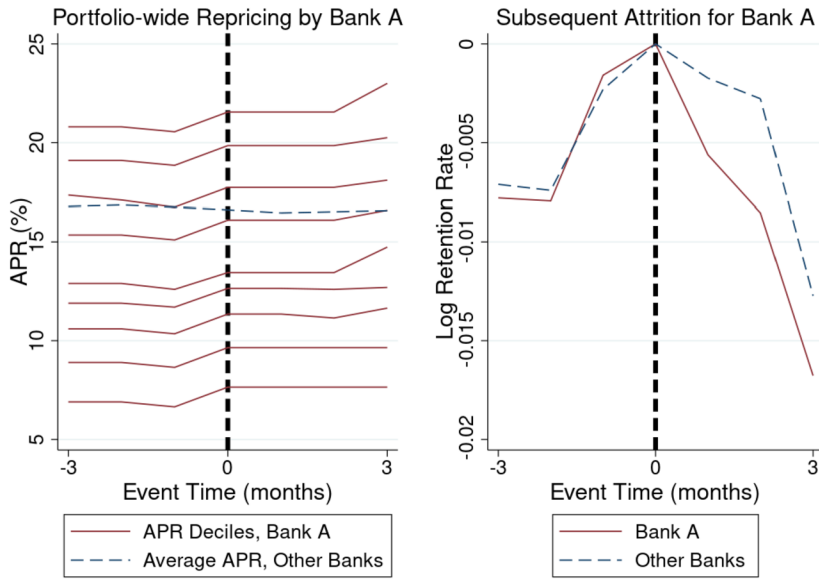


FIGURE 5.—Repricing quasi-experiment and borrowing response. *Notes:* The figure plots the repricing quasi-experiment (left panel) and subsequent attrition from borrowing (right panel) I use to identify price sensitivities. In the left panel, the solid red lines plot deciles of the distribution of annual percentage rates (APRs) on mature, borrowing accounts for one lender in the data, denoted Bank A. All deciles of this distribution rise by 100 basis points in the month labeled event time 0, emphasizing how this repricing campaign affects nearly all accounts in the portfolio. The dotted blue line shows the average APR for all other lenders’ mature, borrowing accounts. In the right panel, log monthly retention rates from borrowing are shown relative to their value in event time 0 for Bank A and for all other lenders. Specifically, the estimates in the right panel of the figure are the $\alpha_{-A,\tau}$ and $\alpha_{A,\tau}$ taken from the regression $\log \sigma_{ld\theta t} = \alpha_{\theta t} + \sum_{\tau \neq 0} \alpha_{-A,\tau} \times \mathbf{1}_{t=\tau} + \sum_{\tau \neq 0} \alpha_{A,\tau} \times \mathbf{1}_{t=\tau} + \beta_1 t + \epsilon_{\theta t}$, where A denotes Bank A-specific terms and $-A$ denotes all other lenders; this specification is discussed further in Appendix A.6.

two econometric issues, both the endogeneity of prices $p_{1t}(\theta)$ with borrower demand, and in time period 0, the appearance of $\sigma_{ld\theta 0}$ on both the right- and left-hand sides. In summary, the first and second stages are then

$$\log P_{\theta t} = a_{\theta t} + a_t + b_l \times t + \pi_{\theta} Z_{1t} \times \mathbf{1}_{\theta} + e_{\theta t}, \tag{4.18}$$

$$\log \sigma_{ld\theta t} = \alpha_{\theta t} + \alpha_t + \beta_l \times t - \gamma_{\theta} \log P_{\theta t} + \epsilon_{\theta t}. \tag{4.19}$$

Given that $P_{\theta t}$ contains the estimated quantity $\sigma_{ld\theta 0}$, I follow Cameron and Miller (2015) in bootstrapping over both individuals and clusters (at the level of lender \times consumer type) to calculate standard errors.

In Table IV, I present estimates of the price coefficient γ in equation (4.19) where, for sake of illustration, this coefficient is restricted to take on a single value for all types θ . In the first column, I show OLS estimates of equation (4.19), while in the second column, I then show corresponding 2SLS estimates. These estimates lend credence to the instrumental variables strategy: the OLS estimate is substantially closer to 0 than is the 2SLS estimate, as would be expected if the instruments overcome price endogeneity.

I further validate this instrumental variables strategy by asking whether the estimated heterogeneity in γ across consumer types is consistent with other reasonable models of consumer preferences. In Appendix A.6, I show that the heterogeneity I estimate across

TABLE IV
BORROWER PRICE SENSITIVITY.

Estimator:	(1)	(2)
	Dependent Variable: Log(Retention Rate)	
	OLS	2SLS
Price Coefficient (γ)	-0.00004 (0.00001)	-0.258 (0.097)
Number of Observations	60,638,012	60,638,012
Number of Clusters	550	550
First-stage F-statistic	.	252.45

Note: The table shows estimates of price coefficients (marginal utilities of income) estimated via OLS and 2SLS using quasi-experimental lender repricing. The variation is presented visually in Figure 5. Parentheses show bootstrapped clustered standard errors at the level of lender \times consumer type, following the procedure in Cameron and Miller (2015). See equations (4.18) and (4.19) in the text.

credit scores (public types) x matches closely what would be expected if all consumers have a constant coefficient of relative risk aversion equal to 3 (Brown and Finkelstein (2008)) and if income varies with credit score as in Albanesi, De Giorgi, and Nosal (2017). My 2SLS estimates therefore suggest consumers place *insurance* value on credit card pricing across states—that is, have different marginal utilities of income given different consumer type realizations—in a manner consistent with relatively standard CRRA preferences, even as I model utility as quasi-linear in pricing *within* any given period or state. I discuss such insurance value further in Section 5.3.2.

Finally, as additional validation of this instrumental variables strategy, Appendix A.6 presents evidence that Bank A's borrowers are comparable to other lenders' credit card borrowers, after controlling for borrower types θ .

4.2.3. Demand Estimation: Taste Parameters

I calibrate firms' discount factor to .98 quarterly and consumers' discount factor to .90 quarterly. Given the above estimates of each consumer's type θ and borrowers' price sensitivities γ , the remaining model parameters to be estimated are then the flow utilities $\delta_{\theta l}$, $\xi_{\theta l}$, $\kappa_{\theta l}$, and $\lambda_{\theta l}$ from Section 4.1.1. I estimate these using a minimum-distance estimator that minimizes squared deviations from several target moments in the data.

Some of these moments are especially informative about particular parameters. First, I target the $|\Theta| \times |L|$ probabilities that borrowers on mature accounts with each lender l in the data choose to borrow again from the same lender the following period; these are the empirical counterparts of the model choice probabilities $\sigma(l, \text{debt}, p_l | s, \bar{p})$. The form of $\sigma(l, \text{debt}, p_l | s, \bar{p})$ shown in expression A.14 in Appendix A.5 makes it clear how this probability depends directly on $\delta_{\theta l}$. I confirm this dependence by simulating the elasticity of these target moments with respect to these parameters at the estimated parameter vector, which I find is always nonzero and averages 1.87 across lenders and consumer types.

Second, I target the $|\Theta| \times |L|$ probabilities that borrowers on mature accounts with each lender l in the data choose to retain their credit card account with l but *not* borrow the following period (i.e., counterparts of the model choice probabilities $\sigma(l, \text{no debt}, p_l | s, \bar{p})$). These probabilities depend directly on liquidity costs $\lambda_{\theta l}$, as in expression A.14 in Appendix A.5. Simulating the elasticities of these moments with respect to these parameters, I find they are always nonzero and average -2.0 . Intuitively, the share of borrowers who

TABLE V
MODEL FIT TO TARGETED AND NONTARGETED MOMENTS.

		Subprime (1)	Prime (2)	Superprime (3)
Borrowing Retention	Model	0.93	0.92	0.85
	Target	0.92	0.91	0.85
Borrowing to Nonborrowing	Model	0.05	0.08	0.15
	Target	0.04	0.07	0.14
Nonborrowing to Nonborrowing	Model	0.62	0.72	0.82
	Target	0.68	0.79	0.86
Entry Rates	Model	0.04	0.05	0.03
	Target	0.04	0.05	0.03
Quarterly Default Rate (p.p.)	$\psi = 1$	3.05	1.73	0.46
	$\psi = 2$	4.75	2.01	0.59
	$\psi = 3$	5.67	2.16	0.68
	$\psi = 4$	6.81	2.53	0.88
	$\psi = 5$	8.24	2.91	1.11
Marginal Cost	$\psi = 1$	2.59	-2.08	-0.99
	$\psi = 2$	5.76	0.84	2.16
	$\psi = 3$	15.52	9.55	7.40
	$\psi = 4$	27.16	13.80	13.24
	$\psi = 5$	59.93	35.34	30.03

Note: The table shows averages of targeted moments across lenders and consumer types together with corresponding model moments. These moments and the parameters they identify most directly are described in Section 4.2.3. In the bottom half of the table, I show evidence of model fit to nontargeted moments by comparing estimated marginal costs to quarterly default rates. While marginal costs capture nondefault costs, such as account maintenance costs and the cost of funds, costs in a well-fitting model should be increasing in default rates. ψ refers to estimated private types. Subprime refers to accounts with FICO scores below 660; prime refers to accounts with FICO scores of 660 or above but below 720; superprime refers to accounts with FICO scores of 720 and above.

choose to pay off their debt (while remaining with the same lender) reveals the costs of paying off that debt.

Third and fourth, I target the $|X| \times |L|$ probabilities that nonborrowers with lender l continue to not borrow with l , and the $|X|$ probabilities that consumers with no credit card account ($l = 0$) in the past period open an account with any lender in a given period. These two groups of target moments can be calculated only at the $X \times L$ -level and at the X -level because, respectively, I do not recover private types in a given period for consumers not borrowing on mature accounts in that period, and the probability of “entering” the market is only observable in the CCP data where, as discussed in Section 2, I do not observe lender identity. These data limitations mean I restrict heterogeneity in the parameters so that account setup costs κ_{0l} vary only at the X -level, and the card-quality parameters ξ_{0l} vary only at the $X \times L$ -level. These final target moments then identify κ_x and ξ_{xl} , as illustrated in expressions A.16 and A.17 in Appendix A.5. Simulating the relevant parameter-moment elasticities, I find they are always nonzero and average -0.38 (for κ) and 0.59 (for ξ).

Although the model is just-identified by these moments, and hence estimating the model could be seen as something of a calibration exercise, the complexity of the de-

mand model means it is still nontrivial to achieve good model fit. Because consumers are forward-looking, all parameters affect all choice probabilities, and the analytic gradient of the L^2 -norm distance between empirical and model moments is difficult to work with directly. Instead, I use a two-step estimator, where I first use gradient-descent with an approximate gradient that ignores consumer's dynamic incentives; this "approximate gradient descent" is a high-speed first step that gets close to a good-fitting estimate for the full model. I then use a quasi-Newton algorithm in a second step that fully accounts for consumers' forward-looking incentives to converge on a best-fitting estimate.

In Table V, I summarize model fit. In general, the fit is quite tight, up to a single percentage point or less for most targeted choice probabilities. Fit is weaker for the $|X| \times |L|$ probabilities of nonborrowing consumers who continue not borrowing, where the model sometimes misses the targets by about 5 percentage points. Fortunately, as I show in the model parameter estimates in Table VI (discussed below in Section 4.3), the card-quality parameters most closely related to these choice probabilities are small enough in magnitude that they play a relatively modest role in the consumer's problem. Finally, in the bottom half of Table V, I present a test of overall model fit to some nontargeted features of the market, which I develop in the next subsection.

4.2.4. *Supply Estimation*

The supply-side estimation is standard: I recover lender marginal costs by minimizing squared violations of the lender's first-order conditions.

To match the dimensions of the lender's pricing problem on new accounts, where lenders set prices at the level of public types x and cannot vary prices by private types ψ , I restrict new-account marginal costs to also be common across private types within the same public type x . This restriction also reflects the type of marketing, underwriting, and account setup costs in the first quarter of an account's life that $c_{0l}(\theta)$ is meant to capture, as discussed in Section 4.1.3, since these costs typically would not vary with consumer unobserved risk. Therefore, there are $|\Theta| \times |L|$ values of $c_{1l}(\theta)$ to estimate and $|X| \times |L|$ values of $c_{0l}(\theta)$ to estimate. Given their linearity, the first-order conditions are easily satisfied exactly at the estimated values.

The marginal cost estimates help confirm the overall fit of the model with respect to nontargeted moments in the data. In particular, as discussed in Section 4.1.3, $c_{1l}(\theta)$ is meant to capture lenders' default costs (in addition to costs such as day-to-day account management), and default is not directly targeted in estimation. A test of untargeted fit therefore is whether these cost estimates are positively correlated with default risk, especially at lower credit scores where default is lenders' primary cost driver. In Table V, I confirm this is the case.

4.3. *Model Parameter Estimates*

In Table VI, I present my estimates of model parameters. Rows correspond to consumers' public types (FICO scores), and each column presents a different set of model parameters, taking averages over private types and over lenders within a row as needed.

In the first column of the table, I show estimates of the flow utility from borrowing δ_{0l} . These are decreasing in FICO score such that lower-risk consumers receive less utility from borrowing; this correlation is consistent with the adverse selectedness of the credit

TABLE VI
MODEL PARAMETER ESTIMATES BY FICO GROUP

FICO	(1) δ	(2) γ	(3) ξ	(4) κ	(5) λ	(6) c_0	(7) c_1	(8) χ
580–599	22.9	0.532	0.02	86.5	6.49	128.8	24.2	0.103
600–619	18.0	0.500	−0.08	89.7	5.66	142.0	24.6	0.050
620–639	14.4	0.420	−0.21	67.1	5.11	145.5	22.9	0.041
640–659	10.6	0.347	−0.66	49.1	4.39	149.6	17.1	0.034
660–679	9.2	0.339	0.49	30.0	3.96	158.1	14.3	0.028
680–699	7.0	0.295	1.16	29.0	3.89	161.6	10.0	0.023
700–719	5.6	0.278	2.04	29.6	3.91	164.0	10.2	0.017
720–739	4.3	0.240	2.54	28.1	4.04	164.6	7.5	0.013
740–759	4.2	0.243	2.04	27.8	3.90	163.3	10.2	0.009
760–779	4.3	0.240	1.67	28.9	3.71	160.8	11.8	0.006
780–799	3.3	0.191	1.71	26.9	3.45	154.3	12.0	0.003

Note: The table shows average model parameter estimates for the FICO score group in each row, where averages are taken across lenders and across private-information types as necessary. The parameters shown in each column respectively are: Flow utilities from borrowing, δ ; disutility from price (marginal utilities of income), γ ; card quality, ξ ; setup costs for opening a new account, κ ; liquidity costs for paying off existing balances, λ ; lender acquisition costs for new accounts, c_0 ; lender marginal costs on mature accounts used for borrowing, c_1 ; and quarterly default probabilities, χ . These default probabilities are estimated through the process illustrated in Figure 3. See the model exposition in Section 4.1, and the demand-side payoff summary in Table II.

card market. In results not shown in the table, I find this adverse selectedness also appears *within* public types: the correlation across private types between default rates and borrowing utility is at least 0.4 or higher in all credit score groups.

In other columns of the table, I show price coefficients γ , card-quality parameters ξ , account setup costs κ , liquidity costs λ , marginal costs c_1 and c_0 for new and mature accounts, and quarterly default rates. Setup costs are highest for consumers with the lowest credit scores, consistent with these consumers receiving fewer direct mail offers to open new credit card accounts (Consumer Financial Protection Bureau (2017)). Liquidity costs are also highest for consumers with the lowest credit scores, consistent with generally high credit constraints in this population (Bhutta, Skiba, and Tobacman (2015)). Setup costs are the more substantial of the two, consistent with other research examining similar adjustment costs in consumer demand for financial products (Handel (2013), Illanes (2017)).

To further quantify what these parameter estimates imply for adverse selection in the credit card market, I use, as a statistical tool, Chiappori and Salanie (2000)'s bivariate probit test for asymmetric information: using borrowing choices and default outcomes simulated in the estimated model, I test whether lenders on new accounts (i.e., when private types are still unobserved) face adverse selection net of the credit scores (public types) that lenders are able to price. I confirm they do; the estimated ρ of unobservables in the bivariate probit is 0.104.

Mature-account costs are positively correlated with actual default rates, as discussed in Section 4.2.4, though estimates show lenders face substantial nondefault costs for the highest-credit score consumers as well. New-account acquisition costs c_0 are increasing in credit score, consistent with lenders needing both greater marketing expenses—for example, more direct mail offers per account opened, as in Grodzicki (2014)—and greater expense on account-opening bonuses—such as a lump sum of airline miles.

5. EQUILIBRIUM UNDER CARD ACT PRICE RESTRICTIONS

I use the estimated model from Section 4 to study the CARD Act's pricing restrictions. I impose these restrictions in the model and I analyze their effects on pricing, borrowing choices, and total welfare after the market converges to a new equilibrium under the new regime. This exercise deliberately holds constant other features of the pre-CARD-Act environment to focus on a precise sense of these restrictions effects: *ceteris paribus* effects that emerge separately from, rather than in conjunction with, other contemporaneous economic and regulatory changes.

5.1. Modeling the CARD Act Price Restrictions

I model the CARD Act price restrictions as a mandate that firms commit to a single long-run price on each credit card contract at the time of origination. Contracts also include a promotional or "teaser" rate for one period before the long-run price takes effect, as such teasers were an important carve-out still permitted under the Act.

A post-CARD-Act contract therefore takes the form of a duple $(p_{0l,\text{post}}(x_0), p_{1l,\text{post}}(x_0))$ for lender l , containing an initial teaser rate and a subsequent long-run rate. This duple depends only on a consumer's public type (FICO score) at origination, x_0 . Pricing no longer depends on private information ψ_t revealed to a lender over time, or on changes in FICO scores x_t over time. Teaser rates continue to depend only on public types at origination, as before.

The choice to include teaser rates in my implementation of the CARD Act price restrictions leads to greater computational difficulty, as it doubles the size of lenders' strategy space. Including teaser rates is important, however, as lenders' ability to effectively set different prices for consumers who are and are not willing to switch accounts frequently has the potential to undo some of the price-pooling effects of the Act that I aim to study.

With lenders now unable to set mature-account prices based on private types and type changes, there are important implications for adverse selection and retention. On mature accounts, lenders face adverse retention similar to that previously documented in Figure 2. On new accounts, adverse selection is exacerbated as the highest-risk private types now anticipate lower pricing (relative to other private types) once the account reaches maturity.

Granted, my characterization of the restrictions also abstracts from some details of the Act, in particular, minor exemptions that would permit discretionary price changes in some circumstances. As discussed in Appendix Sections A.2 and A.4, these exceptions have been rarely used in practice, so abstracting from these seems a reasonable choice for sake of tractability.

I study an equilibrium where each firm can offer only one contract to each public type at origination. I make this restriction in part for sake of realism and in part for tractability, as this restriction avoids the difficulty of solving for an entire menu of contracts for each lender and each public type in an imperfectly competitive environment (Stole (2007)). This "one contract per firm per origination credit score" specification still allows substantial price dispersion at each public type, as differentiated lenders post different prices to each public type. Section 5.3.3 further explores how this one-contract assumption affects strategic interactions among lenders.

5.2. Solving for the Constrained Equilibrium

The firm's problem in the presence of the CARD Act repricing restrictions depends not just on consumers' current types, but also on the public type x_0 a consumer had when

she originated her current contract. Adapting notation from Section 4.1.3, I use $\mu_{\theta x_0/b}$ to denote the stationary distribution of consumers across contracts originated while of type x_0 , with current type θ , at lender l , and past-period borrowing status b . A lender’s total expected discounted profits under the restricted equilibrium can then be written as

$$\begin{aligned} &\Pi_l(p_{l,\text{post}}|\bar{p}_{\text{post}}) \\ &= \sum_{x_0} \underbrace{\Pi_{0l,\text{post}}(p_{l,\text{post}}, x_0|\bar{p}_{\text{post}})}_{\text{newly acquired accounts}} + \sum_{\theta, x_0, b} \underbrace{\mu_{\theta x_0/b} \Pi_{1l,\text{post}}(p_{l,\text{post}}, \theta, l, x_0, b|\bar{p}_{\text{post}})}_{\text{existing accounts}}. \end{aligned} \quad (5.1)$$

The right-hand-side terms $\Pi_{1l,\text{post}}$ and $\Pi_{0l,\text{post}}$ are defined analogously to their counterparts (4.9) and (4.12), updated to include dependence on x_0 ; the full forms of these are shown in expressions A.19 and A.20 in Appendix A.9.

I use successive lender best replies that maximize this profit function to compute the new equilibrium, beginning this process at the pre-CARD-Act equilibrium price vector. At each iteration, each lender computes its best reply to the prior iteration’s market price vector, given consumer behavior determined by the demand side of the model; all of these best replies then together form the market price vector for the next iteration.¹² This process by construction does not depend on the order in which lenders best reply, as each lender’s best reply to the market price vector is computed independently before updating the market price vector and starting the next iteration. Equilibrium convergence is then defined in terms of stability in the market price vector from one iteration to the next.

5.3. Effects of CARD Act Price Restrictions

The new equilibrium shows how prices, quantities, and welfare change under the CARD Act price restrictions.

5.3.1. Prices and Quantities

I find that the Act’s pricing restrictions induce moderate market unraveling for consumers with the lowest credit scores: pooling increases, and pricing newly exceeds willingness to pay for over 30% of safe private types ($\psi = 1$ or 2) with subprime scores who previously borrowed. At higher credit scores, nearly all consumer types face lower prices. Meanwhile, because selection changes the composition of consumers who stay in the market and of contracts that get retained over time, average *traded* prices—that is, prices actually paid by consumers who choose to borrow—decrease at all credit score levels.

I present these results in a series of figures. First, in Figure 6, panels (a), (b), and (c), I show these effects in three FICO score groups across the range of the score distribution:

¹²These best replies serve both as a computational tool to iteratively find the new equilibrium, and as an equilibrium selection device. Similar to some other empirical work that has simulated market equilibrium under a new regulatory regime (e.g., Ryan (2012)), it is difficult to rule out the presence of multiple equilibria in my setting. This process of successive best-replies from the pre-CARD-Act equilibrium seems most plausible as a device to select the post-CARD-Act equilibrium. While an exploration of different starting values, other than the pre-CARD-Act price vector, suggests that other counterfactual equilibria do exist, in general the alternative equilibria I have been able to identify appear to be either: (i) rare and also grossly dissimilar from the observed post-CARD Act equilibrium (e.g., near-total market unraveling), or (ii) essentially similar to the post-CARD-Act equilibrium I study (e.g., with price differences on essentially nontraded contracts). These findings support the empirical relevance of the equilibrium I study. For evidence that firms indeed converge on a new equilibrium gradually by playing best replies to other firms’ most recently observed pricing strategies, see Doraszelski, Lewis, and Pakes (2018).

deep subprime consumers with scores of 580–599; consumers at the border of subprime, with scores of 680–699; and superprime consumers with scores of 780 and above. In the left figure in each panel, I show equilibrium effects on mature-account prices for the five private-information types on the x-axis. These are the annualized fee-inclusive prices denoted $p_{1l,\text{post}}(x_0)$ in the post-CARD-Act equilibrium and $p_{1l}(\theta)$ in the pre-CARD-Act equilibrium; to facilitate comparison between the two equilibria I focus on consumer-contract pairs in both equilibria for which $x(\theta) = x_0$. In the right figure in each panel, I show equilibrium effects on borrowing behavior, or the share of all consumers who choose to borrow on a credit card, for the same credit score groups.

Turning first to deep subprime consumers in panel (a), in the left figure, I show a shift from heterogeneous pricing (a separating equilibrium) across private-information types in pre-CARD-Act data, to nearly complete pooling in the constrained equilibrium. Under this nearly-pooled pricing, all private types with FICO scores of 580 are now estimated to pay a fee-inclusive price of credit of nearly 50% annualized.¹³ However, for the high-risk, high-demand private types this is a lower rate than the average paid in pre-CARD-Act data.


These high prices are an equilibrium outcome driven in part by market unraveling: that is, the safest private-information types exit from borrowing as they are pooled with riskier peers; the cost of lending to the remaining, riskier private-information types then drives prices higher still; these higher prices then induce further exit by relatively safe private-information types; and so on. In the right-side figure of Panel (a), I illustrate this unraveling by showing changes in borrowing behavior for the same credit score group. While the safest private types exit somewhat from borrowing, the riskiest private types borrow more.

Turning to the remaining panels of the figure, other credit score groups do not experience the same degree of unraveling as was seen among deep subprime consumers in panel (a). In panel (b), within the FICO 680 group nearly all private information types experience lower prices as a result of lower markups in the constrained equilibrium; only the safest quintile of private-information types face higher prices while being pooled with their riskier peers. In the right-side figure of panel (b), I then show how these relative price changes affect borrowing shares across types. In panel (c), I show even broader price decreases at higher credit scores: among FICO 780 consumers, all private-information types in fact face either reduced or nearly unchanged loan pricing. Correspondingly, all private-information types in the FICO 780 group have greater borrowing shares in the constrained equilibrium.

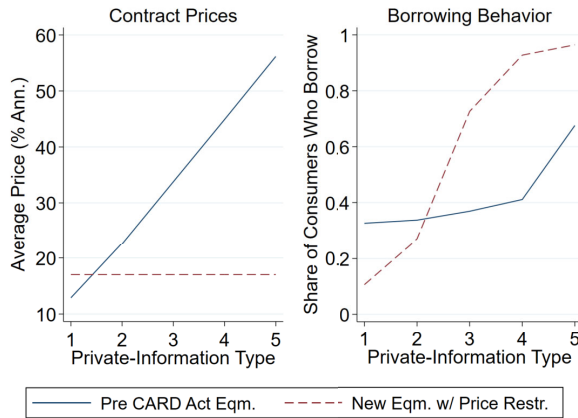
To emphasize, the prices shown in Figure 6 are averages across lenders of the prices available to a consumer with that current credit score. Average *traded* prices may differ from these prices for two reasons. First, in Figure 6, I show that consumers who face price increases tend to exit the market, which changes the composition of which prices are traded. Second, consumer types also change over time—for example, a prime consumer may later become subprime—and the CARD Act restrictions can allow a consumer to retain her earlier contract and its pricing as her type changes.

To summarize the effects of these two compositional changes, I compute averages of the actual prices paid at each FICO score among consumers who choose to borrow. As I show in Appendix Figure 3, these traded prices fall more than the contract prices shown in the prior figure did, reflecting both some consumers' exit from the market and other


¹³Pooling may not be fully complete, to the extent that different lenders set different prices and have different market shares across different ψ types.



(a) FICO 580–599 Consumers



(b) FICO 680–699 Consumers



(c) FICO 780+ Consumers

FIGURE 6.—Equilibrium changes in pricing and borrowing. *Notes:* The figure shows average mature-account prices by credit score and private information type, and the share of each type that uses a credit card for borrowing, in equilibrium with and without the CARD Act price restrictions. The figure is further described in Section 5.3.1 of the text.

consumers' retention of relatively favorable prices over time. Among subprime consumers in particular, this difference between contract and traded prices reflects how relatively few subprime consumers borrow at the (increased) contract prices available to them in the new equilibrium.

In Appendix A.10, I discuss how these estimates of changes in traded prices are reasonably comparable to estimates of the CARD Act's effects in Agarwal et al. (2015b). While my estimates of the Act's effects are somewhat larger, several differences between the time horizons, weighting, and estimands in the two papers appear likely to account for these differences. The same Appendix section also presents other evidence of the model's fit to post-CARD-Act data, including the model's ability to replicate the descriptive patterns shown in Section 3.2 above.

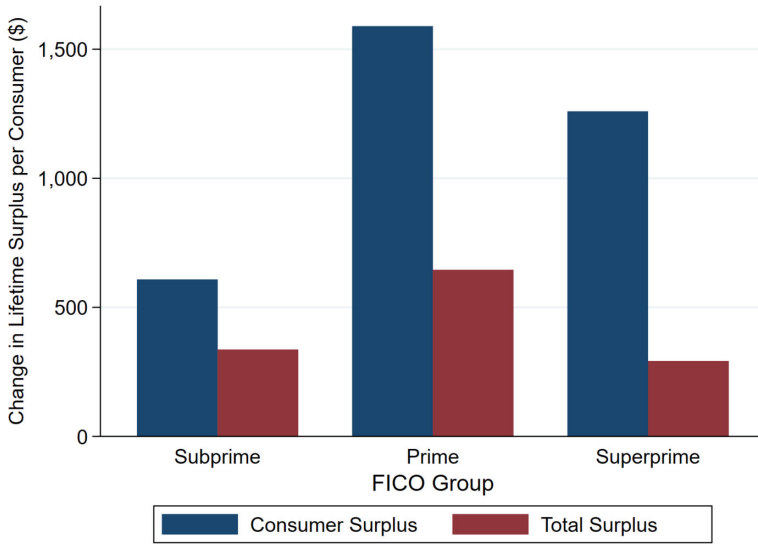
Another measure of market unraveling is change in average costs (Einav, Finkelstein, and Cullen (2010), Crawford, Pavanini, and Schivardi (2018)). In Appendix Figure 6, I show that average costs rose in the same subprime segments where Figure 6, panel (a) illustrated unraveling using choice probabilities. In contrast, average costs fall at other credit score tiers, reflecting the influx of relatively safe private types into borrowing. The figure also illustrates how decreases in teaser rates $p_{0t, \text{post}}$ generally complement the decrease in mature-account prices, though given the small share of consumers on teaser-rate contracts these prices are less consequential for consumer surplus. I turn to quantifying surplus next.

5.3.2. *Welfare: Consumer and Total Surplus*

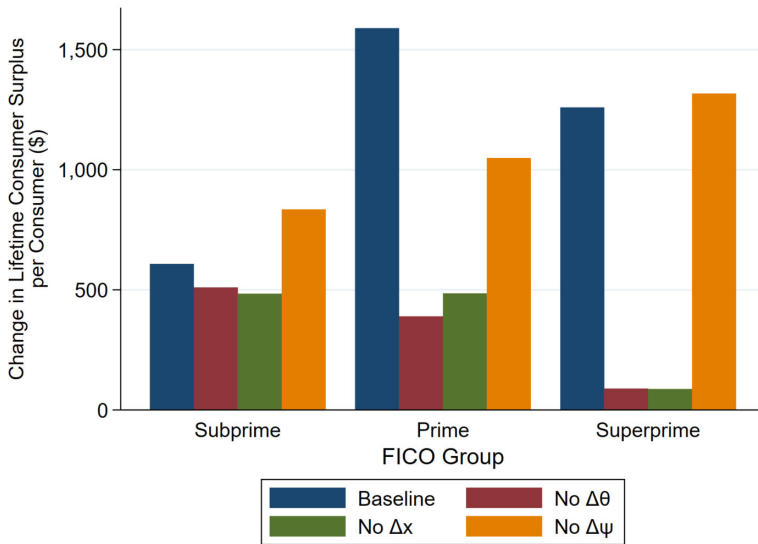
The estimates discussed above suggest that implications for consumer and total surplus could be ambiguous: I find the Act's price restrictions cause prices to rise for some consumer types, who partly exit the market in response; however, I also find these restrictions cause traded prices to fall among the set of consumers who remain in (or newly enter) the market. To quantify the effects of these changes for consumer welfare, I calculate lifetime consumer surplus for each consumer type under both the pre-CARD-Act equilibrium and the constrained equilibrium, and I use each consumer type's marginal utility of income (i.e., the price coefficients estimated with quasi-experimental variation in Section 4.2.2) and the total balance borrowed by the median consumer at each credit-score level in order to dollarize these surplus differences. Because utility is quasi-linear in income when holding consumer types fixed, this yields both a compensating and equivalent variation. Then to quantify overall welfare, I add per-consumer dollarized lender profits to these estimates, yielding total surplus under both equilibria.¹⁴

I find that consumer surplus conditional on credit score in fact rises across all FICO groups as a result of the CARD Act price restrictions, with one-time surplus gains equal to roughly \$600 for subprime consumers and over \$1000 for prime and superprime consumers. These gains are plotted in Figure 7, panel (a). Much of the gain comes at the expense of producer surplus (i.e., lender profits), so as I show in the figure, the net effect on total surplus is more modest even if still positive. In Appendix Figure 11, I decompose the decrease in producer surplus into changes in lifetime revenue and changes in lifetime costs.

¹⁴Note that this strategy for dollarizing consumer surplus has the effect of increasing prime and superprime consumers' surplus more than subprime consumers' surplus, given the higher balances held by these groups. This strategy also treats balances as exogenous; however, if these balances changed endogenously in the post-CARD-Act equilibrium, an envelope theorem argument shows the welfare effects of these unmodeled balance changes would be second order.



(a) Total and Consumer Surplus



(b) Consumer Surplus Decomposition

FIGURE 7.—Consumer and total surplus. *Notes:* Panel (a) shows estimated per-person changes in lifetime consumer surplus and total surplus, from the pre-CARD-Act equilibrium to the new equilibrium found under the CARD Act price restrictions. Total surplus is consumer surplus plus the present value of lender profits. Consumer surplus is dollarized using each type’s marginal utility of income (the price coefficients γ) and using median borrowed balances for a type’s credit score group. Per-person surplus numbers are averaged to coarser credit-score groups using the type probability distribution μ_θ . Panel (b) shows a decomposition of consumer surplus gains under counterfactual scenarios where consumers’ insurance demand is changed by removing their exposure to any consumer-type changes (“no $\Delta\theta$ ”), to changes in public type (“no Δx ”), and to changes in private type (“no $\Delta\psi$ ”). Subprime consumers have FICO scores below 660; prime consumers’ scores are between 660 and 720; superprime consumers’ scores are 720 and above.

There are several sources of these consumer surplus gains. Besides the decrease in markups documented in the previous section, which has been a focus of other recent literature on personalized pricing (Grunewald et al. (2023), Dubé and Misra (2023), Buchholz et al. (2024), Rhodes and Zhou (2024)), I find that another important driver of these gains is the Act's *insurance* value for consumers in cases where default risk deteriorates over time. To quantify how consumers value this protection against price increases, I study a counterfactual environment where consumers have no demand for insurance against price changes because consumer types are perfectly persistent (i.e., $T_{\theta\theta}$ is the identity matrix), and I recompute consumer surplus before and after the CARD Act restrictions within this counterfactual environment. I leave market pricing unchanged from the actual (not counterfactual) environment, in order to focus on mechanisms for surplus gain within equilibria already discussed. Hence, these estimates provide a partial-equilibrium decomposition of the sources of surplus gain.

In Figure 7, panel (b), I show the results of this decomposition. The blue and red bars show starkly that insurance value generates little of the consumer-surplus gains among subprime consumers, and nearly all of the surplus gain among superprime consumers. These superprime consumers are most likely to “lock in” favorable pricing and also to keep their account for long enough to have a substantial probability of migrating to other credit scores. The remaining (green and gold) bars in the figure show analogous results where consumers only face risk to changes in their private type, or only face risk to changes in their public type. The pattern where risk over private types is seen to decrease the insurance value of the CARD Act's restrictions arises because risk over private types tends to be negatively correlated with risk over public types: that is, a transition to a higher-risk public type is often partially offset by a simultaneous transition to a lower-risk private type. Hence the “no- ψ ” scenario in the figure represents a scenario where some consumer types face the greatest repricing risk of all.

To help interpret the magnitude of these consumer surplus gains, recall that these surplus measures are lump-sum equivalent variations for a permanent policy change. The estimates in Figure 7, panel (b) suggest these gains are on the order of a few hundred dollars in the absence of any insurance value. This is closely comparable to present value of other estimates of the Act's price effects for consumers, for example, those in Agarwal et al. (2015b). Interestingly, while these direct pecuniary benefits are sizeable, the insurance value of the Act is of a similar magnitude.

5.3.3. *The Role of Markups*

In markets with adverse selection, informational or pricing restrictions of course do not always translate into surplus gains even when consumers place insurance value on those restrictions. Handel, Hendel, and Whinston (2015), for example, give a prominent example of how similar informational restrictions can induce severe market unraveling in a perfectly competitive insurance setting. What accounts for the relative lack of unraveling here? I show that markups in the pre-CARD-Act period were necessary for the Act's relatively positive effects for welfare.

To show this, I consider two ways of recasting the observed pre-CARD-Act equilibrium as a more competitive one and then implement the CARD Act restrictions within that more competitive setting. In the first approach, I suppose lenders' marginal costs on each account are equal to observed pre-CARD-Act prices on that account; that is, I set $c_{0l}(x) = p_{0l}(x)$ and $c_{1l}(\theta) = p_{1l}(\theta)$. To avoid issues with equilibrium nonexistence in a perfectly competitive selection market (Handel, Hendel, and Whinston (2015)), I then

collapse these lenders to a single representative lender by averaging across all lender-level estimated parameters. This exercise can also be understood as ignoring the finite residual-demand elasticities evident in the repricing analysis in Section 4.2.2 and supposing instead that the lender first-order conditions are satisfied at zero markups, while leaving my demand estimates otherwise unchanged.

In the second approach, I posit instead that the demand response to repricing in Figure 5 might be greater than actually observed, such that the FOC-implied markups would be lower than what I estimate (but nonzero). I then reestimate the rest of the demand model as if this were the case; specifically, I suppose these demand elasticities were sufficiently large as to increase my estimates of price sensitivities γ by one bootstrapped standard error relative to the point estimate in column (2) of Table IV, given how larger demand responses in Figure 5 would map directly to larger price sensitivities.

In the first scenario—supposing zero markups in the pre-CARD-Act equilibrium—I find that the market would completely unravel under the CARD Act restrictions. All prices in the market exceed 150% APR, a vanishingly small share of consumers choose to borrow, and surplus per consumer falls by between \$100 and \$600 depending on the credit score tier. I present these results in Appendix Figure 9.

In the second scenario, I find in Appendix Figure 10 nontrivial unraveling (in the sense of higher average costs) in all credit score tiers, and I find the CARD Act would lead to lower total surplus at all credit score tiers than in my baseline estimates.

A related question is the extent to which average prices fall under the CARD Act price restrictions because of strategic interactions between lenders' pricing strategies. Because prices are strategic complements, and because one lender's pricing of information can serve to make another lender's residual demand curve on new accounts more adversely selected, the CARD Act price restrictions can exert additional downward pressure on prices through strategic channels. To help quantify this effect, I solve for counterfactual pricing strategies in which each lender (i) itself faces the CARD Act's price restrictions but (ii) plays against competitors who continue to play their pre-CARD-Act unrestricted pricing strategy. I find that lender-by-lender best replies in this case indeed reflect higher pricing than in the full post-CARD-Act equilibrium where all lenders face the same restrictions. Interestingly, the differences in best replies vary by FICO score: prices rise for subprime and superprime consumers on average by 1.96 and 0.49 percentage points respectively, while I estimate prices rise by 157 percentage points (i.e., near-total unraveling in these lender-by-lender best replies) for prime consumers. The particularly pronounced unraveling in the prime segment reflects a combination of adverse selection in residual demand on new accounts (which I find is less pronounced for superprime consumers) and a prevalence of switching across lenders (which is less substantial for subprime consumers).

5.3.4. *Limitations*

A notable caveat to these results is that I treat consumers' type changes as exogenous to past and current borrowing decisions. If consumers instead could affect their future types, in particular by borrowing less to decrease the probability of becoming a high-risk type, then policies like the CARD Act's price restrictions that lead to more favorable pricing for high-risk consumers could lead to increased borrowing through a moral hazard channel, potentially attenuating some of the welfare gains that I estimate; this is a distinct notion of moral hazard from the price invariance of default I argued for in Section 4.1.2. Hence, a key question for future research and for the interpretation of my welfare estimates is the extent to which default risk is driven by forces potentially exogenous to borrowing choices, such as job loss, or rather is driven by the endogenous accumulation of balances.

Another important caveat to these results is that I model consumers with rational expectations over their own future types and firms' pricing. This assumption may be in tension with evidence on behavioral frictions in consumer credit markets. While a full exploration of behavioral credit card borrowers is beyond the scope of this paper, I consider robustness to a form of myopia in which consumers at first misunderstand that lenders' pricing will respond to type changes. In particular, I consider a small mass of myopic consumers who enter the market each period, holding no credit card and unaware that credit card pricing responds dynamically to type changes. As soon as they open their first credit card account, they learn that they can face repricing, and they behave like nonmyopic consumers thereafter; for tractability I also suppose the mass of these entrants each period is zero so that they do not affect equilibrium pricing.¹⁵ Under this model of myopia, the CARD Act provides greater surplus gains to myopic consumers than nonmyopic consumers, especially for consumers with higher credit scores, as they incur the largest optimization errors when myopic. For example, in the prime market segment, myopic consumers' lifetime surplus gains from the Act are \$1776, contrasted with gains of \$1589 for nonmyopic consumers.

One more caveat to note is that I assume consumer–lender account relationships are strongly revealing of borrower private types: lender learning is fast. Different modeling tools are needed to study slower learning, and I anticipate nonparametric identification of borrower private types may no longer be possible in such a setting. Understanding the role of imperfect learning in markets with asymmetric information and imperfect competition is an important and difficult question for future work.

6. CONCLUSION

I study regulation that constrains lenders from discretionarily adjusting pricing on their outstanding loans. Under the 2009 Credit CARD Act, US credit card lenders were restricted from such discretionary price increases, suggesting the Act may limit how prices can reflect information learned over the course of lending relationships. I find that the kind of price increases restricted by the Act affected over 50% of borrowing accounts in the pre-CARD-Act period, that these price changes reflected both demand- and risk-relevant information, and that when pricing this information became restricted, price dispersion on newly mature accounts dropped sharply by about one third. Accompanying this shift toward more pooled pricing, I find descriptive patterns consistent with partial market unraveling: some consumers left the market, and this occurred especially for credit scores that saw the greatest price increases in the left (cheap) tail of their price distribution.

I then use a structural model to understand the mechanisms for, distributional patterns in, and welfare consequences of the CARD Act price restrictions' effects. The model quantifies a precise *ceteris paribus* sense of the restrictions' effects wherein I hold constant pre-CARD-Act primitives—including demand, risk, and product differentiation—and then study the effects of regulation that restricts dynamic discretionary pricing. I

¹⁵The mass of myopic types cannot accumulate to a nonzero measure even over infinite horizons, given nonzero probabilities each period of becoming nonmyopic each period via entry into card-holding. A full analysis of equilibrium with a positive mass of both behavioral and nonbehavioral types, for example, in the vein of Gabaix and Laibson (2006), is beyond the scope of this paper. An alternative analysis would consider a market with *only* behavioral types, but this is difficult to implement in the case of myopia about repricing, because lenders' high retention rates among borrowers with high ψ types (see, e.g., Table III) cannot then be rationalized without implausibly high setup costs κ , which then tend to unrealistically depress entry rates into card-holding.

confirm that these restrictions lead to partial market unraveling, especially among subprime consumers, where prices newly exceed willingness to pay for up to 30% of the privately safest borrowers. Consumer surplus nevertheless rises. Investigating mechanisms for these surplus changes, I find that these gains come partly from reduced lender profits, and partly from the restrictions' insurance value for consumers who face deterioration in their risk over time. Although these results are particular to US credit cards, an interesting area for future work is to explore which other markets offer similar conclusions about the welfare benefits of informational pricing restrictions when competition is imperfect, and how financial products or public policy can be optimally designed in response.

REFERENCES

- ADAMS, WILLIAM, LIRAN EINAV, AND JONATHAN LEVIN (2009): "Liquidity Constraints and Imperfect Information in Subprime Lending," *American Economic Review*, 99 (1), 49–84. [1384]
- AGARWAL, NIKHIL (2015): "An Empirical Model of the Medical Match," *American Economic Review*, 105 (7), 1939–1978. [1384]
- AGARWAL, SUMIT, SOUPHALA CHOMSISENGPHET, AND CHUNLIN LIU (2010): "The Importance of Adverse Selection in the Credit Card Market: Evidence From Randomized Trials of Credit Card Solicitations," *Journal of Money, Credit and Banking*, 42 (4), 743–754. [1384]
- AGARWAL, SUMIT, SOUPHALA CHOMSISENGPHET, CHUNLIN LIU, AND NICHOLAS S. SOULELES (2015a): "Do Consumers Choose the Right Credit Contracts?" *The Review of Corporate Finance Studies*, 4 (2), 239–257. [1374]
- AGARWAL, SUMIT, SOUPHALA CHOMSISENGPHET, NEALE MAHONEY, AND JOHANNES STROEBEL (2015b): "Regulating Consumer Financial Products: Evidence From Credit Cards," *The Quarterly Journal of Economics*, 130 (1), 111–164. [1373,1375,1402,1404]
- (2018): "Do Banks Pass Through Credit Expansions to Consumers Who Want to Borrow?" *The Quarterly Journal of Economics*, 133 (1), 129–190. [1374,1375]
- AGARWAL, SUMIT, JOHN C. DRISCOLL, XAVIER GABAIX, AND DAVID LAIBSON (2008): "Learning in the Credit Card Market," Working paper, National Bureau of Economic Research. [1374]
- (2009): "The Age of Reason: Financial Decisions Over the Life Cycle and Implications for Regulation," *Brookings Papers on Economic Activity*, 2009 (2), 51–117. [1374]
- AGUIRREGABIRIA, VICTOR, AND PEDRO MIRA (2010): "Dynamic Discrete Choice Structural Models: A Survey," *Journal of Econometrics*, 156 (1), 38–67. [1381]
- AKERLOF, GEORGE (1970): "He Market for Lemons: Quality Uncertainty and the Market Mechanism," *The Quarterly Journal of Economics*, 84, 488–500. [1372]
- ALBANESI, STEFANIA, GIACOMO DE GIORGI, AND JAROMIR NOSAL (2017): "Credit Growth and the Financial Crisis: A New Narrative," Tech. rep, National Bureau of Economic Research. [1394]
- GRODZICKI, DANIEL, ALEXEI ALEXANDROV, ÖZLEM BEDRE-DEFOLIE, AND SERGEI KOULAYEV (2023): "Consumer Demand for Credit Card Services," *Journal of Financial Services Research*, 63 (3), 273–311. [1374]
- ALLEN, JASON, ROBERT CLARK, AND JEAN-FRANÇOIS HOUE (2014): "The Effect of Mergers in Search Markets: Evidence From the Canadian Mortgage Industry," *American Economic Review*, 104. [1375]
- (2019): "Search Frictions and Market Power in Negotiated-Price Markets," *Journal of Political Economy*, 127. [1375]
- AUSUBEL, LAWRENCE M. (1991): "The Failure of Competition in the Credit Card Market," *The American Economic Review*, 50–81. [1374,1387]
- (1999): "Adverse Selection in the Credit Card Market," Working paper, University of Maryland. [1384]
- BENETTON, MATTEO (2021): "Leverage Regulation and Market Structure: A Structural Model of the U.K. Mortgage Market," *The Journal of Finance*, 76. [1375]
- BERLIN, MITCHELL, AND LORETTA J. MESTER (2004): "Credit Card Rates and Consumer Search," *Review of Financial Economics*, 13 (1), 179–198. [1374]
- BERRY, STEVEN, JAMES LEVINSOHN, AND ARIEL PAKES (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63 (4), 841–890. [1381]
- BERRY, STEVEN T. (1994): "Estimating Discrete-Choice Models of Product Differentiation," *The RAND Journal of Economics*, 242–262. [1381]
- BHATTACHARYA, VIVEK, GASTON ILLANES, AND MANISHA PADI (2025): "Fiduciary Duty and the Market for Financial Advice," Tech. rep, National Bureau of Economic Research. [1375]

- BHUTTA, NEIL, JANE DOKKO, AND HUI SHAN (2017): "Consumer Ruthlessness and Mortgage Default During the 2007 to 2009 Housing Bust," *The Journal of Finance*, 72 (6), 2433–2466. [1384]
- BHUTTA, NEIL, PAIGE MARTA SKIBA, AND JEREMY TOBACMAN (2015): "Payday Loan Choices and Consequences," *Journal of Money, Credit and Banking*, 47 (2–3), 223–260. [1397]
- BRICKER, JESSE, LISA DETTLING, ALICE HENRIQUES, JOANNE HSU, LINDSAY JACOBS, KEVIN MOORE, SARAH PACK, JOHN SABELHAUS, JEFFREY THOMPSON, AND RICHARD WINDLE (2017): Changes in US Family Finances from 2013 to 2016. Federal Reserve Bulletin, (September). [1372,1375]
- BRITO, DAGOBERT L., AND PETER R. HARTLEY (1995): "Consumer Rationality and Credit Cards," *Journal of Political Economy*, 103 (2), 400–433. [1374]
- BROWN, JEFFREY R., AND AMY FINKELSTEIN (2008): "The Interaction of Public and Private Insurance: Medicaid and the Long-Term Care Insurance Market," *American Economic Review*, 98 (3), 1083–1102. [1394]
- BUCHAK, GREG, GREGOR MATVOS, TOMASZ PISKORSKI, AND AMIT SERU (2022): "Why Is Intermediating Houses so Difficult? Evidence From IBuyers," Working Paper. [1375]
- BUCHHOLZ, NICHOLAS, LAURA DOVAL, JAKUB KASTL, FILIP MATĚJKA, AND TOBIAS SALZ (2024): "Personalized Pricing and the Value of Time: Evidence From Auctioned cab Rides," Tech. rep, National Bureau of Economic Research. [1375,1404]
- CAMERON, A. COLIN, AND DOUGLAS L. MILLER (2015): "A Practitioner's Guide to Cluster-Robust Inference," *Journal of Human Resources*, 50 (2), 317–372. [1393,1394]
- CASTELLANOS, SARA G., DIEGO JIMÉNEZ HERNÁNDEZ, APRAJIT MAHAJAN, AND ENRIQUE SEIRA (2025): "Financial Inclusion and Contract Terms: Experimental Evidence From Mexico," Working paper, National Bureau of Economic Research. [1383]
- CHATTERJEE, SATYAJIT, DEAN CORBAE, KYLE DEMPSEY, AND JOSÉ-VÍCTOR RÍOS-RULL (2023): "A Quantitative Theory of the Credit Score," *Econometrica*, 91 (5), 1803–1840. [1375]
- CHIAPPORI, PIERRE-ANDRE, AND BERNARD SALANIE (2000): "Testing for Asymmetric Information in Insurance Markets," *Journal of Political Economy*, 108. [1384,1397]
- CLARK, ROBERT, JEAN-FRANÇOIS HOUDE, AND JAKUB KASTL (2021): "Chapter 15—The Industrial Organization of Financial Markets," in *Handbook of Industrial Organization*, Vol. 5, ed. by Kate Ho, Ali Hortaçsu, and Alessandro Lizzeri. Elsevier. [1375]
- COHEN, ALMA, AND LIRAN EINAV (2007): "Estimating Risk Preferences From Deductible Choice," *The American Economic Review*, 97 (3), 745–788. [1384]
- CONSUMER FINANCIAL PROTECTION BUREAU (2013a): "Consumer Credit Panel," CFPB Confidential Dataset, accessed 2016–24. [1376]
- CONSUMER FINANCIAL PROTECTION BUREAU (2013b): "Credit Card Database," CFPB Confidential Dataset, accessed 2016–24. [1375]
- CONSUMER FINANCIAL PROTECTION BUREAU (2017): "The Consumer Credit Card Market," Tech. rep, Consumer Financial Protection Bureau. [1397]
- CRAWFORD, GREGORY S., NICOLA PAVANINI, AND FABIANO SCHIVARDI (2018): "Asymmetric Information and Imperfect Competition in Lending Markets," *American Economic Review*, 108 (7), 1659–1701. [1381,1384,1402]
- CUESTA, JOSÉ IGNACIO, AND ALBERTO SEPULVEDA (2021): "Price Regulation in Credit Markets: A Trade-off Between Consumer Protection and Credit Access," Working Paper. [1375]
- DEBBAUT, PETER, ANDRA GHENT, AND MARIANNA KUDLYAK (2016): "The CARD Act and Young Borrowers: The Effects and the Affected," *Journal of Money, Credit and Banking*, 48 (7), 1495–1513. [1374]
- DEMPSEY, KYLE, AND MIGUEL FARIA-E-CASTRO (2021): "A Quantitative Theory of Relationship Lending," Working Paper. [1375]
- DORASZELSKI, ULRICH, GREGORY LEWIS, AND ARIEL PAKES (2018): "Just Starting out: Learning and Equilibrium in a New Market," *American Economic Review*, 108 (3), 565–615. [1399]
- DUBÉ, JEAN-PIERRE, AND SANJOG MISRA (2023): "Personalized Pricing and Consumer Welfare," *Journal of Political Economy*, 131 (1), 131–189. [1375,1404]
- EGAN, MARK, ALI HORTAÇSU, AND GREGOR MATVOS (2017): "Deposit Competition and Financial Fragility: Evidence From the US Banking Sector," *American Economic Review*, 107. [1375]
- EINAV, LIRAN, AMY FINKELSTEIN, AND MARK R. CULLEN (2010): "Estimating Welfare in Insurance Markets Using Variation in Prices," *The Quarterly Journal of Economics*, 125 (3), 877–921. [1381,1384,1402]
- EINAV, LIRAN, AMY FINKELSTEIN, AND PAUL SCHRIMPF (2010): "Optimal Mandates and the Welfare Cost of Asymmetric Information: Evidence From the uk Annuity Market," *Econometrica*, 78 (3), 1031–1092. [1381,1384,1388]
- FINKELSTEIN, AMY, AND JAMES POTERBA (2004): "Adverse Selection in Insurance Markets: Policyholder Evidence From the UK Annuity Market," *Journal of Political Economy*, 112 (1), 183–208. [1388]

- FULFORD, SCOTT L. (2015): "How Important Is Variability in Consumer Credit Limits?" *Journal of Monetary Economics*, 72, 42–63. [1374]
- GABAIX, XAVIER, AND DAVID LAIBSON (2006): "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets," *The Quarterly Journal of Economics*, 121 (2), 505–540. [1406]
- GALENIANOS, MANOLIS, AND ALESSANDRO GAVAZZA (2022): "Regulatory Intervention in Consumer Financial Markets: The Case of Credit Cards," *Journal of the European Economic Association*, 20 (5), 1897–1932. [1374,1375]
- GANONG, PETER, AND PASCAL NOEL (2020): "Liquidity vs. Wealth in Household Debt Obligations: Evidence From Housing Policy in the Great Recession," *American Economic Review*, 110 (10), 3100–3138. [1384]
- GATHERGOOD, JOHN, NEALE MAHONEY, NEIL STEWART, AND JÖRG WEBER (2019): "How Do Individuals Repay Their Debt? The Balance-Matching Heuristic," *American Economic Review*, 109 (3), 844–875. [1374]
- GRODZICKI, DANIEL (2012): "The Evolution of Competition in the Credit Card Market," Working paper, Stanford University, Department of Economics, Palo Alto, Calif. [1387]
- (2014): "Competition and Customer Acquisition in the US Credit Card Market," Working paper. [1397]
- GROSS, DAVID B., AND NICHOLAS S. SOULELES (2002): "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence From Credit Card Data," *The Quarterly Journal of Economics*, 117 (1), 149–185. [1374]
- GROSS, TAL, MATTHEW J. NOTOWIDIGDO, AND JIALAN WANG (2020): "The Marginal Propensity to Consume Over the Business Cycle," *American Economic Journal: Macroeconomics*, 12 (2), 351–384. [1374]
- GRUNEWALD, ANDREAS, JONATHAN A. LANNING, DAVID C. LOW, AND TOBIAS SALZ (2023): "Auto Dealer Loan Intermediation: Consumer Behavior and Competitive Effects," Tech. rep, National Bureau of Economic Research. [1375,1404]
- GUISSO, LUIGI, PAOLA SAPIENZA, AND LUIGI ZINGALES (2013): "The Determinants of Attitudes Toward Strategic Default on Mortgages," *The Journal of Finance*, 68 (4), 1473–1515. [1384]
- HAN, SONG, BENJAMIN J. KEYS, AND GENG LI (2017): "Unsecured Credit Supply, Credit Cycles, and Regulation," *The Review of Financial Studies*, 31 (3), 1184–1217. [1374,1385]
- HANDEL, BEN, IGAL HENDEL, AND MICHAEL D. WHINSTON (2015): "Equilibria in Health Exchanges: Adverse Selection versus Reclassification Risk," *Econometrica*, 83 (4), 1261–1313. [1404]
- HANDEL, BENJAMIN R. (2013): "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts," *The American Economic Review*, 103 (7), 2643–2682. [1374,1397]
- HEIDHUES, PAUL, AND BOTON KOSZEGI (2010): "Exploiting Naivete About Self-Control in the Credit Market," *The American Economic Review*, 100 (5), 2279–2303. [1374]
- HEIDHUES, PAUL, AND BOTOND KÓSZEGI (2015): "On the Welfare Costs of Naiveté in the US Credit-Card Market," *Review of Industrial Organization*, 47 (3), 341–354. [1374]
- HERKENHOFF, KYLE F., AND GAJENDRAN RAVEENDRANATHAN (2024): "Who Bears the Welfare Costs of Monopoly? The Case of the Credit Card Industry," *Review of Economic Studies*, rdae098. [1374]
- HONG, SUTING, ROBERT M. HUNT, AND KONSTANTINOS SERFES (2023): "Dynamic Pricing of Credit Cards and the Effects of Regulation," *Journal of Financial Services Research*, 64 (1), 81–131. [1374]
- ILLANES, GASTÓN (2017): "Switching Costs in Pension Plan Choice," Working paper. [1374,1397]
- IMBENS, GUIDO W., AND WHITNEY K. NEWAY (2009): "Identification and Estimation of Triangular Simultaneous Equations Models Without Additivity," *Econometrica*, 77 (5), 1481–1512. [1384]
- INDARTE, SASHA (2023): "Moral Hazard Versus Liquidity in Household Bankruptcy," *The Journal of Finance*, 78 (5), 2421–2464. [1384]
- JAMBULAPATI, VIKRAM, AND JOANNA STAVINS (2014): "Credit CARD Act of 2009: What Did Banks Do?" *Journal of Banking & Finance*, 46, 21–30. [1374]
- JIANG, ERICA XUEWEI (2023): "Financing Competitors: Shadow Banks' Funding and Mortgage Market Competition," *The Review of Financial Studies*, 36 (10), 3861–3905. [1375]
- KEYS, BENJAMIN J., AND JIALAN WANG (2019): "Minimum Payments and Debt Paydown in Consumer Credit Cards," *Journal of Financial Economics*, 131 (3), 528–548. [1374]
- KUCHLER, THERESA, AND MICHAELA PAGEL (2021): "Sticking to Your Plan: The Role of Present Bias for Credit Card Paydown," *Journal of Financial Economics*, 139 (2), 359–388. [1374]
- MAHONEY, NEALE, AND E. GLEN WEYL (2014): "Imperfect Competition in Selection Markets," *Review of Economics and Statistics*. [1384]
- MEIER, STEPHAN, AND CHARLES SPRENGER (2010): "Present-Biased Preferences and Credit Card Borrowing," *American Economic Journal: Applied Economics*, 2 (1), 193–210. [1374]
- MENARD, SCOTT (2000): "Coefficients of Determination for Multiple Logistic Regression Analysis," *The American Statistician*, 54 (1), 17–24. [1389]

- NELSON, SCOTT T. (2025): “Supplement to ‘Private Information and Price Regulation in the US Credit Card Market’,” *Econometrica Supplemental Material*, 93, <https://doi.org/10.3982/ECTA18063>. [1375]
- PETERSEN, MITCHELL A., AND RAGHURAM G. RAJAN (1995): “The Effect of Credit Market Competition on Lending Relationships,” *The Quarterly Journal of Economics*, 110 (2), 407–443. [1371]
- PINHEIRO, TIAGO, JOSHUA RONEN et al. (2016): “Unintended Consequences of the Credit Card Act,” *Journal of Law, Finance, and Accounting*, 1 (1), 93–138. [1374]
- RHODES, ANDREW, AND JIDONG ZHOU (2024): “Personalized Pricing and Competition,” *American Economic Review*, 114 (7), 2141–2170. [1375,1404]
- ROBLES-GARCIA, CLAUDIA (2020): “Competition and Incentives in Mortgage Markets: The Role of Brokers,” Working Paper. [1375]
- RU, HONG, AND ANTOINETTE SCHOAR (2016): “Do Credit Card Companies Screen for Behavioral Biases?” Working paper, National Bureau of Economic Research. [1374]
- RUST, JOHN (1994): “Structural Estimation of Markov Decision Processes,” *Handbook of Econometrics*, 4, 3081–3143. [1381]
- RYAN, STEPHEN P. (2012): “The Costs of Environmental Regulation in a Concentrated Industry,” *Econometrica*, 80 (3), 1019–1061. [1399]
- STAVINS, JOANNA (1996): “Can Demand Elasticities Explain Sticky Credit Card Rates?” *New England Economic Review*, 43–55. [1374]
- STOLE, LARS A. (2007): “Price Discrimination and Competition,” *Handbook of Industrial Organization*, 3, 2221–2299. [1398]
- TRAIN, KENNETH E. (2009): *Discrete Choice Methods With Simulation*. Cambridge University Press. [1384]
- VILLAS-BOAS, J. MIGUEL (1999): “Dynamic Competition With Customer Recognition,” *The Rand Journal of Economics*, 604–631. [1387]

Co-editors Aviv Nevo and Guido Imbens handled this manuscript.

Manuscript received 30 January, 2020; final version accepted 20 December, 2024; available online 3 March, 2025.

The replication package for this paper is available at <https://doi.org/10.5281/zenodo.14025312>. The authors were granted an exemption to publish their data because either access to the data is restricted or the authors do not have the right to republish them. However, the authors included in the package a simulated or synthetic dataset that allows running their codes. The Journal checked the synthetic/simulated data and the codes for their ability to generate all tables and figures in the paper and approved online appendices. However, the synthetic/simulated data are not designed to reproduce the same results.