

Private Information and Price Regulation in the US Credit Card Market

Scott T. Nelson *

January 26, 2018

Job Market Paper
Latest version available [here](#)

Abstract

The 2009 CARD Act limited credit card lenders' ability to raise borrowers' interest rates on the basis of information learned during lending relationships. This paper estimates the efficiency and distributional effects of these restrictions using account-level data from a near-universe of US credit cards. I show that two forces drove these restrictions' effects. First, I show that the Act constrained lenders from adjusting interest rates after learning new information about default risk, which I find exacerbated adverse selection among existing borrowers and caused (partial) market unraveling on new accounts. Second, the Act constrained lenders from adjusting rates in response to new information about demand, which reduced lender rents from inelastic borrowers. To study the net effect of these two forces, I estimate a model that features time-varying consumer risk, flexible correlation between risk and demand characteristics, and differentiated lenders who acquire private information about borrowers over time. When I impose the CARD Act's pricing restrictions in the model, I find that equilibrium market unraveling is most severe for subprime consumers, but lower lender rents are important throughout the market, so that on net, the Act's restrictions allow consumers of all credit scores to capture higher surplus on average. Total surplus inclusive of firm profits rises among prime consumers, whereas gains in subprime consumer surplus are greatest among borrowers who were recently prime.

*Nelson: Department of Economics, MIT (nelsons@mit.edu), and Office of Research, Consumer Financial Protection Bureau (scott.nelson@cfpb.gov). The views expressed herein are those of the author and do not necessarily reflect those of the Consumer Financial Protection Bureau or the United States. For their generous help at all stages of this project, I am deeply indebted to my dissertation committee: 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, Tom Conkling, Amy Finkelstein, Daniel Green, Daniel Grodzicki, Luu Nguyen, Paul Rothstein, Richard Schmalensee, David Silberman, Daniel Waldinger, and Mike Whinston. I am also grateful to Shaista Ahmed, Michelle Kambara, Joe Remy, and Stefano Sciolli for their support and guidance at the CFPB.

1 Introduction

Lenders typically learn new information about their borrowers over time. What are the consequences of regulation that restricts how lenders use such information for loan pricing? 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 lenders' ability to raise credit card borrowers' interest rates over time and also restricted fees that could otherwise substitute for such interest rate increases. Lenders therefore became substantially less able to adjust borrowers' interest rates and other price dimensions in response to learning new information about borrowers.

Understanding the effects of the CARD Act's price restrictions is important both because of these restrictions' economic interest and because of the credit card market's central role in the US consumer credit landscape. Among the estimated 85 million US households with credit cards, roughly 60% use credit cards for at least occasional borrowing, and credit card holders collectively have access to over \$3 trillion in open credit lines. Reliance on credit cards for borrowing is especially pronounced for less credit-worthy consumers, among whom the share of accounts used for at least occasional borrowing exceeds 90%. Credit card regulation is therefore important both for its distributional effects as well as for its implications for the efficient provision of consumer credit more generally.

In this paper, I quantify the distributional and efficiency consequences of the CARD Act's price restrictions. To understand these effects, I analyze two channels through which informational restrictions on pricing can influence credit market outcomes. First, if lenders learn over time about borrower demand, the CARD Act's price restrictions may limit lenders' ability to extract rents from inelastic borrowers. Second, such restrictions may also limit lenders' ability to adjust prices for risk, and the CARD Act may therefore exacerbate adverse selection and induce either partial or complete market unraveling. The interplay of these two channels 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 datasets. The first contains monthly account-level data from the near-universe of US credit card accounts, spanning the period before and after the CARD Act. 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 dataset is a large, randomly sampled panel of US consumer credit reports, also spanning the period before and after the CARD Act. These credit report data reveal patterns that cannot be measured in the account-level data – for example, which consumers are *not* credit card holders at any 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 inter-quartile range of interest rates) on new cohorts of mature accounts dropped immediately by approximately 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.

These results illustrate the complexity of assessing the CARD Act’s equilibrium effects. Restrictions on increasing interest rates may bring lower prices to some borrowers, while other borrowers’ prices will rise as they are pooled with their peers. At the same time, these relative price effects may change the composition of borrowers in the market, further shifting how lenders set rates. Tracing these effects is further complicated by the large number of contemporaneous shocks that were affecting consumer credit markets when the Act took effect in 2009. Moreover, the Act contained many *non*-price regulations as well, including additional disclosure requirements, simpler billing procedures, and “nudges” for borrowers to repay their balances.

With these empirical features in mind, I develop and estimate a detailed structural model of the credit card market to use as a tool for studying the CARD Act’s price restrictions’ effects. I estimate the model on the *pre*-CARD-Act equilibrium observed in the market. I then impose the CARD Act’s price restrictions in the model and analyze their effects for different types of consumers and for total welfare overall. Consequently, this exercise speaks to how the market re-equilibrates in the presence of the CARD Act’s price restrictions in isolation from other coincident shocks in consumer credit markets as well as other, non-price regulations contained in the Act.

In building the model, I begin with a pair of reduced-form analyses of the key forces driving the CARD Act price restrictions’ effects. The first of these analyses shows that the Act changed how the credit card market prices risk, and that these changes led to the adverse retention of risky borrowers over time. I show that prior to the CARD Act, interest rates were strongly responsive to changes in risk, as observed through changes in credit score after origination. In fact, the price gradient of these interest rate changes (as measured in interest rate basis points per point of credit score) was nearly identical to the price gradient of risk observable at the time of origination. In that sense, there was a single average price of risk in the market, which applied equally to risk at origination and risk that emerged over time. In contrast, I find that after the CARD Act, interest rates were less responsive to changes in risk, so that a sizable gap emerged between these two gradients. Newly emergent risk became nearly 75% cheaper for a borrower, per FICO score point, than risk observable at origination. Examining how these relative price effects changed the selection of consumers into and out of borrowing, I estimate that for every one percentage point reduction in interest rates charged to newly risky borrowers, these borrowers responded with a 0.7 percentage point decrease in quarterly attrition rates – a sizable effect given that average attrition rates range from 10 to 15% per quarter.

The second of these reduced-form analyses highlights that the Act also restricted lenders from adjusting interest rates in response to new information about borrowers’ price sensitivity. I find that two of the borrower behaviors that most commonly triggered interest rate increases – late payments of less than 30 days, and transactions in excess of a borrower’s credit limit – were used as signals that borrowers were price inelastic, and that lenders then levied price increases on these accounts to earn *higher* returns than they earned on other, identically risky accounts. In contrast, after the introduction of the CARD Act’s restrictions, lenders’ excess returns on these accounts were either eliminated or sharply reduced, depending on the signal and the credit-worthiness of the borrower. The Act thus made it difficult for lenders to extract rents in response to signals of relative price inelasticity, leading to a decline in rents from inelastic borrowers.

A reduced-form decomposition shows that signals of borrower price inelasticity drove the ma-

jority of interest rate increases on prime accounts, while the majority of interest rate increases on subprime accounts were in response to signals of borrower default risk. A similar decomposition holds for fee revenue. This decomposition suggests the possibility that the CARD Act’s price restrictions mostly led to lower lender rents among prime accounts, and mostly exacerbated information problems through unpriced risk among subprime accounts. However, caution is warranted in relying only on this reduced-form decomposition: since consumers’ credit scores change over time, information asymmetries on subprime accounts can affect prime borrowers’ rates, and even a small amount of unpriced risk can lead to severe market unraveling. This further motivates my use of a model that can predict how the market re-equilibrates overall in order to study these restrictions’ effects.

The structural model features consumers with time-varying risk, differentiated lenders who acquire private information about borrowers over time, and flexible correlation between borrower risk and demand. 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 – occasional, portfolio-wide repricing by certain lenders – to estimate borrowers’ sensitivities to price. I find that riskier borrowers are less price elastic, consistent with the market being adversely selected. I also provide estimates of the extent to which lenders possess private information about their borrowers’ preferences and risk. I find that such private information plays an important role in the credit card market, as my measure of lender private information is nearly as predictive of subsequent default (in per-standard-deviation terms) as borrower credit scores. Other estimates on the demand side of the model indicate that consumers’ set-up costs for opening new credit card accounts are relatively high, consistent with only a subset of consumers taking advantage of promotional or “teaser” interest rates by refinancing balances with new credit cards. Finally, on the supply side of the market, the estimates of lender costs recovered from first-order conditions in the model match closely to industry reports of these costs – for example, the cost of marketing and customer acquisition for new credit card accounts.

After thus estimating the model on the observed pre-CARD-Act equilibrium, I impose the CARD Act’s price restrictions in the model and study how the market responds. Specifically, I study the new equilibrium that emerges when lenders best-reply to each other under a new regulatory regime that does not allow them to change a borrower’s price of borrowing over time, except through promotional or “teaser” rates that were still allowed under the Act.

The results of this exercise reveal a number of interrelated effects of the CARD Act’s price restrictions. First, induced adverse selection is most severe among subprime accounts, while reduced lender rents are evident throughout the market. These patterns are consistent with the results of the earlier reduced-form decomposition. On net, average transacted prices fall throughout the market. This reflects, in part, that many borrowers who previously could access the cheapest credit within their credit score segment face higher prices and exit. Therefore the higher prices they would have faced are censored from transacted prices. This type of partial unraveling is especially pronounced among subprime consumers. Nonetheless, given the importance of lower prices for consumers with the strongest demand for credit, consumer surplus rises throughout the market. Among subprime consumers, the rise in consumer surplus is mostly offset by a fall in lender profits; among prime consumers, total surplus rises. Much of this surplus gain is due to the insurance value of these restrictions for consumers whose credit scores deteriorate over time. While this insurance is most relevant for prime borrowers, it 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 a number of contributions relative to existing literature. Agarwal et al. (2014) also study how the Act affected credit card pricing and find that a measure of the average, fee-inclusive cost of borrowing fell on consumer credit cards after the Act, among the set of consumers who still chose to borrow on credit cards. The primary differences between their study and this paper is that I focus on the mechanisms of how dynamic pricing regulation can affect equilibrium prices by restricting what information is possible to price; I study rich heterogeneity of price effects across borrower types, including borrowers who chose to leave the market after the Act; and I isolate the effects of the Act’s price regulation separately from other omnibus regulation included in the Act, for example restrictions that aimed to make interest charge calculations and due dates more transparent for consumers. Additionally, I use a starkly different empirical strategy: while Agarwal et al. (2014) use business credit cards as a comparison group in a difference-in-differences strategy, I use a combination of reduced-form analyses that leverage cross-account comparisons, and a detailed structural model to help study the Act’s effects in isolation from other contemporaneous shocks in financial markets. Other research on the CARD Act includes Keys and Wang (2016), who study the Act’s “nudges” for borrowers to pay more than their minimum required payment each month, Jambulapati and Stavins (2014) and Santucci (2015), who describe patterns of account closures and credit line changes coinciding with the Act and the Great Recession, Debbaut et al. (2016), who focus on the Act’s particular restrictions to protect young borrowers, and Han et al. (2015), who compare credit cards’ with other financial products’ direct-mail offers 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.¹

This paper also joins a long literature examining the competitiveness of, and sources of market power in, the credit card industry. After seminal work by Ausubel (1991) showed credit card lenders tended not to pass through changes in the cost of funds to their borrowers,² a number of papers explored whether and why the industry may be imperfectly competitive, including for reasons of search costs (Berlin and Mester (2004)), consumer irrationality (Brito and Hartley (1995)),³ and adverse selection for firms that cut prices (Stavins (1996)). My work integrates many of these potential sources of market power in a single model – including switching costs across firms, adverse selection, as well as lender private information – and provides an estimation framework that helps identify the relative importance of each of these. My results on the particular importance of switching costs across firms join a growing recent literature on the importance of switching costs in selection markets, including Handel (2013) and Illanes (2016).

¹There is also a small body of theoretical work focused on the CARD Act’s price restrictions in particular, including Hunt and Serfes (2013) and Pinheiro and Ronen (2015), who present theoretical models of the effects of repricing restrictions, and some research on restrictions to credit card interest rate increases in the law literature (Levitin (2011) and Bar-Gill and Bubb (2011)).

²See Grodzicki (2012) for evidence on how the patterns identified in Ausubel (1991) have become less pronounced in more recent data.

³Research on behavioral consumers in the credit card market has remained quite active, including work by Angeletos et al. (2001), DellaVigna and Malmendier (2004), Grubb (2009), Heidhues and Koszegi (2010), Heidhues and Kőszegi (2015), and Ru and Schoar (2016). Related work focuses on how consumers learn over time how to avoid apparent mistakes with credit cards (Agarwal et al. (2008), Agarwal et al. (2009)), and how the probability of mistakes also falls as consumers face higher stakes, e.g. higher balances borrowed (Agarwal et al. (2015a)).

I also provide new evidence on consumer demand for credit card borrowing and how consumers respond to changes in their terms of credit. 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 et al. (2015b), and Gross et al. (2016)), and how credit limits affect consumers’ holdings of cash on hand (Telyukova and Wright (2008) and Fulford (2015)). In contrast to this work on credit limits, research on how borrowers respond to interest rates and fees has been more limited.⁴ To help fill this gap, I estimate borrower price elasticities across a range of borrower risk types, and also estimate primitives of a rich demand model – including switching costs, liquidity costs, and disutility from price – that predict how price elasticities change non-locally as pricing changes.⁵ Estimates of these primitives help not just for understanding the CARD Act’s price restrictions, but for other applied work in the credit card market as well.

This paper is organized as follows. In Section 2, I provide background on the credit card market, the CARD Act and the two datasets that I use in my analysis. I also present summary statistics from these datasets to highlight key changes in the credit card market around the implementation of the Act. In Section 3, I report reduced-form 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 from 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

2.1 Institutional Background

2.1.1 The Credit Card Industry

Credit cards are well known as a means of transaction. For many households they are also an important source of credit. Credit cards provide over \$3 trillion in open credit lines for unsecured borrowing, and survey estimates suggest that roughly 60% of US households that hold credit cards actively use credit cards to borrow, i.e., do not pay their balance due in full and hence incur interest charges (Bricker et al. (2017)).⁶ The importance of credit cards as a source of credit is especially strong among less credit-worthy consumers, where the prevalence of at least occasional borrowing rises to roughly 85% among accounts held by near-prime consumers and over 95% for subprime consumers.

The credit card market was also relatively unregulated in the period prior to the CARD Act. After US Supreme Court cases in 1978 and 1996 curtailed state regulation of credit card interest rates and fees (Evans and Schmalensee (2005), Hyman (2011)), credit card lending became concentrated among large, national banks that faced few restrictions on pricing strategies or

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 pre-scheduled 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 (Alexandrov et al. (2017)), which, as I detail in Section 3.3, can arise endogenously as lenders respond to shifts in individual borrowers’ risk or demand.

⁵Other modeling work specific to the credit card market includes Drozd and Serrano-Padial (2014).

⁶The account-level administrative data I study in this paper corroborate this survey evidence, as I find that 70% of active credit card accounts are used for borrowing in at least three months of the year.

the terms of credit offered to borrowers (Mandel (1990)). Simultaneously, advances in credit scoring and computing power increased the sophistication of pricing and underwriting, with prices becoming tailored to borrowers' individual risk, price sensitivity, and even shopping habits (Edelberg (2006), FRB (2010)).

Prior to the CARD Act, lenders' pricing strategies rested on two main sources of information. One is consumer credit bureaus, which collect data on consumer borrowing history across a wide range of loan products and then use these data to predict consumers' likelihood of future default. The bureaus transform these predicted default likelihoods to a more familiar credit score on an integer scale with higher numbers corresponding to safer borrowers; one common example is a FICO score.⁷ These scores and the underlying data are sold to credit card issuers to prospect and underwrite new accounts and also to monitor risk on mature accounts.⁸ Because this information is typically available to all firms in the market, this information is best thought of as *public* information for the purposes of studying firm behavior.

The second key source of information for a lender is a consumer's own behavior with a credit card after origination. Much of this information is private for the lender because it is not reported to consumer credit bureaus and is not otherwise observable to competitors, including a consumer's purchase volume, shopping behavior, prevalence of borrowing, repayment rates, and monthly payment timing. For some consumers lenders may receive additional private information as well. For example, consumers may signal their riskiness through interactions with call center representatives – say, explaining an idiosyncratic reason for a late payment when requesting a late fee to be forgiven – or through additional information provided when requesting a credit limit increase, such as updated employment and income information. This private information is generally learned through a relationship with a borrower *after* origination.

Prior to the CARD Act, lenders could use a number of price dimensions to respond to new information learned after origination. First an account's interest rate for borrowing – which in the credit card market is represented as an annual percentage rate (APR)⁹ – could change “at any time for any reason” according to stock language included in nearly all credit card contracts.¹⁰ Credit card contracts also typically delineated a set of “triggers,” such as late payments and over-limit transactions, that would cause the card issuer to consider an interest rate increase. Roughly 52% of borrowers in pre-CARD-Act data experienced a discretionary increase in their card's interest rate over the course of a year, with about half of these increases coinciding with behaviors typically specified as repricing triggers.¹¹ Thus lenders found it optimal to upwardly

⁷Further information on the contents and uses of credit report data is provided in Section 2.2.2.

⁸See Grodzicki (2014) for a discussion of the information that credit card issuers use in prospecting new accounts.

⁹The APR concept was developed by the Truth in Lending Act (TILA) rather than by industry. TILA's implementing regulation specifies that the APR is “determined by multiplying the unit-period rate by the number of unit-periods in a year,” so APRs are annualized without compounding even though credit card interest typically compounds monthly. See 12 CFR Part 1026.

¹⁰See ConsumerAction (2007) for details on the prevalence of these any-time-any-reason terms. Examples include “All terms, including the APRs and fees...may change based on information in your credit report, market conditions, business strategies, or for any reason”, and “We have the right to change the rates, fees, and terms at any time, for any reason...These reasons may also include competitive or market-related factors.”, and “APRs may change to higher APRs, fixed APRs may change to variable APRs, or variable APRs may change to fixed APRs. We may change the terms (including APRs) at any time for any reason.”

¹¹In addition to such discretionary interest rate increases, 36% of borrowers saw a promotional interest rate either introduced or expired during the year. The expiration of a promotional rate differs from a discretionary change in interest rate because it is pre-scheduled at the time the promotion is introduced. Promotional rates are especially common at the time of origination, and hence are often referred to as introductory or “teaser”

reprice the interest rate on many, but not all, borrowers as new information arrived over the course of lending to a consumer.

In addition to these interest rate repricings, credit card pricing also responded to borrower behavior through behavior-contingent fees, such as fees for late payments or over-limit transactions. For an average account prior to the CARD Act, revenue from these fees was 32% as large as interest charges, and on subprime accounts it was 46% of interest charge revenue.¹²

The responsiveness of credit card pricing to borrower behavior became an important motivation for the CARD Act, as consumer advocates and policy-makers both saw an inherent “unfairness”¹³ in price increases that targeted some borrowers rather than others. As I detail in the following section, what emerged from policy debates around the CARD Act were strong restrictions on contingent pricing, i.e. pricing that depended on what lenders learned about borrowers over time, and very limited restrictions on pricing based on information available to lenders at the time of account origination.

2.1.2 The Credit CARD Act

Much of the policy debate around the CARD Act focused on the responsiveness of credit card pricing to borrower behavior. One perspective emphasized that discretionary interest rate repricing and contingent fees could “opportunistically” raise the cost of borrowing for consumers with the most pronounced demand for credit, in effect, extracting rents from those consumers with price-inelastic demand (Levitin (2011)). At the other end of the debate, industry advocates highlighted the importance of raising prices on borrowers revealed to be riskier than expected, so as not to instead make safer borrowers bear the cost of this risk (ABA (2013)).

Ultimately the Act did place strong restrictions on how credit card pricing responds to borrower behavior. First, discretionary increases in interest rates on outstanding balances were almost completely eliminated; the one major exception that was allowed to lenders has, in practice, proved to be an exception lenders rarely choose to use.¹⁴ Second, over-limit fees were one of the most common contingent fees prior to the CARD Act and were likewise almost completely eliminated.¹⁵ Third, the other most commonly used contingent fee, late fees, were effectively

rates. Prior to the CARD Act, 35% of originations included some kind of promotional rate, and among accounts used for borrowing this share reached 71%. These rates were often offered below lenders’ costs; in particular I estimate that 81% of promotional rates were a 0% APR, allowing interest-free borrowing. The profitability of these “teaser” contracts rested on borrowers continuing to borrow after the promotional period ended, and indeed I estimate that 86% of consumers who borrowed during an introductory promotion were still borrowing on their card three months after the end of their promotion. (In principle, promotional rates could also be profitable if borrowers incurred sufficient fees during their interest-period. However, I estimate that fee revenue on cards with promotional balances was only 1.24% annualized as a share of those cards’ balances – insufficient even to cover lenders’ costs of funds for much of the pre-CARD-Act period.)

¹²All major categories of fees were contingent on one or more borrower behaviors revealed after origination, with the exception of annual fees, which made up less than 10% of all fee revenue in pre-CARD-Act data.

¹³Fairness, while not defined in the CARD Act, is invoked both in the Act’s preamble and five separate times in the text of the Act itself.

¹⁴This exception allows for the upward repricing of balances on accounts that are 60 or more days delinquent. In the debates leading up to the CARD Act, industry commentators presented evidence that repricing at this point of delinquency would not be profitable, as such balances are already at high risk of default (FRB (2008)); subsequent experience has borne this out, and lenders today rarely reprice balances that are 60 days late despite being allowed to do so (see Figure 1).

¹⁵While in principle these fees were still allowed if borrowers opt-in to allow these fees, they have virtually disappeared from the market (see Figure 1).

capped by a safe-harbor ceiling of \$25 (or \$39 for subsequent incidences within 6 months).¹⁶ On net, these restrictions strongly restricted lenders from adjusting prices in response to information revealed through borrower behavior over time, while placing little to no restriction on the interest rate set on the account at the time of origination.¹⁷

These interest rate repricing restrictions and over-limit fee restrictions took effect in February 2010 and late fee restrictions take effect in August 2010.¹⁸ These implementation dates followed after a compressed period of policy debate surrounding the Act’s passage. First in December 2008, as a precursor the Act the Federal Reserve issued a rule (originally scheduled to take effect in July 2010) that would have implemented a weaker version of the CARD Act interest rate repricing restrictions and fee restrictions. The CARD Act, introduced in Congress a month later in January 2009, superseded these restrictions and strengthened them to their present form. The Act was then passed and signed into law several months later in May 2009.

Given the Act’s staggered congressional debates, passage, and implementation, I for much of my analysis will focus on a pre-CARD-Act period stretching from July 2008 through June 2009, and a post-CARD-Act period from July 2011 to June 2014. I focus on these full-12-month periods, both beginning in July, in order to avoid overemphasizing any seasonality, such as holiday consumption and subsequent debt repayment timed to the receipt of tax refund payments, that would appear in some months and not in others.

2.2 Data Sources and Summary Statistics

I use two main datasets in my analysis. One dataset contains the near-universe of US credit card accounts in a monthly account-level panel. The second dataset is a large random sample of consumer credit reports, showing all credit cards and other non-credit-card-loans held by a panel of consumers over time. Both are anonymized, administrative datasets furnished by industry and maintained by the Consumer Financial Protection Bureau (CFPB).¹⁹ In this section I introduce both datasets and present summary statistics that highlight key dynamics in the credit card industry before and after the CARD Act.

¹⁶Other pricing restrictions, which affect, for example, the number of fees that can be charged simultaneously or near the time of account origination, are detailed in [CFPB \(2013\)](#).

¹⁷Besides these price restrictions, the CARD Act also included a bevy of restrictions that sought to make credit card borrowing more predictable and transparent for borrowers. Lenders were banned, for example, from changing borrowers’ statement due dates from month to month, or from imposing a cutoff time on due dates that came before 5 PM. Lenders were also required to include additional information on account statements that emphasized how long it would take to pay off a balance at various monthly payment sizes. Changes in account terms were also required to be disclosed to borrowers with 45 days of advance warning rather than the previous 15 day limit. A full review of these restrictions is available in [CFPB \(2013\)](#). However, industry trade association statements suggest the most important part of the Act from industry’s perspective was the restriction on interest rate increases: the American Bankers Association referred to these restrictions as “the core, most important provision of the CARD Act” ([ABA \(2013\)](#)).

¹⁸A limited number of other provisions, including the requirement of earlier disclosure for account changes, took effect soon after the Act’s passage, in mid-2009.

¹⁹Consistent with the CFPB’s confidentiality rules, this paper only presents results that are sufficiently aggregated so as to not identify any specific individuals or institutions. Additionally, the data used contain no direct consumer identifiers.

2.2.1 CCDB Account-Level Dataset

The first dataset I use is the CFPB’s Credit Card Database (CCDB), a near-universe of de-identified credit card account data in a monthly panel from 2008 to present. The data include all open credit card accounts held by 17 to 19 large and midsize credit card issuers under the supervisory authority of either the OCC or the CFPB, which together cover roughly 90% of outstanding general-purpose US credit card balances.²⁰ 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 by type of interest or fee, payment due dates and delinquencies. The dataset also includes some fields that are maintained by the lender but not always included on account statements, such as the consumer’s current FICO score and a flag for whether the account holder keeps other accounts with the same bank, for example a mortgage. These same data fields are typically used by lenders for day-to-day account management.²¹

These data represent a superset of the credit card data used in Agarwal et al. (2014) and Agarwal et al. (2015b), including 9 to 10 additional midsize issuers that cover an additional 17% to 23% of outstanding balances.²² The advantage of using this superset is the inclusion of a more diverse set of firms, especially issuers with relatively concentrated market shares in important submarkets such as subprime or super-prime accounts. While these data are relatively new to academic research, they have been used previously in Keys and Wang (2016), Gross et al. (2016), and Alexandrov et al. (2017),²³ as well as several CFPB market-monitoring publications (CFPB (2013), CFPB (2015)).

More generally, an advantage of using these data is the ability to study an entire *industry’s* behavior under different regulatory regimes using detailed account-level data. Large sample sizes – hundreds of millions of panel observations from credit cards actively used for borrowing in the pre-CARD-Act period, for example – make it possible to estimate rich heterogeneity in borrower demand characteristics and to study how these demand characteristics correlate with default risk, even among borrower types for whom ex-post default is rare. My use of account-level data for this purpose in many ways follows the example of Einav et al. (2012), who illustrate in the related context of subprime auto lending how account-level data can be used to estimate a rich model of credit demand where demand characteristics covary with risk.

For reasons of panel balance and data availability, I restrict my analysis to a subset of CCDB lenders that hold over 88% of all credit card balances observed in the CCDB in 2008-2009. This subset includes all of the issuers studied previously in Agarwal et al. (2014) and several additional issuers, including a large issuer with relative specialization in prime and super-prime lending. Given the presence of some mid-size and regionally-focused issuers in this sample, I also pool data from the smallest issuers into a single “fringe” issuer, as in Somaini (2011), when estimating my model.

²⁰A total of 6 lenders enter or exit at some point in the sample period.

²¹See Trench et al. (2003) for one relevant industry study on this front.

²²The inclusion of detailed data from such a wide range of firms requires considerable effort to regularize the data in cases where firms use different reporting conventions, variable field definitions, or timing conventions. Full details of this regularization work is included in an appendix.

²³Respectively these papers study the CARD Act’s “nudges” for borrowers to pay more than their minimum payments each month, propensities to consume out of changes in credit limits, and the responsiveness of balance size and late payments to interest and fees.

2.2.2 CCP Borrower-Level Dataset

The second database I use is the CFPB’s Consumer Credit Panel (CCP), a large, randomly sampled panel of consumer credit reports showing all credit card accounts and other non-credit-card loans for a set of anonymized consumers over time. The non-credit card loans in these data include mortgages, auto loans, student loans, lines of credit, and installment loans held by a given consumer. The data also include non-loan items such as a measure of past loan applications, defaulted debts in collection, and public records such as bankruptcies.²⁴

The panel is a 1-in-48 random sample, drawn from one of the three nationwide consumer credit reporting agencies.²⁵ This panel is observed quarterly beginning in 2004, with additional observations at an annual frequency from 2001 to 2004.²⁶ The CCP therefore has the advantages of showing a large representative sample of consumers, following these consumers over a longer time frame than is available in the CCDB, and reporting all credit card and non-credit-card accounts for a given consumer. The CFPB CCP data have been used previously in Brevoort and Kambara (2015), Brevoort et al. (2016), and Brevoort et al. (2017).²⁷

In comparison to other credit report data often used in research, in particular the Federal Reserve Bank of New York’s Consumer Credit Panel, the CFPB CCP has the unique feature of being a *loan-level* dataset rather than a borrower-level dataset for credit card accounts. For example, the CFPB CCP shows the quarterly balance on each of a borrower’s credit cards, rather than the total balance summed across all credit cards. The availability of account-level credit report data makes it possible to study how borrowers allocate balances across multiple credit cards and other loans, and how borrower behavior evolves over time across multiple accounts. Additionally, the CCP makes it possible to study borrower entry and exit in the credit card market, as the dataset includes individuals not holding credit cards at any given point in time.

The CCP and CCDB both provide panel data on the credit card market before and after the CARD Act. The CCP has longer panel length and richer borrower-level information, and the CCDB has richer pricing information and lender-level information. Neither accounts nor account-holders can be linked between the CCDB and CCP.

2.2.3 Summary Statistics

In this subsection I use the CCP and CCDB to illustrate the mechanical effects of the Act on three specific price dimensions, to contrast these mechanical effects with the overall changes in the cost of borrowing in equilibrium before and after the Act, and to document changes in borrowing behavior that coincided with these price shifts.

Figure 1 shows the mechanical effects of the Act on three price dimensions that the Act regulated most directly: interest rate repricing, over-limit fees, and late fees. First, Panel A shows the incidence of interest rate increases on current borrowers over time. Forty-eight to fifty-four percent of borrowers experienced a discretionary interest rate increase at least once a

²⁴For further background on data included in consumer credit reports and the uses of these data, see Avery et al. (2003).

²⁵These three are Equifax, Experian and Transunion.

²⁶Additionally, the panel frequency increases to monthly in 2013, although I do not use the monthly data in this paper.

²⁷Respectively these three papers study medical collections’ predictive power for loan default, the prevalence and correlates of not having a credit report file or credit score, and the impact of Medicaid expansions on financial health.

year before the CARD Act.²⁸ The incidence of interest rate increases then dropped sharply, and nearly to zero, when the CARD Act repricing restrictions went into effect. Panel B documents a similar drop in the incidence of over-limit fees, which affected roughly 7% of accounts in an average *month* prior to the CARD Act, and then fell sharply to nearly zero when the Act’s over-limit fee restrictions went into effect. Panel C shows the drop in total late fee revenue at the time the Act’s reasonable-and-proportional late fee restrictions took effect, a decrease of roughly 40%. These three results show that the Act’s restrictions were binding on the price dimensions the Act targeted most directly, and that the Act’s restrictions affected pricing on a sizable majority of accounts.

Figure 2 next shows that these price restrictions’ implementation coincided with an immediate compression in the distribution of interest rates across accounts. The figure shows the inter-quartile range (IQR) of interest rates after controlling for origination FICO score, with one data point presented for each quarterly origination cohort.²⁹ For cohorts reaching maturity before the Act’s repricing restrictions went into effect, these IQRs are consistently equal to nearly 8 percentage points; for cohorts reaching maturity after these restrictions took effect, these IQRs fell sharply to less than 6 percentage points.³⁰ To be clear, this evidence is only an event-study analysis. However, the sharpness of this change around the time of the CARD Act’s implementation suggests that the Act, rather than other coincident changes in the credit card market, caused this fall in price dispersion.

Table 1 presents further evidence on which percentiles of the price distribution compressed and shifted. Each column of the table corresponds to a given statistic of credit card pricing (for example, the 25th percentile of interest costs), and each row highlights a different market segment (for example, borrowers with subprime FICO scores of 620-639). The statistics presented are changes in each measure from pre-CARD-Act data (2008Q3 through 2009Q2) to post-CARD-Act data (2011Q3 through 2014Q2). Effective interest rates³¹ and fee-inclusive borrowing costs³² both compressed from the pre-CARD-Act period to the post-CARD-Act period. For both price measures, the table reveals increases of several hundred basis points in the 25th percentile for

²⁸I focus here on the type of rate increases restricted by the CARD Act, namely rate increases *not* caused by the expiration of a promotional interest rate or by changes in an indexed base rate, and also rate increases not coinciding with a delinquency of 60 days or more.

²⁹I focus here on the age of accounts’ maturity, i.e., the age by which all promotional teaser rates from the time of origination have usually expired, because a substantial amount of price dispersion emerges around the time of promotional rates expiring. In order to focus on within-FICO price dispersion, the IQRs plotted in the figure are for residual borrowing costs after partialling out FICO-score fixed effects.

³⁰For evidence on price dispersion in the credit card market from a slightly earlier time period than is observable in the CCDB data, see [Stango and Zinman \(2015\)](#).

³¹The effective interest rates presented here are calculated by dividing total interest charges by the average amount borrowed, and then annualizing. This is not a fee-inclusive cost or “total” cost, but rather a measure of interest costs. Due to intricacies of how lenders assess interest, these can differ from slightly from the stated APR on the account. Additionally, several APRs may be in effect on an account at any given time, for example, one APR for a promotional balance, one APR applied to a balance accrued through a cash-advance, and another APR applied to non-promotional purchases. This measure of effective interest provides the arguably most representative average of these different APRs.

³²To calculate a measure of the fee-inclusive price of borrowing, I sum interest charges and fee revenue on a given account and divide by the amount borrowed over a given period, such as a month or quarter, and then annualize. I refer to this measure as the fee-inclusive borrowing cost or price, or average borrowing cost. This is the same price measure used previously in research on the credit card market, including by [Agarwal et al. \(2014\)](#), and is equal to the “total cost of credit” as defined by [CFPB \(2013\)](#). Although this is not a *marginal* price for an additional dollar borrowed, it is the relevant marginal cost to consider on the extensive margin of borrowing.

most prime borrowers (FICO scores at or over 660) or in the 10th percentile for most subprime borrowers (FICO scores under 660), while the 75th and 90th percentiles usually fell, sometimes on the order of hundreds of basis points, or at least rose by less than the lower tail rose.

Overall the table shows that most credit scores saw compression in the left tail of the distribution as well as the right tail, and that compression in the price distribution was most pronounced among subprime consumers. Indeed, subprime consumers saw their IQRs of effective interest rates and fee-inclusive borrowing costs both typically fall by over 500 basis points, while the very bottom of the subprime price distribution sometimes rose by over 300 basis points. This compression in the left tail of the distribution cannot be a merely mechanical effect of the CARD Act’s repricing restrictions, which only restricted interest rate *increases* after origination. Rather, this compression is suggestive of an equilibrium outcome whereby borrowers in the left tail of the price distribution faced higher prices as the CARD Act’s repricing restrictions pooled them with their peers.

Figure 3 suggests that these relative price shifts may also have changed borrowing behavior. I focus on the extensive margin of credit card borrowing, both the share of consumers who hold a credit card at all and the share of active credit card accounts used for borrowing instead of transacting.³³ The figure shows that the share of consumers who have any credit card at all fell by up to 10 percentage points in the subprime market, while the share of consumers using cards for borrowing remained broadly unchanged. On net then, there was substantial consumer exit from the credit card market in the same market segments that saw, with the passage of the Act, higher prices in the low-cost left tail of the price distribution. While these patterns are only suggestive, they help motivate my analysis of whether the Act led to partial market unraveling.

I close this section with basic summary statistics that help with understanding the credit card market in the pre-CARD-Act equilibrium. Table 2 shows various statistics of credit card pricing across its columns, while the table’s rows correspond to different market segments and the extent to which these different segments use credit cards for borrowing. The prevalence of borrowing is quite high among active accounts: 96% of credit card accounts with subprime FICO scores of 620-639 are used for borrowing at least three months of the year, and even among prime (resp. super-prime) accounts in the 720-739 (resp. 780+), the prevalence of borrowing at least three months of the year is 67% (resp. 42%). Note also that fee-inclusive prices decrease sharply across the range of FICO scores from roughly 21 percentage points annualized among the subprime accounts shown, on average, down to 10 percentage points among the super-prime accounts shown. There is also a risk gradient in the share of revenue coming from fees: at the subprime end of the market, 5 percentage points out of the total 21% average borrowing cost is generated by contingent fees such as late fees or over-limit fees, while at the super-prime end, less than 1 percentage point out of the total 10% average borrowing cost comes from fee revenue.

3 Reduced Form Evidence

In this section I show in more detail *who* faced relative price changes as a result of the Act. I show that relatively safe borrowers faced higher prices and relatively risky borrowers faced lower prices, and that this engendered a dynamic form of adverse selection whereby lenders retained

³³The share of consumers who hold a credit card at all is taken from the CCP, and the share of active accounts used for borrowing is taken from the CCDB data, both described above. Credit scores in the CCP data are non-FICO scores, but they are presented on the same axis because the two scores are designed to be similarly predictive of default, and because the two scores have the same range.

riskier borrowers over time. Consistent with partial market unraveling, lenders also set higher interest rates on average for all borrowers at origination. However I also show that in some parts of the market – especially prime accounts – the *majority* of the repricing that was restricted by the Act enabled lenders to charge higher markups over the cost risk, not just to adjust prices for risk. Lenders’ excess returns on these marked-up accounts then fell sharply or were reversed after the Act.

3.1 Risk Pricing and Adverse Selection

This subsection examines how credit card lenders price risk that is observable at the time of origination, which I term “origination risk,” and how this compares to the pricing of risk that becomes observable later, which I term “emergent risk.” The CARD Act restricted how lenders price emergent risk but not origination risk, and I show that the Act generated a gap between the pricing of these two types of risk which led to lenders’ adverse retention of riskier borrowers over time.

I first estimate the price gradient of origination risk as a linear relationship between interest rates $r_{i,0}$ 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 plot this gradient in pre-CARD-Act data as the dashed line in Figure 4 against the left and bottom axes, along with an accompanying binscatter.³⁴ There is a consistent relationship between price and risk throughout the FICO distribution: the average price of risk is roughly 32 basis points in annualized interest for every 10 FICO points of expected default risk.

I then estimate the pre-CARD-Act price gradient of emergent risk using a similar linear model, where I estimate the relationship between interest rates and *change* in FICO score since origination,

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

This regression also includes fixed effects α for origination FICO score, $\text{FICO}_{i,0}$, which are included to absorb variation in interest rates $r_{i,0}$ from the time of origination,³⁵ 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 rates expiring over time. Given the presence of 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 the same figure I then plot the estimate of β from this second regression with an accompanying binscatter. These are plotted on the opposite set of axes (right and top axes), which have the same scaling as the main axes for sake of comparability. Both plotted gradients are nearly the same: for both origination risk and emergent risk, borrowers on average face a difference

³⁴A binscatter plots the conditional mean of the dependent variable at each percentile of the regressor, helping illustrate the shape of the relationship between the two across the distribution of the data. This can also be extended to regressions with controls by first partialling out controls from both the dependent variable and the regressor. See [Stepner et al. \(2013\)](#).

³⁵This specification is equivalent to a long-differences specification in price and risk (without controls for origination risk) if the above error terms $e_{i,0}$ and $\epsilon_{i,t}$ are independent. The long-differences specification cannot be estimated directly, as $r_{i,0}$ is typically unobserved in the data for accounts originated prior to 2008. In an appendix I show robustness of this specification’s estimates to an alternative, first-differences specification, which *can* be estimated.

in price of about 30 basis points in annualized interest for every 10 FICO-point difference in risk. This points to the credit card market setting 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.

Figure 5 re-estimates both of these price gradients in post-CARD-Act data. Here there is evidence of the CARD Act’s repricing restrictions causing 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.³⁶

The gap between these gradients leads to weaker incentives for newly risky borrowers to attrite from borrowing, and likewise gives newly safe borrowers stronger incentives to attrite. I look for evidence of this type of dynamic adverse selection by estimating the relationship between borrower retention and changes in FICO score since origination, using a specification similar to equation 3.2,

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

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, $\text{FICO}_{i,0}$. The equation is again estimated at a quarterly frequency. The β coefficient therefore captures how quarterly linear-probability hazards from borrowing to non-borrowing 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 binscatters in Figure 6. The gap between the two plotted relationships shows the difference between attrition hazards at each credit score. The gaps show that borrowers who become safer over time become more likely to attrite from borrowing after the Act relative to before. Similarly, borrowers who become riskier over time become less likely to attrite than before the Act. The estimates imply that for every one percentage point by which emergent risk is mispriced relative to origination risk, borrowers respond with a 0.7 percentage point change in the quarterly hazard of attrition from borrowing.

In an appendix, I show that these two core results – the divergence between emergent and origination risk and the ensuing adverse retention of risky borrowers – are robust to a number of different specifications. These specifications include the following cases: if fees are included in addition to interest rates in the definition of the “price” of borrowing; if only very young (i.e., recently originated) accounts are included to estimate the origination price-risk gradient; if the sample only includes accounts old enough that all originated *prior* to the CARD Act; if a short-differences specification is used to relate quarterly changes in interest rates to quarterly changes in FICO score; if attrition from accounts is extended to include charge-off; if accounts with promotional rates are included in the sample used to estimate origination price-risk gradients; and if a Cox proportional hazard model is used instead of a linear probability model to estimate

³⁶One intriguing question is why the post-CARD-Act price gradient of emergent risk is sloped at all, and furthermore, why it is not kinked at zero, seeing as the Act did not restrict interest rate decreases for borrowers who became safer over time. The likely answers to this particular questions are related. First, the Act still allowed several channels through which lenders are able to update interest rates as borrower risk evolves: lenders could change interest rates on future balances, albeit not on current balances; lenders could pass through base rate increases to borrowers but could also selectively choose to cancel these increases; and lenders could still offer promotional rates to borrowers, even on mature accounts. However, with the exception of a scheduled expiration of such a promotional rate, the Act provided no means for a lender to “claw back” any rate decrease for a borrower after offering that decrease, so lenders’ incentive to offer rate decreases to newly safe borrowers was blunted by dynamic considerations.

these attrition hazards.

3.2 Price Elasticity Signals and Lender Rents

Consumer behavior on credit cards may reveal information not just about risk, but also about borrowing demand characteristics. In this section I provide evidence on which consumer behaviors reveal price elasticities of borrowing demand – behaviors that I term “price elasticity signals.” To do so, I analyze heterogeneity in lender returns across accounts that exhibit different behaviors in pre-CARD-Act data, and I identify which behaviors predict *higher* returns relative to returns on other, equally risky accounts that exhibit no such particular behavior. These higher returns are evidence of lenders’ learning about borrowers’ price elasticities, as they indicate lenders are able to adjust loan pricing in excess of what would compensate the lender for any change in borrower risk.

My core finding in this exercise is that two of the most common causes of interest rate repricing the pre-CARD-Act data – transactions exceeding an account’s credit limit, and delinquencies of less than 30 days – were in fact price elasticity signals in many FICO-score segments. In particular, delinquencies of less than 30 days predicted excess returns as high as 500 basis points at some FICO scores. I also confirm that, for accounts exhibiting either of these two behaviors, lenders’ excess returns were either sharply reduced or eliminated after the Act. In contrast, all other behaviors that were typically denoted as potential causes for repricing in pre-CARD-Act credit card contracts predict greater default rates and (often sharply) lower returns in the pre-CARD-Act period.

Using ex-post returns to identify price elasticity signals is an appealing approach because such signals are otherwise inherently difficult for a researcher to identify in the CCDB data. This is true for at least two reasons. First, there is no analog of a FICO score that can be used to track changing demand, rather than risk, over time. Second, lenders’ endogenous price responses to such signals can make the borrowers in question appear less, not more, likely to borrow than their peers. However even when these endogenous price changes lead to higher attrition, they still lead to higher ex-post returns if lenders are profit-maximizing and if a behavior is indeed revealing of higher price inelasticity.

To categorize borrower behaviors as price elasticity signals, I calculate the expected present value of lender revenues minus default losses among accounts that exhibit a certain behavior s in period $t = 0$, relative to the expected present value of balances lent on the same accounts, and I compare this measure of returns to the corresponding returns on accounts that do not exhibit any such particular behavior. Concretely this measure of expected returns is,

$$\hat{\mathbb{E}}[Y|s] = \sum_{t=0}^T \frac{\sum_{i:b_0(i)=s} R_{it} - L_{it}}{\sum_{i:b_0(i)=s} B_{it}/12} \quad (3.4)$$

where $b_t(i)$ is the behavior exhibited by consumer i in period t , and respectively R_{it} , L_{it} , and B_{it} are revenues, default losses, and revolved balances for that consumer. I then classify s as a price elasticity signal if,

$$\hat{\mathbb{E}}[Y|s] > \hat{\mathbb{E}}[Y|0] \quad (3.5)$$

where the behavior “0” on the right-hand-side of the inequality signifies that an account displayed “normal” behavior in that period, or more precisely, exhibited none of the signals I study.

I conduct this exercise for all behaviors that were typically included in pre-CARD-Act credit

card contracts as causes for either a penalty fee of some kind or a potential change in interest rate: over limit transactions, delinquencies in paying a monthly bill of various severity (less than 30 days, 30 to 60 days, and over 60 days), as well as changes in FICO score or other credit report information. I also consider several interactions of these behaviors, for example late payment that coincides with an over-limit transaction in the same billing cycle.

Note that I do not require an account to never exhibit behavior s in order to be included in the sum over $\{i : b_0(i) = s\}$; I only require that the account not exhibit s in period 0. In my baseline results, I take $T = 24$ to correspond to a 2-year horizon, which is a standard horizon over which to evaluate outcomes in consumer credit; results are also robust to taking $T = 12$. Given the front-loading of revenue relative to losses, the shorter-horizon specification leads to additional behaviors being classified as price elasticity signals as well.

Figure 7 shows the difference in expected returns, $\hat{E}[Y|s] - \hat{E}[Y|0]$, for two primary signals that I identify as price elasticity signals: over-limit transactions not coinciding with delinquencies, and delinquent payments that are late by less than 30 days. Over-limit transactions are generally price elasticity signals on subprime accounts, while late payments of less than 30 days are generally price elasticity signals on prime accounts.³⁷ Such late payments may be indicative of less price-elastic demand for a number of reasons, including credit constraints, a higher cost of time, or borrower inattention.³⁸

Table 3 then shows the results of this exercise for all other behaviors not classified as price elasticity signals. As shown in the table, each of these other signals predicts greater lender losses over the next two years. For example, among near-prime accounts with credit scores of 660-679, a quarterly FICO score drop of 30 to 59 points predicts lower annual returns by 3.66 percentage points off a baseline return of 5.09%, whereas late payments of 60 to 89 days predict lower returns by 42 percentage points.

3.3 Decomposition of Contingent Pricing

In this section I find that such price elasticity signals drove the majority of repricing on prime accounts, but not subprime accounts. I find that this result holds whether one considers interest rate repricing in response to contract-specified triggers, or any-time-any-reason interest rate repricing,³⁹ or pricing through fees rather than interest rates. This decomposition suggests, and my model results later confirm, that the CARD Act price restrictions' primary effect for *prime* borrowers is to restrict lenders from pricing information about borrower demand characteristics. In contrast, among subprime consumers the Act restricts the pricing of more risk-relevant information.

Figure 8 decomposes the share of interest rate increases in the pre-CARD-Act period that coincide with various contract-specified repricing triggers, for example transactions in excess of an account's credit limit. This decomposition is done separately for subprime accounts in the left

³⁷This difference between prime and subprime accounts comports with some basic institutional features of the credit card market: credit limits on prime accounts are typically high enough that an over-limit transaction for a prime consumer would suggest severe liquidity needs, likely predictive of substantial risk; in contrast late payments of less than 30 days on prime accounts may signal inattention and hence lower price sensitivity, whereas any late payment on subprime accounts may signal a liquidity shortfall.

³⁸As further evidence that late payments of less than 30 days may indicate inattention among some borrowers, I find in CCP data that these payments are positively correlated with borrowers reporting having a credit card misplaced or stolen.

³⁹See Section 2.1.1 for more details on any-time-any-reason repricing.

panel and prime accounts in the right panel, and each trigger is colored to emphasize whether I identified it as a price elasticity signal in section 3.2 above. Price elasticity signals (colored in green) are by far the dominant cause of interest rate increases on prime accounts; in contrast, other triggers (colored in red) dominate on subprime accounts.

To investigate whether this basic pattern also appears in *fee* revenue rather than interest rate increases, Table 4 next shows the share of fee revenue coming from various signals across various FICO score groups. The share of fee revenue attributable to price elasticity signals again depends on FICO score. Among prime accounts, over 70% of all contingent fee revenue comes from a behavior I find to be a price elasticity signal, delinquencies of less than 30 days. Among subprime accounts, only about 20% of fee revenue comes from the behavior I find to be a price elasticity signal in this market segment, over-limit transactions not coinciding with delinquencies. These patterns suggest that, for fee revenue just as for interest rate increases, the CARD Act price restrictions primarily restricted the pricing of risk-relevant information in the subprime market, whereas they primarily restricted the pricing of demand-relevant information in the prime market.

In Figure 9 I next ask whether a similar pattern holds for non-triggered interest rate increases, i.e. any-time-any-reason repricing.⁴⁰ The figure shows that non-triggered price increases primarily occur at times when a borrower’s FICO score is in fact *increasing*, i.e. default risk is falling. Hence roughly the same decomposition seen in triggered interest rate increases and in fee revenue can also be seen in non-triggered interest rate increases. That is, in the subprime market, the types of pricing restricted by the CARD Act had been primarily responsive to signals of borrower risk, while in the prime market, these pricing restrictions primarily limited price response to signals of borrower demand.

3.3.1 The Need for a Model

The results in the preceding subsections 3.1 and 3.2 point to a key tradeoff emerging from the CARD Act’s pricing restrictions. On the one hand, restricting lenders’ ability to raise prices on borrowers in response to a signal of borrowers’ price elasticity can lower markups on some borrowers, bringing prices closer in line with marginal costs and reducing the deadweight loss associated with these markups. On the other hand, restricting lenders’ ability to raise prices in response to risk information can engender adverse selection (at any price), which brings deadweight loss of its own. For consumers, the net effect of the Act on pricing depends on which of these two forces dominates in equilibrium, and for total surplus in the market, the net effect of the Act depends on the relative sizes of these two deadweight losses.

Empirically assessing the relative sizes of these effects empirically is difficult for two reasons. First, the CARD Act substantially changed the composition of borrowers in the credit card market. This makes the Act’s price effects difficult to measure for borrowers who were induced to leave the market. Second, the implementation of the Act coincided with a number of other credit market reforms⁴¹ and with a time of unique turbulence in consumer financial markets, and

⁴⁰The data do not actually assign an interest rate increase to a particular cause, so these non-triggered repricings are inferred based on which behaviors or potential causes are seen to coincide with an interest rate increase within the preceding quarter.

⁴¹Particularly relevant for credit card lending is the Federal Accounting Standards Board’s release of FAS 166/167 in June 2009, which made securitization of credit card loans more costly for lenders. See [Tian and Zhang \(2016\)](#), who use a difference-in-differences strategy between securitizing and non-securitizing credit card lenders to estimate that these accounting changes led to a 40% reduction in loan balances by the most affected banks.

the Act itself contained a number of policy changes unrelated to the repricing restrictions that I focus on here. For all of these reasons, it is a fraught exercise to isolate the effects of the Act’s repricing restrictions per se in any data taken from after the Act’s implementation.

These empirical difficulties notwithstanding, it *is* still an empirical question whether exacerbated information problems or lower lender markups were dominant when the Act’s pricing restrictions took effect. Intuitively, the key issue underlying this question is the whether the information restricted by the Act resolved more uncertainty about borrower risk or demand. The more this information was relevant for borrower demand, then the greater were the Act’s effects on markups. I formalize this intuition through a graphical example in an appendix, where I stylize the Act’s pricing restrictions as requiring two borrower types who previously could be priced differently to instead be pooled. The more these two borrower types differed in terms of their demand elasticities, the more overall prices in the market fall as a result of the Act’s restrictions, and the more does total surplus increase; conversely, the more these two borrower types differed in terms of their default risk, the worse is the resulting adverse selection problem and the more does total surplus fall.

In the following sections, I extend the intuition from that two-borrower example into a more realistic model of the credit card market, including multiple firms, private information, and a dynamic setting where lenders attempt to poach profitable borrowers from each other while borrower types also change over time. As I emphasize in section 4.1, each of these features plays a crucial role in a model designed to predict the CARD Act repricing restrictions’ effects; for example, private information and dynamic borrower types are important in light of how the CARD Act restricted the pricing of information that either changes or is revealed privately over time.

4 A Model of the Credit Card Market

In this section I develop and estimate a model of the credit card market. I estimate the model on the equilibrium observed in pre-CARD-Act data, so that I can later, in Section 5, use the model as a tool to study the effects of introducing the CARD Act’s price restrictions into this equilibrium. The model incorporates two features of the credit card market highlighted in the preceding section: lenders learn new information over time about both risk and demand, and lenders respond to this information in the pre-CARD-Act regulatory regime by changing loan pricing. The model also has three other prominent features – heterogeneous price sensitivities among borrowers, adjustment costs for consumers who switch lenders or pay off their balances, and private information among lenders about borrowers. In subsection 4.1 I motivate these three model features and illustrate how these features are identified by the data. I then formally introduce the model in subsection 4.2, discuss estimation in subsection 4.3 and present model parameter estimates in subsection 4.4.

4.1 Credit Card Demand: Three Key Facts

4.1.1 Fact 1: Price Sensitivity of Demand

This subsection establishes that credit card borrowers are sensitive to price and illustrates how it is possible to identify heterogeneous price sensitivities in the data. This heterogeneity will play a key role when I later use the model to study the equilibrium effects of the CARD Act’s price

restrictions, because this heterogeneity affects which types of consumers change their borrowing behavior in response to different relative price changes.

I estimate these price sensitivities by exploiting a novel source of price variation in the credit card market: occasional, idiosyncratic repricing campaigns in the pre-CARD-Act data in which banks change interest rates on entire extant credit card portfolios simultaneously.⁴² These campaigns come in two varieties. Occasionally, a credit card lender will reprice nearly all of its accounts at once, across all credit card types issued by that lender. In other cases, lenders will focus such repricing on all accounts in a single portfolio, such as a portfolio of airline credit cards. Such portfolios can be identified in the available data through clerical “tags” that lenders use to differentiate their accounts. It is plausible that these repricing campaigns are motivated by factors exogenous to consumer credit demand, such as changes to lenders’ internal cost of funds, changes in individual portfolio managers’ taste for risk, or a desire to shrink loan portfolios in advance of other institutional changes such as a merger or acquisition.

As an example of such repricing campaigns, Figure 10’s left panel illustrates a campaign in which one lender, referred to as “Bank A,” raised the APR on nearly all extant accounts by exactly 100 basis points in a month labeled as event time 0. The nine red lines show that all APR deciles of Bank A’s accounts rose simultaneously, after a preceding period with minimal price change. This campaign occurred more than a year before the passage of the CARD Act, and occurred at a time when, as shown by the figure’s dashed blue line, other lenders’ pricing was on average unchanged.

This change in Bank A’s pricing relative to its competitors facilitates a difference-in-difference analysis of borrower retention. The right panel of Figure 10 presents the standard difference-in-difference event-study plot for these two retention rates. Specifically, the right panel shows event-time-specific estimates from the equation,

$$\log Q_{jt} = \alpha_j + \alpha_t + \beta_j t + \alpha_{A,t} + \epsilon_{jt} \quad (4.1)$$

where Q_{jt} denotes retention rates among existing borrowers for lender j in month t , i.e. the share of borrowers who continue to borrow. The first two α terms in this equation implement a standard difference-in-differences design, while the $\alpha_{t,A}$ terms capture differences between Bank A and other, non-campaign banks. For sake of presentation, the β term is included to account for different time trends among the included banks, though as I show later this does not substantially affect the model parameters ultimately estimated off of this variation.

As can be seen in the right panel of Figure 10, the retention rate for Bank A’s borrowers falls relative to other banks’ borrowers immediately after the repricing campaign, with the greatest difference in the first month and a sustained but lesser gap in subsequent months. 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.

When estimating the demand side of the model, I use such price variation to estimate heterogeneous price sensitivities across different borrower types. This heterogeneity plays a key role in determining the equilibrium effects of the CARD Act’s price restrictions, as it affects which types of borrowers – for example, high or low risk borrowers – are most likely to enter or exit the market in response to relative price changes.

⁴²Firm-level price variation has also been used elsewhere in consumer finance research, for example by Cox (2017) in the context of student loan refinancing.

4.1.2 Fact 2: Persistence and Adjustment Costs

The previous subsection showed that price elasticities of borrowing demand are nonzero; this subsection considers reasons why elasticities are also not infinite. In particular I posit two kinds of adjustment costs faced by credit card users and I show evidence for these costs in pre-CARD-Act data. These adjustment costs will play an important role when I use the model to study the CARD Act price restrictions' effects, as they affect both the intensity of competition between lenders for different borrower types, and also the degree to which different borrowers substitute toward accounts with promotional pricing if other prices rise.

I present evidence for these adjustment costs by showing persistence in two dimensions of consumer behavior. One dimension is borrowing choices: consumers who use a card for borrowing in one month are highly likely to continue borrowing in the next month, while consumers who do not borrow are highly likely to continue not borrowing. A second is firm choice: regardless of whether they are borrowing or not, consumers persist in holding a card from a given bank, despite sometimes strong incentives for switching to another bank's credit card. These two types of persistence suggest adjustment costs both in paying off balances and in switching to a new credit card issuer.

Table 5 presents evidence that consumers face some kind of adjustment cost when paying off credit card balances: throughout the FICO score distribution, consumers are substantially more likely to borrow on a credit card in a given month if they also borrowed in the preceding month (columns 1 and 3) than if they did not borrow in the preceding month (columns 2 and 4). In the first half of the table, columns (1) and (2) make this point in a subsample of consumers with a demonstrated preference for borrowing – those consumers who borrowed on their credit card at least once in the past six months. As an illustrative example, note that FICO 720 consumers in this subsample who were borrowers in the preceding month have an 87% chance of continuing to borrow in the current month, whereas their non-borrower counterparts in the preceding month have only a 9% chance of borrowing. Columns (3) and (4) then extend this analysis to the whole population of credit card holders, not just those who borrowed at some time in the past six months. There is strong persistence in this broader population too: to again consider the example of FICO 720 consumers, the probability of continuing to borrow is 70%, while the probability of new borrowing is only 2%. This persistence is suggestive of some kind of adjustment cost in paying off credit card balances, which I term a “liquidity cost” to reflect the opportunity cost of using other funds to repay a credit card balance.

I next show that borrowers often face strong incentives to switch credit cards but nevertheless switch cards infrequently. To illustrate these strong incentives to switch cards, Table 6 follows a format similar to Table 2, here showing *introductory* rates on newly originated accounts in the pre-CARD-Act period. Here prices are shown for newly originated accounts to which a borrower transferred a previous balance at a promotional interest rate. Discounts relative to mature accounts appear throughout the FICO score distribution. For example, among FICO 740 consumers, the average cost of borrowing is roughly 600 basis points lower on newly originated accounts with promotional balance transfers, relative to mature accounts. Next, in Figure 11 I examine how frequently borrowers switch cards in the presence of these price incentives. To estimate these switch rates, I calculate the total number of balance transfers with promotional rates per quarter in the pre-CARD-Act period, and I compare this flow to the stock of consumers borrowing on mature accounts at non-promotional rates.⁴³ The figure shows this rate, along with

⁴³This ratio differs from the true balance transfer rate insofar as a single consumer may account for multiple

the total count of balance transfers, across a range of FICO groups. Even on a quarterly basis, only 16% of prime consumers and less than 5% of subprime consumers respond to the price incentives shown previously in Table 6 by transferring balances to a new credit card, indicating that many consumers face some kind of adjustment cost in setting up accounts with new issuers.

When estimating the model, I use these differences in switch rates and retention rates across borrower types to identify two corresponding sets of adjustment cost parameters – liquidity costs for paying off a balance, and set-up costs for opening a new account with a new lender. These adjustment cost parameters then determine which borrowers are most likely to substitute to promotional pricing and which borrowers are most likely to switch lenders when the CARD Act price restrictions are introduced.

4.1.3 Fact 3: Asymmetric Information

This subsection illustrates that lenders possess a substantial amount of private information about their ongoing borrowers. I also find that such private information was reflected in pre-CARD-Act loan pricing. These facts suggest that the CARD Act’s pricing restrictions – which make it difficult for lenders to adjust prices when they acquire private information about borrowers over time – have different price effects across different consumers depending on these consumers’ privately revealed types. Incorporating such private information in my model therefore becomes important in anticipation of using the model to study the Act.

I recover such private information from observed lender pricing in pre-CARD-Act data. This information is indeed private, because interest rates and fees in the credit card market are typically not observable to a lender’s competitors.⁴⁴ Equilibrium pricing therefore reveals lenders’ private information so long as distinct prices are assigned to distinct consumer types; I formalize some conditions sufficient for such pricing later.

To study the importance of this private information formally, I assign each borrower an index of private information corresponding to that borrower’s location in the distribution of prices charged by their lender to other borrowers at their FICO score. I will develop this index in detail in section 4.3.1. This index has the properties that borrowers with the same index value and the same FICO score have the same expected default rate regardless of which lender they borrow from (despite different lenders pricing different risk levels differently); indexes are, by sign convention, increasing in risk; and, when indexes are discretized, they are discretized such that an equal share of borrowers in the market is assigned to each index.

I use these indexes in Table 7, where I present linear-probability estimates of default rates by quintile of this private default-risk index. In this analysis I control flexibly for 20-point bins of FICO score in order to measure the predictive power of private information *within* observably similar borrowers. Formally, I estimate these effects in the following equation,

$$\text{Default}_{i,t:t+12} = \alpha_{j(i),x(i)} + \alpha_t + \sum_{n=1}^5 \beta_n 1_{\psi_{i,t}=n} + \epsilon_{it} \quad (4.2)$$

Here the dependent variable is an indicator for any instance of default by borrower i in the subsequent 12 months after period t , and the key coefficients β_n capture differences in default

balance transfers in the same quarter, for example when closing two cards and transferring both cards’ balances to the same new card. It is impossible to quantify the number of such instances using the CCDB.

⁴⁴The unobservability of competitors’ prices stems from the issue I discussed when introducing the CCP credit report data, that credit reports contain no data on prices paid for each loan.

rates across five quintiles of the private information index, which I denote by ψ . Meanwhile the fixed effects for borrower i 's firm j , FICO score x , and time period t help ensure that these risk comparisons are made within otherwise observably similar borrowers.

Estimates of β_n are presented in Table 7, first for all credit card borrowers in column (1), and then separately for prime and subprime borrowers in columns (2) and (3) respectively. The first (lowest risk) quintile is omitted, so that all other coefficients are relative to this group. The table shows that private information has substantial predictive power for default risk, especially in the subprime market. Overall, the fifth quintile of private information has 9 percentage points higher probability of default than the lowest quintile, and in the subprime market this gap grows to 20 percentage points.

To help benchmark these estimates against median default rates at various FICO scores, Table 8 then presents default rates across the FICO score distribution at the top-quintile, bottom-quintile, and median of such private information. Strikingly, the top quintile among FICO 720 borrowers has roughly the same expected default rate as the median borrower with a FICO 680 score, while the bottom quintile among these FICO 720 borrowers has roughly the same expected default rate as a median borrower with a FICO 740 score. Further perspective on these gaps can come from the overall distribution of FICO scores among credit card holders: I find that moving from the first to the fifth quintile of privately-known default risk is, on average across all FICO scores, roughly equivalent to a 2 standard deviation (174 point) decrease in FICO score in the overall distribution of scores; likewise, one standard deviation of privately-known default risk is just as predictive of future risk as 0.74 standard deviations of borrower credit score.

These results highlight the importance of incorporating private information in the model in order to study the CARD Act price restrictions' effects. As the Act's restrictions limit lenders' ability to adjust loan pricing when they learn such private information over time, borrowers with different privately revealed types, and hence default risk, will experience different relative price effects and face different incentives to either continue or attrite from borrowing.

4.2 Model Exposition

This section presents my model of the credit card market. The backbone of the demand model is a finite mixture of consumer types, each of whom has logit demand over credit card lenders and over the choice of whether to use his credit card for borrowing or not. Precisely, in a market with J banks there are $2J + 1$ discrete choices available to each consumer each period: two choices per bank (i.e. borrowing, or holding a credit card from that bank without borrowing) and one outside good, which is the option to hold no credit card at all. Consumers choose at most one bank at any point in time, and with this bank consumers choose only whether or not to borrow – that is, I model only the extensive margin of borrowing, not the choice of how much to borrow.⁴⁵

⁴⁵These two modeling decisions – that consumers single-home over banks and choose extensive rather than intensive-margin borrowing – are primarily made for sake of tractability. However, these decisions also do not depart much from realism in the credit card market. First, using CCP data, I find that a large majority of consumers hold only one “primary” credit card, where primary is defined as carrying the majority of a consumer's credit card balances. Depending on FICO score, this share ranges from at least 80% to over 90%. Hence a single-homing model can in many respects be thought of as a model of a consumer's choice of primary card. Additionally, a majority of deep subprime consumers and a large minority of prime consumers indeed hold only one credit card in CCP data. Second, there is some evidence that firms compete more on the extensive margin using price, and then use credit limits as their preferred instrument on the intensive margin (Trench et al. (2003), Agarwal et al. (2015b)). In fact, many credit limits are not disclosed until after a borrower has made the extensive margin choice of whether to open a credit card or not, whereas prices are advertised heavily to consumers considering a new

Each type has different tastes for each choice.

I denote types by θ . I specify several taste parameters to be estimated for each type. First, each type enjoys a flow utility $d_{j\theta}$ from borrowing with bank j and a flow utility $n_{j\theta}$ from transacting (rather than borrowing) with bank j ; meanwhile the utility of the outside good (holding no credit card at all) is normalized to zero. Additionally, in order to capture the adjustment costs documented earlier in this section, each type pays a setup cost $s_{j\theta}$ for opening a new account with bank j and a liquidity cost $l_{j\theta}$ or paying off a balance and transitioning to transacting (non-borrowing) status after borrowing with bank j in the past period. Additionally, types have heterogeneous marginal utilities of income γ_θ (i.e., the price coefficient in logit demand). The parameters $\{d_{j\theta}, n_{j\theta}, s_{j\theta}, l_{j\theta}, \gamma_\theta\}_{(\theta,j) \in \Theta \times J}$ are the key demand parameters to be estimated in the model, along with a probability distribution μ_θ over types.

This parameterization allows a type's preferences each period to depend on what bank he held a credit card from in the *previous* period, and also on whether he borrowed or not in the previous period. Because this is a model of industry-wide dynamics with differentiated firms, the total number of choice probabilities modeled is large ($|\Theta| \cdot (2J + 1)^2$). I therefore use Table 9 to summarize which parameters enter different borrowers' flow utilities for each choice. The three rows of the table correspond to the consumer's circumstances at the end of the preceding period: a consumer either (i) has an open credit card from some bank j that he used for borrowing, (ii) has a credit card from j that he did not use for borrowing, or (iii) holds no credit card at all. The five columns of the table then correspond to the consumer's choice in the current period: a consumer either keeps his credit card from the same bank j (columns 1 and 2), or opens a new card with some other bank $j' \neq j$ (columns 3 and 4), or chooses the outside good of no credit card at all (column 5). When holding a credit card, a consumer chooses either to use it for borrowing (columns 1 and 3) or not (columns 2 and 4).

In reading the table, note that these banks j and j' can be any bank in the set of banks J , so there are $|J|$ distinct values of each parameter subscripted by j or j' . An important pattern to note in the table is that consumers only pay setup costs s when transitioning from some bank j to a new bank $j' \neq j$, and only pay liquidity costs l when transitioning from borrowing to transacting.

Meanwhile, as shown in the table, prices differ for consumers who are newly opening a credit card with a bank and consumers who held a credit card with that bank in the past period. These two prices are denoted $p_{\theta j}^0$ and $p_{\theta j}^1$. Allowing these prices to differ between new and mature accounts helps pin down consumers' switching costs across accounts when estimating the model, which then is helpful in predicting how consumers respond when such new-account discounts ("teaser" rates) change after I impose the CARD Act price restrictions. Note also that these prices are one-dimensional, so in practice I use the fee-inclusive borrowing cost introduced in Section 2.2.3 when I estimate these prices in the data; these are also the appropriate marginal prices to use when modeling the extensive margin of borrowing.

The presence of adjustment costs makes the consumer's problem dynamic. Therefore the total expected payoff for a given choice is the sum of the relevant flow utility from Table 9 and also a discounted expectation of continuation values (plus also, given logit demand, the realization of an extreme value type-1 i.i.d. taste shock). To describe these continuation values, let $k \in \{\text{borrow}, \text{transact}\} \equiv \{\text{b}, \text{n}\}$ denote a consumer's choice of how to use his credit card

card. Incorporating the intensive margin in the model would therefore seem to require including both prices and credit limits, which would expand the firms' strategy space to the point of intractability.

and $j \in J$ again denote a consumer's choice of card.⁴⁶ I then write these continuation values as $V(\theta', j, k)$. Note that θ' is a consumer's type in the *next* period while j and k correspond to the current period. For example, a consumer i 's total expected payoff for choosing to borrow ("b") with bank j in the current period after having also borrowed with bank j in the past period is,

$$\underbrace{d_{j\theta} - \gamma_{\theta} p_{\theta j}^1}_{\text{flow utility}} + \underbrace{\beta \mathbb{E}_{\theta} [V(\theta', j, b)]}_{\text{exp. cont. value}} + \epsilon_{ijb} \quad (4.3)$$

Integrating over taste shocks ϵ for each choice yields the standard Bellman equation for continuation values V ,

$$V(\theta, j, k) = \log \left(\sum_{j', k'} \exp(v(j', k' | j, k, \theta)) \right) \quad (4.4)$$

where the lower-case v term denotes total expected payoffs for a given choice exclusive of taste shocks. The value of v depends on consumers' past-period and current-period choices as described previously in Table 9. For example, in the case of a consumer who chooses (as in equation 4.3) to borrow ("b") with bank j in the current period after having also borrowed with bank j in the past period, the value of v is,

$$v(j, b | j, b, \theta) = d_{j\theta} - \gamma_{\theta} p_{\theta j}^1 + \beta \mathbb{E}_{\theta} [V(\theta', j, b)] \quad (4.5)$$

Besides determining flow utilities as above, consumer types θ additionally govern heterogeneity in default rates. Specifically each type defaults at exogenous rate $\delta(\theta)$ in periods when he chooses to borrow. Default occurs after all flow utilities are realized in that period. I later discuss how these default rates determine firms' costs, but here I emphasize how default rates also matter for consumer payoffs. In particular, a consumer who defaults has his credit card account "closed" and is reassigned to the outside good (holding no credit card at all) for purposes of computing adjustment costs in the next period. Hence default rates affect expected payoffs only through the expectation over future continuation values.

To tractably model expectations over continuation values, I follow the standard approach in the dynamic discrete choice literature and suppose types evolve according to a Markov process,⁴⁷ with a transition matrix that I denote $T_{\theta\theta'}$. Transitions occur independently of default, consumer choices, and taste shocks. Hence, for consumers who use their credit card for borrowing, the expectation \mathbb{E}_{θ} can be decomposed as,

$$\mathbb{E}_{\theta} [V(\theta', j, b)] = \underbrace{(1 - \delta(\theta)) T_{\theta\theta'}(\theta) V(\theta', j, b)}_{\text{no default}} + \underbrace{\delta(\theta) T_{\theta\theta'}(\theta) V(\theta', 0, 0)}_{\text{default}} \quad (4.6)$$

where the θ argument in $T_{\theta\theta'}(\theta)$ selects the relevant row of the matrix $T_{\theta\theta'}$. In the second term on the right-hand-side, recall that I use $(j, k) = (0, 0)$ to denote the outside good.

In contrast, for consumers who do not choose to borrow (i.e., who choose $k = n$ or $k = 0$), the expectation \mathbb{E}_{θ} does not depend directly on default rates and takes the form,

$$\mathbb{E}_{\theta} [V(\theta', j, k)] = T_{\theta\theta'}(\theta) V(\theta', j, k) \quad (4.7)$$

⁴⁶In this notation I also represent the outside good as $(j, k) = (0, 0)$.

⁴⁷See Rust (1994) for a review of this literature and a taxonomy of assumptions typically used to help make such models tractable.

The above exposition makes clear how the demand side of the model captures two of the three stylized facts I highlighted – price sensitivity and adjustment frictions. To capture the third stylized fact – the importance of private information – I now describe how the model parameterizes consumer types. Specifically 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. Note that the public type x is best thought of as a FICO score, as FICO scores are expressly designed to be a composite of public information about a consumer, and are indeed observable to all firms in the market.⁴⁸ The joint of these two components is then a consumer’s overall type, $\theta \equiv (x, \psi)$.

Two assumptions on borrower types will prove useful in estimating the model. One assumption, which is arguably the stronger of the two, is that borrower default rates depend only on types, and in particular do not depend on prices p_j^0 and p_j^1 or on bank j . This can be thought of as a “no moral hazard” assumption and I will refer to it as **Assumption 1**:

$$\delta = \delta(\theta) \quad \forall j, p_j^0, p_j^1 \quad (4.8)$$

In an appendix, I collect several pieces of evidence in support of this assumption, including evidence on how observed default rates respond to two sources of plausibly exogenous price variation in the CCDB data, an argument using CCP data that examines how credit card price changes affect most consumers’ overall cost of debt service (a budget constraint argument), and a review of related research in consumer finance suggesting little to no moral hazard channel through which prices affect default rates. 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 et al. \(2010\)](#).⁴⁹

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. I order private types ψ at each public type x such that default is increasing in ψ ,

$$\psi' > \psi \implies \delta(x, \psi') > \delta(x, \psi) \quad \forall x \quad (4.9)$$

A second assumption, which I view as the weaker of the two, is a “non-advantageous selection” assumption. This assumption is supported by extensive evidence showing the credit card market is not merely non-advantageously selected, but is indeed adversely selected ([Ausubel \(1999\)](#), [Agarwal et al. \(2010\)](#)).⁵⁰ Formally the assumption is that higher-risk private types do not have less demand for borrowing from any given lender than do lower-risk private types, at a given FICO score. I express this assumption in terms of resultant choice probabilities, which I term **Assumption 2**:

$$\psi' > \psi \implies \Pr(j, b | j, b, x, \psi') \geq \Pr(j, b | j, b, x, \psi) \quad \forall x \in X, j \in J \quad (4.10)$$

Note that this assumption embeds some restrictions on the competitive environment, namely that one lender’s relative quality advantage over competing lenders (as expressed in differences

⁴⁸See Section 2.2.2 for further information on the contents and availability of credit report data.

⁴⁹Additionally, I highlight in section 4.3 where the estimation procedure could be adapted, albeit at considerable computational expense, should this assumption fail.

⁵⁰To clarify these terms, advantageous selection is the case where higher prices induce the composition of borrowers to become less risky; adverse selection is the more familiar opposite of this case. Non-advantageous selection includes adverse selection as well as the intermediate case where the composition of borrower risk is unchanged with price.

across j in demand parameters such as the flow utility from borrowing, $d_{j\theta}$) does not change so drastically with ψ , the private dimension of θ , such that lenders in fact face lower demand as private risk rises. That is, residual demand curves and not just aggregate demand curves are non-advantageously selected in the pre-CARD-Act equilibrium.⁵¹

The precise timing of the demand side of the model is as follows. First, borrower types θ are realized at the start of the period. Banks then post prices p^0 and p^1 for each type. Consumers choose a bank and a borrowing status, and enjoy flow utility from their choice. Default then arrives exogenously. Borrowers who default are forced into the outside good (no account with any bank) for purposes of determining their adjustment costs in the following period. Borrowers who do not default continue on to the next period with their chosen bank.

On the supply side of the model, a credit card lender's price-setting problem has two parts: what price of borrowing to offer on existing accounts, and what promotional or "teaser" price to offer for new customers. As in the consumer's choice problem, these two sets of prices are denoted $p_{\theta j}^1$ and $p_{\theta j}^0$ respectively, where subscripts denote bank j and consumer type θ .

Corresponding to these two types of prices, credit card lenders' costs can also readily be grouped into two types: acquisition costs related to originating a new account, which include underwriting costs, account set-up costs, and marketing expenses; and account maintenance and charge-off costs on existing accounts, which include day-to-day account management plus costs of default net of recoveries. I denote these costs $c_{\theta j}^0$ and $c_{\theta j}^1$ respectively.

My model focuses on the extensive margin of borrowing,⁵² so lender flow profits for consumers who choose to borrow are the difference between the relevant price and cost: that is, flow profits for lender j are $p_{\theta j}^1 - c_{\theta j}^1$ for existing borrowers and $p_{\theta j}^0 - c_{\theta j}^0$ for borrowers opening a new account. I suppose acquisition costs must also be paid for new accounts even if consumers choose not to borrow, given that new-account costs are primarily driven by set-up and marketing expenses rather than default cost. This cost structure implies that expected discounted lifetime profits for a new consumer, Π^0 , take the form,

$$\begin{aligned} \Pi^0(p_j, p_{-j}, \theta, k) = & \underbrace{\Pr_j^0(b|\theta, p, k)p_{\theta j}^0 - c_{\theta j}^0}_{\text{flow profit}} + \\ & \underbrace{\Pr_j^0(b|\theta, p, k)\beta(1 - \delta(\theta))T_{\theta\theta'}(\theta)\Pi^1(p_j, p_{-j}, \theta', b)}_{\text{exp. cont. profit | borrow}} + \\ & \underbrace{\Pr_j^0(n|\theta, p, k)\beta T_{\theta\theta'}(\theta)\Pi^1(p_j, p_{-j}, \theta', n)}_{\text{exp. cont. profit | not borrow}} \end{aligned} \quad (4.11)$$

Here the notation $\Pr_j^0(b|\theta, p, k)$ denotes the probability of consumer type θ choosing to borrow conditional on having opened a new account with lender j in the current period, and conditional on having chosen $k \in \{\text{borrow, transact, out}\} \equiv \{b, n, 0\}$ in the preceding period. Similarly $\Pr_j^0(b|\theta, p,)$ denotes the probability of choosing to transact (i.e., hold a credit card without borrowing). The dependence on k is a result of consumers facing different adjustment costs depending on whether they borrowed in the previous period, and hence exhibiting different choice probabilities in the current period. As in the demand side of the model, $\delta(\theta)$ denotes

⁵¹While it would be preferable to express this assumption in terms of primitives, it appears quite cumbersome to do so, and the assumption as expressed in choice probabilities helps clarify the essential – and also most likely directly testable – content of the assumption.

⁵²See footnote 45.

borrower default probabilities, and the notation $T_{\theta\theta'}(\theta)$ selects the appropriate θ -specific row of the consumer type transition matrix. Also note that $p = (p_j, p_{-j})$ denotes the market price vector (including both existing-account prices and teaser prices). The final piece of new notation to introduce is $\Pi^1(p_j, p_{-j}, \theta', k)$, which is lenders' continuation profits on existing accounts, as a function of the consumer's choice $k \in \{\text{borrow}, \text{transact}\} \equiv \{b, n\}$ in the current period. These profits on existing accounts are defined further below.

Some intuition about issuers' dynamic incentives in the previous expression may be helpful. These continuation profits are the sum of two objects: first, the probability that a consumer chooses to borrow on a card, times the sum of both a one-period payoff and a discounted expected continuation value given that choice; and second, the probability that a consumer chooses to use a card only for transactional purposes (not for borrowing), times a corresponding payoff and continuation value. Accounts have higher continuation values the more likely these choices are, and the higher lenders' payoffs are given these choices. Account holders may also choose to close their account, which yields zero payoff and continuation value for the firm.⁵³

Profits on existing accounts take a similar form to profits on new accounts,

$$\begin{aligned} \Pi^1(p_j, p_{-j}, \theta, k) = & \underbrace{\Pr_j^1(b|\theta, p, k) (p_{\theta j}^1 - c_{\theta j}^1)}_{\text{flow profit}} + \\ & \underbrace{\Pr_j^1(b|\theta, p, k) \beta (1 - \delta(\theta)) T_{\theta\theta'}(\theta) \Pi^1(p_j, p_{-j}, \theta', b)}_{\text{exp. cont. profit | borrow}} + \\ & \underbrace{\Pr_j^1(n|\theta, p, k) \beta T_{\theta\theta'}(\theta) \Pi^1(p_j, p_{-j}, \theta', n)}_{\text{exp. cont. profit | not borrow}} \end{aligned} \quad (4.12)$$

Here the primary difference between existing account profits and new account profits is that expected costs $c_{\theta j}^1$ are only paid if a consumer chooses to borrow, reflecting how existing-account costs primarily depend on loan default. Additionally, firms earn existing-account prices $p_{\theta j}^1$ and incur existing-account costs $c_{\theta j}^1$ rather than the new-account terms $p_{\theta j}^0$ and $c_{\theta j}^0$.

Notwithstanding the apparent similarity in these two profit functions, lenders' pricing problem on new accounts is starkly different from the pricing problem on existing accounts. This is because of the different types of information available to lenders on new and existing accounts. As discussed earlier, lenders must make new account pricing decisions on the basis of "public" information available in credit reports, whereas pricing on existing accounts can depend on private information that a lender learns over the course of a lending relationship. I express these constraints in the following informational assumption: lenders observe only a borrower's public type x on a newly originated account, whereas lenders observe a consumer's full type $\theta = (x, \psi)$, including the private type ψ , on existing accounts.⁵⁴

Given this informational assumption, I impose the natural restriction that lender pricing strategies on new accounts must be the same for all types θ that have the same public type x ,

$$p_{\theta j}^0 = p_{x(\theta)j}^0 \quad \forall \theta$$

⁵³Furthermore, lenders also lose any continuation value (but still receive flow profits) if an account used for borrowing goes into default at the end of the period; as described previously in the demand model, accounts in default are closed permanently at the end of the period.

⁵⁴This assumption precludes borrowers behaving strategically in a way that prevents lenders from observing their true type, although it does allow for a signal-jamming behavior in which all consumers try to appear safer or more price sensitive than they truly are, so that lenders nonetheless infer their type.

where $x(\theta)$ selects the public component of types $\theta = (x, \psi)$. To be consistent with this restriction, I also suppose acquisition costs take the form $c_{\theta j}^0 = c_{x(\theta)j}^0 \forall \theta$.

In choosing prices p_{xj}^0 a lender therefore takes into consideration its expectation of which private types ψ it acquires as new customers at any given price level, expressed below as a sum over types θ that share a given FICO score x , competing lenders j' , and borrowers' past-period choices k ,

$$\Pi^0(p_j, p_{-j}, x) = \sum_{j' \neq j} \sum_{\theta: x(\theta)=x} \sum_{k \in \{b, n, 0\}} \mu_{j', \theta, k}(p) \Pr(j|p, j', k, \theta) \Pi^0(p_j, p_{-j}, \theta, k) \quad (4.13)$$

Here the weights $\mu_{j', \theta, k}$ are the share of consumers who are of type θ , who held a credit card from lender j' in the prior period (or held no card in the case of $j' = 0$), and who used that card for $k \in \{\text{borrow}, \text{transact}, \text{out}\} \equiv \{b, n, 0\}$, as a function of the market price vector p .

Given the above expressions for $\Pi^0(p_j, p_{-j}, x)$ and $\Pi^1(p_j, p_{-j}, \theta, k)$, the lender's pricing problem can now be written as,

$$\max_{p_j} \sum_x \Pi^0(p_j, p_{-j}, x) + \sum_{\theta} [\mu_{j, \theta, b}(p) \Pi^1(p_j, p_{-j}, \theta, b) + \mu_{j, \theta, n}(p) \Pi^1(p_j, p_{-j}, \theta, n)] \quad (4.14)$$

In the following subsection I describe how I estimate the supply side of the model using the first-order conditions of this optimization problem. I also describe three distinct steps in estimating the demand side of the model: recovering borrower types θ and the probability distribution over types μ_{θ} ; estimating the parameters γ_{θ} that govern consumers' price elasticities, conditional on types; and finally estimating all remaining demand parameters, conditional on both types and estimated elasticities.

4.3 Model Estimation

4.3.1 Demand Estimation: Borrower Private Types

The first step in demand estimation is recovering a type θ for each borrower in the data. To emphasize, rather than estimating a *parametric* mixture model of types, in which the key objects to be estimated would be parameters of the type mixture distribution, I instead recover a single type for each consumer in the data, and allow the distribution over types to remain flexible.

Recall types θ are the joint of public and private types, $\theta = (x, \psi)$. Finding borrowers' public types x is straightforward: I allow each borrower's public type to be a binned version of his FICO score. I make this choice because FICO scores are expressly designed to be a one-dimensional composite of all publicly available information predicting default, and because FICO scores are readily observable in the data. I use 20-point FICO score bins, which are a standard set of bins, or "breaks," the credit card industry uses to group borrowers for account management purposes. Additionally I pool all FICO scores of 599 or below into a single bin and all FICO scores of 780 or above into a single bin. This yields a total of 11 distinct public types x .

With these public types so defined, the remaining part of this exercise is to recover private types ψ . Empirically, my approach here 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 et al. (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 develop an empirical strategy that recovers these private types from the observed pricing that each borrower faces in each period.

Here I make use of Assumptions 1 and 2 developed in the previous section. As I show in an appendix, a straightforward implication of these two assumptions is that equilibrium prices p^* are increasing in private types ψ for all public types x and all lenders j ,⁵⁵

$$\psi' > \psi \implies p_{j,x,\psi'}^{*1} > p_{j,x,\psi}^{*1} \quad \forall x, j \quad (4.15)$$

Recall also from equation 4.9 that default rates δ are also increasing in private types ψ for all FICO scores x and all banks j . So, default rates and prices p^{*1} are increasing with respect to each other,

$$\hat{\delta}_{jx} (p_j^1) \nearrow p_j^1 \quad (4.16)$$

where $\hat{\delta}_{jx}$ is the default rate as an indirect function of prices in equilibrium, among borrowers with FICO score x for lender j . Finally, using the inverse of δ implied by equation 4.9, private types can be recovered by inverting default rates observed at each price level,

$$(x, \psi) = \delta_x^{-1}(\hat{\delta}_{jx}(p_j^1(x, \psi))) \quad \forall x \quad (4.17)$$

Note that equilibrium price schedules p_j^1 are lender-specific, as are the indirect functions δ_{jx} relating these prices to realized default rates. However the inverse δ_x^{-1} maps default rates, which are *common* for all borrowers of a given type, back to types. So in estimating the model, $\hat{\delta}_{jx}$ is estimated separately by lender and by FICO score x , while δ_x^{-1} is estimated across all lenders – i.e., for the market as a whole – within each FICO group.

To do this inversion in practice, I first use isotonic regression to estimate $\hat{\delta}_{jx}$ for each lender j and FICO score group x . The default measure I use is delinquencies of 90+ days within the following two 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 $\hat{\delta}_{jx}$ being a consistent estimate of actual default rates at each price level, given Assumptions 1 and 2.

To define the inverse $\delta_x^{-1}(\cdot)$, I use the fact that private types ψ are an index of default risk (see equation 4.9), and I therefore specify $\delta_x^{-1}(\cdot)$ 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 . I then also bin each lender's pricing functions $p_j^1(x, \psi)$ to that lender's average price at each bin.

This process is illustrated for two actual lenders in the data in the three panels of Figure 12. As can be seen, a borrower of a given type shares a common default rate regardless of his current bank, while the price faced by each borrower is different depending on the bank he chooses. 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.

The consumer types estimated in this process make it straightforward to study the dynam-

⁵⁵This result also makes use of an informational assumption I develop on the supply side of the model, that lenders observe ψ after having a relationship with a consumer in the preceding period.

ics of how types change over time. In particular, the transition matrix $T_{\theta\theta'}$ can be estimated non-parametrically 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 borrower choices and default outcomes: type transitions do not depend on borrower choices or realized default, 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 13. Here, the integer-labeled type indices correspond to the 11 different 20-point FICO score groups described earlier, while the sub-ticks 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 persistent, in both public and private dimensions. The rippling pattern evident in the plot shows the same phenomenon seen previously in Table 8, whereby 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 . Recovering this distribution is necessary even though ψ was taken to be quintiles of a default rate distribution. This is because these default rates are only observed for consumers who choose to *borrow*; hence, while there is a uniform distribution (within FICO score) of types among borrowers, the overall distribution of types may not be uniform, if different types have different probabilities of borrowing. To overcome this difficulty, I simply use the fact that type transition matrix $T_{\theta\theta'}$ operates independently of consumers' choices of whether to hold a credit card and whether to borrow, so ergodicity implies a unique steady state μ_θ that satisfies the equation $\mu_\theta = T_{\theta\theta'}\mu_\theta$.

4.3.2 Demand Estimation: Demand Elasticities

The next demand parameters to estimate are price elasticities of borrowing demand η_{ij} , across consumers i and credit card issuers j . In general demand elasticities change as prices change, so it is helpful to estimate primitives that determine these elasticities rather than merely estimate the elasticities themselves, as the latter only are local to a particular equilibrium. I therefore use the well-known relationship between demand elasticities and marginal utilities of income in logit demand,

$$\eta_{ij} = -\gamma_i p_{ij}(1 - Q_{ij}) \quad (4.18)$$

Here γ_i is consumer i 's marginal utility of income, p_{ij} is consumer i 's price of borrowing from lender j , and Q_{ij} is consumer i 's probability of choosing to borrow from lender j . In particular I use η_{ij} to denote the elasticity of *continued* borrowing among current borrowers, so that the price on the right-hand side denotes lender j 's pricing on mature credit card accounts, and Q_{ij} denotes a retention probability for current borrowers. Intuitively in this expression higher marginal utilities of income make borrowers more price elastic.

To derive an estimating equation for γ_i that uses price variation such as the repricing campaign illustrated previously in Figure 10, I first substitute for η_{ij} using the definition of an elasticity,

$$\text{dlog}(Q_{ij}) = -\gamma_i p_{ij}(1 - Q_{ij})\text{dlog}(p_{ij}) \quad (4.19)$$

I then draw on the form of borrower heterogeneity specified in section 4.2, and I take this equation from the level of individual consumers i to the level of consumer types θ . This leverages in particular the assumption that borrower types x and ψ capture all relevant borrower heterogeneity in the model (with $\theta = (x, \psi)$). This simply changes i subscripts to θ subscripts in the

above, and substitutes observed type-level retention rates $Q_{\theta j}$ in lieu of individual retention probabilities Q_{ij} .

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

$$\log Q_{\theta jt} = \alpha_{\theta j} + \alpha_t + \beta_j t - \gamma_{\theta} \log P_{\theta jt} + \epsilon_{\theta jt} \quad (4.20)$$

Here the fixed effects denoted by α implement difference-in-differences, and the term $P_{\theta jt}$ is a price term scaled as in equation 4.19 above, with scalars taken from the period immediately prior to a repricing denoted here by $t = 0$,⁵⁶

$$\log P_{\theta jt} = (1 - Q_{\theta j0}) p_{\theta j0} \log(p_{\theta jt}) \quad (4.21)$$

Meanwhile the β term is included to account for different trends among the included banks; I explore robustness to excluding this term below. This equation differs from the earlier event-study version shown in equation 4.1 and Figure 10 only through the regressor $P_{\theta jt}$, which, following the above derivation, makes it possible to recover the primitives γ_{θ} rather than just a local elasticity.

I estimate γ_{θ} using both limited-information maximum likelihood and two-stage least squares, with instruments that isolate the type of repricing variation highlighted in Figure 10. Specifically, I instrument for the endogenous price term $P_{\theta jt}$ with a dummy instrument Z_{jt} equal to unity in all periods t following a repricing campaign by lender j . As is standard in a model that is fully interacted with consumer types θ , these instruments are also interacted with indicators for borrower types θ , so that there are $|\Theta|$ instruments corresponding to the $|\Theta|$ endogenous regressors $P_{\theta jt}$.⁵⁷ Note that these instrumental variables address two econometric issues, both the endogeneity of prices $p_{\theta j}$ with borrowers' marginal utilities γ_{θ} , and, in time period 0, the appearance of $Q_{\theta j0}$ on both the right- and left-hand sides. In summary, the first and second stage equations are then,

$$\log P_{\theta jt} = a_{\theta j} + a_t + b_j \times t + \pi_{\theta} Z_{jt} \times 1_{\theta} + e_{\theta jt} \quad (4.22)$$

$$\log Q_{\theta jt} = \alpha_{\theta j} + \alpha_t + \beta_j \times t - \gamma_{\theta} \log P_{\theta jt} + \epsilon_{\theta jt} \quad (4.23)$$

Given that $P_{\theta jt}$ contains the estimated quantity $Q_{\theta j0}$, I bootstrap to calculate standard errors.

Table 10 presents estimates corresponding to the repricing quasi-experiment shown in Figure 10. The first column shows OLS estimates of equation 4.23, while the second column then shows corresponding 2SLS estimates that use variation from the first-stage equation 4.22. Comparing these two estimates lends credence to the instrumental variables strategy: the OLS estimate of γ is substantially closer to 0 than is the 2SLS estimate, as would be expected if the instruments overcome the standard endogeneity problem whereby higher prices are charged to less price-sensitive borrowers in equilibrium.

The next column of the table then examines how estimates change with the exclusion of bank-specific time trends β_j , and the final column of the table explores heterogeneity in marginal utilities γ across borrower types. As can be seen, the inclusion of bank-specific trends changes the resulting estimates of γ slightly, with estimates falling from .106 to .0696 when these trends are excluded. Furthermore, heterogeneous estimates of γ across borrower types vary in a sensible

⁵⁶These base-period values are chosen because they correspond to demand elasticities at the time of the repricing, as in equation (4.18).

⁵⁷The high number of interaction terms motivates using limited-information maximum likelihood estimates in lieu of two-stage least squares estimates, to help overcome finite-sample bias in a setting with many instruments.

way, with higher marginal utilities of income for lower-credit score borrowers who are, on average, lower income. Further evidence on this heterogeneity is presented in an appendix. Note that the estimates of γ I ultimately use in solving the model are presented in Figure 14, where I allow γ to vary flexibly across consumers' public types x .

For the estimates I present in this table, the set of instrumental variables I use are drawn from the repricing quasi-experiment illustrated previously in Figure 10. This particular quasi-experiment has the advantage that I have been able to verify important background details that help validate the exogeneity of this repricing campaign vis-a-vis existing borrowers' demand: documents from this lender's investor relations materials emphasize that the lender was seeking to consolidate its credit card portfolio at this time in advance of an upcoming merger or acquisition – as would be rationalized by the bank's internal cost of capital changing in anticipation of such an acquisition. This merger or acquisition was not consummated until several quarters after the repricing event in question, so it likely did not substantially change the competitive environment in the event-time months immediately following the repricing quasi-experiment I use here. The lack of a detectable change in other competitors' pricing strategies at this time, as evidenced by the blue dashed line in Figure 10, also support the exogeneity of the repricing event.

4.3.3 Demand Estimation: Taste Parameters

Given the above estimates of each consumer's type θ and borrowers' price sensitivities corresponding to γ_θ , the remaining model parameters to be estimated are the flow utilities $d_{j\theta}$, $n_{j\theta}$, $s_{j\theta}$, and $l_{j\theta}$. Recall these terms are, respectively, flow utilities from borrowing, flow utilities from transacting (rather than borrowing), setup costs for opening an account with a new lender, and liquidity costs for paying off a balance in order to transition from borrowing to transacting. These are estimated by matching key moments of the data corresponding to the three key facts outlined in section 4.1, each moment being an observed probability that is matched to a corresponding likelihood predicted by the model. In particular, these moments are: borrowers' persistence in borrowing behavior; non-borrowing consumers' persistence in *non*-borrowing behavior; account closure rates for borrowers; and account opening rates for consumers not holding credit cards.

To help illustrate how such moments identify the remaining model parameters, Figure 15 shows the example of how borrowers' persistence in borrowing behavior (i.e., lenders' retention rates among borrowers) identify flow utilities from borrowing, $d_{j\theta}$. The figure shows, for each FICO score group on the x-axis, the highest and lowest borrower retention rates across all lenders in solid lines; these lines are simply the upper and lower envelopes of retention rates in the market. The figure also shows in dashed lines the fee-inclusive prices⁵⁸ charged by the lenders in these upper and lower envelopes.⁵⁹ Reading across the FICO score distribution from low to high, note that at the bottom of the distribution the lender with the highest retention rate also charges a relatively high price of 45 percentage points annualized, relative to 20 percentage points for the lowest-retention lender; meanwhile in the middle of the FICO distribution, the price gap between high- and low-retention lenders converges to nearly zero, and at the top of the price distribution, the highest-retention lender instead charges lower prices than the lowest-retention lender. This pattern identifies differences in $d_{j\theta}$ for these high- and low-retention lenders across the FICO score distribution. At the bottom of the distribution, the gap in these two values of $d_{j\theta}$

⁵⁸Recall I use the *fee-inclusive* borrowing cost introduced in Section 2.2.3 when I estimate these prices in the data; these are also the appropriate marginal prices to use when modeling the *extensive* margin of borrowing.

⁵⁹The figure is designed this way, using upper and lower envelopes rather than just showing two example lenders, so as to protect firms' confidentiality and avoid displaying the full price schedule for any single lender.

must be large enough to rationalize borrowers having relatively high retention rates on a lender that is also relatively high priced; intuitively, this high-retention lender is a bank whose credit cards are particularly attractive to borrowers at the bottom of the FICO score distribution. In the middle of the FICO score distribution, the smaller price gap at comparable retention rates identifies a smaller gap between these two values of $d_{j\theta}$, and at the top of the distribution, the reversal of the price gap points to this $d_{j\theta}$ gap being even smaller. In brief, the patterns in Figure 15 point to credit card product differentiation being a relatively important determinant of borrowing demand at the bottom of the credit score distribution, and a less important factor at higher credit scores.

4.3.4 Supply Estimation

The lender's maximization problem in equation 4.14 has tractable first-order conditions because many pricing decisions are made independently. This independence follows from lenders' lack of commitment power in the pre-CARD-Act regulatory regime, which implies a deviation in $p_{j\theta}^1$ only affects profits earned on existing accounts for consumers of type θ , and likewise a deviation in p_{jx}^0 only affects profits earned on new accounts among consumers of public type x . Furthermore continuation profits are unaffected by these one-period deviations. The first-order condition for $p_{j\theta}^1$ at the equilibrium price vector p^* is thus, for a given θ ,

$$\begin{aligned} \sum_{k \in \{b, n\}} \Pr_j^1(b|\theta, p^*, k) &= \sum_{k \in \{b, n\}} \gamma_{\theta} \mu_{b, \theta, k}(p^*) \Pr_j^1(b|\theta, p^*, k) (\Pr_j^1(b|\theta, p^*, k)) \times \\ &\quad [p_{\theta j}^1 - c_{\theta j}^1 + \beta(1 - \delta(\theta))T_{\theta\theta'}(\theta)\Pi^1(p_j, p_{-j}, \theta', b)] \\ &\quad - \gamma_{\theta} \mu_{n, \theta, k}(p^*) \Pr_j^1(b|\theta, p^*, k) (\Pr_j^1(n|\theta, p^*, k)) \times \\ &\quad [\beta(1 - \delta(\theta))T_{\theta\theta'}(\theta)\Pi^1(p_j, p_{-j}, \theta', b)] \end{aligned} \quad (4.24)$$

First-order conditions for prices on newly originated accounts p_{jx}^0 are similarly, for a given x ,

$$\begin{aligned} \sum_{j' \neq j} \sum_{\theta: x(\theta)=x} \sum_{k \in \{b, n, 0\}} \mu_{j', \theta, k}(p) \Pr(j|p, j', k, \theta) &= \\ \sum_{j' \neq j} \sum_{\theta: x(\theta)=x} \sum_{k \in \{b, n, 0\}} \gamma_{\theta} \mu_{b, \theta, k}(p^*) \Pr_j^0(b|\theta, p^*, k) (1 - \Pr_j^0(b|\theta, p^*, k)) \times \\ &\quad [p_{\theta j}^1 - c_{\theta j}^1 + \beta(1 - \delta(\theta))T_{\theta\theta'}(\theta)\Pi^1(p_j, p_{-j}, \theta', b)] \\ &\quad - \gamma_{\theta} \mu_{b, \theta, k}(p^*) \Pr_j^0(b|\theta, p^*, k) \Pr_j^0(n|\theta, p^*, k) \times \\ &\quad [\beta T_{\theta\theta'}(\theta)\Pi^1(p_j, p_{-j}, \theta', b)] \end{aligned} \quad (4.25)$$

The number of free supply parameters $\{c_{xj}^0, c_{\theta j}^1\}$ is equal to the number of prices set for all lenders j , and hence equal to the number of first-order conditions. The parameters are therefore just-identified and quickly converge in a procedure that minimizes squared violations of these FOCs.

4.4 Model Parameter Estimates

This subsection presents my estimates of model parameters. I emphasize two key results on the demand side of the model and two key results on the supply side.

First, on the demand side, my estimates of consumers' utility from borrowing (the parameters $d_{j\theta}$) correlate strongly with default rates across borrower types; this confirms a basic adverse selection property, that the highest-risk borrowers are also the borrowers with the greatest demand for credit. In Figure 16 I plot estimates of these flow utilities and also borrowers' average default rates, by type and by lender. The three panels of the figure correspond to three representative FICO scores, while the x-axis of each figure shows different borrower private types ψ . The evident pattern in these figures is that, across the FICO score distribution, borrower default rates are strongly correlated with demand for borrowing, with the highest-risk types also exhibiting the highest credit demand. The correlation between these two quantities across types θ ranges from .44 to as high as .88, depending on the lender. This correlation emerges mostly from the strong correlation between price and risk in the pre-CARD-Act data used to estimate the model, as these high demand parameters are revealed by consumers' willingness to borrow at those high rates.

Second, my estimates of the remaining demand parameters indicate that account set-up costs are a substantial friction for consumers, in particular limiting the degree to which many borrowers are able to refinance balances with competing lenders. Other parameters, including exit costs from borrowing and flow utilities from transacting, are only modestly important for determining consumer behavior. In Figure 17 I plot average account setup costs s , liquidity costs l , and utilities from transacting n across lenders and across borrower private types. The x-axis shows consumer FICO score groups, and the y-axis plots dollarized values of these utility parameters. For sake of comparability, these utilities are dollarized using the homogeneous marginal utilities of income (logit price coefficients) estimated in Table 10, not the heterogeneous marginal utilities presented in Figure 14. These estimates indicate that account setup costs are a substantial friction for consumers looking to switch credit cards or refinance their credit card balance with another lender, while exit costs from borrowing and demand for credit cards as transactional products are less important in driving consumer behavior. In particular, I estimate that for a FICO 700 borrower, the dollarized switch cost for setting up a new credit card account is roughly \$900, on par with the total pecuniary benefit from a typical new credit card's teaser interest rate spread over 4.5 years.⁶⁰

On the supply side of the model, I first highlight that my estimated marginal cost parameters for borrowing correspond with default rates – the primary driver of lender costs – across consumer types. I also highlight how my estimates of lenders' costs for originating new accounts correspond to industry reports of average marketing, underwriting, and processing costs associated with account origination. I present these results in Figures 18 and 19. Figure 18 follows a similar format to Figure 16 above: the panels of the figure correspond to three different FICO score groups, while the x-axis of each figure shows different borrower private types ψ . The striking pattern from the figure is that these costs are strongly correlated with default, but also that the cost estimates are not a consistent scalar multiple of default rates. On the one hand, this strong correlation indicates that the model first-order conditions are able to recover lender costs that closely follow the primary driver of actual costs, as reported in administrative data in the CCDB; this is an important validation for the model. On the other hand, the ratio between these estimated costs and data on default rates suggests a roughly 0% recovery rate on defaulted

⁶⁰This back-of-the-envelope calculation draws on the average annualized price gap between mature and new credit card accounts for FICO 700 borrowers shown in Table 6, and also the average credit card balance for FICO 700 borrowers of \$4000 dollars. These switch costs are several times larger than but still the same order of magnitude as the switch costs estimated in Shui and Ausubel (2004).

loans for the riskiest borrowers, and a closer to 100% recovery rate on defaulted loans for the safest borrowers. With industry average recovery rates around 10% and the majority of defaults generated by the riskiest borrowers, the scaling on these marginal costs therefore also appears consistent with industry benchmarks. Second, in Figure 19 I plot my estimates of lenders' costs for originating new accounts, separately by lender and across the FICO score distribution. The clear pattern in the plot is that lenders' acquisition costs are steadily increasing in FICO score; this is consistent with the extra incentives, for example airline miles, that lenders often use to encourage opening of new credit card accounts for higher FICO-score consumers. These estimates are roughly on par with industry estimates of the average cost of marketing, underwriting, and processing new accounts, which average roughly \$200 per account.⁶¹

5 Equilibrium Effects of CARD Act Price Restrictions

I now use the model developed in the previous section as a tool to study the CARD Act's pricing restrictions. I impose the Act's restrictions in the model while otherwise leaving the pre-CARD-Act environment unchanged, and I analyze these restrictions' effects on pricing, borrowing choices, and total welfare after the model converges to a new equilibrium. This exercise is informative in three ways. First, this exercise makes it possible to analyze the mechanisms behind the effects of CARD-Act-like pricing regulation. Second, I use this exercise to assess the CARD Act pricing restrictions' effects across a range of consumer types, including borrowers who choose to exit the market after the restrictions take effect. Finally, this exercise helps identify the CARD Act pricing restrictions' effects in isolation from other non-price regulation included in the Act and other contemporaneous shocks to consumer credit markets.

5.1 Modeling 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 credit card contract under the new restrictions therefore takes the form of a duple (p_j^0, p_j^1) for lender j , containing an initial teaser rate and a subsequent long-run rate.

The prices on a contract depend only on a consumer's public type (FICO score) at origination, x_0 . In particular, a contract's long run price can no longer depend on private information ψ_t revealed to a lender over the course of an account-holding relationship, as these private types are unobservable at origination. A contract's long run price also can no longer depend on updated FICO scores x_t over time. That is,

$$\begin{aligned} \text{Pre-CARD-Act: } p_1^j &= p_1^j(x_t, \psi_t) \\ \text{Post-CARD-Act: } p_1^j &= p_1^j(x_0) \end{aligned} \tag{5.1}$$

However teaser rates continue to depend only on public types at origination, as they did in the pre-CARD-Act regime.

⁶¹While industry contacts emphasize the high cost of new account acquisition, it is also plausible for the model to estimate these costs to be negative, especially on subprime accounts, reflecting fee revenue at the time of origination such as application fees that are not otherwise reflected in lender revenues in the model.

It is important to consider such teaser rates when imposing the CARD Act’s pricing restrictions in the model, because the availability of these rates implies that the Act’s price effects may differ substantially for consumers with different propensities to switch credit cards frequently. Consumer types who bear low setup costs (the demand parameter $s_{\theta j}$) on new accounts might serially transfer balances across cards to take advantage of promotional rates repeatedly, whereas consumers who bear higher setup costs are less likely to do so. Additionally, the Act may lead to less generous terms on new accounts by reducing the rents lenders are able to extract on these accounts in later periods (as in Petersen and Rajan (1995)), and including teaser rates p^0 when imposing the CARD Act’s pricing restrictions in the model provides a means to study such effects.

To emphasize, the prices set at origination are only in effect for as long as a consumer keeps a given contract. Once the consumer closes a given credit card account and opens another, the new account’s pricing reflects the consumer’s public type at the time the new contract is originated. A basic intuition explains switching behavior in this environment: all else equal, a consumer becomes more likely to switch accounts as the gap increases between (1) his current contract’s long-run price, $p_j^1(x_0)$, which was determined by his *past* public type at the time he originated this contract, and (2) a competing lender’s teaser rate on a new contract, $p_{j'}^0(x_t)$, which is determined by the consumer’s *current* public type.

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. It is in practice rare for credit card lenders to offer a *menu* of contracts to the same borrower at the same point in time, and this restriction also 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)). As my model results later confirm, this “one contract per firm per origination credit score” specification still allows substantial price dispersion at each public type, as differentiated lenders post different price duples (p^0, p^1) to each public type.

Specifying the firm’s problem in the presence of these repricing restrictions requires keeping track of the share of consumers of each type θ who hold a contract that they originated when they were of type x_0 , where x_0 is potentially different from the current public type $x(\theta)$. This requires a slight update to the notation I used in the original model exposition. Previously I used $\mu_{j,\theta,k}(p)$ to denote the share of consumers of each type θ who hold a credit card with bank j , who use that card the purpose $k \in \{\text{borrow}, \text{transact}\} \equiv \{\text{b}, \text{n}\}$, in a market where banks offer the price vector p . I now additionally keep track of the share of consumers who make each of those choices *while* holding a contract they originated at public type x_0 , which I denote $\mu_{j,\theta,x_0,k}(p)$. As before, this vector denotes the (unique) long-run distribution of consumers across contracts and choices for a given price vectpr p , where flows into a given component of $\mu_{j,\theta,x_0,k}(p)$ are equal to flows out of that component.

Bank j ’s expected discounted lifetime profits on a mature account can then be written in a

form similar to equation 4.12,

$$\begin{aligned} \Pi^1(p_j, p_{-j}, \theta, x_0, k) = & \underbrace{\Pr_j^1(b|\theta, p, x_0, k) (p_{x_0j}^1 - c_{\theta j}^1)}_{\text{flow profit for type } \theta, \text{ contract } x_0} + \\ & \underbrace{\Pr_j^1(b|\theta, p, x_0, k)\beta(1 - \delta(\theta))T_{\theta\theta'}(\theta)\Pi^1(p_j, p_{-j}, \theta', x_0, b)}_{\text{exp. cont. profit | borrow}} + \\ & \underbrace{\Pr_j^1(n|\theta, p, x_0, k)\beta T_{\theta\theta'}(\theta)\Pi^1(p_j, p_{-j}, \theta', x_0, n)}_{\text{exp. cont. profit | not borrow}} \end{aligned} \quad (5.2)$$

Expected discounted lifetime profits on new accounts are defined analogously, by making slight revisions to equation 4.11 to show continuation profits' dependence on a consumer's origination type $x_0(\theta)$,

$$\begin{aligned} \Pi^0(p_j, p_{-j}, \theta, k) = & \underbrace{\Pr_j^0(b|\theta, p, k)p_{x_0(\theta),j}^0 - c_{x_0(\theta),j}^0}_{\text{flow profit}} + \\ & \underbrace{\Pr_j^0(b|\theta, p, k)\beta(1 - \delta(\theta))T_{\theta\theta'}(\theta)\Pi^1(p_j, p_{-j}, \theta', x_0(\theta), b)}_{\text{exp. cont. profit | borrow}} + \\ & \underbrace{\Pr_j^0(n|\theta, p, k)\beta T_{\theta\theta'}(\theta)\Pi^1(p_j, p_{-j}, \theta', x_0(\theta), n)}_{\text{exp. cont. profit | not borrow}} \end{aligned} \quad (5.3)$$

and likewise by revising equation 4.13 to reflect a sum over inflows from competitors' contracts originated at types x_0 ,

$$\Pi^0(p_j, p_{-j}, x) = \sum_{j' \neq j} \sum_{\theta: x(\theta)=x} \sum_{k \in \{b, n, 0\}} \sum_{x_0} \mu_{j', \theta, x_0, k}(p) \Pr(j|p, j', \theta, x_0, k) \Pi^0(p_j, p_{-j}, \theta, k) \quad (5.4)$$

With this notation in hand, a lender's total expected discounted profits across both new and mature accounts under the restricted equilibrium can be written as,

$$\begin{aligned} \Pi_j(p_j, p_{-j}) = & \underbrace{\sum_x \Pi^0(p_j, p_{-j}, x)}_{\text{newly acquired accounts}} + \\ & \underbrace{\sum_{\theta} \sum_{x_0} \sum_{k \in \{b, n, 0\}} \mu_{j, \theta, x_0, k}(p) \times \Pi^1(p_j, p_{-j}, \theta, x_0, k)}_{\text{existing accounts}} \end{aligned} \quad (5.5)$$

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.⁶² In practice,

⁶²These 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 a new 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, which here can be seen most clearly from the fact that some bank j may find it profitable to originate contracts to a given FICO score if bank j' does not, whereas bank j' may find doing so profitable only if bank j does not. These two putative equilibria are also not unique up to the labeling of j and j' , as each bank in the data faces different costs and provides different flow utilities for each consumer type. This process of successive

I find that a market equilibrium gets close to convergence after 5 to 8 iterations of updating lenders’ best replies, while subsequent iterations are mostly needed to pin down prices on thinly traded contracts that few consumers choose in equilibrium. For market-level aggregate statistics, such as the average price paid at origination by a consumer of type x_0 , model runs therefore exhibit substantial stability quite early in this iteration process.

5.2 Equilibrium Effects of CARD Act Price Restrictions

The estimated post-CARD-Act equilibrium reveals how the two forces of market power and adverse selection trade off in different parts of the credit card market. I find that market unraveling due to adverse selection after the Act is moderately severe among the most subprime of consumers, whereas the benefits of reduced markups are dominant at higher credit scores. Nevertheless, consumer surplus conditional on credit score rises at *all* credit scores, even in credit score segments where unraveling is relatively severe, reflecting the relative importance of price decreases for the riskiest and most inelastic of borrowers. Total surplus as well as consumer surplus rises in the highest credit score segments, where surplus lost due to adverse selection is lowest.

To illustrate these effects, Figures 20 and 21 respectively show contract prices and shares of consumers who borrow on credit cards, in pre-CARD-Act data and in the estimated post-CARD-Act equilibrium. The figures are divided into three panels for three representative FICO score segments in the deep subprime part of the market (FICO 580), in the near-prime segment (FICO 680, at the cusp between subprime and prime), and in the superprime segment (FICO 780). Each panel shows prices or borrowing shares across different private-information types within the relevant FICO score group.

Turning first to panel (a) of Figure 20, there is a shift from heterogeneous pricing (a separating equilibrium) across private-information types in pre-CARD-Act data, to nearly complete pooling in the estimated post-CARD-Act equilibrium. Under this pooled pricing, all private types are now estimated to pay a fee-inclusive cost of credit in excess of 50% annualized.⁶³ Only for the very riskiest and most inelastic of private types is this a lower rate than the average they paid in pre-CARD-Act data, and all other types face higher prices than they faced before.

These high prices are an equilibrium outcome driven in part by partial unraveling, whereby the safest private-information types exit from borrowing as prices rise, and the cost of lending to only the riskiest private-information types then drives prices higher still. Turning from panel (a) of Figure 20 to panel (a) of Figure 21, the data show these exit patterns sharply. In the pre-CARD-Act data, at least 30% of all private-information types used credit cards for borrowing; among all but the highest (riskiest) quintile, the shares who borrowed were roughly equal.⁶⁴ In contrast, in the estimated post-CARD-Act equilibrium the figure shows that the safest private-information types exit almost entirely from borrowing, and even the median private-information

best-replies from the pre-CARD-Act equilibrium seems most plausible as a device to select the post-CARD-Act equilibrium (as opposed to, for example, a starting price vector where all firms charge prices of zero). For evidence that firms indeed may converge on a new equilibrium gradually after a regulatory change by playing best replies to other firms’ most recently observed pricing strategies, see Doraszelski et al. (2016).

⁶³These borrowing rates in fact track closely to some APRs seen among deep subprime credit cards in recent years, for example a 79.9% APR subprime credit card marketed in 2010 (Prater, 2010).

⁶⁴In fact, this lack of correlation between prices and borrowing share is related to the near-zero correlation between borrowing probability and prices seen previously in the OLS estimates on pre-CARD-Act data in Table 10.

type has its borrowing share fall by over two-thirds. Meanwhile, the riskiest private-information types increase their borrowing share in response to the lower prices they face, as they are pooled with the safer of their peers who remain borrowers.

Turning to panels (b) and (c) of both figures, other credit score segments do not experience the same degree of unraveling as was seen among deep subprime consumers in panel (a). First in panel (b) of Figure 20, in the FICO 680 group nearly all private information types experience lower prices as a result of lower markups in the estimated post-CARD-Act equilibrium; only the safest quintile of private-information types face higher prices while being pooled with their riskier peers. Panel (b) of Figure 21 then shows how these relative price changes affect borrowing shares across types. While the very safest private types exit somewhat from borrowing, in response to the higher prices they face in the estimated post-CARD-Act equilibrium, they do not exit to the same degree that analogously safe private types exited in panel (a). Meanwhile a greater share of all other private types borrow, reflecting these types' price decreases in the estimated post-CARD-Act equilibrium.

Panel (c) of both figures shows that the effects of reduced markups are even more pronounced at higher credit scores. In the example shown in Figure 20, in the FICO 780 group at the superprime end of the credit score distribution, all private-information types in fact face either reduced or nearly unchanged loan pricing. Correspondingly, in Figure 21, all private-information types in the FICO 780 group have greater borrowing shares in the estimated post-CARD-Act equilibrium.

While the price changes shown in Figure 20 give a sense of the long-run contract pricing faced by a consumer who *originates* a credit card at a given FICO score, some consumers hold contracts they originated in earlier periods when their credit scores differed. In particular, in the estimated post-CARD-Act equilibrium consumers can “lock in” relatively favorable rates by retaining a contract they originated at a higher credit score, as in equation (5.1). Given that borrowers are more likely to retain favorable contracts than unfavorable ones – the same adverse retention phenomenon documented previously in the reduced form results in Figure 6 – the average of all transacted prices among borrowers with a given credit score will generally be lower than contract prices for consumers who originated a contract *at* that FICO score. This phenomenon becomes clear in Figure 22, which shows average transacted prices on mature contracts at each FICO score, averaged across all private types, both in pre-CARD-Act data and in the estimated post-CARD-Act equilibrium. The average of transacted prices in the post-CARD-Act equilibrium is indeed lower than the contract prices shown in the previous Figure 20; for example, FICO 580 consumers' average prices are fully a third lower on average. Furthermore, the plot makes it clear that average *transacted* prices fall throughout the credit score distribution, reflecting both the attrition of borrowers who face price increases and the greater shares of borrowing among consumers who paid the highest prices conditional on their FICO score in pre-CARD-Act data.

Notwithstanding these decreases in transacted prices, the entry and exit patterns in Figure 21 suggest that the CARD Act's effects on consumer as well as total surplus could be ambiguous: quantities rise in some parts of the market and fall elsewhere.⁶⁵ In Figures 23 and 24, I show changes to consumer and total surplus resulting from the CARD Act price restrictions at each FICO group as well as in the market overall. As can be seen, consumer surplus rises across *all*

⁶⁵In an adversely selected market with market power, equilibrium quantities are necessarily lower than efficient levels (Mahoney and Weyl (2014)), hence an increase in quantities *may* indicate a rise in total surplus. However, total surplus can still fall when quantities rise, depending on the composition of borrowers selecting into the market.

FICO groups, reflecting the importance of markups on inelastic borrowers in determining pre-CARD-Act prices for many consumers. However, reflecting the relative importance of adverse selection among subprime consumers, the rise in subprime consumer surplus is mostly offset by a fall in lender profits on subprime accounts. In contrast, in the prime segment of the market both consumer and total surplus rise, reflecting the relative importance of pre-CARD-Act markups rather than risk adjustment among these accounts.

These consumer surplus estimates also in part reflect the insurance value of the CARD Act’s restrictions. Given that I estimated marginal utilities of income to generally rise as credit scores fall (see Figure 14), consumers prefer ex ante to shift high prices away from future states of the world where their credit scores are lower, and toward states of the world where their credit scores are higher.⁶⁶ This result also suggests analyzing the redistributive effects of the Act’s pricing restrictions relative to other policies with more explicit redistributive goals (Hendren (2017)). To illustrate these redistributive effects more concretely, Figure 25 shows ZIP-imputed income for each consumer in the data as a function of that consumer’s change in contract long-run prices under the CARD Act price restrictions.⁶⁷ Even though individual incomes can vary importantly from ZIP-code-level incomes, the figure suggests that the greatest price decreases as a result of the Act were incident on consumers who also had relatively low incomes, which emphasizes the value of exploring the Act’s insurance value and redistributive effects in future work.

Despite the insurance value of the Act’s restrictions, not all consumers with lower credit scores necessarily benefit. In part this was already seen in Figures 20 and 22, which show that although the average consumer who *holds* a credit card contract at low FICO scores benefits substantially from lower prices, consumers who wish to originate a *new* contract while holding a low credit score often face higher long-run prices on those contracts, especially if these borrowers are relatively safe private-information types. Figure 26 extends this finding to look at how lenders’ total outlays on acquiring new accounts – including the cost of both promotional teaser rates and the direct costs paid for new account acquisition⁶⁸ – differ before and after implementing the Act’s price restrictions. Consistent with the fall in long-run profitability among subprime accounts, but not among prime accounts, lender outlays for acquiring new subprime accounts fall modestly, while outlays to acquire new prime consumers increase substantially. Interestingly, this increase in outlays for new prime account acquisition matches recent trends in the post-CARD-Act credit market, where credit card issuers have invested heavily in new prime account acquisition (Kerr (2017)).

⁶⁶This demand for insurance is in spite of a weak but opposite force that can be shown analytically to result from logit demand, whereby consumers prefer to shift high prices from states of the world where borrowing has a relatively choice probability (i.e., higher credit scores) to states of the world where borrowing has a higher choice probability.

⁶⁷Whereas income is irregularly reported in the CCDB (and is drawn from credit card applications, where income is typically only self-reported), I use the availability of borrower ZIP code in the data to impute an average income at each ZIP code, using IRS Statistics on Income public data. I use IRS SOI data from the 2008 tax year, corresponding to the pre-CARD-Act equilibrium estimated in the model.

⁶⁸Because the model captures the total outlay that credit card issuers invest in customer acquisition in two distinct parameters – both the acquisition cost of new accounts, and the teaser price provided to mature accounts – I present results on how the Act changes the sum of both parameters average across new accounts in different market segments. This sum in part reflects changing market shares across lenders with different acquisition costs for new accounts.

6 Conclusion

In this paper I study the consequences of restricting lenders from adjusting borrowers' interest rates in response to information acquired over the course of lending relationships. I focus on such restrictions in the 2009 CARD Act, which I find limited lenders' ability to adjust loan pricing in response to information about risk, but also in response to information about borrower demand characteristics. Building on reduced-form evidence, I develop and estimate a model that assesses how this policy caused partial market unraveling through unpriced risk, but also reduced lenders' rents on inelastic borrowers, and I use the model to study how this tradeoff affected pricing, borrowing choices, and total welfare in the market. Model estimates also uncover new facts about the credit card market, including the correlation between demand characteristics and risk, and the importance of lenders' private information in predicting borrower default. When I impose the CARD Act's price restrictions in the model, I find that the credit card market's new equilibrium involves partial unraveling, especially on subprime accounts, but sufficiently lower rents are extracted from most borrowers, such that consumer surplus rises and, in the prime credit card market, total surplus rises as well.

One important mechanism driving these results is that the CARD Act's price restrictions effectively provide price insurance for borrowers with deteriorating risk over time. Hence even though credit cards are not insurance products per se, they involve a tradeoff between insurance value and adverse selection similar to many insurance products. [Handel et al. \(2015\)](#) and [Handel et al. \(2016\)](#) evaluate this tradeoff empirically in a simulated health insurance exchange, and they find that the insurance value of restricting firms from pricing certain types of health information can be greater than the resulting welfare costs due to adverse selection. My results reach a similar conclusion in a very different setting, where I also consider issues of lender market power due, in part, to private information that lenders learn about consumers over time. Additionally, [Handel et al. \(2015\)](#) also find that restrictions on the pricing of health status lead to more severe unraveling than I estimate in the credit card market with CARD Act pricing restrictions, perhaps reflecting the nontrivial amount of risk-based pricing still allowed under the Act.

Promising areas for future work include studying the optimality of the CARD Act's price restrictions in a broader class of possible restrictions, potentially generalized through a tax on lenders' price changes that can be designed to balance the key forces I study here. Other alternative policies that can be evaluated in my modeling framework include a weaker version of the CARD Act's pricing restrictions that would allow lenders to adjust prices in response to changes in FICO score – but not other signals from borrowers – over time, and a stronger version of the Act's restrictions that would ban promotional teaser rates in addition to the Act's other price restrictions. In the credit card market more generally, my results also motivate additional analyses on what drives the dimension of consumer risk that appears through private-information types – for example, unanticipated income shocks versus permanently heterogeneous preferences – and how lenders differentially invest in screening such private information under different regulatory regimes.

References

ABA (2013): “Response to Notice and Request for Information Regarding the Credit Card Market,” *Docket Number CFPB 2012 0048*. [2.1.2](#), [17](#)

- AGARWAL, S., S. CHOMSISENGPHET, AND C. LIU (2010): “The importance of adverse selection in the credit card market: Evidence from randomized trials of credit card solicitations,” *Journal of Money, Credit and Banking*, 42, 743–754. [4.2](#)
- AGARWAL, S., S. CHOMSISENGPHET, C. LIU, AND N. S. SOULELES (2015a): “Do consumers choose the right credit contracts?” *The Review of Corporate Finance Studies*, 4, 239–257. [3](#)
- AGARWAL, S., S. CHOMSISENGPHET, N. MAHONEY, AND J. STROEBEL (2014): “Regulating consumer financial products: Evidence from credit cards,” *The Quarterly Journal of Economics*, 130, 111–164. [1](#), [2.2.1](#), [32](#)
- (2015b): “Do banks pass through credit expansions? The marginal profitability of consumer lending during the great recession,” *Working Paper*. [1](#), [2.2.1](#), [45](#)
- AGARWAL, S., J. C. DRISCOLL, X. GABAIX, AND D. LAIBSON (2008): “Learning in the credit card market,” Tech. rep., National Bureau of Economic Research. [3](#)
- (2009): “The age of reason: Financial decisions over the life cycle and implications for regulation,” *Brookings Papers on Economic Activity*, 2009, 51–117. [3](#)
- ALEXANDROV, A., O. BEDRE-DEFOLIE, AND D. GRODZICKI (2017): “Consumer Demand for Credit Card Services,” *Working Paper*. [4](#), [2.2.1](#)
- ANGELETOS, G.-M., D. LAIBSON, A. REPETTO, J. TOBACMAN, AND S. WEINBERG (2001): “The hyperbolic consumption model: Calibration, simulation, and empirical evaluation,” *The Journal of Economic Perspectives*, 15, 47–68. [3](#)
- AUSUBEL, L. M. (1991): “The failure of competition in the credit card market,” *The American Economic Review*, 50–81. [1](#), [2](#)
- (1999): “Adverse selection in the credit card market,” Tech. rep., working paper, University of Maryland. [4.2](#)
- AVERY, R. B., P. S. CALEM, G. B. CANNER, AND R. W. BOSTIC (2003): “An overview of consumer data and credit reporting,” *Fed. Res. Bull.*, 89, 47. [24](#)
- BAR-GILL, O. AND R. BUBB (2011): “Credit card pricing: The card act and beyond,” *Cornell L. Rev.*, 97, 967. [1](#)
- BERLIN, M. AND L. J. MESTER (2004): “Credit card rates and consumer search,” *Review of Financial Economics*, 13, 179–198. [1](#)
- BREVOORT, K., D. GRODZICKI, AND M. HACKMANN (2017): “Medicaid and Financial Health,” *Unpublished mimeo*. [2.2.2](#)
- BREVOORT, K. P., P. GRIMM, AND M. KAMBARA (2016): “Credit Invisibles and the Unscored,” *Working Paper*. [2.2.2](#)
- BREVOORT, K. P. AND M. KAMBARA (2015): “Are all collections equal? The case of medical debt,” *Working Paper*. [2.2.2](#)

- BRICKER, J., L. DETTLING, A. HENRIQUES, J. HSU, L. JACOBS, K. MOORE, S. PACK, J. SABELHAUS, J. THOMPSON, AND R. WINDLE (2017): “Changes in US Family Finances from 2013 to 2016: Evidence from the Survey of Consumer Finances,” *Federal Reserve Bulletin*. [2.1.1](#)
- BRITO, D. L. AND P. R. HARTLEY (1995): “Consumer rationality and credit cards,” *Journal of Political Economy*, 103, 400–433. [1](#)
- CFPB (2013): “CARD Act Report,” Tech. rep., Consumer Financial Protection Bureau. [16](#), [17](#), [2.2.1](#), [32](#)
- (2015): “The Consumer Credit Card Market,” Tech. rep., Consumer Financial Protection Bureau. [2.2.1](#)
- COHEN, A. AND L. EINAV (2007): “Estimating risk preferences from deductible choice,” *The American economic review*, 97, 745–788. [4.2](#)
- CONSUMERACTION (2007): “2007 Credit Card Survey,” Tech. rep., Consumer Action News. [10](#)
- COX, N. (2017): “Pricing, Selection, and Welfare in the Student Loan Market: Evidence from Borrower Repayment Decisions,” . [42](#)
- DEBBAUT, P., A. GHENT, AND M. KUDLYAK (2016): “The CARD Act and Young Borrowers: The Effects and the Affected,” *Journal of Money, Credit and Banking*, 48, 1495–1513. [1](#)
- DELLAVIGNA, S. AND U. MALMENDIER (2004): “Contract design and self-control: Theory and evidence,” *The Quarterly Journal of Economics*, 119, 353–402. [3](#)
- DORASZELSKI, U., G. LEWIS, AND A. PAKES (2016): “Just starting out: Learning and equilibrium in a new market,” Tech. rep., National Bureau of Economic Research. [62](#)
- DROZD, L. AND R. SERRANO-PADIAL (2014): “On Prepayment and Rollover Risk in the US Credit Card Market,” in *2014 Meeting Papers*, Society for Economic Dynamics, 215. [5](#)
- EDELBERG, W. (2006): “Risk-based pricing of interest rates for consumer loans,” *Journal of Monetary Economics*, 53, 2283–2298. [2.1.1](#)
- EINAV, L., A. FINKELSTEIN, AND P. SCHRIMPF (2010): “Optimal mandates and the welfare cost of asymmetric information: Evidence from the uk annuity market,” *Econometrica*, 78, 1031–1092. [4.2](#), [4.3.1](#)
- EINAV, L., M. JENKINS, AND J. LEVIN (2012): “Contract pricing in consumer credit markets,” *Econometrica*, 80, 1387–1432. [2.2.1](#)
- EVANS, D. S. AND R. SCHMALENSEE (2005): *Paying with plastic: the digital revolution in buying and borrowing*, Mit Press. [2.1.1](#)
- FINKELSTEIN, A. AND J. POTERBA (2004): “Adverse selection in insurance markets: Policyholder evidence from the UK annuity market,” *Journal of Political Economy*, 112, 183–208. [4.3.1](#)

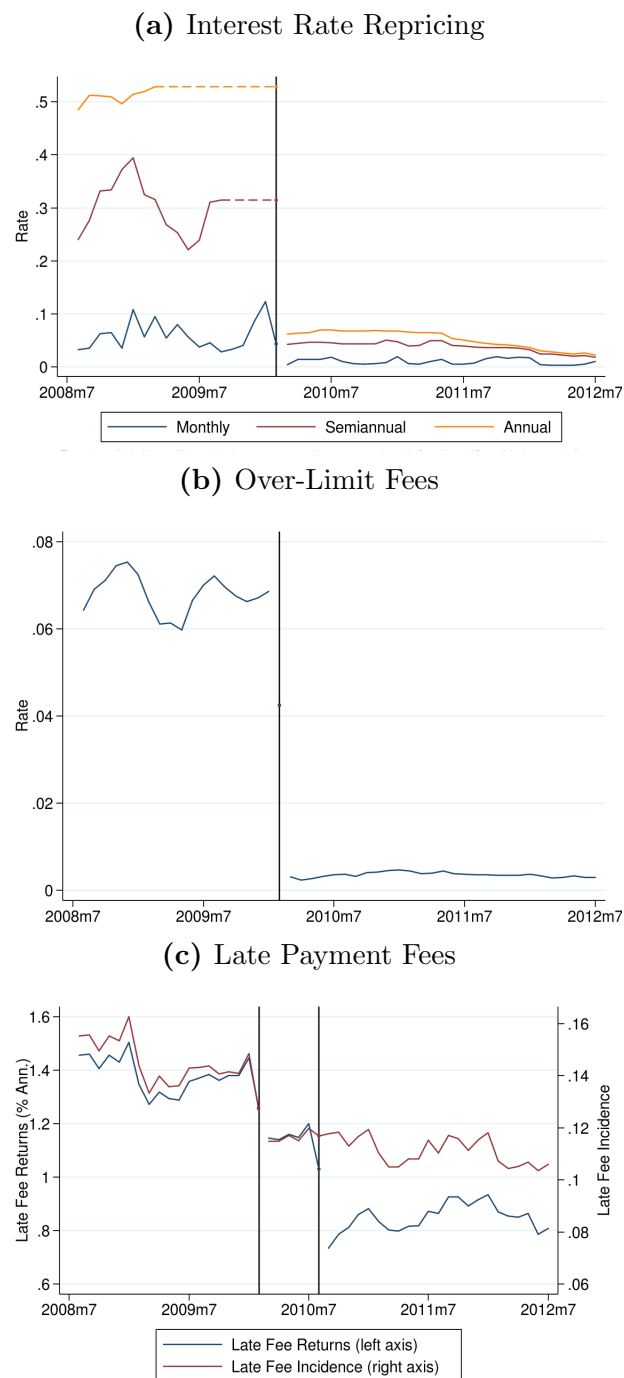
- FRB (2008): “Comment from Oliver I. Ireland,” Tech. rep., Federal Reserve Board of Governors. [14](#)
- (2010): “Report to the Congress on Reductions of Consumer Credit Limits Based on Certain Information as to Experience or Transactions of the Consumer,” Tech. rep., Federal Reserve Board of Governors. [2.1.1](#)
- FULFORD, S. L. (2015): “How important is variability in consumer credit limits?” *Journal of Monetary Economics*, 72, 42–63. [1](#)
- GRODZICKI, D. (2012): “The evolution of competition in the credit card market,” Tech. rep., Working paper. Stanford University, Department of Economics, Palo Alto, Calif. [2](#)
- (2014): “Competition and Customer Acquisition in the US Credit Card Market,” *Working Paper*. [8](#)
- GROSS, D. B. AND N. S. SOULELES (2002): “Do liquidity constraints and interest rates matter for consumer behavior? Evidence from credit card data,” *The Quarterly journal of economics*, 117, 149–185. [1](#), [4](#)
- GROSS, T., M. J. NOTOWIDIGDO, AND J. WANG (2016): “The marginal propensity to consume over the business cycle,” Tech. rep., National Bureau of Economic Research. [1](#), [2.2.1](#)
- GRUBB, M. D. (2009): “Selling to overconfident consumers,” *The American Economic Review*, 99, 1770–1807. [3](#)
- HAN, S., B. J. KEYS, AND G. LI (2015): “Information, contract design, and unsecured credit supply: Evidence from credit card mailings,” . [1](#)
- HANDEL, B., I. HENDEL, AND M. WHINSTON (2016): “The welfare impact of long-term contracts,” Tech. rep., Working Paper. [6](#)
- HANDEL, B., I. HENDEL, AND M. D. WHINSTON (2015): “Equilibria in health exchanges: Adverse selection versus reclassification risk,” *Econometrica*, 83, 1261–1313. [6](#)
- HANDEL, B. R. (2013): “Adverse selection and inertia in health insurance markets: When nudging hurts,” *The American Economic Review*, 103, 2643–2682. [1](#)
- HEIDHUES, P. AND B. KOSZEGI (2010): “Exploiting naivete about self-control in the credit market,” *The American Economic Review*, 100, 2279–2303. [3](#)
- HEIDHUES, P. AND B. KÓSZEGI (2015): “On the welfare costs of naiveté in the US credit-card market,” *Review of Industrial Organization*, 47, 341–354. [3](#)
- HENDREN, N. (2017): “Efficient Welfare Weights,” . [5.2](#)
- HUNT, R. M. AND K. SERFES (2013): “Dynamic Pricing of Credit Cards and the Effects of Regulation,” *Working Paper*. [1](#)
- HYMAN, L. (2011): *Debtor nation: The history of America in red ink*, Princeton University Press. [2.1.1](#)

- ILLANES, G. (2016): “Switching Costs in Pension Plan Choice,” Tech. rep., Working Paper. 1
- JAMBULAPATI, V. AND J. STAVINS (2014): “Credit CARD Act of 2009: What did banks do?” *Journal of Banking & Finance*, 46, 21–30. 1
- KERR, R. (2017): “Why US Credit Card Bonuses Are So High – And Why That May Change,” *The Points Guy*. 5.2
- KEYS, B. J. AND J. WANG (2016): “Minimum Payments and Debt Paydown in Consumer Credit Cards,” Tech. rep., National Bureau of Economic Research. 1, 2.2.1
- LEVITIN, A. J. (2011): “Rate-Jacking: Risk-Based & Opportunistic Pricing in Credit Cards,” *Utah L. Rev.*, 339. 1, 2.1.2
- MAHONEY, N. AND E. G. WEYL (2014): “Imperfect competition in selection markets,” *Review of Economics and Statistics*. 65
- MANDEL, L. (1990): “The Credit Card Industry,” *A History*. Boston, Twayne Publishers. 2.1.1
- PETERSEN, M. A. AND R. G. RAJAN (1995): “The effect of credit market competition on lending relationships,” *The Quarterly Journal of Economics*, 110, 407–443. 5.1
- PINHEIRO, T. AND J. RONEN (2015): “Unintended Consequences of the Credit Card Act,” *Working Paper*. 1
- RU, H. AND A. SCHOAR (2016): “Do credit card companies screen for behavioral biases?” Tech. rep., National Bureau of Economic Research. 3
- RUST, J. (1994): “Structural estimation of Markov decision processes,” *Handbook of econometrics*, 4, 3081–3143. 47
- RYAN, S. P. (2012): “The costs of environmental regulation in a concentrated industry,” *Econometrica*, 80, 1019–1061. 62
- SANTUCCI, L. (2015): “A tale of two vintages: credit limit management before and after the CARD act and Great Recession,” *Working Paper*. 1
- SHUI, H. AND L. M. AUSUBEL (2004): “Time inconsistency in the credit card market,” . 60
- SOMAINI, P. (2011): “Competition and interdependent costs in highway procurement,” *Unpublished manuscript*. [394, 395]. 2.2.1
- STANGO, V. AND J. ZINMAN (2015): “Borrowing high versus borrowing higher: price dispersion and shopping behavior in the US credit card market,” *The Review of Financial Studies*, 29, 979–1006. 30
- STAVINS, J. (1996): “Can demand elasticities explain sticky credit card rates?” *New England Economic Review*, 43–55. 1
- STEPNER, M. ET AL. (2013): “BINSCTTER: Stata module to generate binned scatterplots,” *Statistical Software Components*. 34

- STOLE, L. A. (2007): “Price discrimination and competition,” *Handbook of industrial organization*, 3, 2221–2299. [5.1](#)
- TELYUKOVA, I. A. AND R. WRIGHT (2008): “A model of money and credit, with application to the credit card debt puzzle,” *The Review of Economic Studies*, 75, 629–647. [1](#)
- TIAN, X. S. AND H. ZHANG (2016): “Impact of FAS 166/167 on credit card securitization,” *Working Paper*. [41](#)
- TRENCH, M. S., S. P. PEDERSON, E. T. LAU, L. MA, H. WANG, AND S. K. NAIR (2003): “Managing credit lines and prices for bank one credit cards,” *Interfaces*, 33, 4–21. [21](#), [45](#)

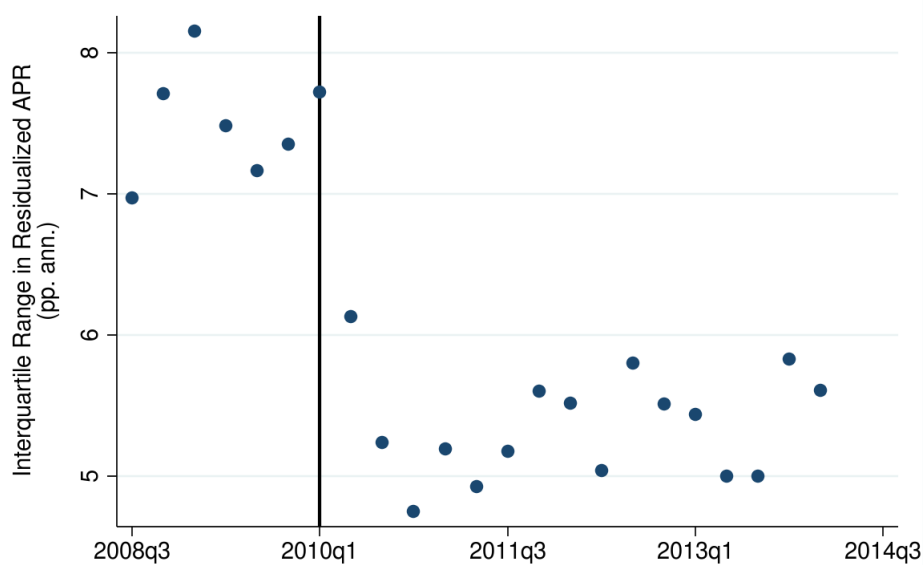
7 Figures

Figure 1: Direct Price Effects at CARD Act Implementation



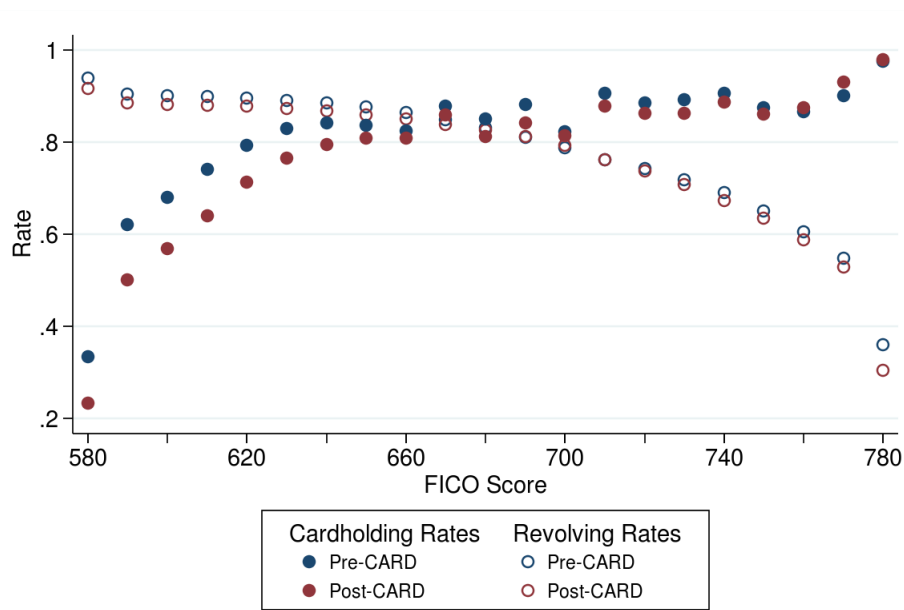
Notes: Panel (a) shows the incidence of interest rate increases on current borrowers over 1-month, 6-month, and 12-month horizons, excluding interest rate increases permitted by the CARD Act (i.e., increases coinciding with the expiration of a promotional rate, with changes in an index rate, or with 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 interest rate repricing restrictions in February 2010, which is marked by the vertical black line. Panel (b) shows the monthly incidence of over-limit fees on current borrowers, excluding any fees subsequently reversed. The implementation date of the CARD Act's over-limit fee restrictions in February 2010 is marked by the vertical black line. Panel (c) shows annualized lender returns from late fees relative to total outstanding balances on borrowing accounts (left axis) and the average incidence of late fees across accounts (right axis). The vertical black lines show the CARD Act's implementation dates for restrictions on interest-rate increases and over-limit fees in February 2010 and for restrictions on late fee amounts in August 2010.

Figure 2: Interquartile Ranges in Credit Card APRs by Vintage



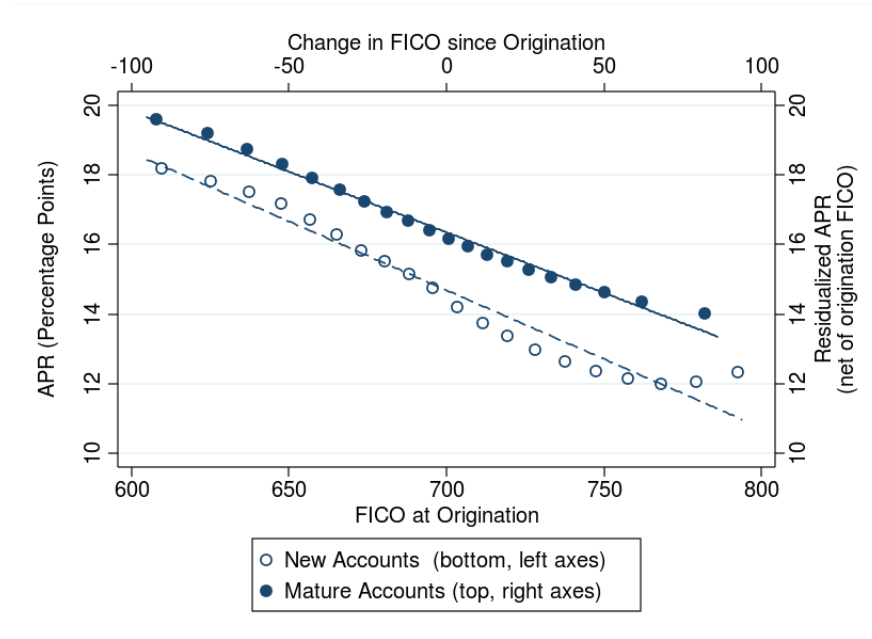
Notes: The figure shows the interquartile range (IQR) of annual percentage rates on borrowing accounts by origination cohort, after partialling out origination credit score and origination month. The date shown for each cohort is its age of maturity (18 months), by which point introductory promotional 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 line shows the date of implementation for the CARD Act's restrictions on interest rate increases, in February 2010.

Figure 3: Prevalence of Cardholding and Borrowing
Pre- and Post-CARD Act



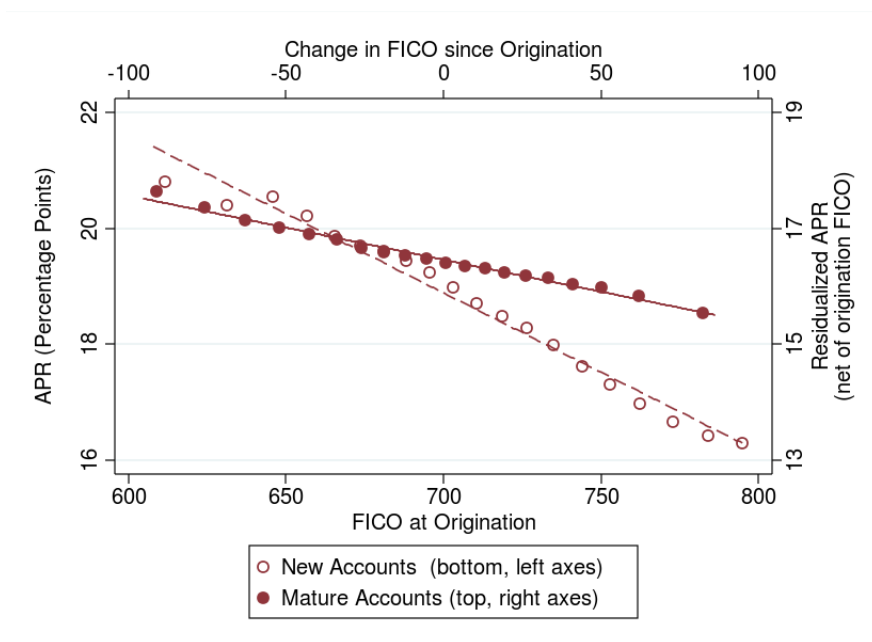
Notes: The figure shows the rate of credit card-holding among individuals in each credit score bin (CCP data) and the share of active credit card accounts used for borrowing (CCDB data), in pre- and post-CARD-Act periods (2008Q3 to 2009Q2, and 2011Q3 to 2014Q2, respectively). Borrowing is defined as not paying a balance in full for two successive billing cycles.

Figure 4: Pre-CARD-Act Price Gradients
for Origination Risk and Emergent Risk



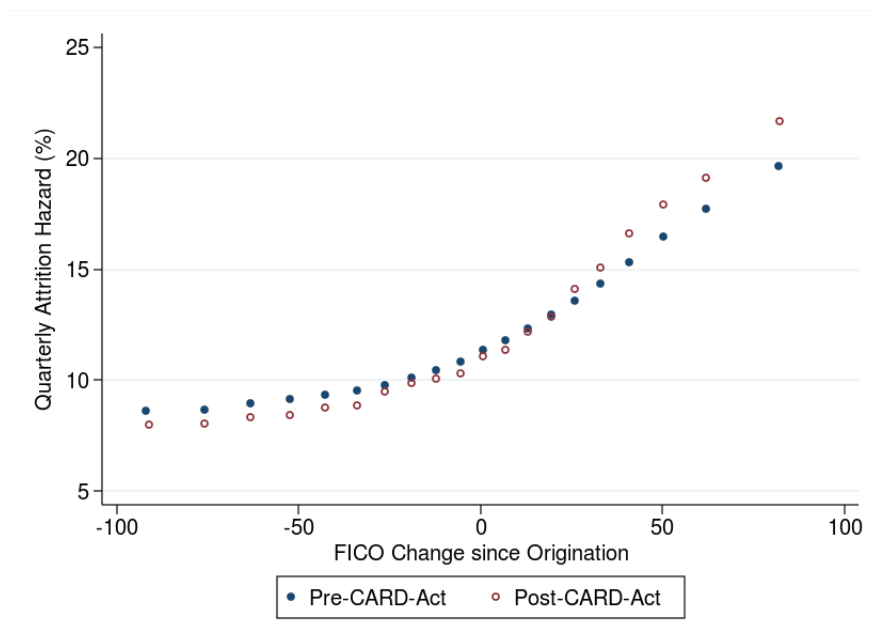
Notes: The figure shows two different gradients of risk in the pre-CARD-Act era (2008Q3 to 2009Q2) on two pairs of axes. On the left, bottom axes, the figure plots the average annual percentage rate (APR) on newly originated accounts across quantiles of the credit score distribution, together with a line of best fit. On the right, 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.

Figure 5: Post-CARD-Act Price Gradients for Origination Risk and Emergent Risk



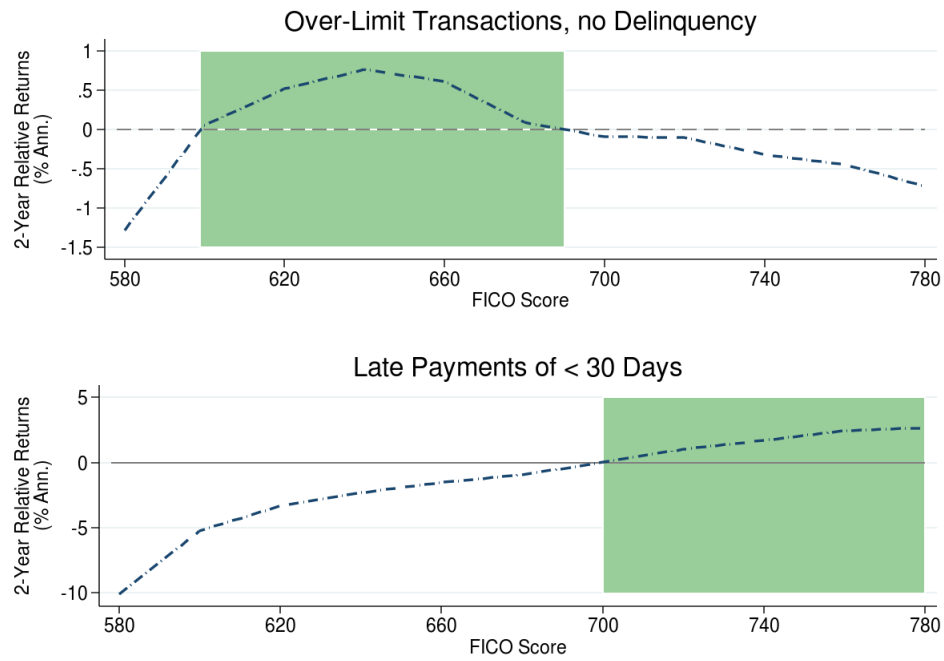
Notes: The figure presents the same price-risk gradients as in Figure 4 but in post-CARD-Act data (2011Q3 to 2014Q2). The two y-axes have the same axis scale, but the axis ranges are shifted to facilitate comparison of the two gradients. See notes to Figure 4 for further detail.

Figure 6: Adverse Retention in Response to Risk Mispricing



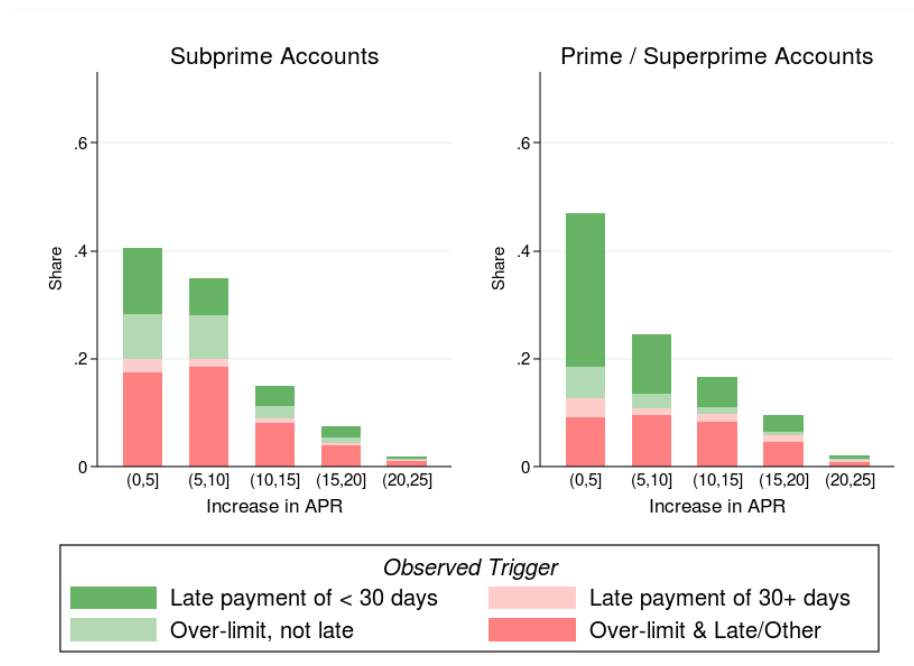
Notes: The figure plots quarterly attrition rates from borrowing (including both attrition through account closure and also attrition through paying off a credit card's balance) across quantiles of borrowing accounts' changes in FICO score since origination, separately in pre-CARD-Act data and post-CARD-Act data (2008Q3 to 2009Q2 and 2011Q3 to 2014Q2, respectively). See equation 3.3 in the text.

Figure 7: Pre-CARD Act Price Elasticity Signals



Notes: The figure highlights two commonly used triggers for interest rate increases that I identify as price elasticity signals (see equation 3.5 in the text): over-limit transactions not coinciding with delinquency, and late payments of less than thirty days. The plotted line shows the change in lenders' expected returns after observing the relevant signal on an account, relative to expected returns on accounts that send no particular signal (behavior "0" in equation 3.5), as a function of accounts' credit score. Green shading emphasizes the credit score segments where behaviors are identified as price elasticity signals.

Figure 8: Causes of Triggered Repricing



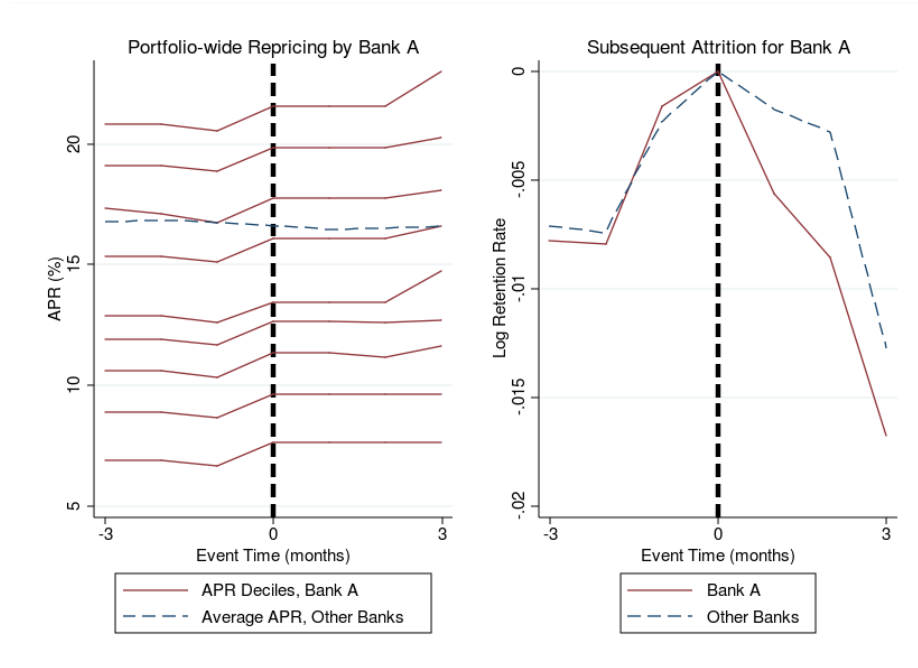
Notes: The figure shows a decomposition of interest rate increases in pre-CARD-Act data (2008Q3 - 2009Q2) across various standard triggers that may coincide with an interest rate increase. This decomposition is shown separately for subprime and prime accounts (left and right panels) and separately by the size of the APR increase (grouped across the x-axes). Color shading emphasizes which triggers are behaviors that predict higher vs. lower lender returns, with the darkest green showing the highest future returns and the darkest red showing the most negative future returns on average across accounts. See Figure 8 and Table 3 for evidence on which signals predict higher and lower future returns.

Figure 9: Causes of Discretionary Repricing



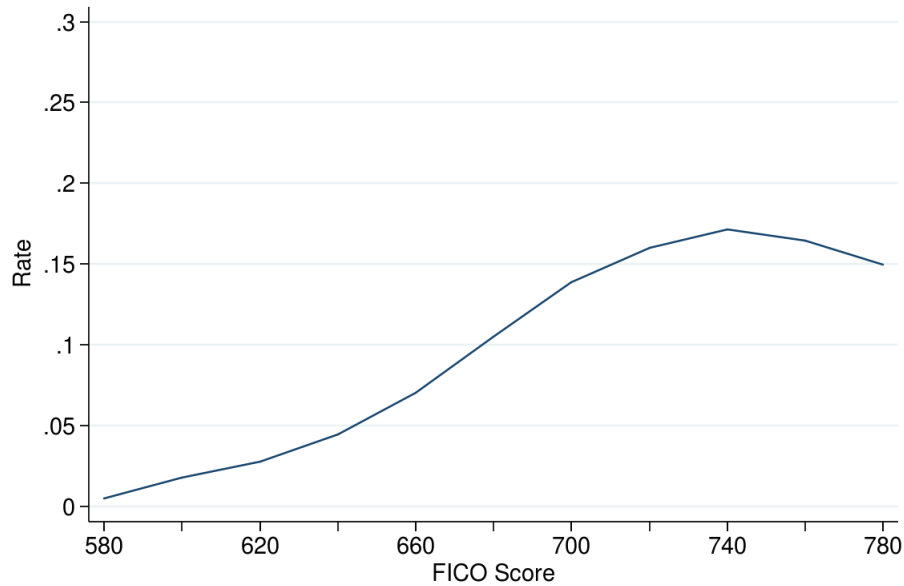
Notes: The figure presents evidence on how credit scores are evolving at the time of non-triggered interest rate increases. Whereas some interest rate increases are “triggered” by a behavior typically delineated in credit card contracts as a justification for raising interest rates, such as over-limit transactions or late payments, lenders can also impose “non-triggered” interest rate increases through contract provisions that allow interest rates to be increased “at any time, for any reason” in the absence of any particular trigger. These interest rate increases are shown separately by size, in percentage points, across the x-axis. The shading illustrates the one-quarter change in credit score preceding the interest rate increase, with darker green indicating declining risk for lenders and darker red shades indicating rising risk.

Figure 10: Example of Repricing Quasi-Experiment



Notes: The figure plots an example of a repricing quasi-experiment (left panel) and subsequent attrition from borrowing (right panel) from the pre-CARD-Act data. 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 attrition rates from borrowing are shown relative to their value in event time 0 for Bank A and for all other banks. Here attrition includes attrition through paying off a balance, through refinancing with another lender, or through closing a card. See equation 4.1 in the text.

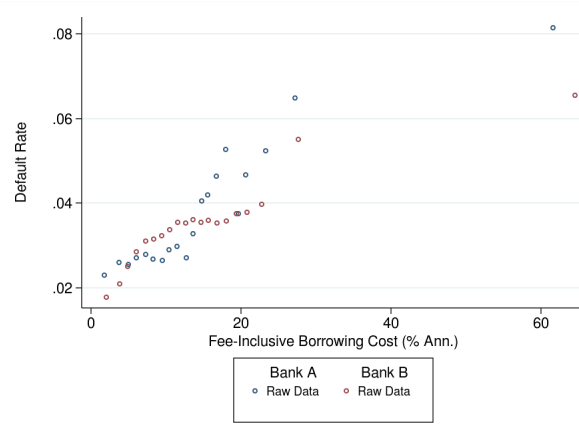
Figure 11: Prevalence of Balance Transfer Activity by FICO Score Pre-CARD Act



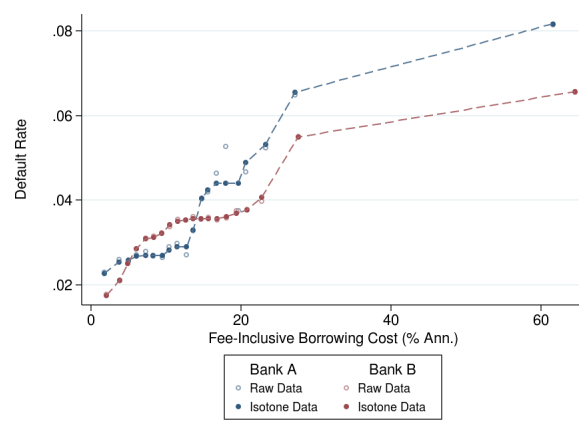
Notes: The figure shows the rate of balance transfers by credit score, calculated as the ratio of incoming balance transfers at promotional rates or on newly originated accounts, to the number of mature borrowing accounts without promotional rates in effect. Borrowing is defined as not paying a balance in full for two subsequent billing cycles.

Figure 12: Recovering Private-Information Types from Equilibrium Pricing

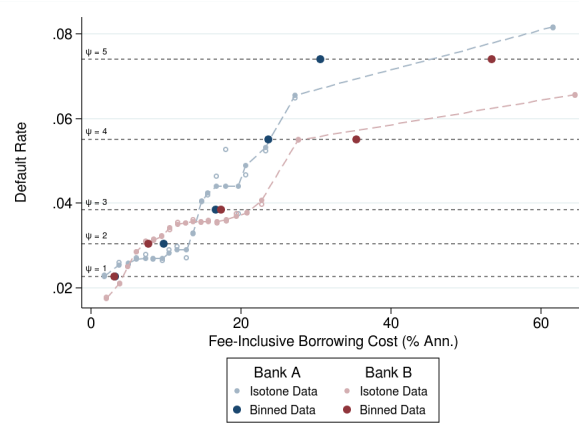
(a) Step 1: Inverse Pricing Functions for Ex-Post Default



(b) Step 2: Isotonic Inverse Pricing Functions

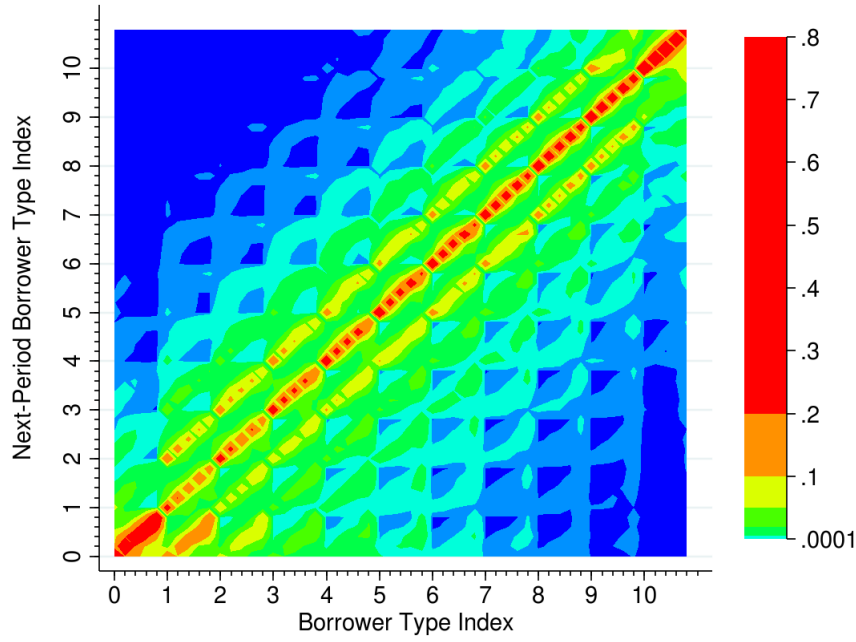


(c) Step 3: Discretizing Private Types ψ from Pricing Functions



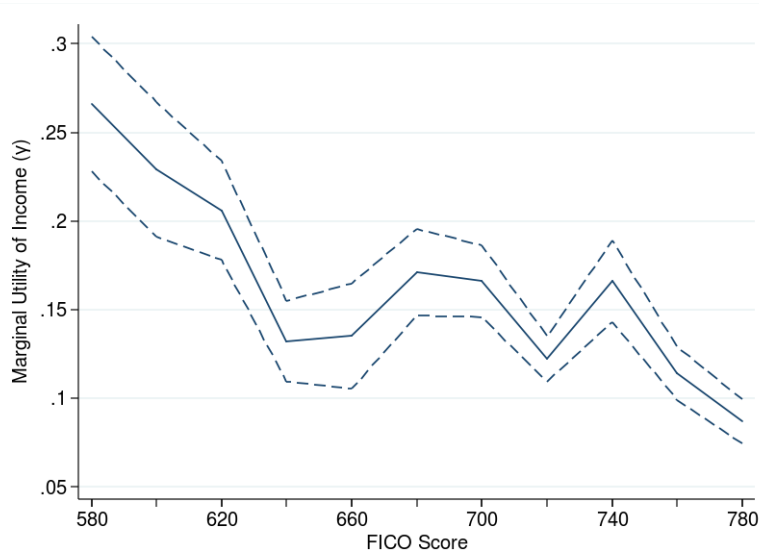
Notes: The figure illustrates the process of recovering private-information types from observed equilibrium pricing in pre-CARD-Act data, as described in equations 4.16 and 4.17 in the text. This example is taken from the market segment defined by the credit score range 720-739. Panel (a) shows raw data on observed default rates at quantiles of price levels on two different banks, labeled Bank A and Bank B. 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) for sake of comparison. 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.

Figure 13: Transition Rates Among Public and Private Types



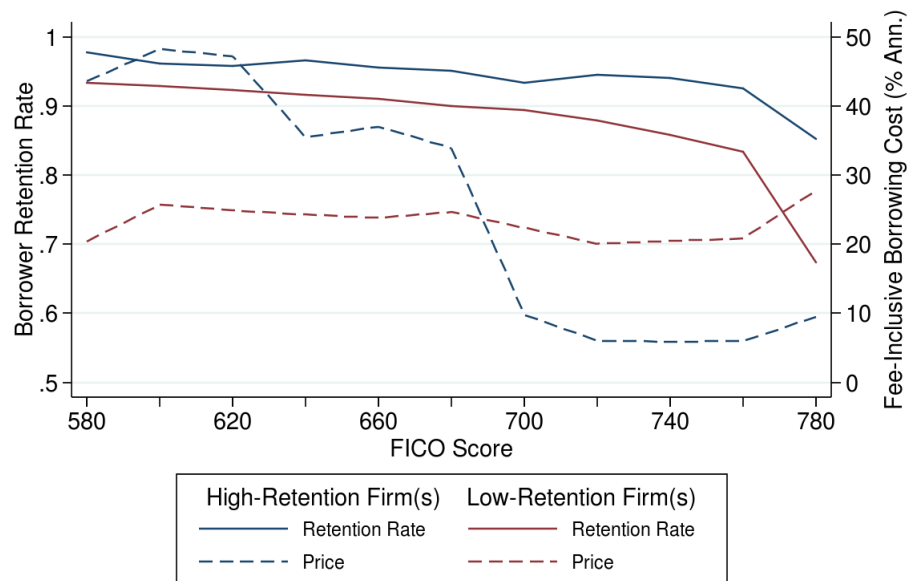
Notes: The figure displays a contour plot of period-to-period transition probabilities among consumer types. These probabilities are estimated quarterly among borrowers observed for two subsequent quarters, using the joint of public and private types recovered through the process illustrated in Figure 12. 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, the range $[1,2)$ corresponds to the 600-619 FICO score group, and so-on. Within integers, the sub-ticks correspond to the five private-information types recovered at each FICO score level, in order of increasing risk.

Figure 14: Heterogeneity in Price Coefficients



Notes: The figure displays estimates of heterogeneous price coefficients (marginal utilities of income γ_x) across FICO score, estimated via equation 4.23. Dotted lines display 95% confidence bands.

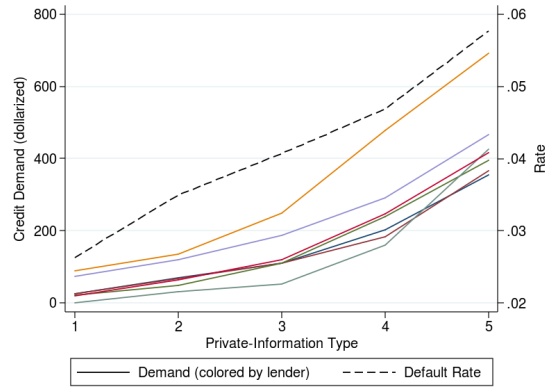
Figure 15: Identification of Demand Parameters



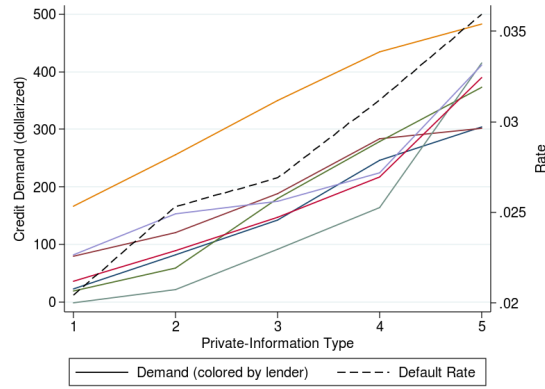
Notes: The figure shows borrower retention rates for the highest-retention and lowest-retention credit card issuers at each 20-point credit score group. Hence the retention lines are upper and lower envelopes across the market, not the set of retention rates for any single firm. For each firm included in these envelopes, corresponding prices are shown in the dotted lines. Results are shown for the median private-information type in each FICO score group.

Figure 16: Borrowing Demand and Default Rates by Consumer Type

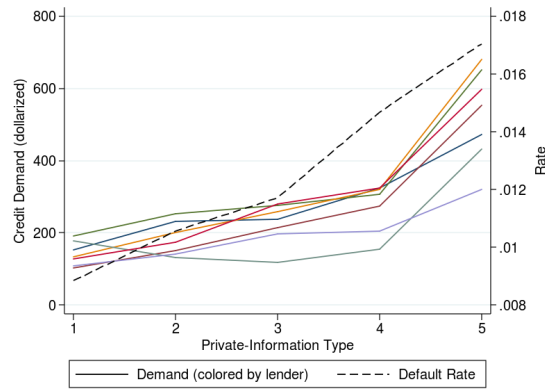
(a) FICO 620-639 Consumers



(b) FICO 660-679 Consumers

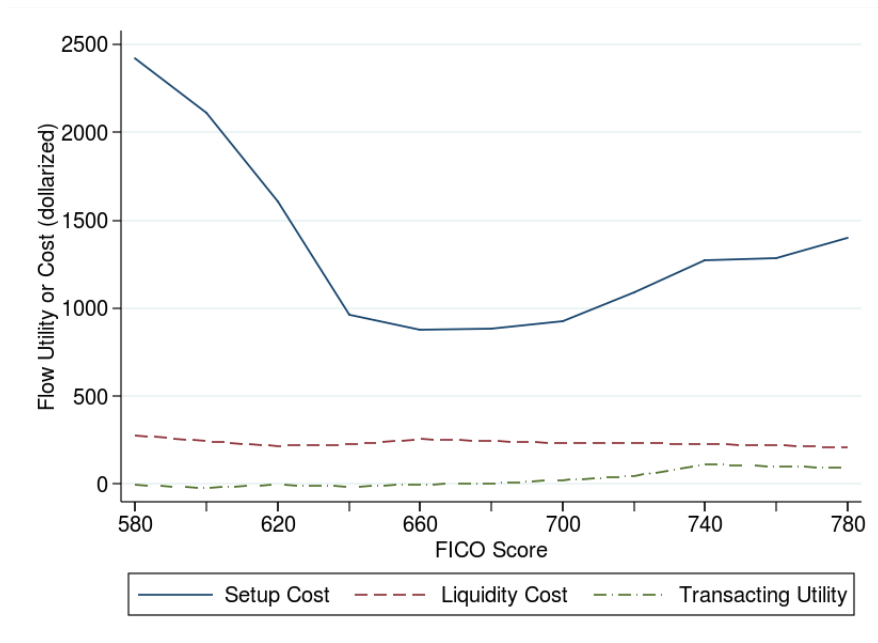


(c) FICO 720-739 Consumers



Notes: The figure shows estimates of consumer types' flow utilities from borrowing, together with these types' default rates. Consumer types $\theta = (x, \psi)$ are shown separately by private-information type ψ (across the x-axes) and by public type x , i.e., credit score group (three selected groups are shown separately in the three panels). Flow utilities (the parameter $d_{\theta j}$) are plotted separately by lender j in solid lines. These flow utilities are dollarized using each type's marginal utility of income (the price coefficient γ_{θ}) and using average borrowed balances for that credit score group. Default rates measure the probability of being 90+ days delinquent at a quarterly horizon.

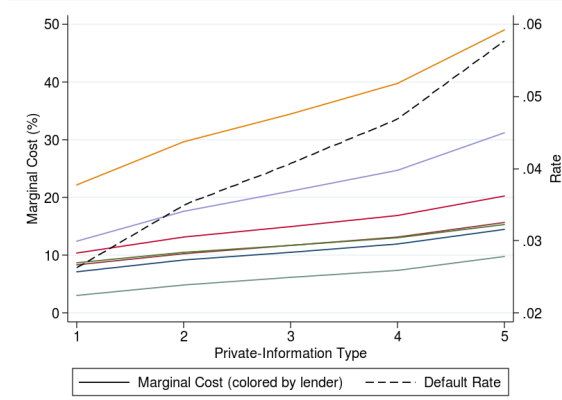
Figure 17: Setup Costs, Exit Costs and Transacting Demand



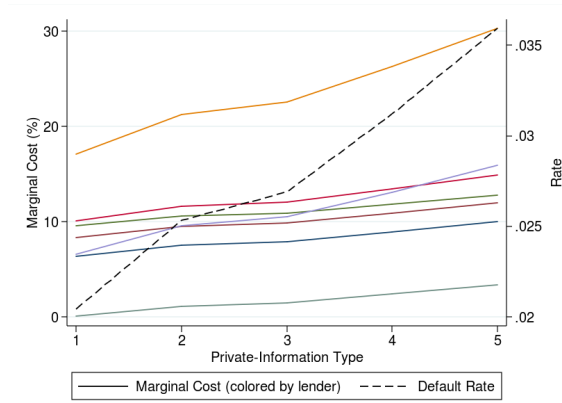
Notes: The figure shows estimates of flow utilities from transacting, liquidity costs to paying off a balance, and set-up costs for opening a new account, separately by 20-point bin of credit score. Parameters that are estimated separately by lender and by private-information type are averaged within credit-score group, using pre-CARD-Act market share weighting by lender and the probability distribution μ_θ across private types. Parameters are dollarized using a population-average marginal utility of income, estimated in column (1) of Table 10, and using average borrowed balances for each credit score group.

Figure 18: Marginal Costs and Default Rates by Consumer Type

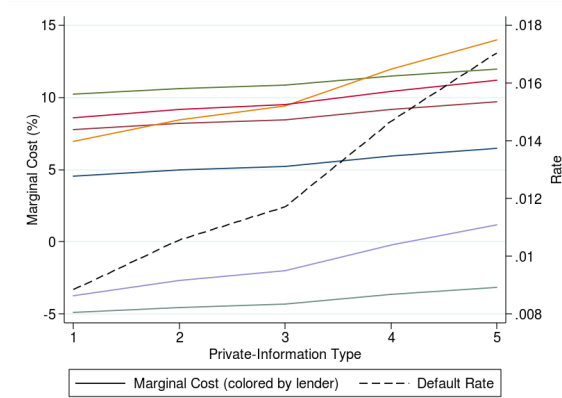
(a) FICO 620-639 Consumers



(b) FICO 660-679 Consumers

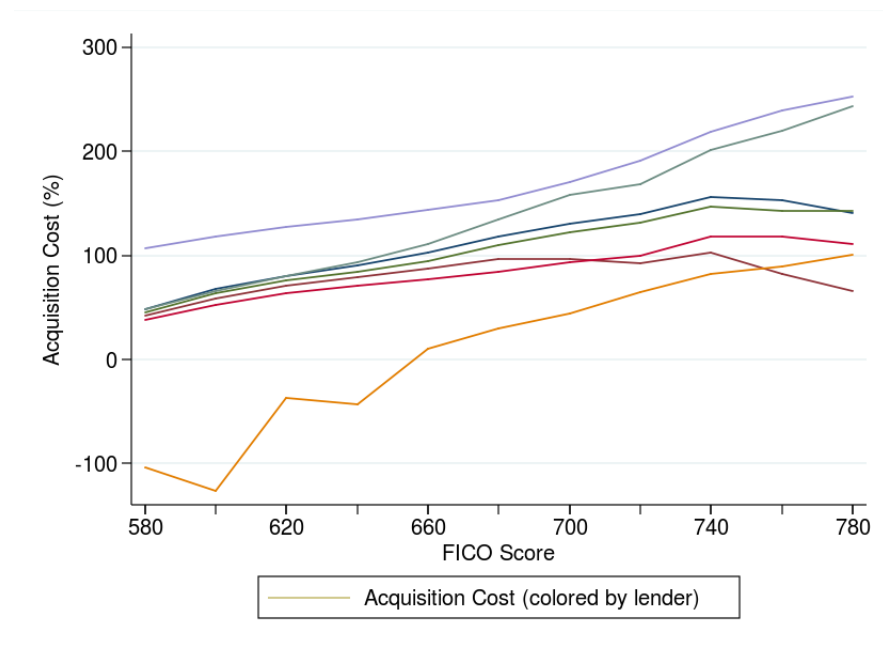


(c) FICO 720-739 Consumers



Notes: The figure shows estimates of firms' marginal cost of lending to each consumer type in three selected credit score groups, together with these types' default rates. Marginal costs are expressed as an annualized percentage of average borrowed balances, and default rates measure the probability of being 90+ days delinquent at a quarterly horizon.

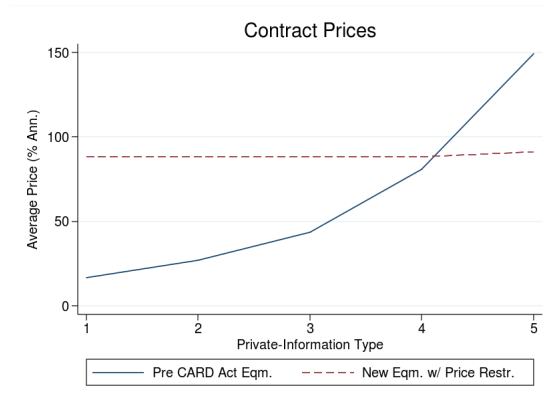
Figure 19: Consumer Acquisition Costs



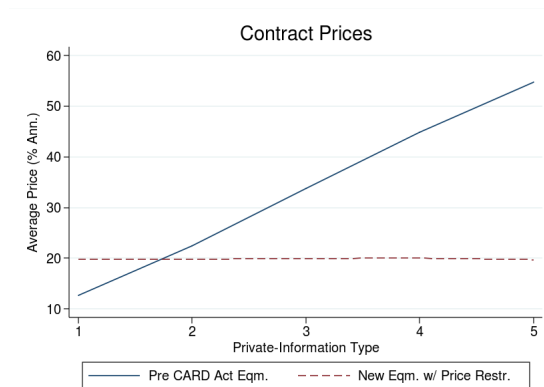
Notes: The figure shows estimates of firms' per-account acquisition cost for consumers in each 20-point credit score group, expressed as an annualized percentage of average borrowed balances in that credit score group. Occasional estimates of negative acquisition costs may reflect fee revenue at the time of account origination, such as application fees, as discussed in footnote 61.

Figure 20: Equilibrium Changes in Contract Pricing with CARD-Act Pricing Restrictions

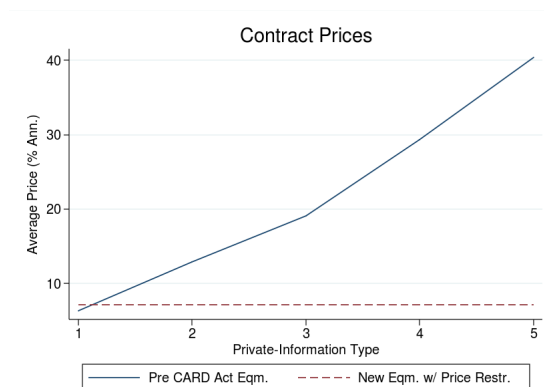
(a) FICO 580-599 Consumers



(b) FICO 680-699 Consumers



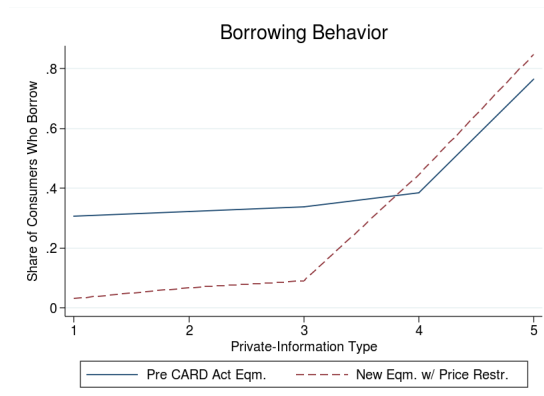
(c) FICO 780+ Consumers



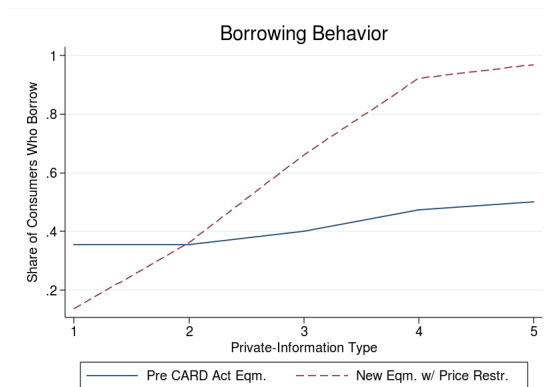
Notes: The figure shows observed average contract prices for each consumer type in three selected credit score groups in the pre-CARD-Act equilibrium, together with model results for these types' equilibrium contract prices after imposing the CARD Act price restrictions. The prices shown are annualized, account-level averages at a quarterly frequency inclusive of both interest charges and fees, normalized by the amount borrowed. This price measure is described in Section 2.2.3 of the text.

Figure 21: Equilibrium Entry/Exit from Borrowing with CARD-Act Pricing Restrictions

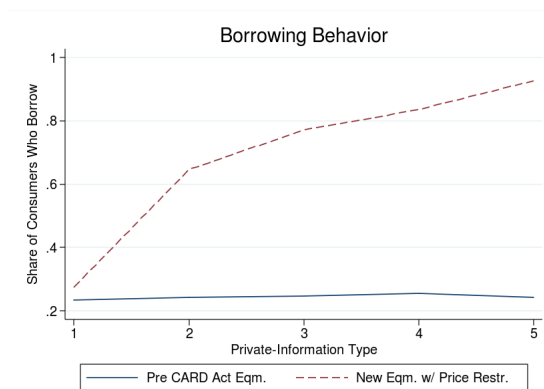
(a) FICO 580-599 Consumers



(b) FICO 680-699 Consumers

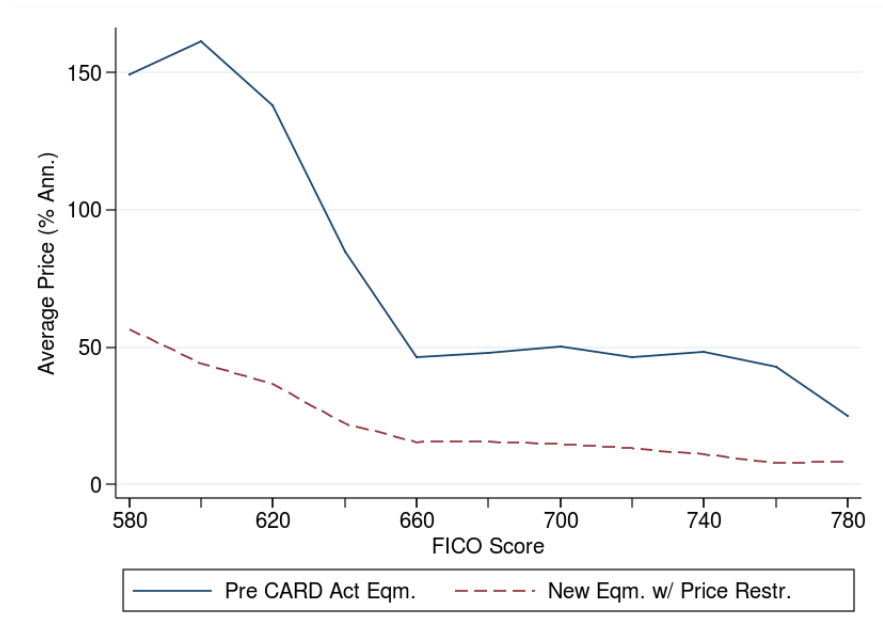


(c) FICO 780+ Consumers



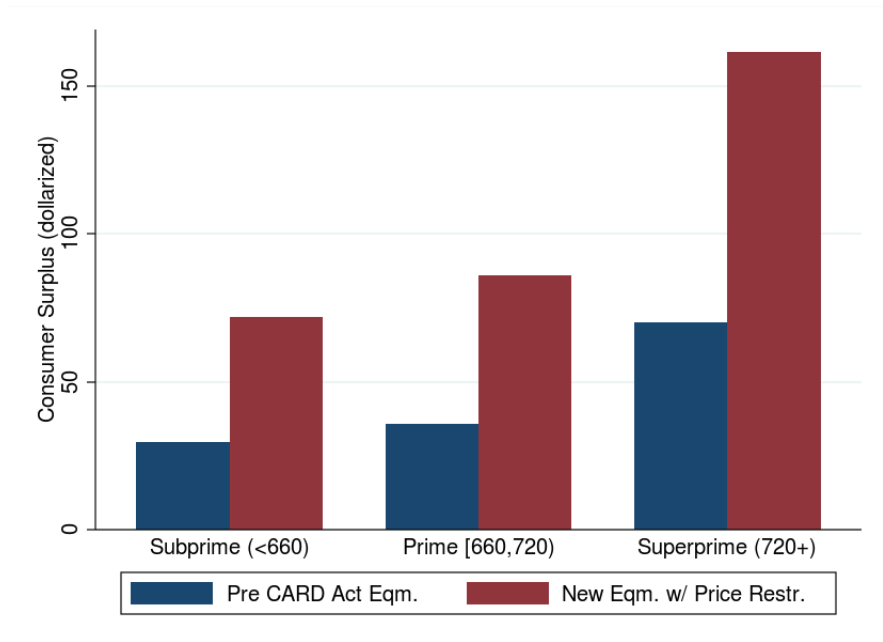
Notes: The figure shows the share of consumers who use a credit card for borrowing among various consumer types. Shares range from 0 to 1. Shares for the new equilibrium with price restrictions reflect the effect of CARD Act price restrictions when implemented in the model, holding constant other parameter estimates from the *pre*-CARD-Act equilibrium. Private-information types are shown across the x-axis of each panel and the three panels show three selected public information (credit-score) groups.

Figure 22: Changes in Transacted Contract Prices



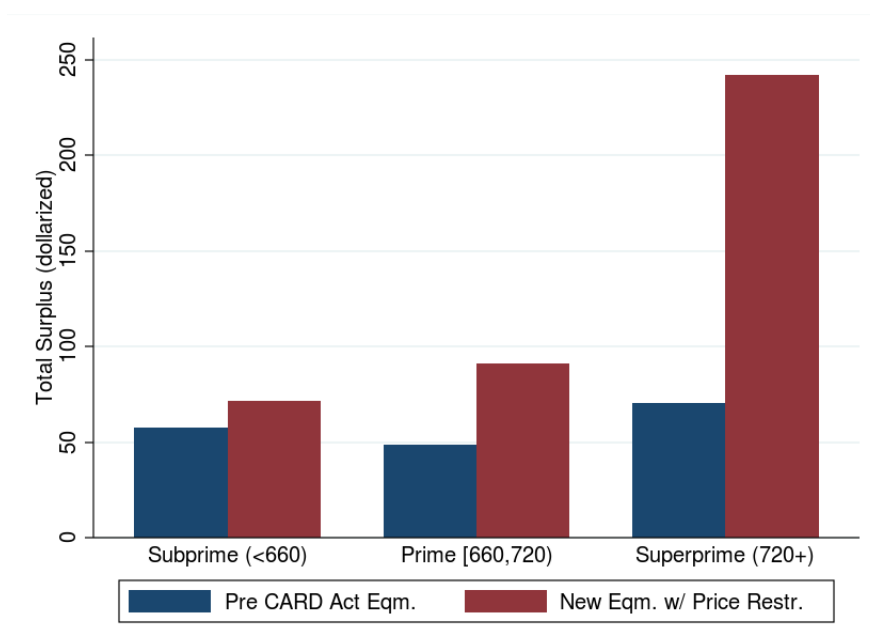
Notes: The figure shows changes in *transacted* long-run contract prices across FICO scores on the x-axis. Consumers who exit the market are therefore not counted in the new equilibrium with price restrictions. Prices shown are weighted averages across private types and across lenders.

Figure 23: Changes in Consumer Surplus



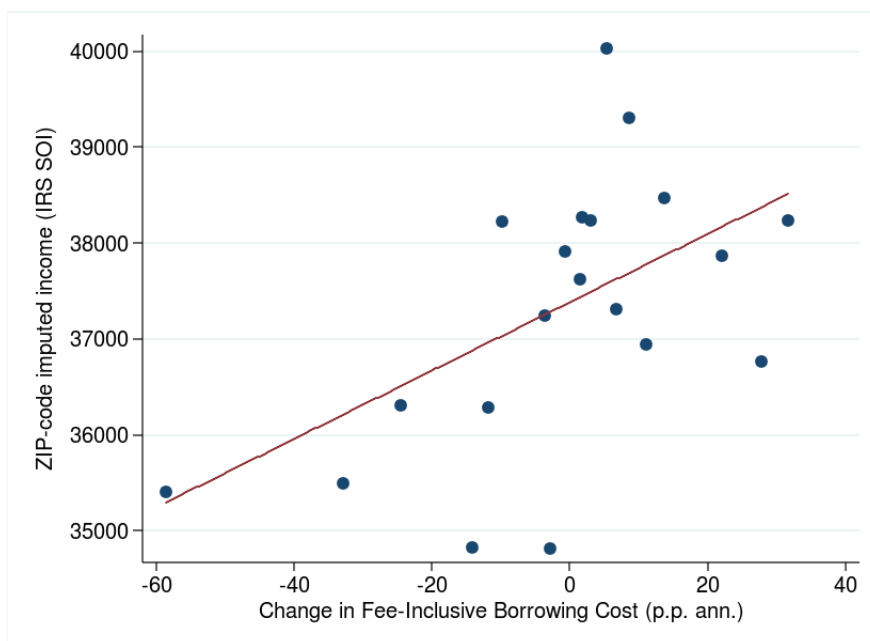
Notes: The figure shows estimated per-person consumer surplus (including both borrowers and non-borrowers) in the pre-CARD-Act equilibrium and also in the new equilibrium found in the model after imposing the CARD Act price restrictions. Surplus is dollarized using each type's marginal utility of income (the price coefficient γ_θ) and using average 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 .

Figure 24: Changes in Total Surplus



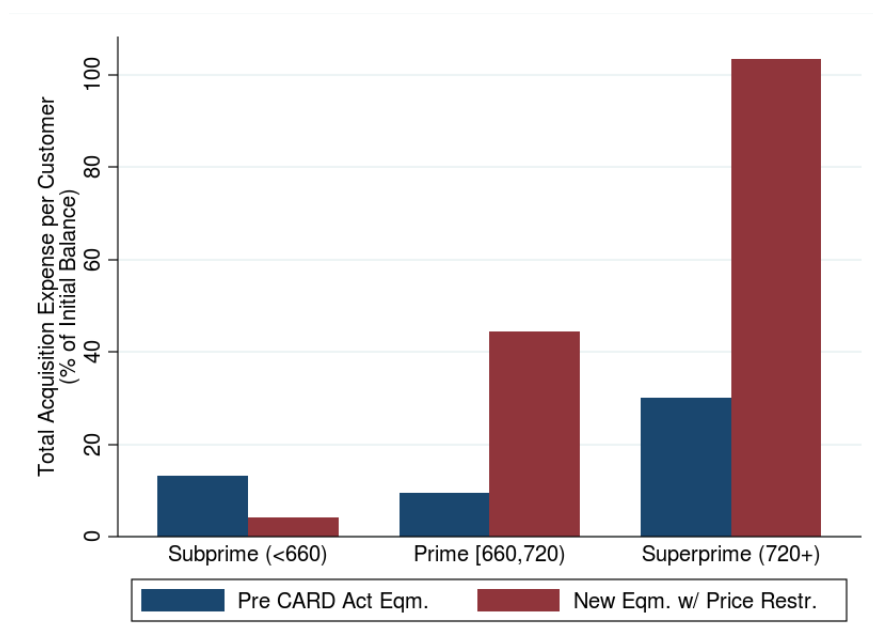
Notes: The figure shows estimated per-person total surplus (including both borrowers and non-borrowers' consumer surplus as well as firm profits) in the pre-CARD-Act equilibrium and also in the new equilibrium found in the model after imposing the CARD Act price restrictions. Consumer surplus is dollarized using each type's marginal utility of income (the price coefficient γ_θ) and using average 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 .

Figure 25: Incidence of CARD Act Price Changes across Income



Notes: The figure plots annual incomes imputed at the ZIP-code level using IRS Statistics of Income data against the predicted change in the contract price of borrowing. This price change is from the pre-CARD-Act equilibrium to the new equilibrium found in the model after imposing the CARD Act price restrictions. See Figure 20 for further discussion of this price measure.

Figure 26: Changes in Total Outlay for New Account Acquisition



The figure shows firms' total per-account outlay for new account acquisition in the pre-CARD-Act equilibrium and also in the new equilibrium found in the model after imposing the CARD Act price restrictions. Outlay is defined as account acquisition costs (a model parameter) minus introductory prices (p^0) offered on new accounts (a variable chosen by firms in the model). Outlay is averaged across firms using equilibrium market share, so changes in outlay reflect both changing market shares across firms with different acquisition costs and also changes in introductory prices.

8 Tables

Table 1: Observed Price Changes from Pre- to Post-CARD Act

Panel A: Changes in Interest Charges (% Ann.)					
FICO	P10	P25	Mean	P75	P90
580 - 599	2.46	-0.03	-2.52	-4.62	-2.83
600 - 619	2.16	0.89	-1.54	-4.32	-2.28
620 - 639	2.66	1.70	-0.75	-3.66	-1.91
640 - 659	3.03	2.49	0.12	-2.69	-2.11
660 - 679	3.01	2.95	0.88	-1.06	-2.15
680 - 699	2.67	3.15	1.38	0.05	-1.50
700 - 719	1.44	3.22	1.59	0.99	-0.49
720 - 739	0.44	3.18	1.56	1.33	0.44
740 - 759	-0.99	2.68	1.45	1.44	0.28
760 - 779	-2.55	1.91	1.07	1.44	-0.04
780 - 799	-2.54	-0.02	0.82	1.41	1.07

Panel B: Changes in Fee-Inclusive Charges (% Ann.)					
FICO	P10	P25	Mean	P75	P90
580 - 599	3.14	-0.06	-6.10	-7.39	-10.60
600 - 619	2.27	0.83	-3.43	-5.61	-6.31
620 - 639	2.76	1.64	-2.22	-4.87	-4.71
640 - 659	3.21	2.50	-0.90	-3.41	-3.49
660 - 679	3.14	3.04	0.20	-1.70	-2.86
680 - 699	2.78	3.25	0.90	-0.23	-2.37
700 - 719	1.50	3.23	1.25	0.36	-1.32
720 - 739	0.63	3.27	1.31	1.20	-0.35
740 - 759	-0.88	2.73	1.25	1.24	0.06
760 - 779	-2.35	1.97	0.88	1.30	-0.23
780 - 799	-2.74	0.10	0.68	1.42	0.76

Notes: The table shows percentage point changes in two price measures across the FICO score distribution from before the CARD Act to after (2008Q3 to 2009Q2 and 2011Q3 to 2014Q2 respectively). The first price measure, shown in Panel A, is an account's annualized percentage interest charges, defined as annualized monthly interest charges divided by borrowed balances. The second price measure, shown in Panel B, adds fee charges to the numerator of the first price measure.

Table 2: Pre-CARD Act Price Distribution on Mature Accounts

FICO Group	Cum. Months of Borrowing	Share within FICO Group	Interest Charges (% Ann.)			Fee-Inclusive Charges (% Ann.)		
			P25	Mean	P75	P25	Mean	P75
620 - 639	0	1.81%
	1-2	2.13%	10.23	17.90	25.73	11.03	25.75	29.03
	3-5	4.10%	8.31	16.14	24.91	8.86	21.35	27.98
	6-11	20.79%	9.50	16.73	25.12	9.98	21.19	27.92
	12	71.16%	11.62	18.29	26.00	12.18	21.15	27.99
680 - 699	0	5.33%
	1-2	4.23%	4.78	12.89	19.34	4.94	16.14	21.21
	3-5	6.33%	2.87	11.28	17.79	2.96	13.34	19.21
	6-11	23.33%	4.33	12.02	18.13	4.61	13.58	19.34
	12	60.77%	8.35	14.36	19.46	8.57	15.36	20.38
740 - 759	0	15.86%
	1-2	8.01%	2.11	9.56	14.61	2.16	11.52	15.65
	3-5	9.27%	1.23	8.56	13.41	1.68	9.84	14.29
	6-11	24.03%	3.10	9.32	13.66	3.17	10.24	14.36
	12	42.83%	6.13	11.07	14.50	6.20	11.59	14.98
800 - 819	0	44.68%
	1-2	14.22%	0.00	7.59	12.71	0.00	9.82	13.41
	3-5	10.97%	0.27	8.26	12.86	0.47	9.66	13.40
	6-11	16.24%	3.79	8.82	12.72	3.90	9.69	13.15
	12	13.89%	5.46	9.71	12.72	5.51	10.15	13.01

Notes: The table shows price quartiles and means at selected FICO score groups and across accounts with different cumulative months of borrowing over the course of the year in the pre-CARD-Act period (2008Q3 to 2009Q2). This sample includes only mature accounts (observed at 18 or more months since origination). The two price measures shown are, first, an account's annualized percentage interest charges, defined as annualized monthly interest charges divided by borrowed balances, and second, a price measure that adds fees charged to the numerator of the first price

Table 3: Lender Returns after Borrower Risk Signals

FICO Group	Baseline (% Ann.)	Over-Limit and Delinquent	Late by 90+ Days	Late by 60-89 Days	Late by 30 - 59 Days	FICO Drop of 60+ Points	FICO Drop of 30-59 Points
580 - 599	0.89	-36.65	-40.77	-34.25	-27.34	-12.65	-6.26
600 - 619	2.99	-25.36	-41.97	-35.77	-25.69	-9.55	-6.17
620 - 639	3.30	-21.90	-43.67	-37.73	-24.20	-7.96	-5.82
640 - 659	3.69	-19.95	-45.26	-38.92	-23.20	-6.16	-5.30
660 - 679	4.35	-19.04	-47.23	-40.17	-23.39	-5.22	-4.59
680 - 699	5.09	-18.70	-48.28	-42.01	-23.00	-4.21	-3.66
700 - 719	6.02	-17.94	-49.06	-42.51	-22.06	-3.87	-2.89
720 - 739	6.99	-16.88	-51.65	-44.83	-19.53	-3.87	-2.37
740 - 759	7.92	-16.07	-53.05	-45.12	-16.27	-3.18	-2.10
760 - 779	8.82	-15.78	-52.16	-43.26	-11.98	-2.79	-1.81
780 - 799	9.24	-17.47	-50.81	-42.11	-8.19	-2.31	-1.43

Notes: The table shows baseline annual percent returns on accounts in each FICO score group (column 1) in the pre-CARD-Act period (2008Q3 to 2009Q2), and differences from these baseline returns that are predicted in the pre-CARD-act period by the risk signals in each column. Returns are calculated by dividing finance revenue less default cost by borrowed balances.

Table 4: Fee Revenue Shares by Signal Type

FICO Group	Late by <30 Days	Over-Limit not Delinquent	Over-Limit and Delinquent	Late by 30+ Days	FICO Drop of 30+ Points
580 - 599	11.49	9.85	72.42	6.15	0.10
600 - 619	27.11	18.20	47.57	6.78	0.35
620 - 639	32.15	20.33	41.04	6.01	0.47
640 - 659	38.71	20.63	34.25	5.76	0.64
660 - 679	47.20	19.00	27.18	5.70	0.92
680 - 699	56.19	16.38	20.38	5.88	1.18
700 - 719	64.78	13.51	13.98	6.25	1.47
720 - 739	71.26	11.02	9.60	6.59	1.53
740 - 759	77.00	8.40	6.34	7.06	1.19
760 - 779	82.71	5.13	3.62	7.80	0.74
780 - 799	85.03	2.63	2.11	9.97	0.26

Notes: The table shows the share of fee revenue in each FICO score group generated by the fee categories in each column in the pre-CARD-Act period (2008Q3 to 2009Q2). Late fees are shown separately by delinquency status and by whether they coincided with an over-limit fee.

Table 5: Persistence in Consumer Revolving Behavior

FICO Group	<u>Recent Borrowers</u>		<u>All Accounts</u>	
	Transactor	Borrower	Transactor	Borrower
580	0.16	0.85	0.05	0.84
600	0.14	0.89	0.05	0.80
620	0.13	0.89	0.05	0.79
640	0.12	0.89	0.04	0.81
660	0.12	0.89	0.03	0.77
680	0.11	0.88	0.03	0.79
700	0.10	0.88	0.02	0.75
720	0.09	0.87	0.02	0.72
740	0.08	0.87	0.02	0.70
760	0.08	0.86	0.01	0.65
780	0.08	0.82	0.01	0.49

Notes: The table shows probabilities of next-quarter borrowing in the pre-CARD-Act period (2008Q3-2009Q2) for consumers who are either transactors or borrowers in the current period. The first two columns restrict the sample to consumers who have borrowed at least once in the past 6 months (recent borrowers), and the latter two columns extend these results to the full sample of active credit-card holders.

Table 6: Pre-CARD Act Price Distribution on New Accounts

FICO Group	Cum. Months of Borrowing	Share within FICO Group	Interest Charges (% Ann.)			Fee-Inclusive Charges (% Ann.)		
			P25	Mean	P75	P25	Mean	P75
539 - 639	0	2.03%
	1-2	2.49%	0.00	11.45	18.86	0.00	20.36	24.58
	3-5	4.90%	0.00	10.99	17.85	0.00	17.40	21.16
	6-11	38.15%	0.84	11.78	18.24	1.94	17.52	22.01
	12	52.43%	4.77	12.41	18.15	5.52	16.76	21.83
640 - 699	0	6.47%
	1-2	5.44%	0.00	6.90	12.62	0.00	10.17	14.56
	3-5	8.27%	0.00	6.58	12.35	0.00	8.87	13.88
	6-11	38.84%	0.00	6.81	12.44	0.00	8.64	13.62
	12	40.98%	0.00	8.16	12.88	0.00	9.57	13.91
700 - 759	0	15.00%
	1-2	8.83%	0.00	3.94	8.14	0.00	5.38	8.74
	3-5	11.36%	0.00	3.49	5.93	0.00	4.48	6.82
	6-11	36.24%	0.00	3.46	6.34	0.00	4.25	7.10
	12	28.58%	0.00	5.36	9.47	0.00	6.03	9.86
760 - 819	0	28.17%
	1-2	14.11%	0.00	3.37	5.08	0.00	4.97	6.48
	3-5	14.66%	0.00	2.74	0.61	0.00	3.67	2.12
	6-11	29.25%	0.00	2.38	1.54	0.00	2.89	2.87
	12	13.81%	0.00	3.62	7.17	0.00	4.12	7.56

Notes: The table shows price quartiles and means at selected FICO score groups and across accounts with different cumulative months of borrowing over the course of the year in the pre-CARD-Act period (2008Q3 to 2009Q2). This sample includes only young accounts (observed at 12 or fewer months since origination). The two price measures shown are, first, an account's annualized percentage interest charges, defined as annualized monthly interest charges divided by borrowed balances, and second, a price measure that adds fees charged to the numerator of the first price measure. Borrowing is defined as not repaying a balance in full at the end of a given month.

Table 7: Default Rates by Private-Information Type
(relative to lowest quintile)

	(1)	(2)	(3)
Dependent Variable	One-Year Default Rate		
Sample	All Accounts	Subprime	Prime
Estimator	OLS	OLS	OLS
2nd Quintile	0.0317*** (0.0000460)	0.0902*** (0.000116)	0.00176*** (0.0000310)
3rd Quintile	0.0585*** (0.0000503)	0.147*** (0.000118)	0.00502*** (0.0000355)
4th Quintile	0.0780*** (0.0000535)	0.191*** (0.000131)	0.0129*** (0.0000367)
5th Quintile	0.0904*** (0.0000627)	0.198*** (0.000150)	0.0257*** (0.0000437)
Quarter FEs	YES	YES	YES
Bank x FICO FEs	YES	YES	YES
Observations	243734158	88264172	155469986

Notes: The table shows regression estimates for a model using private information types as well as public types (FICO scores) to predict 1-year default. Private information types are presented as quintiles of the distribution of lender private information; estimates are relative to the lowest quintile of the private information distribution.

Table 8: Average Default Rates by Public and Private-Information Type

FICO Group	One-Year Default Rate by Quintile of Private-Information Type (%)				
	1st	2nd	3rd	4th	5th
580 - 599	14.92	31.14	39.14	45.75	45.73
600 - 619	5.93	9.37	13.75	16.78	20.47
620 - 639	5.02	7.12	10.23	12.35	15.47
640 - 659	4.18	5.25	7.17	9.20	11.54
660 - 679	3.34	4.08	5.13	6.80	8.75
680 - 699	2.66	3.08	3.41	4.58	6.72
700 - 719	1.80	1.97	2.21	3.40	4.76
720 - 739	1.05	1.29	1.59	2.18	3.27
740 - 759	0.64	0.76	0.99	1.40	2.45
760 - 779	0.42	0.48	0.64	0.90	1.77
780 - 799	0.29	0.30	0.43	0.58	1.22

Notes: The table shows one-year default rates by private information types (quintiles of the private information distribution) in the pre-CARD-Act period (2008Q3 to 2009Q2), for the FICO score group in each row. Default is defined as any instance of delinquency of over 90 days. Private information types are constructed to be weakly increasing in default risk, but the relative predictiveness of private vs. public information (FICO scores) remains flexible.

Table 9: Demand Model: Consumers' One-Period Payoffs by State

<i>Prior Period:</i>	Current Period:	Same Bank j		New Bank j'		No Credit Card with Any Bank
		Borrower	Non-Borrower	Borrower	Non-Borrower	
<i>Borrower, on Credit Card with Bank j</i>		$d_{\theta j} - \gamma_{\theta} p_{1\theta j}$	$n_{\theta j} - l_{\theta j}$	$d_{\theta j'} - s_{\theta j'} - \gamma_{\theta} p_{\theta\theta j'}$	$n_{\theta j'} - s_{\theta j'} - l_{\theta j}$	$-l_{\theta j}$
<i>Non-Borrower, on Credit Card with Bank j</i>		$d_{\theta j} - \gamma_{\theta} p_{1\theta j}$	$n_{\theta j}$	$d_{\theta j'} - s_{\theta j'} - \gamma_{\theta} p_{\theta\theta j'}$	$n_{\theta j'} - s_{\theta j'}$	0
<i>No Credit Card with Any Bank</i>		<i>(choice not available)</i>	<i>(choice not available)</i>	$d_{\theta j'} - s_{\theta j'} - \gamma_{\theta} p_{\theta\theta j'}$	$n_{\theta j'} - s_{\theta j'}$	0

Notes: The table shows a consumer's one-period flow payoffs depending on the consumer's circumstances at the end of the previous period (by row) and the consumer's choice in the current period (by column). The parameters shown include the flow utility from borrowing, $d_{\theta j}$, and the flow utility from holding a credit card without borrowing, $n_{\theta j}$, as well as disutility from price (marginal utilities of income), γ_{θ} , and two adjustment costs, including setup costs for opening new accounts, $s_{\theta j'}$, and liquidity costs for paying off existing balances, $l_{\theta j}$. The subscripts j and j' can refer to any bank in the set of banks J , while subscripts θ refer to consumer types.

Table 10: Demand Model: Marginal Utilities of Income

	(1)	(2)	(3)	(4)
Dependent Variable	<u>Log(Retention Rate)</u>		<u>Log(Retention Rate)</u>	
Estimator	OLS	2SLS	2SLS	2SLS
Gamma	-0.0000339*** (0.0000118)	-0.106*** (0.0129)	.	-0.0696*** (0.00664)
Gamma Subprime			-0.187*** (0.0281)	
Gamma Prime			-0.141*** (0.0108)	
Gamma Superprime			-0.104*** (0.0104)	
Bank-Specific Trends	YES	YES	YES	NO
Observations	60638012	60638012	60638012	60638012
1st-Stage F-Statistic		54.26	47.759	51.31
Clusters	550	550	550	550

Notes: The table shows estimates of price coefficients (marginal utilities of income) estimated via OLS and 2SLS using quasi-experimental lender repricing. Subprime, prime, and superprime accounts in column (3) are defined as FICO scores less than 660, from 660 to 719, and 720 or above respectively. 2SLS estimators use a total of 55 instruments from repricing event dummies interacted with consumer types.