Risk-Based Borrowing Limits in Credit Card Markets*

William Matcham§

March 10, 2025

Abstract

I use novel statement-level data on the 2010—2015 UK credit card market to show that lenders individualize contracts through risk-based credit limits. Though shared with other European credit markets, this feature contrasts with the US counterpart, where interest rates are also individualized. To quantify the implications of this distinction, I estimate a structural model relating individualized interest rates and credit limits to lender-specific credit scores. I evaluate a counterfactual where lenders can freely tailor prices and credit limits, which the UK environment precludes. Lenders control default risk with credit limits and use prices to extract surplus from inelastic borrowers.

Keywords: Risk-based credit limits, risk-based pricing, adverse selection, credit cards
JEL Classification: D22, D82, E51, G21, G51, L13, L50

*This paper is a revision of the first two chapters of my Ph.D. dissertation. I am greatly indebted to Alessandro Gavazza for guidance and support on this project. I thank John Gathergood, Scott Nelson, Matteo Benetton, Johannes Stroebel, Amit Seru, Lu Liu, Dirk Jenter, Daniel Paravisini, Darrell Duffie, and Jonathan Shaw for helpful comments. For the same reason, I thank seminar participants at Tel Aviv University, Oxford Saïd, Erasmus University Rotterdam, Bank of England, Duke Fuqua, Boston University Questrom, Boston College Carroll, EIEF, NYU Stern, the Stanford Institute for Theoretical Economics (SITE), and Stanford GSB. In this paper, I use data provided by the Financial Conduct Authority (FCA) under a data-access agreement for research purposes. The views expressed are my own. Therefore, this paper should not be reported as representing the views of the FCA. I thank Jeroen Nieboer, Cherryl Ng, and Ann Sanders for their support and assistance at the FCA.

§Royal Holloway University of London. william.matcham@rhul.ac.uk, willmatcham.com

1 Introduction

US credit card lenders individualize interest rates and credit limits according to assessments of customers' default risk. This paper shows that the leading UK lenders tailor credit limits but keep interest rates constant across customers of varying default risk. Other European credit markets follow suit, in line with EU-wide regulations limiting credit lenders' ability to price discriminate. How does this striking institutional difference affect consumers' and lenders' outcomes? Further, why do US lenders tailor both credit limits and interest rates in equilibrium?

The central contribution of this paper is to estimate a structural model of the credit card market to shed light on these hitherto unanswered questions. The model includes two interrelated consumer characteristics on which lenders could tailor contracts: default risk and price elasticity. My model and counterfactual results imply that lenders tailor credit limits to mitigate potential default risk associated with adverse selection, and tailor interest rates to maximize interest revenues. The use of risk-based interest rates benefits lenders and price-elastic borrowers at the expense of higher interest rates for inelastic borrowers, who are typically most at risk of default.

Beyond the standalone interest in the functioning of consumer credit markets, two factors underscore the importance of answering these questions. The first is the longstanding interest in credit constraints and credit rationing. An extensive literature has proposed credit constraints as the factor reconciling discrepancies between the empirical facts on borrowing and the theoretical predictions of traditional lifecycle models (Attanasio and Weber, 2010). On credit rationing, several theoretical papers study why similar borrowers have, in the past at least, varied in their ability to obtain credit (Stiglitz and Weiss, 1981). Nowadays, credit products are widely accessible, but credit rationing occurs through credit limits.[1] Despite theoretical interest, there have been almost no empirical attempts to explain how lenders ration credit on the intensive margin. This gap in the literature persists despite the relevance of rationing to information economics (Akerlof, 2001), the macroeconomy (Blinder and Stiglitz, 1983), and economic development (McKinnon, 1973).

This topic is also important because of its policy implications, which extend beyond credit markets. The academic literature and policy discourse focus on regulating price discrimination.[2] Recently, interest in tailored prices has increased in response to the tension between the development of AI (used for algorithmic pricing) and regulatory interventions aimed at limiting its scope. Firms, though, set multidimensional contracts with many levers for discrimination. Effective regulation requires an understanding of how firms discriminate multiple variables and how multidimensional individualization adjusts in response to regulation. This paper is a first attempt to study credit

[1] In July 2023, approximately 80% of US adults owned at least one of the 578 million credit cards in circulation (see https://www.newyorkfed.org/newsevents/news/research/2023/20230808).

[2] A leading example of a US Act that limits price discrimination is the Affordable Care Act, which prohibits health insurers from tailoring premiums based on gender, health status, medical history, or occupation.

market quantity *and* price discrimination in tandem.

At least two challenges have stymied efforts to establish the implications of interest rate and credit limit discrimination in credit card markets. First, studies focus on the US, where lenders can and do tailor both interest rates and credit limits. This collinearity limits the variation available to disentangle their impacts. I overcome this by studying the European context, in which credit market regulation inhibits lenders in tailoring interest rates. In EU credit markets, lenders must present a representative interest rate for each product and use this price (or lower) for the majority of customers. While this improves the transparency of credit products, it may exacerbate adverse selection issues if the safest individuals are not willing to borrow at the prevailing advertised rate.

The second challenge is the scarcity of administrative, statement-level panel data that includes credit scores, interest rates, and credit limits. To address this, I use a new source of statement-level data on approximately 80% of UK credit cards active between 2010 and 2015. The data include cardholder demographics and card characteristics for every card, along with monthly spending, repayment, default decisions, and lenders' funding (marginal) costs. Among other advantages, the data contain lenders' proprietary credit scores for every card originated. Hence, I can check whether lenders tailor interest rates and credit limits to their own predictions of customers' risk. The data reveal widespread intensive margin credit rationing: around 40% of individuals use over 90% of their credit limit on at least one occasion in the first two years of owning the card.

In my formal empirical analysis, I document significant variation in credit scores and credit limits, both within and across lenders. This variation persists within credit card products, with higher credit scores corresponding to larger credit limits. In contrast, interest rates vary minimally across cards, remain almost constant at the card-month level, and are not risk-based. This choice by lenders exceeds the aforementioned EU regulatory requirements. I also report heterogeneity in the shape and scale of credit limit distributions across lenders. Finally, the results of Chiappori and Salanie (2000) probit regressions imply that lenders control unobserved default risk through credit limits and show that unobservables driving unpaid balances and default are positively correlated, consistent with adverse selection.

To investigate the welfare impacts of individualized interest rates and credit limits on cardholders and lenders, I develop and estimate a structural model of the credit card market. My primary modeling novelty relates to the supply side. I endow each lender with a screening technology that generates a noisy signal of each individual's private type, which is their risk. Optimal credit limits trade off maximizing interest revenues from larger balances when the cardholder does not default, with adverse selection coming from the positive correlation between unobservables driving the desired balance and default probability. The model includes potential costs of individualizing a customer's interest rate above that which they advertise. Examples of these costs include infrastructure requirements to operationalize tailored interest rates, and the reputational damage from customer negativity. The model represents the first quantitative analysis of credit card lenders' screening technologies, credit limits, and individualized interest rates.

My supply-side estimates reveal substantial variation in lenders' screening technologies, which aligns with observed differences in credit limit distributions. Lenders with precise screening technologies have fewer statements on which the customer repays their entire balance. This finding is consistent with a segmentation of credit card lenders, in which those with the most precise screening technologies serve a riskier but more profitable market segment. Lenders with precise screening technologies are more willing to serve customers who will borrow but may default because they can more accurately set lower credit limits for customers they perceive to be riskier.

The estimated costs of individualizing interest rates are substantial: each lender faces an average shadow cost of £87 (around $110) per customer for deviating five percentage points above that advertised. These costs vary across lenders and correlate positively with income and credit score. This latter finding is as expected, if individualized interest rates are shrouded to lower income, higher risk customers.

The demand model explains borrowers' credit card choices, balances unpaid, and default decisions, incorporating observed and unobserved heterogeneity in all endogenous demand-side variables. For credit card and borrowing choices, preferences over interest rates vary with individuals' incomes. To identify demand parameters, I leverage a novel source of quasi-experimental price variation: the cost shock resulting from the 2011 High Court ruling on the mis-selling of payment protection insurance (PPI). Banks were forced to compensate thousands of consumers after the court deemed they had mis-sold PPI, which led to rises in interest rates on some credit cards.

My demand estimates show that individuals with the lowest income have the most inelastic demand, both in their borrowing and credit card choices. Consequently, individuals with a high default probability also have inelastic demand. They are therefore susceptible to high interest rates due to their inelasticity, not just their increased cost.

Finally, counterfactual simulations illustrate how the option of fully individualized prices affects consumers' and lenders' welfare. Similar to the US, in the counterfactual, lenders face no costs or constraints in individualizing interest rates and credit limits. Interest rate and credit limit discrimination emerges as a result. Low-income, inelastic individuals experience increases in interest rates and thus reductions in consumer surplus, but consumer surplus increases for elastic borrowers. Tailoring interest rates increases lenders' profits by 23%. The profit increases in the counterfactual suggest that the costs to lenders of tailoring rates in the EU environment limit lenders in extracting surplus from inelastic consumers. Though it is not the question I answer in this paper, understanding the sources of these costs is an important endeavor. I conclude by further discussing their potential sources, focusing on reputational risk.

Relative to the US, the EU context tilts the market in favor of low-income individuals and away from lenders, whose profits are lower because of constrained abilities to price discriminate. My counterfactual reveals that in consumer credit markets, tailoring prices and tailoring quantities are complementary tools rather than substitutes. The punchline of my counterfactual analysis is that in unregulated environments, interest rates are risk-based to maximize interest revenues, and credit limits are risk-based to cover downside default risk.

A central and non-obvious insight of my analysis is that those most likely to default receive high interest rates because they have the most inelastic demand, not because they have the highest costs. Indeed, individuals in the counterfactual with high default risk signals, but elastic demand obtain *lower* interest rates in the counterfactual relative to the baseline. Further, there is no difference in the relationship between interest rates and elasticities for those with low and high risk signals. This suggests that the negative correlation observed in the US between FICO scores and interest rates may be driven as much by standard price discrimination as by firms pricing in default risk.

2 Related Literature

Credit Rationing

Credit limits are a device for credit rationing, and since the seminal work of Stiglitz and Weiss (1981), there has been a longstanding interest in credit rationing in credit markets (Calomiris, Longhofer, and Jaffee, 2017). In the Stiglitz and Weiss model, market-level interest rates do not rise to clear the market because higher interest rates attract riskier borrowers (adverse selection effect) and may directly induce more defaults. As a result, among similarly risky projects, some receive a loan, and others do not.

Motivated by recent empirical findings, which I detail in Sections 3 and 5, my model of credit rationing differs from Stiglitz and Weiss (1981) along three main dimensions. First, I argue that the credit card market is not perfectly competitive, so lenders set their own prices. Second, several recent papers have documented the price-invariance of default in credit markets, so I exclude a direct effect of interest rates on default. Finally, in my framework, lenders obtain signals on borrowers' risk, whereas lenders in Stiglitz and Weiss (1981) infer default risk based on borrowers' willingness to accept higher rates. In prior work, credit rationing occurs as a result of the aversion of lenders to raise rates to clear the market. I contribute by producing a model that generates intensive margin (rather than extensive margin) credit rationing through lenders *choosing* optimally to set credit limits that may bind for some individuals.

The limited existing empirical work focuses on the effect of credit limits on borrowers' outcomes. Gross and Souleles (2002a) and Agarwal, Chomsisengphet, Mahoney, and Stroebel (2017) estimate the causal effect of credit limits on borrowing and default. Aydin (2022) presents an interesting experiment randomizing credit limit shocks across credit card in Turkey. I am the first to estimate a model that explains credit limit distributions as a function of lenders' risk signals.

Risk-Based Pricing

My work also contributes to the literature on risk-based pricing. Papers have documented the presence of risk-based pricing in some financial markets (e.g., Edelberg, 2006 for US consumer loans) and its absence in others (e.g., Benetton, 2021 for UK mortgages). Livshits, Mac Gee, and Tertilt (2016) is a key contribution to this literature. The authors argue that credit has become widely available, or "democratized," in response to financial innovations such as credit scoring and risk-based pricing. Their empirical work shows that in the US, the availability of credit to riskier borrowers coincided with a significant rise in interest rate dispersion, consistent with their model's predictions. I show that the functioning of the UK credit card market up to 2015 differs: credit has become widely available at almost all levels of risk despite limited variation in interest rates across the distribution. My work shows that in the UK market, the widespread availability of credit occurs alongside risk-based credit limits and not risk-based interest rates.

Regulation of Credit Markets

My final primary contribution is to the literature on regulating credit markets. Much of the existing work focuses on the effects of the 2009 US Credit Card Accountability, Responsibility, and Disclosure (CARD) Act.[3] Agarwal, Chomsisengphet, Mahoney, and Stroebel (2014) documents substantial consumer savings due to the Act. Nelson (2022) focuses on how the Act limited lenders' abilities to reprice credit card customers after origination. I focus entirely on pricing and credit limits at origination since risk-based repricing has limited application, and no regulatory restriction, in the UK (see Section 4). I contribute to this literature by arguing that the limited risk-based pricing in the EU credit market context tilts the market in favor of low-income consumers and away from lenders and their profits.

3 Data and Context

My primary data source is the Financial Conduct Authority (FCA) Credit Card Market Study (CCMS) Database (FCA, 2015b). Using its regulatory authority, the FCA collected data from 14 lenders on all credit cards active between 2010 and 2015.[4] This dataset covers approximately 80% of all UK cards active during this period, amounting to around 74 million cards. The FCA's data collection resulted in three main datasets, which have yet to be used for research.

The first dataset includes cardholder and card information at origination, such as demographics (age, income, employment, and homeownership status, etc.), acquisition channel (whether in branch, online, by post, via telephone, etc.), and initial interest rate, credit limit, and 0% promo-

[3] Another related context is the Chilean credit market, studied among others by Cuesta and Sepulveda (2021). The paper shows that tighter interest rate caps decrease surplus, with the welfare costs from loss of credit access outweighing the lower equilibrium prices. Related to my work, they show that risk-based interest rate caps mitigate welfare decreases.

[4] The FCA chose 11 firms (split into 14 separate lending entities) to be representative of the entire credit card market. For confidentiality reasons, I cannot reveal their identity. In the main analysis, I omit store cards.

tional deal length. Notably, this dataset includes lender-specific credit scores, a unique feature I analyze comprehensively in Section 4.1. Since there is no single equivalent of the FICO score in the UK, this addition is valuable because without data on the credit scores that lenders actually use, I cannot accurately assess whether interest rates and credit limits are risk-based. Table A.1 provides detailed summary statistics, revealing stark differences in credit limit and interest rate variation. For instance, the coefficient of variation (the ratio of standard deviation to mean) for credit limits is 93%, compared to 36% for APRs. Sections 4.2 and 4.3 further decompose these variations by lenders and cards in what is a central element of my descriptive findings.

The second dataset is a monthly panel of statement data for credit cards between January 2010 and January 2015. It includes balances (opening and closing), repayments, the number and value of transactions, fees, interest, and failures to repay. These data allow me to measure both *utilization* (balance used before the closing date) and *revolving* (balance left on the account after repayment). Further, it provides the evolution of credit scores, interest rates, and credit limits, which is often lacking in existing ones used for research on the US market. Following industry practice, I use observations on failures to repay to classify customer default as 90+ days without a repayment.

I use the statement data to assess the extent of credit rationing from binding credit limits. Credit is considered rationed through credit limits if individuals' spending on the card is at or close to the credit limit.[5] I define credit limit utilization by calculating the closing balance as a percentage of the credit limit. Across all statements with a positive closing balance, 23% close with a balance over 90% of the credit limit, and 27% have a closing balance exceeding 85% utilization. Thus, the closing balance is at or close to the credit limit on over one in five statements on which the card is used. At the individual level, in the first two years post-origination, approximately 40% use over 90% of their credit limit on at least one occasion. So near- or full-utilization is a prominent feature of the UK credit card market.

The decision of a customer to repay the entire balance—also known as transacting—is another prevalent feature of the statement data. Repayment covers the whole balance on approximately 50% of statements, and approximately 25% of cardholders repay in full every month in the first 12 months post-origination. Further, there is substantial variation in the proportion of transacting statements across lenders, ranging from 22% to 85% (see Figure A.1 for all values). These findings motivate including an extensive margin transaction decision for consumers in my model.

The third dataset is a monthly panel of card characteristics covering January 2010 through January 2015, capturing annual fees, rewards, income thresholds, and advertised APRs. Additionally, the dataset includes rarely observed lender funding costs. The statistics on funding rates in Table A.1

[5] Individuals using, say 90% of their credit limit are considered as credit constrained because they might not make a purchase on credit because it would push them over their credit limit, which incurs a fee and will lead to a decrease in credit score.

suggest that credit card lenders enjoy substantial markups, as also shown for US business credit cards in Benetton and Buchak (2024). Mean yearly funding rates are 2.28%, and less than 3% of borrowers default each year. Even under the conservative assumption that lenders cannot recover any part of a defaulted balance and ignoring interchange revenue, this implies a markup (price to marginal cost ratio) of approximately four. Since the marginal lender prices above marginal cost, I am inclined to consider an alternative to a perfectly competitive model. Furthermore, the Herfindahl-Hirschman Index (based on the value of borrowing) is 1,496, where a value exceeding 1,000 implies a concentrated industry (FCA, 2015a).

Finally, the CCMS data package also includes a credit reference agency (CRA) dataset that matches cards to individuals. Also, I occasionally complement my analysis and motivate modeling choices with an FCA survey of cardholders, detailed in FCA (2015c).

UK Credit Card Market and EU Regulation

The UK credit market differs from the US equivalent in several ways. First, the UK market is more passive regarding rewards, fees, and purchase promotional deals. In the UK, cashback and air miles are scant, present in only 11% and 6% of card product-months, respectively, and annual fees are zero in 88% of card product-months. In understanding these discrepancies, it is worth noting that EU regulation limits interchange fees to 0.3%, as compared to typical values of 2% in the US (see Wang, 2023 for more details on the positive relationship between interchange fees and rewards).

Second, US individuals own more credit cards than UK individuals. My CRA data confirm that most UK individuals have only one card each (see Figure A.2). The 2015 US mean number of cards per person was 2.24.

Third, regulatory differences exist. All promotional material and documentation for a credit card product in an EU credit market must include a "representative" ("advertised") APR. Before February 2011, at least 66% of customers each month had to obtain the advertised APR or lower. The regulation changed in February 2011 when the UK harmonized with the EU to reduce the threshold to 51%, and it has not changed since. Also, UK customers cannot discover their personal interest rate or credit limit until after they are accepted. The US has no such regulation or practices. The 1974 Consumer Credit Act mandates a "cooling-off period" during which consumers can freely cancel their card, though my data show that this option is exercised in only 0.2% of originations. These shopping periods are also standard in the US, though the law does not mandate them. Finally, the UK has no laws that cap interest rates on credit cards, and while some US states have usury laws, since national banks are regulated at the federal level, these caps tend not to apply.[6]

Like in the US, UK lenders securitize credit card receivables (the future principal and interest payments) into Special Purpose Vehicles, which issue these as Asset-Backed Securities to investors

[6] Matcham (2024) offers more detail on the exact methods lenders use to charge interest on credit cards, and more specifics of US and UK credit card regulation.

(Fitch, 2013). While this allows lenders to transfer default risk to outside investors, lenders must retain sufficient exposure to comply with regulations introduced after the 2008 Financial Crisis.

Finally, in Section 4.2, I describe two other notable features of the UK credit card market (limited lender product portfolios and limited ex-post repricing) as they pertain to ruling out alternative mechanisms through which lenders could implement risk-based pricing.

4 Descriptive Evidence

The main aim of this section is to present robust evidence that up to 2015, the leading UK credit card lenders individualized credit card contracts through risk-based credit limits rather than interest rates. I organize the first three subsections around the three relevant variables: credit scores, interest rates, and credit limits. Though not detailed in below, rewards and fees are almost always constant at the card-month level. At the end of each subsection, I summarize the descriptive facts presented and explain their implications for a model of the UK credit card market.

4.1 Credit Scores

I start by highlighting two features of the credit score data: (i) credit scores differ across lenders and are based on proprietary information, and (ii) there is substantial *within-card* variation in credit scores across cardholders at origination.

First, I analyze the distributions of origination credit scores across different lenders (see Figure A.3). The figures highlight that lenders' scores differ in their numerical scales and distribution shape. These features suggest that lenders construct their own scores, but this could result from a mere rescaling process or customers with differing default risks selecting across different lenders.

To explore potential discrepancies in lenders' credit scores further, I regress each lender's proprietary credit scores on a fine set of demographics collected by lenders during the application process, including percentile bins for income and age, employment and homeownership status dummies, and month fixed effects. In these regressions, the proportion of variation in private credit scores explained is 21% on average, varying from 7% to 34% (see Figure A.4 for the range of values).[7] These findings imply that most lenders' proprietary credit scores are based on much more than the readily available customer demographics, though some lenders rely on these demographics more

[7] Similar findings emerge when I perform the same exercise, replacing demographics with the main publicly available UK credit score, though I only have data on the public credit score for a limited set of months. In these regressions, the mean R-squared across lenders is around 22%, and it varies from 6% to 36%. These findings imply that public credit scores only explain a moderate proportion of the variation in each lender's proprietary credit scores: lenders either input private data sources into their algorithms or use alternative algorithms to the public credit score providers.

than others. The fact that the R-squared values vary greatly across lenders suggests that lenders' proprietary credit scoring algorithms differ.

The use of proprietary credit scores in the UK contrasts with the US, where FICO scores provide a standardized measure of customer creditworthiness that many banks use as part of their lending decisions (Agarwal, Chomsisengphet, Mahoney, and Stroebel, 2017). Recent academic research justifies why lenders might create their own risk scores. For example, Albanesi and Vamossy (2019) shows that machine learning (specifically deep learning) methods consistently outperform standard credit scoring models, even when trained on the same data source.[8]

Next, I investigate whether lenders sort customers of similar risk onto different cards. For each lender-month, I perform a one-way Analysis of Variance (ANOVA), decomposing the variation in proprietary credit scores into a within-card and between-card component (see Online Appendix Subsection B.1 for a mathematical formulation). Table A.2 column (1) contains the results across lenders. When averaging over lenders and months, the within-card variation accounts for 87% of the total variation, indicating substantial variation in customers' credit scores *within* each card product. This fact implies that lenders do not sort customers of varying risk onto separate cards. As described in the next section, lenders do not offer many products in their portfolio anyway, and their cards differ in alternative features such as their network (e.g., Visa/Mastercard), rewards, and branding. To summarize this subsection:

Empirical Finding 1 (Credit Score Variation) *Lenders construct their own credit scores, which differ from publicly available scores and are not well-explained by typical demographics such as income or age. Additionally, substantial within-card variation in credit scores indicates that lenders do not sort customers by risk onto separate cards.*

Model Implication 1 (Screening Technologies) *Lender-specific credit scores, termed "screening technologies," should be a lender-specific variable in a model of the supply side of the UK credit card market. Credit scores observed are a combination of demographics such as income, and private signals on customers' risk.*

4.2 Interest Rates

4.2.1 Limited Total and Within-Card Variation in Lenders' Rates

Next, I document limited variation in interest rates over customers at each lender. Table A.2 column (2) reports the average (over months) of lenders' interest rate coefficient of variation. The values are below 23%, and the average across prime and superprime lenders (weighted by

[8] See FCA (2023) for a recent report on the UK credit information market and credit reference agencies. The report confirms that the credit information market is concentrated, and points to "several areas where it could be working better."

originations) is 14%. This finding implies that the standard deviation in the interest rate is, on average, one-seventh of the mean at a lender in a given month. Further, as detailed in Table A.2 columns (3) and (4), the across-lender weighted average of the ratio of the 75^{th} to 25^{th} percentile (respectively 90^{th} to 10^{th}) for interest rates is 1.19 (respectively 1.38), further illustrating limited UK variation in interest rates within lenders. In contrast, Galenianos and Gavazza (2022) show that for US interest rates, the ratio of 90^{th} percentile to 10^{th} percentile is as large as 3, even after controlling for borrower and card characteristics. Finally, the UK coefficient of variation in interest rates is 0.36 when calculated across *all lenders*. This is over 2.5 times larger than the within-lender average, indicating some differences in lenders' average interest rates.

For the leading UK credit card lenders, a modest proportion of the already small total variation in interest rates is found across the individuals on a given credit card product. To show this feature, I perform the same one-way ANOVA as in Section 4.1, but this time for interest rates. The within variation for prime and superprime lenders is, on average, 24% of the total variation.[9]

4.2.2 High Proportion of Customers Obtaining Advertised APR

To explain the lack of within-card variation in interest rates, I calculate the percentage of customers each month that obtain their card's advertised APR (recall that regulation forces lenders to advertise an APR for each card product). I plot the time series in Figure A.5. The proportion of customers receiving the advertised APR across all credit cards in the sample remains consistently high at 80 to 90%, and this stability persists even after the regulatory change in February 2011 that relaxed the requirements for advertised APRs. Even though regulation required lenders to give the advertised APR (or lower) to only 51% of their customers after February 2011, most lenders still gave almost all their customers the advertised APR. Further, in 77% of card-months, over 90% of originations obtain the advertised APR. This statistic confirms that most *cards*, not just *lenders*, give most of their holders the advertised APR. The results suggest that either the benefits of individualizing interest rates is small, or that lenders face significant shadow costs of individualizing rates. Indeed, one aim of the quantitative model I estimate is to separately identify these benefits and costs.

4.2.3 Ruling Out Alternative Forms of Risk-Based Pricing

Lenders could employ risk-based pricing by adjusting interest rates after origination, *repricing* customers according to their evolving risk and behavior. However, limited repricing occurs in the UK credit card markets. Outside promotional deals, I calculate that lenders reprice 3% of

[9] The weighted average including subprime lenders is 33%. I discuss subprime lenders separately in Online Appendix B.2. Table A.2 column (5) reports the values of the percentage of within-card variation for all lenders. In the extreme case, one lender gives all customers on a given credit card the same interest rate in *all* months. In that case, all the variation in interest rates at origination is at the card level.

individuals within nine months of origination and 5% within one year of origination.

Furthermore, lenders could employ risk-based pricing by offering multiple distinct cards and sorting customers of differing risk onto different cards. Empirical Finding 1 rules out that customers are sorted onto cards by credit score, and the above evidence shows that there is limited between-card variation in interest rates at a lender. Furthermore, as shown through the statistics in Table A.3, lenders offer a small set of distinct card products each month. For example, the share of originations on the top two cards at each lender is 86% at the mean and 91% at the median.

Finally, lenders may refrain from using risk-based pricing because they collude on interest rates in a cartel. I provide brief empirical evidence inconsistent with this notion in Appendix B.3.

Empirical Finding 2 (Interest Rate Variation) *At each lender, interest rates exhibit limited total and minimal within-card variation. Nearly 90% of customers obtain the advertised APR at origination each month, corroborating the limited within-card variation in interest rates. Interest rates are not risk-based within credit card, and alternative means by which lenders could employ risk-based pricing (through sorting or dynamics) are also absent.*

Model Implication 2 (Card-level interest rates) *The model should allow for some potential friction (regulatory or otherwise) associated with individualizing each customer's interest rate.*

4.3 Credit Limits

4.3.1 Substantial Variation in Lenders' Credit Limits Across and Within Cards

Having confirmed the lack of variation (particularly within-card variation) in interest rates, I turn to credit limits, which, unlike interest rates, exhibit substantial variation. The coefficient of variation in credit limits across lenders is 78% on average (when weighted by originations). This value is over five times larger than that for interest rates. Columns (7) and (8) of Table A.2 report the across-lender weighted average of the 75th to 25th percentile and the 90th to 10th percentile credit limit ratios, which are 3.34 and 9.18, respectively, indicating significant variation in credit limits within each lender. These ratios are not well-documented in the literature for the US.

Like before, I perform a within-card and between-card decomposition of credit limit variation. Eighty percent of total variation is found within product. Like credit scores, the dominance of within variation suggests that lenders do not sort individuals onto cards with varying average credit limits. Instead, there is considerable variation in credit limits, even within a credit card product and month. The section thus far has presented a set of descriptive results. Figure A.6 summarizes the main findings on total and within-card variation in credit scores, interest rates, and credit limits.

FIGURE 1. EMPIRICAL CDFS OF TWO PARTICULAR LENDERS' CREDIT LIMITS

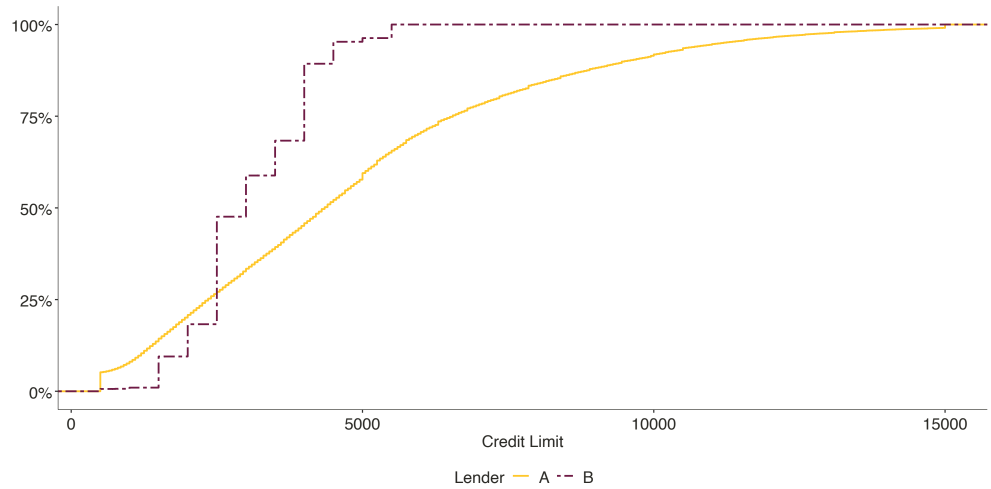

Notes: Monetary values here and everywhere that follows are quoted in 2015 Great British Pounds (GBP).

4.3.2 Variation in the Shape and Scale of Lenders' Credit Limit Distributions

The distribution of credit limits varies substantially across all lenders, both in shape and scale.[10] I illustrate this in Figure 1, by plotting the empirical cumulative distribution function (CDF) of credit limits for two contrasting lenders, lenders A and B. Two substantial differences are evident. The first relates to the *shape* of the credit limit distributions. Lender B's curve is step-like, implying a coarse process of assigning credit limits to individuals, where groups of consumers obtain the same credit limit. Lender A's smooth curve is consistent with a more finely tuned allocation mechanism for origination credit limits. The second difference relates to the *scale* of the credit limit distributions. Lender A has lower values of credit limits than lender B for the first 25 percentiles; however, all percentiles after the 25^{th} are larger. The range of lender A's credit limit distribution is indeed much larger.

Other lenders' credit limit CDFs, plotted in Figure A.7, lie between the two lenders in Figure 1. This range in the shape and scale of distributions is consistent with lenders who vary in the coarseness of their credit limit assignment. Some lenders offer large groups of customers the same credit limit, while others with smoother CDFs adjust their credit limits more precisely.

As expected, lenders link each individual's credit limit to an assessment of their default risk. In Figure 2, I plot the mean of the origination credit limit along application credit scores for two

[10]To confirm differences between lenders' credit limit distributions formally, I conduct multiple distribution "Kolmogorov-Smirnov" hypothesis tests. I strongly reject the equality of empirical CDFs across lenders at lower than 0.5% significance levels in all tests. Details are available on request.

Figure 2. Mean credit limits across credit scores for two particular lenders

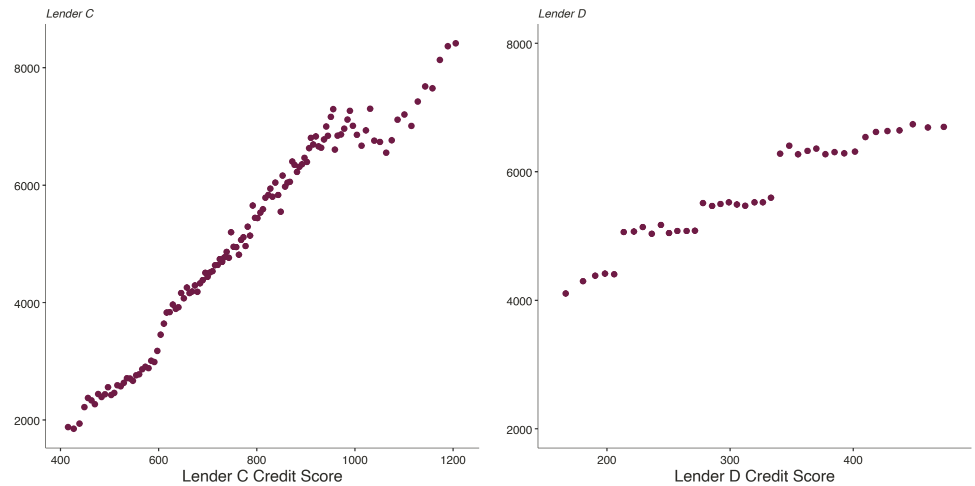

Notes: Credit score scales differ across lenders so cannot be compared.

contrasting lenders for 2013.[11] Both curves are upward-sloping, consistent with risk-based credit limits. Further, the right-hand lender has discontinuities in credit scores at credit score thresholds similar to those exploited in Agarwal, Chomsisengphet, Mahoney, and Stroebel (2017). Accordingly, my model aims to rationalize discreteness and discontinuities in lenders' credit limit distributions through coarse (discrete) assessments of customers' risk.

Empirical Finding 3 (Risk-Based Credit Limit Distributions) *There is substantial within-card variation in credit limits across lenders. The distributions of credit limits differ in shape and scale across lenders. Credit limits vary with lender-specific credit scores, and heterogeneity exists in how lenders map their credit scores into credit limits.*

Model Implication 3 (Risk-Based Credit Limits) *Lenders should choose credit limits optimally according to their predictions of customers' risk. Differences in their screening technologies should deliver lender-specific distributions of credit limits that vary in scale and coarseness.*

4.4 Default and Asymmetric Information

I follow the longstanding empirical literature in credit and insurance markets and examine the relationship between credit demand and default. My focus is on the intensive margin of demand, looking at the relationship between the *amount* of credit and failures to repay.

[11] In Figure A.8, I plot the mean of origination credit limit for each lender, along application credit scores. All curves are upward-sloping, consistent with risk-based credit limits. In unreported plots, the same patterns emerge when produced by month and by card.

I estimate the correlation between unobservables determining credit card usage and serious delinquencies by following the approach first suggested in Chiappori and Salanie (2000) and subsequently used in credit markets by Crawford, Pavanini, and Schivardi (2018), Nelson (2022) and Benetton and Buchak (2024). As mentioned in Section 2, a positive correlation in unobservables driving credit demand and default can occur because of "adverse selection" (those with ex-ante unobservables driving high credit demand have unobservables driving increased likelihood of failure to repay) or positively correlated ex-post shocks to borrowing and ability to repay. While I am not attempting to separately identify these forces, any positive correlation will represent for the extent of combined information asymmetries in this market.

The two estimating equations are $y_{it}^{v*} = x_{it}\beta^v + \epsilon_{it}^v$, for $v \in \{b, d\}$. Across all specifications, the variable y_{it}^{d*} represents the latent net utility from failing to repay the credit card. Consistent with standard practice by UK credit card lenders, the observable $y_{it}^d = 1(y_{it}^{d*} > 0)$ is equal to one if consumer i that originated in month t went 90 days without making a credit card repayment in the first 18 months after t. These serious *delinquencies* occur for around 3% of customers on average. Depending on the specification, y_{it}^b is a dummy variable equal to one for individuals with an above-median credit limit or a dummy variable equal to one depending on the individual's level of revolving.[12] The variable x_{it} corresponds to a flexible set of individual characteristics (logged income, origination month, card product dummies, distribution channel, etc.).

I assume that ϵ_{it}^b and ϵ_{it}^d are standard normal variables to yield the Chiappori and Salanie "Pair of Probits." In Table 1, I report the correlation between the generalized residuals $\hat{\epsilon}_{it}^b$ and $\hat{\epsilon}_{it}^d$ along with the Chiappori and Salanie (2000) W statistic, which is asymptotically $\chi^2(1)$ under the null of $\text{corr}(\epsilon_{it}^b, \epsilon_{it}^d) = 0$.[13] Results are robust to the inclusion of lender dummies, rather than credit card products, and the choice to include dummies for the distribution channel of the card (e.g., originated online, in-store, etc.).

In the first column of Table 1, y_{it}^b is a dummy variable equal to one if the individual has an above-median credit limit. In this case, the correlation between the credit limit equation residual and the default equation is negative, with a tight confidence interval around -0.03. We would expect such a negative correlation if lenders successfully screen riskier individuals and consequently give them lower credit limits.

Next, I discuss revolving levels, measured as the balance left on the account after repayment. In

[12] I also estimate models where y_{it}^{*b} is a continuous variable representing credit limit (or revolving) and the results are in line with the findings presented.

[13] Chiappori and Salanie (2000) specify weights w_{it} and write

$$W = \frac{\left(\sum w_{it}\hat{\epsilon}_{it}^b\hat{\epsilon}_{it}^d\right)^2}{\sum \left(w_{it}\hat{\epsilon}_{it}^b\hat{\epsilon}_{it}^d\right)^2}.$$

I use an unweighted version of the statistic.

Table 1. Correlation between demand and default unobservables

	Credit Limit	**Any Revolving**	**Full Revolving**
	(1)	(2)	(3)
Correlation Between Unobservables	-0.030	0.108	0.151
	[-0.031,-0.0282]	[0.107,0.110]	[0.150,0.153]
CS (2000) W Statistic	1,855	25,214	14,347

Notes: All specifications include logged income along with card product, month, and distribution channel dummies. 95% confidence intervals are provided below estimates. CS (2000) W statistic is given in Footnote 13.

column (2), y_{it}^b is a dummy variable equal to one if the individual revolves any balance in the first 18 months post-origination. The correlation is now positive, around 0.10. The size of the W statistic shows evidence of a clear rejection of the null of no correlation. Finally, in column (3), y_{it}^b is taken as a dummy variable equal to one if the individual revolves their entire credit limit. In this case, the correlation increases 50% to 0.15, suggesting that the level of revolving is even more strongly correlated with default than the mere decision to revolve a balance. It is worth noting that these estimates are strikingly similar to the findings in Benetton and Buchak (2024), which calculates similar statistics in the case of US business credit cards.

Empirical Finding 4 (Default and Correlation with Revolving) *The average default rate for consumer credit cards is about 3%. Unobservable factors driving credit card utilization positively correlate with credit card default, providing suggestive evidence of adverse selection or correlated ex-post shocks to borrowing and default.*

Model Implication 4 (Positive Correlation of Borrowing and Default Unobservables) *The demand-side model should allow for correlation between unobservables driving revolving and default. The supply-side model should have a mechanism through which lenders' credit limit choices act as a screening mechanism against potential default risk.*

4.5 Implications of Descriptive Findings

This section reveals that UK credit card lenders individualize credit limits based on assessments of customer risk, but do not individualize interest rates. The next step is understanding how lender heterogeneity and regulatory environments impact consumers' and lenders' outcomes and welfare. The empirical setting is not insightful on the magnitude of the costs of individualizing interest rates, nor can it explain how lenders would set interest rates if they were not required to set and advertise a card-level APR. In the absence of meaningful exogenous variation in the regulatory environment or the makeup of lenders, the best—and perhaps only—way to achieve this aim is to build an economic model of the credit card market. This model follows in the next section.

5 Model of the Credit Card Market

In this section, I develop a model of lenders' interest rate and credit limit choices and individuals' credit card, revolving, and default choices. The model is designed to explain how lenders set credit limits using their own credit scores and quantify the costs lenders face in individualizing interest rates. The setup allows me to measure the causal effect of restrictions (regulatory or otherwise) that inhibit lenders in tailoring interest rates. Before providing the details, I clarify what the model is *not* intended for. First, while I am able to estimate the size of costs of individualizing interest rates, the model and the data cannot identify the specific sources of these costs. I discuss some possibilities, focusing on infrastructure investments and reputational concerns, and provide further details in Online Appendix E.2. While unpacking this black box is a worthwhile endeavor, it is beyond the scope of this paper.

Second, as my descriptive work reveals, lenders do not offer an extensive menu of cards with a broad spectrum of advertised interest rates. Lenders also do not sort consumers of similar risk across their limited set of cards. I am not trying to rationalize these features with my model either. I take lenders' menus of cards as given, and instead, I estimate a model of lenders' credit limits and interest rates on a given product portfolio. My model is equivalent one in which lenders initially have a product portfolio and face, at least in the short run, large fixed costs of changing that set of products. Given the importance of branding and advertising, this assumption is not so unreasonable. Further, the model is fit for the purpose of simulating how lenders would set interest rates (admittedly, on the same set of cards) in the absence of *any* limitations on how they are individualized – I develop that model in the counterfactual section. This modeling setup is sufficient to shed light on the two aims of the paper, which are to understand (1) the effect of restrictions limiting the tailoring of interest rates on lenders' profits and borrowers' outcomes and (2) the respective roles of tailored interest rates and credit limits, as observed in the US context.

Third, I am not modeling how consumers search over credit card choices; I will assume that consumers consider all products they qualify for. Galenianos and Gavazza (2022) provide evidence that US individuals between 2006 and 2008 do not consider all credit cards they are eligible for. However, the UK data between 2010 and 2015 tell a different story, with lenders reporting that several customers acquire their cards through price comparison websites (PCW), which instantly reveals the full set of card products a customer could choose. While many consumers still do not use PCWs, abstracting from consumer search is a profitable abstraction in the context I study.

5.1 Preliminaries and the Credit Card Product

The market is a pair (m, t), where t represents an origination month between January 2010 and June 2013, and m represents one of three distribution channels: branch, online, or other channels such as telephone or post.[14] I describe the model through its three features: the credit card $j \in J_{mt}$, consumers currently without a credit card $i \in I_{mt}$ (demand), and lenders $\ell \in L_{mt}$ (supply).

I focus on the preferences of those currently without a credit card for two reasons. First, as discussed in Section 3, most UK adults hold only one credit card, making it a relevant subset for empirical analysis. Second, estimating my model on the sample currently without a credit card circumvents complications arising from (i) balance transfers and (ii) balance-matching heuristics in repayment across multiple cards (Gathergood, Mahoney, Stewart, and Weber, 2019).

My demand model of card origination, revolving, and default can be microfounded in a typical consumption-savings lifecycle setup. However, since my focus is lenders' credit limit and interest rate choices, I prefer to specify demand-side estimating equations as a set of linearized equations agnostic to the behavior that generates them. Matcham (2024) discusses the costs and benefits of this approach—chiefly, the ability to abstract from the myriad behavioral biases potentially present—and other approaches in other credit market models.

Following the tradition of Lancaster (1966), I model a credit card product as a bundle of features. There are four components. The first is the *advertised* interest rate r_{jmt}. The second is the income threshold $\underline{Y}_{jmt}$, explained in Section 5.2. The third and fourth are characteristics: those I observe, denoted X_{jmt} (e.g., rewards such as air miles), and those I do not, ξ_{jmt} (e.g., prestige and loyalty).

5.2 Consumer

Like Crawford, Pavanini, and Schivardi (2018), I model three primary endogenous demand-side variables: card choice, initial revolving level, and default.[15] I detail each of these in turn.

5.2.1 Card Choice

Consumers choose a card and decide whether to use the card *at first* for transacting or revolving. Choosing to transact, denoted $j = 0$, involves paying off the balance in full. Revolvers leave some

[14] I stop at June 2013 to ensure that I observe sufficient borrowing and default data on each individual.

[15] Benetton and Buchak (2024) use a similar approach to model the demand-side of the US business credit card market though their second equation explains utilization rather than revolving.

balance unpaid, accruing interest.[16] The consumer's utility from revolving on card j is

$$V^E_{ijmt} = \bar{V}^E(X^E_{jmt}, \xi^E_{jmt}, r, \eta^E_{mt}, y_i; \theta^E_{mt}) + \nu_{ijmt}.$$

Throughout the model, superscript E represents the ***E***xtensive margin. The term X^E_{jmt} denotes the elements of observed card characteristics X_{jmt} that affect card choice, and the same convention applies to ξ. The term ν_{ijmt} represents a random taste shock. I model ν_{ijmt} as generalized type-1 extreme value distributed taste shocks. These random taste shocks are independent and identically distributed (iid) across customers and correlated across choices. The final components of credit card utility currently undefined are η^E_{mt}, a card-utility market fixed effect; y_i, which denotes logged income; and θ^E_{mt}, which denotes market-specific parameters that govern indirect utility.

To justify my choice of components for $\bar{V}^E$, I draw on the results of a question from a cardholder survey (FCA, 2015c). Participants were asked, "Which of the following applied when you took out your credit card?" The most common response is rewards, which 33% of respondents provide. Hence, I include X^E_{jmt} in $\bar{V}^E$. Twelve percent of customers mention the card's interest rate. Hence, I include the interest rate in $\bar{V}^E$. I deal with the issue of whether this is the advertised rate r_{jmt} or the individualized rate r_{ijmt} when it comes to estimating the model in Section 6.

Other non-price, non-reward, and non-promotional deal responses comprise some remaining survey responses, implying the importance of ξ^E_{jmt}. Such responses include "use abroad" (15%), "low fees (4%), and "good deal offered" (13%), all of which are examples of unobserved characteristics contained in ξ^E_{jmt}. Finally, there is little to no mention of individualized credit limits $\bar{b}_{ijmt}$, which I omit from $\bar{V}^E$ directly. Despite this survey finding, if credit limits do affect card choice, the card-average credit limit would be part of ξ^E_{jmt}, so the model does allow individuals to prefer certain cards because they know (or think) these cards have higher *average* credit limits.[17]

I follow the literature (e.g., Berry, Levinsohn, and Pakes, 1995) and linearize $\bar{V}^E$ so that

$$V^E_{ijmt} = \beta^{E'} X^E_{jmt} + \xi^E_{jmt} + \nu_{ijmt} + \alpha^E_{imt} r + \eta^E_{mt}. \quad (1)$$

[16] That consumers choose whether they will use the card initially for revolving or transacting is one of few substantive assumptions on consumer behavior I impose. I impose it as it simplifies the lender's problem. Though not all consumers commit to transacting or revolving, consumers' use of direct debits (automatic transfers) suggests that many consumers have decided how they intend to use their credit card at origination. In the first three months of originating the card, 25% have set up a direct debit, rising to 31% by six months. Of those who set up a direct debit at origination, around 40% set up a direct debit to automatically pay off their entire balance each month, suggesting they intend to be a transactor. Of the remaining 60% who set up a direct debit for an amount less than the full balance, 76% set up a direct debit to pay the *minimum repayment*, which is usually the maximum of (i) 1-2.5% of the balance, and (ii) £5 (around $6).

[17] In ongoing work, I estimate the probabilities of switching card or originating an extra card for individuals either side of the credit limit discontinuities I described in Subsection 4.3.2. The findings from this work will shed light on whether the indifference to credit limits in card choices as implied by the survey aligns with individuals' revealed preferences.

The random coefficient α^E_{imt} represents individual-specific preferences over interest rates. Since my counterfactual scenarios explore how lenders may choose individualized interest rates, I must allow for the possibility that preferences over interest rates differ across individuals. Heterogeneous preferences over interest rates read

$$\alpha^E_{imt} = \alpha^E + \Omega^{E,r}_{mt}\tilde{y}_{imt}. \tag{2}$$

The term $\tilde{y}_{imt} = y_i - \bar{y}_{mt}$ denotes log income centered around the market average, where the market average is given by $\bar{y}_{mt} = I^{-1}_{mt}\sum_{i\in I_{mt}} y_i$. I center logged income around the market average so that α^E represents the mean interest rate sensitivity in the card choice equation.

I generate choice sets for individuals by comparing their income at origination to the card's income threshold. Individuals qualify for a card if their income Y_i exceeds the income threshold $\underline{Y}_{jmt}$. Consequently, the set of cards available to customer i is

$$J_{imt} = \{j \in J_{mt} | Y_i > \underline{Y}_{jmt}\}.$$

I discuss the rationale for lenders' use of income thresholds in Subsection 5.3. The utility from transacting, also linearized, is given by $V^E_{i0mt} = \delta_{0mt} + \nu_{i0mt} + \Omega^{E,\text{cons}}_{mt}\tilde{y}_{imt}$, where δ_{0mt} is a market-level constant of transacting utility. Individuals choose the card j^* in their choice set corresponding to the maximal value of V^E_{ijmt}, and individuals transact if V^E_{i0mt} exceeds $V^E_{ij^*mt}$.

5.2.2 Borrowing

Revolvers choose the initial amount of spending to leave unpaid on their card. I refer to this as the borrowing and revolving level interchangeably. This is not the level of spending; it is the expenditure that remains *after repayment*. Denote by b^*_{ijmt} the *desired* level of revolving, which represents the individual's level of balance unpaid in the absence of any credit limit. The word "desired" reflects that individuals may wish to borrow more than their credit limit $\bar{b}_{ijmt}$ allows. The value of b^*_{ijmt} satisfies the following borrowing function

$$b^*_{ijmt} = b(X^B_{jmt}, \xi^B_{jmt}, r, \eta^B_{mt}, y_i, \varepsilon^B_{imt}; \theta^B_{mt})$$

and as in card choice utility, the log of borrowing is linear in parameters:

$$\log(b^*_{ijmt}) = \beta^{B'} X^B_{jmt} + \xi^B_{jmt} + \alpha^B_{imt} r + \eta^B_{mt} + \Omega^{B,\text{cons}}_{mt}\tilde{y}_{imt} + \varepsilon^B_{imt}. \tag{3}$$

The terms X^B_{jmt}, ξ^B_{jmt}, α^B_{imt}, and η^B_{mt} in (3) have analogous definitions to those in (1) and (2), swapping E for ***B***orrowing. The random variable ε^B_{imt} reflects a revolver's unobserved demand for borrowing. For example, ε^B_{imt} would be high if an individual has an unreported health issue that requires them to quit their job. Both the lender and I do not observe ε^B_{imt} perfectly. I define its distribution in Subsection 5.2.4.

In the data, revolvers are likely to make monthly borrowing choices, such as those implied by the solution to an intertemporal consumption-savings problem. However, this paper concerns lenders' choices of *origination* credit limits and interest rates. When choosing origination credit limits, what matters to lenders are consumers' overall borrowing over the immediate period they use the card and less so the dynamics of borrowing within that period. After all, lenders can reprice at any point after origination if they desire. Hence, my setup does not require a model of multiple borrowing values across periods, as a consumption-savings problem implies. Modeling static borrowing is a clear-cut profitable abstraction for my context.

5.2.3 Default

Finally, revolvers choose whether to default on their balance. The net utility from defaulting reads

$$V_{imt}^{D} = V^{D}(\eta_{mt}^{D}, r, y_i, \varepsilon_{imt}^{D}; \theta_{mt}^{D}),$$

where, again, all terms are analogous to those defined in (1) and (3), swapping E for ***D***efault. The individual defaults if $V_{imt}^{D} > 0$. Once again, I linearize V_{imt}^{D}, implying

$$V_{imt}^{D} = \eta_{mt}^{D} + \alpha_i^{D} r + \Omega_{mt}^{D} \tilde{y}_{imt} + \varepsilon_{imt}^{D}. \tag{4}$$

Following the empirical literature on credit markets, I omit the interest rate from default utility and so set $\alpha_i^D = 0$. Nelson (2022) and Castellanos, Jiménez Hernández, Mahajan, Alcaraz Prous, and Seira (2018) provide empirical evidence of an insignificant effect of price on default in credit markets. Collier, Hartley, Keys, and Ng (2024) exploits a discontinuity in the interest rate of disaster loans and finds no difference in delinquency rates either side of the discontinuity. Much of the research on default implies that short-run liquidity drives default rather than the long-run value of a loan (Ganong and Noel, 2020; Indarte, 2023). Also, making default invariant to price follows other structural models of selection markets without direct interest rate effects on default, for example, Cohen and Einav (2007).

Further, as in Nelson (2022), default is not a direct function of the credit limit. Empirical work supporting this choice includes (i) Gross and Souleles (2002b), which shows that increases in credit limit do not explain increases in default, and (ii) Agarwal, Chomsisengphet, Mahoney, and Stroebel (2017), which provides evidence that credit limits do not affect default rates for prime customers. In Online Appendix C.1, I show that the association between credit limit and default in my data is consistent with there being no positive causal effect of credit limit on default.

5.2.4 Private Information Structure

I decompose private characteristics $(\varepsilon_{imt}^{B}, \varepsilon_{imt}^{D})$ into a common component $\tilde{\varepsilon}_i$ and an idiosyncratic component $\tilde{\varepsilon}_i^h$ so that

$$\varepsilon_{imt}^{h} = \sigma_{mt}^{h} \tilde{\varepsilon}_i + \tilde{\varepsilon}_i^{h}$$

for $h \in \{B, D\}$. The common component simplifies the lender signal structure (following in Subsection 5.3) and generates a correlation among unobserved private characteristics for each individual.

I simplify by setting $\tilde{\varepsilon}_i^B$ to zero and letting $(\tilde{\varepsilon}_i, \tilde{\varepsilon}_i^D)$ be independently standard normal distributed. This approach greatly reduces the complexity of estimation while still allowing for arbitrary correlation between ε_{imt}^B and ε_{imt}^D, which I term as intensive margin selection. Henceforth, I simplify the notation, writing ε_i instead of $\tilde{\varepsilon}_i$.

I do not include a term ε_{imt}^E in the card choice equation for the following reason. Because the estimation is conditional on choosing a card, and the first nest is between transacting and revolving, a term ε_{imt}^E would not represent unobserved preferences over having a credit card, but instead unobserved preferences that drive the choice to revolve. Since I include ε_{imt}^B to allow for unobserved preferences over the level of revolving, it is superfluous to also include unobserved preference over whether to revolve. As well, estimating the demand and supply models would become more complex with the inclusion of ε_{imt}^E, beyond what it would add to the model.

5.3 Lender

This part presents my model of lenders' exogenous screening technologies and endogenous interest rates and credit limits. Existing approaches, which are theoretical or calibrated against limited data, focus on how lenders *choose* the coarseness of their screening technology in the context of fixed costs of creating "scorecards" (e.g., Livshits, Mac Gee, and Tertilt, 2016 and Agarwal, Chomsisengphet, Mahoney, and Stroebel, 2017). Instead, I treat screening technologies as exogenous within a given month, which is reasonable, since investing in higher-quality proprietary data and setting up new credit scoring algorithms is costly. My novelty is to design a model that can be used to estimate lenders' screening technologies off optimal credit limit choices, and lenders' costs of individualizing prices off optimal interest rate choices.

5.3.1 Preliminaries and Timing

Lenders observe individuals' incomes Y_i and take characteristics X_{jmt}, ξ_{jmt}, and income thresholds $\underline{Y}_{jmt}$ as given in each market. Card characteristics are exogenous for three reasons. First, in the data, lenders do not individualize rewards, which are sticky and rarely change over the entire five-year period on which I have data. Second, lenders cannot adjust many of the unobserved characteristics, such as loyalty, in a given month. Third, contract pricing introduces issues in equilibrium existence and uniqueness that are profitable to abstract from where justified.

Income thresholds determine the set of individuals qualifying for a given card. UK lenders use income thresholds partly because they must be able to inform consumers of the information used to reject them if they source data from a credit reference agency (Department for Business Innovation and Skills, 2010). Consequently, lenders base decisions on *eligibility*, at least in part on income.

The timing, which matches the institutional environment, is as follows. The regulatory climate mandates that lenders set advertised APRs at the card-month level at the start of each month. Then, individuals choose their card. Finally, *after individuals have originated a card*, lenders choose credit limits and how, if at all, to individualize interest rates, ensuring their choices comply with the regulatory requirements. While the advertised interest rates are chosen competitively, the credit limit and individualized interest rate choices happen after individuals select their card, and are thus chosen non-strategically by lenders to maximize profits.

5.3.2 Screening Technology

Each lender employs its own screening technology. The screening technology takes in data available to the lender on a customer and provides the lender with a tailored prediction of possible values of the customer's common risk component ε_i. Without a screening technology, for each customer, the lender would take expectation over a standard normal, which is the population distribution of ε_i. The screening technology intends to provide a distribution with a mean close to each individual's realization of ε_i and a variance less than one, that of the population distribution.

Two features characterize the lender-specific, tailored distributions that the screening technology delivers. The first is the set of signals, or central points, around which the tailored distributions can be based. I denote these as $e_{i\ell t}$, which can take a finite number of lender-specific values $\{e_{\ell t1}, \ldots, e_{\ell t L_{\ell t}}\}$. The second feature is the precision of the distribution it generates. The distribution generated by the screening technology accounts for the fact that the signal may not be a correct prediction of a customer's risk, i.e., it allows for error. For an individual who generated the signal $e_{i\ell t}$, the distribution provided by the screening technology is normal with mean $e_{i\ell t}$ and variance $\sigma_{\ell t}^2 \leq 1$, and I call $\sigma_{\ell t}$ the precision parameter. Equivalently, given the value of $e_{i\ell t}$, the screening technology models ε_i as $\hat{\varepsilon}_i = e_{i\ell t} + w_{i\ell t}$, where $w_{i\ell t} \sim \mathcal{N}(0, \sigma_{\ell t}^2)$. When setting profits, the lender takes expectations using the distribution $\mathcal{N}(e_{i\ell t}, \sigma_{\ell t}^2)$, as the screening technology provides.

Figure 3 depicts distributions of ε_i and $\hat{\varepsilon}_i$ for two fictitious lenders. The risk distribution provided by lender 1's screening technology for customer i is $\mathcal{N}(e_{i1t}, 0.95)$. The mean of the conditional distribution e_{i1t} is far from the customer i's actual realization of $\varepsilon_i = \breve{\varepsilon}$. Lender 2 has a better screening technology. Its screening technology gives the signal e_{i2t}, which is closer to $\breve{\varepsilon}$. Furthermore, since σ_2 is smaller than σ_1, the signal errors at lender 2 are less dispersed around the signal than at lender 1. When setting credit limits for customer i, lender 2 will put more weight (relative to lender 1) on potential values close to $\breve{\varepsilon}$ and less on incorrect values, such as those near zero.

5.3.3 Advertised Interest Rates

First, lenders set advertised interest rates. Following Crawford, Pavanini, and Schivardi (2018) and Benetton and Buchak (2024), lenders strategically choose r_{jmt} to maximize expected profits so that interest rates form a Nash-Bertrand equilibrium. Let $\boldsymbol{r}^*_{-\ell mt}$ denote the equilibrium interest rates on cards at lenders other than ℓ. Then for lender ℓ, their vector of advertised rates $\boldsymbol{r}^*_{\ell mt}$ solves

$$\max_{\boldsymbol{r}_{\ell mt}} \sum_{i \in I_{mt}} \sum_{j \in J_{i\ell mt}} s^E_{ijmt} \Pi_{ijmt}.$$

Figure 3. Distribution of ε (solid) and $\hat{\varepsilon}_i$ (dashed) across two lenders for a customer with unknown value $\varepsilon_i = \breve{\varepsilon}$

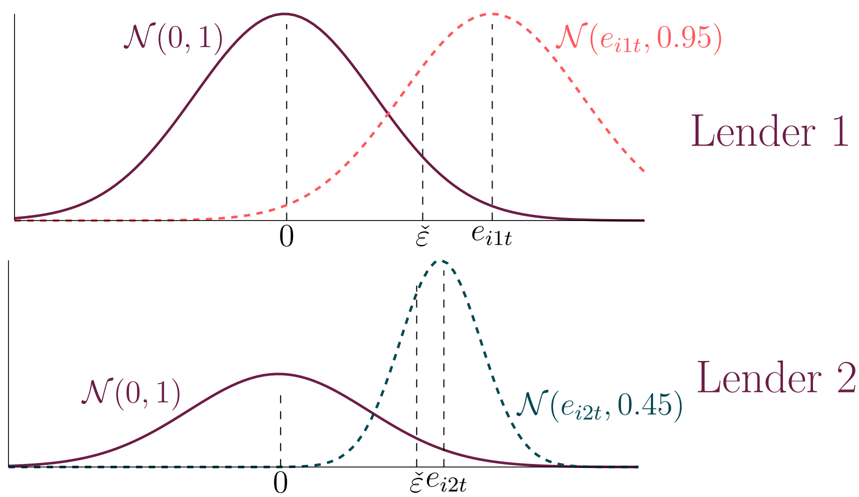

Notes: The bottom lender's screening technology, which delivers the signal e_{i2t}, outperforms the top lender's signal of e_{i1t} for this individual. The standard normal $\mathcal{N}(0,1)$ is the population distribution of ε_i.

The term $J_{i\ell mt} = J_{imt} \cap J_{\ell mt}$ is the set of cards that individual i qualifies for at lender ℓ. That is, it is the set of cards offered by lender ℓ with income thresholds lower than Y_i. The term s^E_{ijmt} denotes the probability of individual i originating card j as a borrower, implied by the logit model of card choice. Equation (19) provides the exact functional form for this term. The profit from individual i from borrowing on card j, denoted Π_{ijmt}, is defined in equation (6) of the following subsection.

5.3.4 Individualized Interest Rates and Credit Limits

Having set the advertised market level interest rates r_{jmt}, lenders set individual-specific credit limits $\bar{b}_{ijmt}$ and determine individual-specific interest rates r_{ijmt}. The individualized rate satisfies $r_{ijmt} = r_{jmt} + z_{ijmt}$ where z_{ijmt} is defined as the deviation for individual i from the advertised rate. With r_{jmt} fixed, choosing z_{ijmt} determines r_{ijmt}. As reported in the descriptive work, z_{ijmt} is zero for around 90% of originations. The origination data show that for prime and superprime lenders, when z_{ijmt} is not equal to zero, it is positive for 97% of originations, so I focus on the empirically relevant case of $z_{ijmt} \geq 0$, noting that the regulation requires that z_{ijmt} is non-positive for at least 51% of originations on a given card within a month (66% before February 2011).

Modeling lenders' credit limit and interest rate choices requires an expression of their profits. Regarding costs, I focus on the cost of funds, denoted c, charge-off (default) costs, and shadow costs of individualizing interest rates. According to statistics from US credit card lenders, the first two here account for over two-thirds of lenders' total cost of issuing credit cards (Evans and Schmalensee, 2005). The remaining third comprises mainly fixed costs (overhead and operational), which I can ignore since they are sunk at the point of choosing rates and limits.

Regarding revenue, I focus entirely on finance charges coming from interest. For US lenders in 2001, this accounted for 70% of their card revenue (Evans and Schmalensee, 2005). The remaining 30% comes from three main factors: interchange revenue, fees, and cash advances. Online Appendix C.2 describes the three factors in more detail and explains why they are less relevant in the UK credit card market than in the US.

Each lender's profit from a transacting customer, denoted Π_{i0mt}, is unrelated to the credit limit and interest rate.[18] Therefore, the credit limit and individualized interest rate decision is unaffected by whether the individual originating card j is a transactor or a borrower.

Let Δ_{imt} denote the probability that borrower i defaults, and c_{jmt} denote the funding rate. Then the profit per unit of credit borrowed from individual i is the interest rate minus the funding cost if the customer does not default, and $-(1-\psi)-c_{jmt}$ if they do, where ψ is the proportion of the balance that debt collectors can recover, which I set to zero in my empirical specification.[19] Hence, the expected profit per unit credit (gross of any costs of individualizing interest rates) for individual i on card j is

$$\pi_{ijmt} = (1-\Delta_{imt})(r_{ijmt}-c_{jmt}) + \Delta_{imt}(-1-c_{jmt}). \tag{5}$$

Then the expected profit from borrower i choosing card j in market (m,t), *gross* of any costs of individualizing interest rates, is

$$\Pi_{ijmt} = \mathbb{E}\left[\min\{b^*_{ijmt}, \bar{b}_{ijmt}\}\pi_{ijmt}\right]. \tag{6}$$

The observed borrowing $b_{ijmt} = \min\{b^*_{ijmt}, \bar{b}_{ijmt}\}$ reflects the fact that the individual cannot borrow more than the credit limit.

Let I_{jmt} denote the set of revolvers choosing card j in distribution channel m and month t. Then, for card j and all customers $i \in I_{jmt}$, the lender decides customer i's credit limit, $\bar{b}_{ijmt}$, and their deviation from the advertised rate, z_{ijmt}, to maximize expected profits, subject to the regulatory

[18] The revenue and costs from transactors do not relate to the interest rate, since they do not revolve a balance on which interest accrues. Lenders' variable cost from non-defaulting customers is per-unit credit, and therefore lenders' costs from transactors are unrelated to the credit limit. The credit limit may affect interchange revenue, but I abstract from interchange revenue for revolvers and do so for transactors for the same reason. Resultantly, profits from transactors are not related to credit limit and interest rate choices.

[19] When cardholders default, payment card issuers start collection procedures. These cardholders will often have other debts, which may be collected before credit card debt. Debt collection procedures are very costly relative to the size of the loan for credit card lenders. Further, in the US in 2002, 50% of all charge-offs resulted from bankruptcy, where debt collection is often futile (Evans and Schmalensee, 2005). These factors considered together, $\psi = 0$ is a reasonable abstraction.

constraints. Hence, for each card j, the lender solves

$$\max_{\bar{b}_{ijmt}, z_{ijmt}} \sum_{i \in I_{jmt}} \Pi_{ijmt}(\bar{b}_{ijmt}, z_{ijmt}) - C(z_{ijmt}, \kappa_{ijmt}),$$

subject to $z_{ijmt} \geq 0$ and the regulatory constraint

$$\frac{1}{I_{jmt}} \sum_{i \in I_{jmt}} 1(z_{ijmt} > 0) \leq \chi_t, \tag{7}$$

where $1(A)$ denotes the indicator function, equal to 1 if A is true and 0 otherwise, and χ_t is the regulatory requirement on the proportion receiving the advertised rate, equal to 49% after February 2011 and 34% before.

The term $C(z_{ijmt}, \kappa_{ijmt})$ represents the cost to the lender of choosing to deviate by an amount z_{ijmt} from the advertised rate r_{jmt} when setting individual i's interest rate, r_{ijmt}. The term κ_{ijmt} denotes the parameters governing the costs of individualizing interest rates for customer i. The function C is increasing and weakly convex in z_{ijmt} and satisfies $C(0, \kappa_{ijmt}) = 0$, so there is no cost to the lender from customer i if they give customer i the advertised rate.

My context does not enable me to identify the exact sources of C. Rather, the exercise is to quantify its magnitude, given lenders' optimal choices in individualizing interest rates. That said, it may help the reader to have some idea of what C might be. I discuss these in more detail in Section 8.2.3 and Online Appendix E.2. A brief summary of the main factors mentioned in interviews with industry experts is as follows. First, and most importantly, lenders might encounter reputational costs if they advertise a particular APR but then provide customers with a different, individualized, APR, especially since the individualized rate is set after the individual signs the contract. The larger the deviation from the advertised rate, the more egregious this might feel to the consumer, explaining the increasing and convex shape of C.

Second, C could include the administrative expenses of constructing the infrastructure and software to set optimal individualized prices. Extending pricing teams to provide individualized rates alongside advertised rates will be expensive for lenders. Also, by law, lenders must be able to explain to the customer how they chose their individualized rate, which can also create administrative costs, especially if lenders use advanced analytics such as machine learning models (which can be hard to explain in layman's terms) to set individualized interest rates optimally.

Individualized Interest Rate Optimality

In the data, the regulatory constraint, that is, equation (7), is not close to binding for any lender, so its Lagrange multiplier will be zero and the constraint can be ignored at the solution. Hence, the optimality condition for z_{ijmt} is

$$\frac{\partial \Pi_{ijmt}}{\partial z_{ijmt}} = \frac{\partial C}{\partial z_{ijmt}} - \lambda_{ijmt}, \tag{8}$$

where the Lagrange multiplier λ_{ijmt} on the constraint $z_{ijmt} \geq 0$ is non-negative itself and satisfies $\lambda_{ijmt} z_{ijmt} = 0$. Equation (8) will allow me to estimate the parameters κ_{ijmt} of C for those individuals with $z_{ijmt} > 0$. For those individuals, as a standard variational argument would imply, the marginal increase in profit from changing z_{ijmt}, given by $\partial \Pi_{ijmt} / \partial z_{ijmt}$, must equal the marginal cost of individualizing interest rates $\partial C_{ijmt} / \partial z_{ijmt}$.

Credit Limit Optimality

As derived in Online Appendix C.3, the first order condition with respect to the credit limit is

$$\mathbb{E}\left[\pi_{ijmt} | b^*_{ijmt} \geq \bar{b}_{ijmt}\right] = 0. \tag{9}$$

The intuition behind the first order condition is that at the optimal credit limit, the expected profit per unit credit, over those with unobservables that drive them to use their full credit line, is zero. If, for instance, the expected profit per unit credit were positive, the lender should raise the credit limit because the expected benefit of safer types using the entire credit limit exceeds the expected costs of riskier types using the whole credit limit. The converse is true if the expected profit per unit credit were negative. Notably, the first order condition here is not a zero-profit condition. Expected profit per unit credit over those infra-marginal individuals with unobservables that drive them to use less than their full credit line is positive.

This intuition and the technical conditions behind the first order condition rely on a positive correlation between the unobservables driving borrowing and default, which I term intensive margin adverse selection. Existing studies of credit markets estimate significant adverse selection (Nelson, 2022; Crawford, Pavanini, and Schivardi, 2018). My demand estimates, which I estimate entirely *independently* of the supply side, verify this assumption (see Sections 6 and 7 for details).

Because default probability is not a direct function of borrowing, the lender's optimization problem would be piecewise linear in credit limit if there were no adverse selection. In this case, the corner solution would make lenders give zero credit to those with negative expected profit per unit credit and unboundedly large limits to those with positive expected profit per unit credit. However, in the presence of adverse selection, an interior solution arises because, conditional on observables, sufficiently large credit limits will only be utilized by those with the highest unobserved default risk, that is, those with negative expected profit per unit credit. Of course, the lender does not want to offer extra credit to these individuals. To summarize, in the context of adverse selection, the choice of credit limit must be made with consideration of the distribution of default risk *among those* utilizing that amount of credit. This insight reveals the effect of adverse selection on lenders' optimal credit limit choices.

My descriptive findings on the differences in lenders' credit limit distributions motivate the tight relationship between lenders' screening technologies and the shape of the distribution of credit limits. Each unique signal implies a different choice of credit limit for the lender, and therefore, given income, there is a mapping between the number of unique credit limits at each lender and the number of unique signals provided by their screening technology. Lenders who give observably identical consumers (to the econometrician) a wide range of credit limits must have a wide range of different signals of these consumers' unobserved risk. I leverage this link between credit limits and signals to estimate the distribution of signals from each of the unique credit limit values. I detail this estimation process in the following section.

6 Estimation

In this section, I outline my method for estimating model parameters. My approach to demand estimation shares features with Benetton (2021) and Benetton, Gavazza, and Surico (2025). Figure A.9 displays the five steps of the estimation procedure.

6.1 Demand

6.1.1 Log-likelihood Conditional on Borrowing

I start by estimating the parameters of the demand model (Step 1 in Figure A.9). My demand model for those who revolve a balance consists of equations for card choice (equation 1), revolving (equation 3), and default (equation 4). The equations map cardholders' demographics along with lenders' interest rates, credit limits, and card characteristics into card choice, revolving level, and default. Together with stochastic assumptions on unobservables, the three equations imply a log-likelihood function for observed decisions.

The log-likelihood $\log \mathcal{L}_{mt}$ comprises two parts: the log-likelihood for card choice $\log \mathcal{L}_{mt,E}$, and the joint log-likelihood for borrowing and default choices $\log \mathcal{L}_{mt,BD}$. This form follows from the fact that unobservables for card choice are uninformative about the unobservables driving borrowing and default. Online Appendix D.1 provides detailed expressions for the terms of the likelihood. In the text below, I focus on how the estimation approach overcomes two primary challenges and discuss the exogenous variation I exploit to identify the parameters.

The truncation in borrowing is the first of three primary challenges in estimating the parameters of the likelihood function. Specifically, I observe the *constrained* level of borrowing $b_{ijmt} = \min\{b^*_{ijmt}, \bar{b}_{ijmt}\}$, rather than the *desired* level b^*_{ijmt}. As a result, I do not observe the desired borrowing for any revolvers who borrow their entire credit limit. Revolvers either use their entire credit line (full utilization) or not (interior utilization) and do or do not default. This creates four possible outcomes for revolver i:

1. $i \in I_1$: Interior utilization and default
2. $i \in I_2$: Interior utilization and no default
3. $i \in I_3$: Full utilization and default
4. $i \in I_4$: Full utilization and no default

Let $s^{(g)}_{ijmt}$ denote the likelihood of individual i being in group I_g. Then

$$\log \mathcal{L}_{mt,BD} = \sum_{i \in I_{mt}} \sum_{j \in J_{imt}} \sum_{g=1}^{4} 1^{(g)}_{ijmt} \log(s^{(g)}_{ijmt}), \tag{10}$$

where $1^{(g)}_{ijmt}$ is a dummy equal to one if individual i chooses card j and is in group I_g. I provide the expressions for $s^{(g)}_{ijmt}$ in Online Appendix D.1.

Individuals exhibiting full utilization create the most complication. Since their desired borrowing is not observed, their contribution to the likelihood is an integral with no closed form. Hence, I use simulated maximum likelihood (Gouriéroux and Monfort, 1996) with Halton draws.

The second challenge in estimating the demand model comes in specifying demand-side responses to interest rates. To make estimation feasible, preferences over cards and the initial borrowing choice must depend on the advertised interest rate. In what follows, I justify this choice on both economic and econometric grounds.

Regarding economic justifications, first, since individualized rates are determined after the contract has been signed, it is reasonable to think that individuals choose their card and *initial* borrowing based on the advertised rate, and then adjust their long-term borrowing after they have comprehended their individualized rate. Further, regarding borrowing, individuals might commit to some spending on the card before their individualized rate is determined.

Second, this modeling choice fits the framework of Gabaix and Laibson (2006) with individualized rates as the shrouded object that lenders are trying to conceal from their customers, particularly in the event that they differ from the advertised rate. If individuals choose their card based on comparisons of advertised interest rates, then the advertised rate will be the salient variable. And if individuals commit to their initial revolving before getting the card, they might well choose that variable with consideration of the advertised rate.

The choice can also be defended on econometric grounds. Since $r_{jmt} = r_{ijmt} - z_{ijmt}$, using advertised rates is equivalent to having additive measurement error, where the advertised rate is used as a mismeasured version of the true individualized rate. Because 90% of customers obtain the advertised rate, only 10% of observations are measured with error. Hence, even if the reader believes that the correct model of demand would have card choices and borrowing depend on individualized interest rates, then the quantitative impact of using the advertised rate is likely to be small. A final point to make on this matter is that an alternative option would be to estimate the demand model on the subset of originations obtaining the advertised rate, though, of course, this may comprise a selected sample of observations.

A third challenge for demand-side estimation is the endogeneity of advertised interest rates in the card choice and borrowing level equations. Advertised rates are likely to correlate with unobserved card characteristics ξ_{jmt}. For example, interest rates may be high on a given card because its unobserved card characteristics imply high demand for the card. Without addressing this issue, estimates might suggest that individuals prefer higher interest rates when, in fact, they prefer products with attractive unobservable features that are resultantly priced higher. To deal with this, I estimate a full set of product-channel-month fixed effects δ_{jmt} in the card choice and borrowing equations that subsumes the endogeneity between r_{jmt} and ξ_{jmt}.

Formally, I rewrite equations (1) and (3) respectively as

$$\begin{aligned} V^E_{ijmt} &= \delta^E_{jmt} + \nu_{ijmt} + u^E_{ijmt}, \\ \delta^E_{jmt} &= \beta^{E\prime} X^E_{jmt} + \xi^E_{jmt} + \eta^E_{mt} + \alpha^E r_{jmt}, \\ u^E_{ijmt} &= \Omega^{E,r}_{mt} \tilde{y}_{imt} r_{jmt}, \end{aligned} \tag{11}$$

and

$$\begin{aligned} \log(b^*_{ijmt}) &= \delta^B_{jmt} + \varepsilon^B_{imt} + u^B_{ijmt}, \\ \delta^B_{jmt} &= \beta^{B\prime} X^B_{jmt} + \xi^B_{jmt} + \alpha^B r_{jmt} + \eta^B_{mt}, \\ u^B_{ijmt} &= \Omega^{B,\text{cons}}_{mt} \tilde{y}_{imt} + \Omega^{B,r}_{mt} \tilde{y}_{imt} r_{jmt}, \end{aligned} \tag{12}$$

where δ^E_{jmt} and δ^B_{jmt} are the card-channel-month fixed effects to be estimated. Because of the typical identification issue in discrete choice models, I normalize $\delta^E_{0mt} = 0$ and take interest rates and card characteristics in (11) and (12) as differences from the outside option.

The term in the log-likelihood containing the card choice parameters is

$$\log \mathcal{L}_{mt,E} = \sum_{i \in I_{mt}} \sum_{j \in J_{imt}} 1^E_{ijmt} \log(s^E_{ijmt|j \in J_{imt}}), \tag{13}$$

where $1^E_{ijmt} = 1(j^*_{imt} = j)$ is a dummy equal to one if individual i chooses card j in their choice set J_{imt} and $s^E_{ijmt|j\in J_{imt}}$ are logit shares, derived in Online Appendix D.1. The term $s^E_{ijmt|j\in J_{imt}}$ reflects the probability that individual i chooses card j in channel m and origination month t, *conditional* on individual i choosing to revolve a credit card balance.

To summarize, the first step of demand estimation involves market-by-market simulated maximum likelihood estimation on the log-likelihood for card choice, borrowing, and default, conditional on borrowing, to estimate product-market fixed effects (δ^E_{jmt} and δ^B_{jmt}). Estimating the fixed effects sidesteps the endogeneity problem for the moment. This step also estimates the variance-covariance matrix of private characteristics ($\varepsilon^B_{imt}, \varepsilon^D_{imt}$) (specifically σ^B_{mt} and σ^D_{mt}) and the demographic coefficients ($\Omega^{E,r}_{mt}$, $\Omega^{B,r}_{mt}$, and $\Omega^{B,\text{cons}}_{mt}$).

6.1.2 Log-likelihood for Borrowing and Transacting

In the second step of demand estimation (Step 2 in Figure A.9), I maximize a log-likelihood for the choice between transacting and revolving, which estimates δ_{0mt} and outside option utility term $\Omega^{E,\text{cons}}_{mt}$, along with the correlation coefficient for the extreme value shocks, ϱ_{mt}. I provide details and an expression for the log-likelihood of revolving/transacting in Online Appendix D.2.

6.1.3 Constant Demand Parameters

In the third and final step of demand estimation (Step 3 in Figure A.9), I estimate the constant parameters of the card-choice and borrowing equations by projecting the estimates of card-channel-month fixed effects ($\delta^E_{jmt}, \delta^B_{jmt}$) onto market fixed effects, interest rates, and observed characteristics as in (11) and (12). The same endogeneity problem persists, so I estimate the equation using an instrumental variable.

As an instrument for interest rates, I exploit a cost shock to lenders in 2011 relating to the mis-selling of payment protection insurance (PPI). PPI is a form of insurance designed to cover a loan if an individual cannot make repayments due to adverse events such as unemployment, illness, or disability. In the late 20^{th} Century, UK lenders started bundling PPI with credit products such as credit cards. In the mid-2000s, it was claimed that lenders were mis-selling PPI to borrowers. For example, lenders were selling PPI to self-employed individuals who could not use it because of their employment status. In 2006, the Financial Services Authority (FSA) started imposing fines on financial institutions for mis-selling PPI.

A significant development in the case came in January 2011, when the British Bankers' Association (BBA) took the FSA to court over its decision to impose standards on the correct selling of PPI *retrospectively*.[20] The British Bankers' Association lost the case, and in mid-2011, banks informed the BBA that they were withdrawing their support for an appeal of the decision. The ruling forced banks to reopen thousands of claims for PPI mis-selling. Around 64 million policies were mis-sold between the 1970s and 2000s, with over £33bn repaid to individuals who complained about a PPI sale. The court case loss in mid-2011 and the reopening of PPI claims led to cost increases, which were spread unevenly among banks according to how frequently they mis-sold PPI. Shortly after, some credit card lenders increased advertised interest rates for some cards in their portfolios.[21]

I create an instrument for interest rates from this cost shock by interacting lender fixed effects with a post treatment dummy.[22] The assumption is that the only channel through which the court case ruling affects individuals' card choice and borrowing is the impact of cost increases on cards' interest rates. I know no other events in the same period that affected lenders' unobservable card characteristics. I confirm the instrument's relevance in Table A.4.

[20] See *R (on the application of the British Bankers' Association) v Financial Services Authority and another [2011] EWHC 999*.

[21] Previous work argued that in the US, credit card rates are sticky relative to the cost of funds (Ausubel, 1991). However, this doesn't seem to be the case in the time series or cross section in the UK, at least in the period for which I have data. Regarding the former, monthly average funding rates are almost exactly always 10% of interest rates, with both rising by 14% between 2010 and mid-2013. And they are not merely trending together: when I regress logged interest rates on logged funding costs along with distribution channel and *month* fixed effects, the elasticity of interest rates with respect to funding costs is 0.16 and significant at lower than 1% significance level.

[22] Unfortunately, I do not have data on the proportion of PPI repayments made by each lender over time. Otherwise, I could construct a measure of lenders' exposure to the court case decision as an instrument.

6.2 Supply

6.2.1 Screening Technologies

The fourth step of estimation relates to the screening technologies. The parameters to estimate are the screening technology signals $e_{i\ell t}$ and the standard deviation of the signal noise, $\sigma_{\ell t}$. I estimate these by minimizing the residual sum of squares from the first order condition of the credit limit optimization problem. As derived in Appendix C.3, for each unique observed credit limit $\bar{b}_{ijmt}$ on card j at lender ℓ in month t, the corresponding signal $e_{i\ell t}$ satisfies

$$\int_{\omega_{i\ell t}(\bar{b}_{ijmt}, e_{i\ell t})}^{\infty} \pi_{ijmt}(e_{i\ell t}, w_{i\ell t})\phi\left(\frac{w_{i\ell t}}{\sigma_{\ell t}}\right) dw_{i\ell t} = 0, \tag{14}$$

where I define π_{ijmt} in (5). Towards an estimation strategy, note that under the distributional assumptions on private characteristics, the probability of default, as featured in π_{ijmt} is given by

$$\Delta_{imt} = \Phi\left(\eta_{mt}^D + \Omega_{mt}^D \tilde{y}_{imt} + \sigma_{mt}^D(e_{i\ell t} + w_{i\ell t})\right).$$

From this expression, I can calculate Δ_{imt} and hence π_{ijmt}, and therefore the integrand, as a function of the model parameters and the signal error.

For each observed credit limit and income, equation (14) provides an equation in which the only unknowns are the screening technology $e_{i\ell t}$ and $\sigma_{\ell t}$ (once demand parameters have been replaced with their estimates). The basis of the strategy is to estimate the screening technologies as the values that minimize the sum of squared deviations (over individuals) from the integral in (14). As in parts of demand estimation, the integral in (14) has no closed form. Therefore, for each lender-month, I simulate the integral using H Halton draws $\omega_{i\ell t}^h$, and solve

$$\min_{\{e_{i\ell t}\}, \sigma_{\ell t}} \sum_{i \in I_{\ell t}} \left(\frac{1}{H}\sum_{h=1}^{H} 1\left(\sigma_{\ell t}\omega_{i\ell t}^h > \omega_{i\ell t}(\bar{b}_{ijmt}, e_{i\ell t})\right) \pi_{ijmt}(e_{i\ell t}, \sigma_{\ell t}\omega_{i\ell t}^h)\right)^2,$$

While I could estimate the model at the lender-month level, I prefer more parsimonious models that either (i) pool months within a year or (ii) pool over all months.

6.2.2 Costs of Individualizing Interest Rates

The final step of estimation involves estimating the parameters of the function governing the costs of individualizing interest rates, $C(z_{ijmt}, \kappa_{ijmt})$. For this, I specify $C = \kappa_{ijmt} z_{ijmt}$ so that the costs of individualizing interest rates are linear in the deviation from the advertised rate, and κ_{ijmt} defines the marginal cost of deviating from the advertised rate for individual i. From equation (8), for those with $z_{ijmt} > 0$,

$$\kappa_{ijmt} = \frac{\partial \Pi_{ijmt}}{\partial z_{ijmt}}$$

so that the individual-specific marginal costs of individualizing interest rates are point identified from the derivative of Π_{ijmt}. Equation (22) in Online Appendix D.3 provides the expression for

TABLE 2. FIRST AND SECOND STEP DEMAND ESTIMATES

Variable	**Interpretation**	**Parameter**	**SE**
η^D	Default Constant	-1.90	0.02
Ω^D	Default-Income Gradient	-0.15	0.02
σ^D	S.D. in Default Unobservables	0.48	0.02
$\Omega^{B,\text{cons}}$	Revolving-Income Gradient	0.24	0.02
$\Omega^{B,r}$	Income Gradient for Revolving Elasticity	-1.16	0.02
σ^B	S.D. in Revolving Unobservable	3.70	0.06
$\text{Corr}(\varepsilon^B, \varepsilon^D)$	Correlation in Unobservables	0.38	0.02
$\Omega^{E,r}$	Income Gradient for Card-Choice Elasticity	-0.22	0.00
$\Omega^{E,\text{cons}}$	Transacting-Income Gradient	-0.11	0.01
ϱ	Transact/Revolve Substitution Parameter	0.29	0.00

$\partial\Pi_{ijmt}/\partial z_{ijmt}$, for which all the parameters have been estimated in previous steps. While this allows me to estimate a distribution of κ for the individuals not receiving the advertised rate, I cannot identify the full distribution of κ because it is not identified for those that receive the advertised rate. Alternatively, I could assume a constant marginal cost at each lender, so that $C = \kappa_{\ell mt} z_{ijmt}$ and estimate $\kappa_{\ell mt}$ for each lender and market, but as the results in the next section show, the data reject a constant value of κ across all individuals at a lender.

7 Model Estimates and Findings

7.1 Estimates of Demand Parameters

Table 2 presents the demand estimates from the first stage (log-likelihood of card choice, borrowing, and default) and the second stage (log-likelihood for transacting/revolving) of demand estimation. I report means (over markets) of parameter estimates and standard errors. Standard errors are asymptotic, coming from the inverse of the corresponding Hessian matrices.

First, I consider gradients of default utility, transacting utility, and the level of borrowing with respect to income. The negative value for Ω^D implies that higher-income borrowers are less likely to default. The positive estimate of $\Omega^{B,\text{cons}}$ means that higher-income individuals desire to borrow more. And the negative value for $\Omega^{E,\text{cons}}$ in the transaction utility indicates that higher-income individuals are less likely to transact. These findings are consistent with the Relative Income Hypothesis (Duesenberry, 1949), which posits that higher-income individuals are influenced by the consumption patterns of peers, leading to increased spending and use of credit.

The most notable insight comes from the gradients of interest rate sensitivities (i.e., α_i^E and α_i^B).

I estimate negative values of $\Omega^{B,r}$ and $\Omega^{E,r}$, implying that lower-income individuals who decide to borrow are less sensitive to interest rates (i.e., more inelastic) when choosing their card and how much to borrow. By a standard price discrimination argument, this implies that lenders have an incentive to set higher interest rates for lower-income individuals. Since low-income individuals are more likely to default, the negative correlation between credit scores (e.g., FICO) and interest rates, as observed in other countries and markets, may result from standard price discrimination instead of/alongside the pricing of default risk. I elaborate on this finding when discussing the counterfactual results in Section 8.

Regarding the distribution of private characteristics, the mean value of 0.48 for σ^D indicates unobserved heterogeneity in default, underscoring the importance of lenders' screening technologies. The correlation between unobserved preferences for borrowing and default is 0.38, implying that revolvers with a higher unobserved preference to borrow have a higher unobserved preference to default. This finding is strong evidence of adverse selection along the intensive borrowing margin. The presence of adverse selection justifies the model intuition for how lenders set optimal credit limits. The estimate is larger than the estimate of 0.14 obtained by Crawford, Pavanini, and Schivardi (2018), whose context is the Italian market for small business loans between 1988 and 1998. The estimate of the correlation between utilization and default in Benetton and Buchak (2024) is also smaller, but it is intuitive that the correlation between *utilization* and default would be smaller than that between *revolving* and default. Finally, the parameter ϱ, estimated at 0.29, indicates a reasonable substitution between the choice to transact or revolve.

Next, I consider model fit. Figure A.10 displays the model-implied and data distributions of market shares and borrowing. The fit is good, indicating that the model captures the heterogeneity in borrower behavior, adding credence to the subsequent policy simulations.

Table A.4 reports estimates and standard errors of the demand parameters recovered in the third stage of demand estimation. OLS coefficients on interest rates in both card choice (α^E) and borrowing (α^B) equations are positive, whereas IV estimates are negative. This finding indicates the severity of interest rate endogeneity. Coefficients on dummies for most rewards (i.e., air miles and purchase protection) in the card choice equation are positive across specifications, except for cashback. Cashback rewards are rare in the UK and the cashback rate tends to be lower than in the US, due to lower interchange fees in the UK.

Finally, I turn to interest rate elasticities (see equations (18) and (21) in Online Appendix C.4). Figure A.11 plots the distribution of elasticities over individuals. Three noteworthy features emerge. First, individuals are much more elastic to the interest rate in their card choice relative to their borrowing choice: this suggests that individuals do shop over advertised interest rates, but their choice of how much revolve is not so sensitive to them. Second, there is a large degree of dispersion in both elasticities. The coefficient of variation of card choice and borrowing elasticity is over one. This implies substantial heterogeneity in responsiveness to changes in interest rates across individuals. Third, both distributions are skewed. The distribution of card choice elasticities has a long tail, and the distribution of borrowing elasticities has a large mass close to zero, implying several consumers who are completely inelastic to the interest rate. Finally, elasticities are similar to other estimates of interest rate elasticities in credit markets (Karlan and Zinman, 2018) and US business credit cards (Benetton and Buchak, 2024).

7.2 Estimates of Screening Technologies

Estimation of the screening technologies delivers two sets of parameter estimates, where the first is the variation in signal mismeasurement across lenders, σ_ℓ. Table 3 reports summary statistics in the values of σ_ℓ across lenders. The coefficient of variation is 1.46, indicating substantial differences in the precision of lenders' screening technologies.

The second set of parameter estimates from supply estimation are the lenders' screening technology signals, denoted e_ℓ. Figure 4 shows the estimated screening technologies for two contrasting lenders superimposed onto a standard normal distribution. Each vertical line represents one of the lender's possible signals. I superimpose the values onto a standard normal distribution since the signals partition the standard normal distribution of ε. The left lender (lender E) has a screening technology that produces several possible signals. It is a sophisticated screening technology that provides sharp signals of borrowers' type. Lender F on the right side has a screening technology that offers only a few values, implying less precise signals on borrowers' unobservables. Figure A.12 shows the screening partitions for other lenders. Like with the values of σ_ℓ, there is substantial variation in the values and the coarseness of the screening technology across lenders.

In Section 3, I described the substantial variation across lenders in the proportion of transactors. Having estimated the structural model, I can check whether this variation correlates with the quality of lenders' screening technologies. Indeed, the correlation between σ_ℓ and the proportion of periods in which individuals repay the entire balance is 0.25. This estimate is consistent with a segmentation of credit card lenders in which lenders with the most precise screening technologies serve a riskier, but more profitable, market segment, on average (see Agarwal, Chomsisengphet, Mahoney, and Stroebel (2014) and Benetton and Buchak (2024) for credit card evidence on the profitability of various segments of the risk distribution). Lenders with more precise screening technologies are willing to serve customers who will revolve but may default, because they can accurately set lower credit limits for customers they perceive to be riskier.

In the context of endogenous screening technologies, it would be insightful for further work to assess the direction of causality between the quality of a lender's screening technologies and the risk profile of their customers. One possibility is that lenders face external factors that drive them to serve safer customers and, consequently, set high-income thresholds and do not invest in screening technologies. Another possibility is that external factors drive lenders to invest in (or be endowed with) higher-quality screening technologies, encouraging them to set lower income thresholds and

Figure 4. Screening technology at two lenders

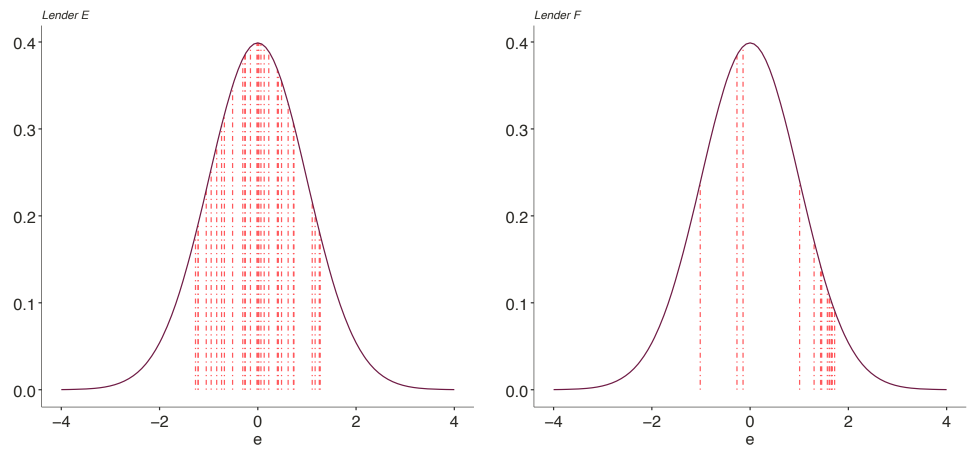

Table 3. Summary statistics for variation in signal mismeasurement

Variable	**Mean**	**SD**	**10%**	**25%**	**50%**	**75%**	**90%**
Variation in Signal Mismeasurement σ_ℓ	0.13	0.19	0.02	0.02	0.06	0.16	0.23

accept more profitable borrowers.

7.3 Estimates of the Costs of Individualizing Rates

Finally, I present and analyze the estimates of the marginal costs of individualizing interest rates, κ_{ijmt}. To start, I describe the distribution of estimates. I estimate the mean value of κ to be $1,158$, implying that increasing a single individual's interest rate five percentage points above the advertised rate (a typical value in the data) costs the lender £87 (around \$110) in reputational/computational costs on average. Figure A.13 depicts the distribution of κ_{ijmt} estimates. The plot reveals significant variation across individuals, ranging from essentially 0 to values as large as 3000 for some individuals. Note that these estimates represent the marginal costs of individualizing interest rates for those that do not receive the advertised interest rate. I hypothesize that this distribution is a lower bound on the full distribution of costs including those that do receive the advertised rate, because those that receive the advertised rate should, all else equal, be the ones with higher values of κ_{ijmt}. To summarize, these estimates are substantial, and hence explain why lenders do not exhaust the full extent of the regulation limit on individualizing interest rates.

Next, there are differences across lenders in their customers' marginal cost of individualizing interest rates. Estimates of the lender-average of κ vary from 918 to 1,627. The averages are largest at lenders associated with more sophisticated, higher-education, and higher-income customer bases. These differences are statistically significant, as confirmed by regressions of κ_{ijmt} on lender and distribution-month dummies.

Finally, I relate the values of κ_{ijmt} to other model variables. In the left side of Figure 5, I present binscatter plots of the log of differences between individualized and advertised interest rates, $\log(z_{ijmt})$, against κ_{ijmt}. As expected, among those with $z_{ijmt} > 0$, those with the largest values of κ_{ijmt} receive the smallest values of z_{ijmt}. On the right side of Figure 5, I present binscatter plots of the values of κ_{ijmt} against logged income, y_i. Those with the highest income have the largest value of κ_{ijmt}. To analyze this finding further, I regress the values of κ_{ijmt} against logged income and risk signals, along with dummies for lenders and distribution-month pairs. The estimates are provided in Table A.5. The coefficient on logged income implies that those with a 1% higher income have a 3.20 higher marginal cost of individualizing the interest rate. The coefficient on the risk signal is negative, implying that riskier individuals have a lower value of κ. These results align with what we would expect if higher-income, lower-risk individuals are either more aware of—or more aggrieved by—receiving an interest rate higher than advertised. Indeed, we may expect that individualized interest rates are *shrouded* to higher-risk, lower-income individuals.

8 Counterfactual Analysis

This paper's central empirical finding is that the credit limit is the primary contractual variable individualized by lenders in the UK credit card market. Related to this empirical fact is the regulatory environment, which requires lenders to advertise an interest rate for each credit card product offered. Despite the requirement to advertise a card-level interest rate, lenders have the flexibility to individualize interest rates that goes almost entirely unused. My model explains the lack of individualized interest rates through shadow costs of individualizing interest rates, κ. In this section, I use the estimated model to understand what would happen to lender and borrower outcomes if these costs and regulatory constraints were removed.

8.1 Implementation

I design a counterfactual to compare the two extremes of interest rate individualizing. In the baseline, close to the EU context, lenders give all customers the advertised market interest rates. The counterfactual, which mimics the US environment, has lenders setting interest rates and credit limits without any requirements to promote an interest rate as a regulatory requirement and no costs of individualizing interest rates. Mathematically, this means removing equation (7) and setting $\kappa_{ijmt} = 0$ for all customers. After simulating the baseline and counterfactual environments, I analyze the resulting differences in interest rates and credit limits, among other outcomes.

To implement these scenarios, I simulate the February 2013 in-branch market under the new regime, taking income thresholds and card characteristics from the data. In the counterfactual, for customer i, lender ℓ now solves simultaneously for all interest rates and credit limits across their

Figure 5. Costs of individualizing interest rates

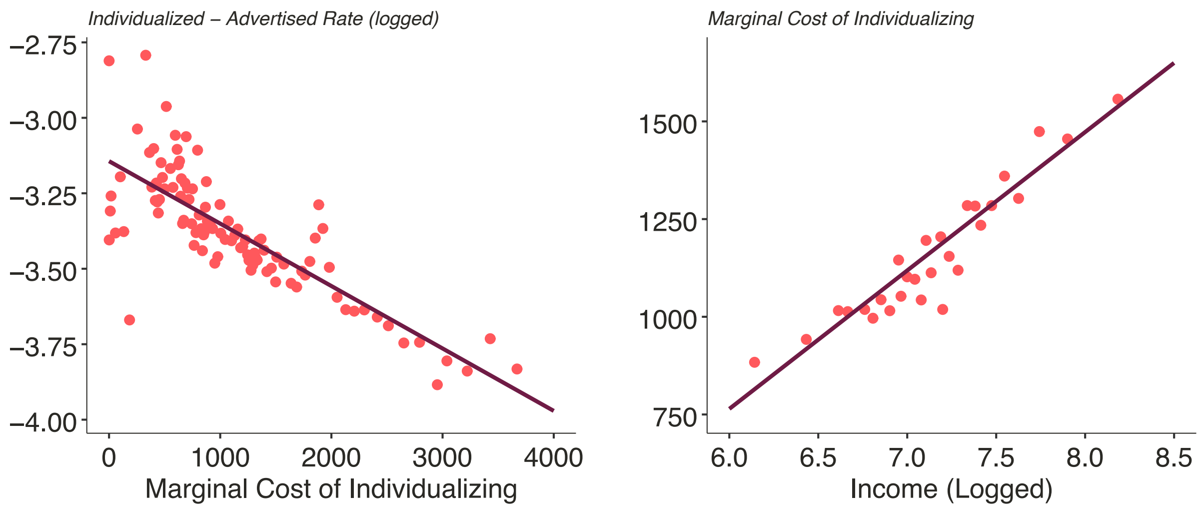

cards $J_{i\ell}$ that consumer i is eligible for. Formally, given other lenders' optimal interest rate choices $\boldsymbol{r}^*_{-i\ell mt}$, for customer i, lender ℓ solves the unconstrained problem

$$\max_{\boldsymbol{r}_{i\ell},\bar{\boldsymbol{b}}_{i\ell}} \sum_{j\in J_{i\ell}} s^E_{ij}(\boldsymbol{r}_{i\ell},\boldsymbol{r}^*_{-i\ell})\mathbb{E}\left[\min\{b^*_{ij},\bar{b}_{ij}\}\pi_{ij}\right], \tag{15}$$

where s^E_{ij} represents card choice probability, $\min\{b^*_{ij},\bar{b}_{ij}\}$ reflects the constrained revolving level, and π_{ij} denotes profit per-unit credit. Customer i chooses their card, unpaid balance, and whether to default.[23] Like supply estimation, I minimize the residual from the first order conditions to equation (15) to calculate $\boldsymbol{r}_{i\ell}$ and $\bar{\boldsymbol{b}}_{i\ell}$ for all individuals i. Appendix E.1 provides the first order conditions I use to calculate counterfactual interest rates and credit limits. This implementation is computationally intensive because I have to solve a separate optimization problem for each consumer. Consequently, I draw a representative sample of approximately 20% of the market.

I measure three sets of endogenous variables in the baseline and counterfactual. The first is the set of equilibrium origination interest rates and credit limits. The second set—the demand-side variables—includes consumers' card choice, borrowing level, and consumer surplus. The latter is

$$CS_i = \frac{1}{\alpha_i}\log\left(\sum_{j\in J_i}\exp\left(\bar{U}^E_{ij}\right)\right), \tag{16}$$

[23]In the counterfactual, I follow the baseline model by assuming that individuals know their potential interest rate at each lender when choosing their card; results from the case in which consumers do not know interest rates are available on request. I maintain the assumption that consumers do not know their credit limit to ensure that I am only changing one object at a time and also due to the absence of any credible method to measure what individuals' preferences over credit limits would be, were they known to the consumer.

Figure 6. Distribution of interest rates in baseline and counterfactual

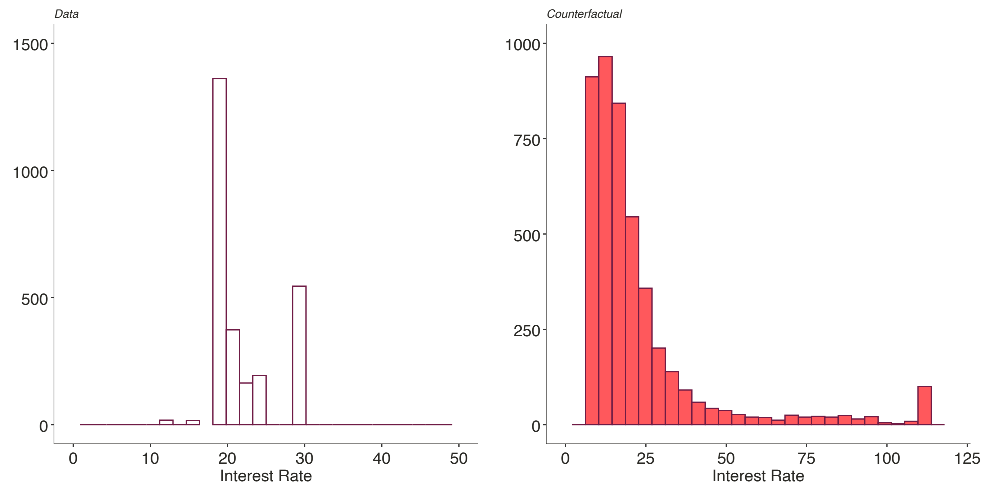

where $\bar{U}_{ij}^{E}$ is equal to $\bar{V}_{ij}^{E}/\varrho$, a scaled version of indirect utility. Note that this expresses the value of having access to credit cards in terms of percentage points of interest. This is because the price variable in the demand model is the interest rate r, rather than the total monetary cost in dollars of borrowing, as is usually the case in standard demand estimation. This is purely a matter of interpretation: the expression in equation (16) can be rescaled by a representative revolving balance to convert the value into the typical version of consumer surplus, denominated in dollars. Furthermore, the following results present consumer surplus as percentage changes relative to the baseline, so this rescaling factor cancels.

The third set of endogenous variables includes supply-side variables. I focus on ex-post profit, which for a revolver i is given by

$$\pi_{ij}^{\text{post}} = b_{ij}\Bigg[(1-\mathcal{D}_i)(r_j - c_j) + \mathcal{D}_i(-1-c_j)\Bigg],$$

where $\mathcal{D}_i$ is equal to 1 if borrower i defaults.

8.2 Counterfactual Results

8.2.1 Interest Rates and Credit Limits

The main variable driving changes in the counterfactual is the interest rate. Figure 6 shows the distribution of interest rates in the baseline and the counterfactual scenarios. The distribution becomes highly individualized in the counterfactual, with thousands of unique interest rate values compared to a few distinct values in the baseline. The coefficient of variation in interest rates increases from 20% in the baseline to 92% in the counterfactual, and the standard deviation increases from 0.04 to 0.21. Together, these imply a significant increase in interest rate dispersion.

The net directional effect on the values of interest rates is *ex ante* ambiguous. Average interest rates may increase because lenders can now price discriminate, but they could also decrease because lenders need not pool interest rates across risk types. The former dominates in the counterfactual, with interest rates increasing by three percentage points, equivalent to a 13% increase.

The net increase in interest rates in the counterfactual masks vast heterogeneity in interest rate changes across borrowers. In the counterfactual, lenders practice traditional price discrimination. Individuals with below median elasticity of demand (the most inelastic customers) receive an average interest rate increase of 12 percentage points. On the contrary, interest rates fall by six percentage points for the most elastic individuals.

Since the most inelastic individuals are those with the lowest incomes and thus the highest average default risk, interest rate and default probability are positively correlated in the counterfactual. However, this is not caused by lenders pricing in default risk. In the counterfactual, individuals with high overall default risk, but elastic demand obtain a lower interest rate than in the baseline.

To shed further light on this finding, in Figure 7, I plot baseline and counterfactual interest rates as a function of interest rate sensitivity $|\alpha_i^E|$. I do this separately for those with small and large risk signals $\breve{\varepsilon}$. Three features are noteworthy. First, as expected, there is no relationship between rates and sensitivity (nor rates and risk signals) in the baseline. This finding follows from the absence of within-card and across-card variation in rates. Second, the counterfactual gradient is negative for both low and high default risk signals, consistent with the price discrimination mentioned previously. Third, the counterfactual curves and patterns are almost identical for those with high and low risk signals, confirming that interest rates do not price in this default risk.

The lenders' second screening instrument is the credit limit. Figure A.14 displays the distribution of credit limits in the data and the counterfactual scenario. Credit limits remain individualized and become more dispersed, with the coefficient of variation in credit limits increasing by 38% and the standard deviation rising by 24%. Credit limits fall by 10% on average in the counterfactual. The coincidence of rising interest rates and falling credit limits follows the intuition of downward sloping cost curves in Einav and Finkelstein (2011).

The intuition for why counterfactual lenders combine individualized interest rates and credit limits is as follows. Credit limits do not affect individuals' card choices or desired borrowing and *only* serve to manage downside risk from default. Interest rates, however, affect an individual's choice of card and level of desired borrowing through the terms s_{ij}^E and b_{ij}^* in the profit function for individual i. Individualized prices are thus a device for standard price discrimination. As an example, lenders set lower interest rates for individuals with elastic demand to encourage them to choose their cards and generate interest revenue. Then, among those with elastic demand, they set low (respectively, high) credit limits if the unobserved risk signal is large (respectively, small).

Figure 7. Counterfactual interest rates and price elasticities by risk type

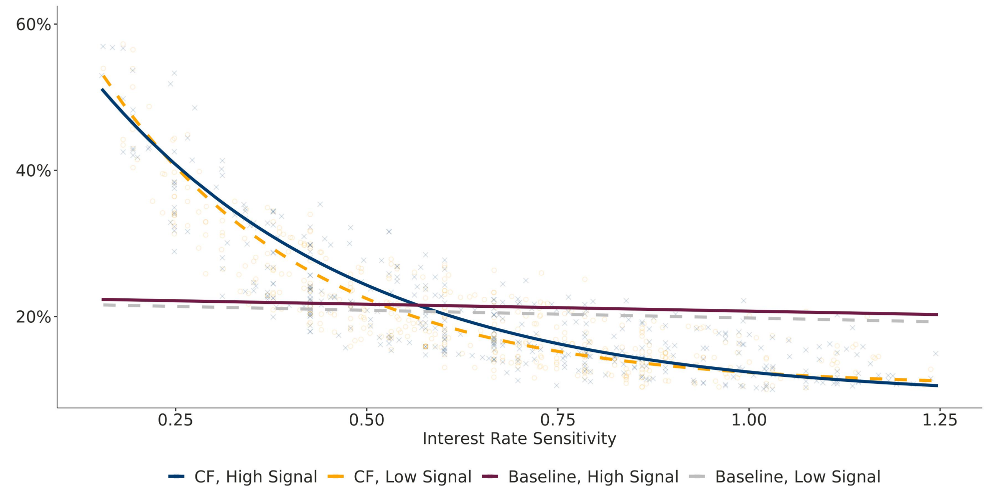

Notes: The solid downward-sloping curve represents the least squares fit between observations of interest rates and interest rate sensitivity in the counterfactual scenario for those with default risk signal above the 85^{th} percentile; the dashed downward-sloping curve represents the analog of the solid curve except estimated on those with default risk below the 15^{th} percentile. The solid horizontal line represents the best fit between interest rate and interest rate sensitivity in the baseline for those with default risk signal above the 85^{th} percentile; the dashed horizontal line represents the analog except estimated on those with default risk below the 15^{th} percentile.

8.2.2 Demand-Side Variables

Next, I explore changes to borrowers' outcomes. In the counterfactual, 10% of customers switch cards and 8% switch lenders. That the vast majority of customers choose the same card in the context of changing interest rates highlights the importance of card characteristics (X_{jmt} and ξ_{jmt}) along with the relative price-inelasticity of demand. Though the borrowing level does not adjust materially on net, it increases by approximately 11% for those with above-median income and decreases by 15% for those with below-median income. This result is consistent with lenders' motives to price-discriminate: lenders decrease interest rates for safe, elastic individuals to incentivize them to borrow/borrow more.

In the counterfactual, consumer surplus falls by 1% on average relative to the baseline. However, as with interest rates, this decrease in the average masks heterogeneity across borrowers. In Figure A.15, I plot the distribution of percentage changes in consumer surplus for high-elasticity and low-elasticity individuals. Consumer surplus generally increases in the counterfactual for the high-elasticity group—a 15% increase on average—because they benefit from lower interest rates. Consumer surplus falls by 2% on average for the low-elasticity group.

8.2.3 Lenders' Profits

By price discriminating, lenders' ex-post profits increase by 23% on average. Such increases are as expected given the size of the costs of individualizing interest rates that previously justified why lenders only deviated from the advertised rate for 10% of originations. Identifying the exact sources of these costs is not the focus of this paper. I briefly discussed some possibilities in Section 5.3.4, and I add further details here and in Online Appendix E.2. Interviews with policy and industry experts highlighted two factors: (i) infrastructure investments, and (ii) reputational concerns. While important in practice, the former is less nuanced, so I focus here on reputational risk.

Industry and policy documents allude to reputational risk to lenders from adopting risk-based pricing in the EU regulatory context (House of Commons Treasury Committee, 2003). For instance, there could be significant backlash, were lenders to advertise a low interest rate and then set high individualized APRs after individuals have chosen their card. Indeed, in 2003, the UK Government Treasury Committee described this kind of risk-based pricing as an "unacceptable practice," raising "serious transparency issues."[24] In the aftermath of a global financial crisis and a "PPI scandal," any further erosion of trust in UK banks may have come at significant reputational cost.

8.2.4 Assumptions Underlying the Counterfactual Exercise

Before I conclude, it is worth acknowledging the assumptions underlying the results of this counterfactual exercise. As is standard in counterfactual analysis, the maintained assumption is that all exogenous model elements remain unchanged. However, we may expect variables such as product characteristics to change in response to the removal of restrictions on individualizing interest rates. Further, the counterfactual relies on the assumption that consumers know their individualized rate at each lender and that lenders can price discriminate because elasticities depend only on observable income. Further, I cannot account in my framework for endogenous responses in alternative credit markets (for example, payday loans) that might follow changes to the credit card market. And, the strong intuitions on how lenders take advantage of being able to individualize interest rates and credit limits may not follow through to other markets if product choices depend on the quantity of credit the individual plans to use. Finally, since I cannot include the extensive margin of not holding a credit card at all, I cannot capture the transitions in and out of the credit card

[24] In April 2022, the UK Chancellor of the Exchequer stated that it was "important that advertised APRs reflect the rate the consumer is likely to receive." `https://on.ft.com/3uKGZ92` last accessed 9 March 2025. The Chancellor's statement was made in response to a report on advertised APRs by the largest UK consumer website, MoneySavingExpert.com (`https://www.moneysavingexpert.com/news/2022/03/chancellor-ask-regulator-credit-card-loan-aprs-martin-lewis/` last accessed 9 March 2025.) As part of their report, the website conducted two nationally representative surveys of over 2,000 British adults. The findings revealed that 35% of customers who were offered a higher rate than advertised stated that it had a negative effect on their financial well-being, and the same percentage claimed the higher rate had a detrimental impact on their emotional well-being.

market that would result from changes to interest rates. These extensions would all need to be acknowledged in a more complete assessment of the net effect of altering interest rate regulations.

9 Concluding Remarks

In this paper, I investigate how credit card lenders manage customers' unobserved default risk by individualizing contracts. Using novel, statement-level microdata, I estimate a structural model of the credit card market. The model features credit rationing on the *intensive margin* as lenders set binding credit limits. The central modeling innovation is the lender screening technology that provides noisy signals on borrowers' unobserved types. Lenders make credit limits contingent on these signals, and the coarseness of the set of signals offered by the screening technology corresponds to the coarseness of their equilibrium credit limit distribution. Indeed, supply-side estimates imply marked differences in lenders' screening technologies. Regarding their pricing choices, UK lenders face heterogeneous and non-trivial costs of individualizing a customer's interest rates. As a result, lenders do not frequently deviate from the rate they must advertise by law.

On the other side of the market, the demand estimates imply that lower-income borrowers have more inelastic demand. This finding highlights an alternative motive for "risk-based" pricing in line with price discrimination. Lenders use credit limits to manage default risk, whether risk-based pricing is allowed or not.

I use the estimated model to evaluate a counterfactual where lenders can fully and freely individualize interest rates and credit limits. The corresponding interest rate discrimination results in consumer surplus gains for high-income individuals and losses for low-income individuals, and lenders' profits increase on average. My findings imply that lenders' costs of individualizing interest rates protect high-risk consumers at the expense of lenders' profits and low-risk consumers.

Important questions remain regarding why UK credit card lenders do not base interest rates on risk, be it through (i) exploiting the full extent that the regulation allows, (ii) offering a broader menu of cards of varying advertised rates, or (iii) repricing customers after origination. I describe some relevant considerations in Online Appendix E.2, but a more thorough investigation is warranted. Given that the estimated costs of individualizing interest rates are substantial, finding the drivers of these shadow costs is a fruitful avenue for further work.

There are other possible extensions of this paper. First, future work could analyze counterfactuals that change lenders' screening technologies. One example would be a scenario in which lenders must share their screening technologies. This counterfactual would create a setting closer to the US, where many lenders use FICO scores.[25] Second, building on the empirical work of Panetta,

[25] One of the remedies in the FCA study of the UK credit information market was to mandate data sharing amongst credit reference agencies (FCA, 2023). Guttman-Kenney and Shahidinejad (2024) argues that mandating

Schivardi, and Shum (2009), my model can analyze the welfare effects of mergers in which the merging lenders combine their screening technologies. Along with the typical trade-off between cost synergies and increased concentration, mergers would benefit from improved screening technologies. The model can measure this element of merger synergies, which is typically challenging.

Regarding external validity, consumer credit markets across different countries and time periods have used varying combinations of individualized prices and quantities. No general theory exists to explain how product features and regulatory environments interact to influence how lenders individualize contracts. Deriving conditions on the market that deliver tailored prices or quantities is another natural sequel to this work.

Appendices

Supplemental appendices are available **here**.

References

The numbers at the end of every reference link to the pages citing the reference.

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. 5, 34

——— (2017): "Do Banks Pass through Credit Expansions to Consumers Who want to Borrow?" *The Quarterly Journal of Economics*, 133, 129–190. 4, 9, 13, 20, 21

AKERLOF, G. A. (2001): "Behavioral Macroeconomics and Macroeconomic Behavior," Nobel Prize Committee, Nobel Prize Lecture. 1

ALBANESI, S. AND D. F. VAMOSSY (2019): "Predicting Consumer Default: A Deep Learning Approach," Working Paper 26165, National Bureau of Economic Research, https://doi.org/10.3386/w26165. 9

ATTANASIO, O. P. AND G. WEBER (2010): "Consumption and Saving: Models of Intertemporal Allocation and Their Implications for Public Policy," *Journal of Economic Literature*, 48, 693—-751. 1

AUSUBEL, L. M. (1991): "The Failure of Competition in the Credit Card Market," *The American Economic Review*, 81, 50–81. 30

information sharing in credit card markets increases competition.

Aydin, D. (2022): "Consumption Response to Credit Expansions: Evidence from Experimental Assignment of 45,307 Credit Lines," *American Economic Review*, 112, 1–40. 4

Benetton, M. (2021): "Leverage Regulation and Market Structure: A Structural Model of the U.K. Mortgage Market," *The Journal of Finance*, 76, 2997–3053. 5, 27

Benetton, M. and G. Buchak (2024): "Revolving Credit to SMEs: The Role of Business Credit Cards," *Working Paper*, http://dx.doi.org/10.2139/ssrn.4997456. 7, 14, 15, 17, 22, 33, 34

Benetton, M., A. Gavazza, and P. Surico (2025): "Mortgage Pricing and Monetary Policy," *American Economic Review*, 115, 823–863. 27

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63, 841–890. 18

Blinder, A. S. and J. E. Stiglitz (1983): "Money, Credit Constraints, and Economic Activity," *The American Economic Review*, 73, 297–302. 1

Calomiris, C. W., S. D. Longhofer, and D. M. Jaffee (2017): *Credit Rationing*, London: Palgrave Macmillan UK, 1–10. 4

Castellanos, S. G., D. Jiménez Hernández, A. Mahajan, E. Alcaraz Prous, and E. Seira (2018): "Contract Terms, Employment Shocks, and Default in Credit Cards," Working Paper 24849, National Bureau of Economic Research, https://doi.org/10.3386/w24849. 20

Chiappori, P. and B. Salanie (2000): "Testing for Asymmetric Information in Insurance Markets," *Journal of Political Economy*, 108, 56–78. 2, 14

Cohen, A. and L. Einav (2007): "Estimating Risk Preferences from Deductible Choice," *American Economic Review*, 97, 745–788. 20

Collier, B. L., D. A. Hartley, B. J. Keys, and J. X. Ng (2024): "Credit When You Need It," Working Paper 32845, National Bureau of Economic Research, https://doi.org/10.3386/w32845. 20

Crawford, G. S., N. Pavanini, and F. Schivardi (2018): "Asymmetric Information and Imperfect Competition in Lending Markets," *American Economic Review*, 108, 1659–1701. 14, 17, 22, 26, 33

Cuesta, J. I. and A. Sepulveda (2021): "Price Regulation in Credit Markets: A Trade-Off between Consumer Protection and Credit Access," *Working Paper*, http://dx.doi.org/10.2139/ssrn.3282910. 5

Department for Business Innovation and Skills (2010): "Guidance on the regulations implementing the Consumer Credit Directive updated for EU Commission Directive," Guidance Note. 21

Duesenberry, J. (1949): *Income, Saving, and the Theory of Consumer Behavior*, Harvard University Press. 32

Edelberg, W. (2006): "Risk-based pricing of interest rates for consumer loans," *Journal of Monetary Economics*, 53, 2283–2298. 5

Einav, L. and A. Finkelstein (2011): "Selection in Insurance Markets: Theory and Empirics in Pictures," *The Journal of Economic Perspectives*, 25, 115–138. 39

Evans, D. and R. Schmalensee (2005): *Paying with Plastic: The Digital Revolution in Buying and Borrowing*, MIT Press Books, The MIT Press, 2nd ed. 23, 24

FCA (2015a): "Credit Card Market Study: Final Findings Report," *Financial Conduct Authority Reserach Publications*. 7

——— (2015b): "Credit Card Market Study Interim Report: Annex 10 - Account Level Data," *Financial Conduct Authority Research Publications*. 5

——— (2015c): "Credit Card Market Study Interim Report: Annex 3 - Results from the consumer survey," *Financial Conduct Authority Research Publications*. 7, 18

——— (2023): "Credit Information Market Study: Final Report," *Financial Conduct Authority Research Publications*. 9, 42

Fitch (2013): "Credit Card Index — UK (Q1 2013)," *Fitch Ratings*. 8

Gabaix, X. and D. Laibson (2006): "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets," *The Quarterly Journal of Economics*, 121, 505–540. 28

Galenianos, M. and A. Gavazza (2022): "Regulatory Interventions in Consumer Financial Markets: The Case of Credit Cards," *Journal of the European Economic Association*, 1897–1932. 10, 16

Ganong, P. and P. Noel (2020): "Liquidity versus Wealth in Household Debt Obligations: Evidence from Housing Policy in the Great Recession," *American Economic Review*, 110, 3100–3138. 20

Gathergood, J., N. Mahoney, N. Stewart, and J. Weber (2019): "How Do Individuals Repay Their Debt? The Balance-Matching Heuristic," *American Economic Review*, 109, 844–75. 17

Gouriéroux, C. and A. Monfort (1996): *Simulation-based Econometric Methods*, Oxford University Press. 28

Gross, D. B. and N. S. Souleles (2002a): "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data," *The Quarterly Journal of Economics*, 117, 149–185. 4

——— (2002b): "An Empirical Analysis of Personal Bankruptcy and Delinquency," *The Review of Financial Studies*, 15, 319–347. 20

Guttman-Kenney, B. and A. Shahidinejad (2024): "Unraveling Information Sharing in Consumer Credit Markets," *Working Paper*, `http://dx.doi.org/10.2139/ssrn.4629496`. 42

House of Commons Treasury Committee (2003): "Transparency of Credit Card Charges," *Stationery Office by Order of the House*. 41

Indarte, S. (2023): "Moral Hazard versus Liquidity in Household Bankruptcy," *The Journal of Finance*, 78, 2421–2464. 20

Karlan, D. and J. Zinman (2018): "Long-Run Price Elasticities of Demand for Credit: Evidence from a Countrywide Field Experiment in Mexico," *The Review of Economic Studies*, 86, 1704–1746. 34

Lancaster, K. J. (1966): "A New Approach to Consumer Theory," *Journal of Political Economy*, 74, 132–157. 17

Livshits, I., J. C. Mac Gee, and M. Tertilt (2016): "The Democratization of Credit and the Rise in Consumer Bankruptcies," *The Review of Economic Studies*, 83, 1673–1710. 5, 21

Matcham, W. (2024): "A Note on Broader Credit Card Literature and Credit Card Regulation," *Mimeo*. 7, 17

McKinnon, R. I. (1973): *Money and Capital in Economic Development*, Washington, D.C.: Brookings Institution. 1

Nelson, S. (2022): "Private Information and Price Regulation in the US Credit Card Market," *Working Paper*. 5, 14, 20, 26

Panetta, F., F. Schivardi, and M. Shum (2009): "Do Mergers Improve Information? Evidence from the Loan Market," *Journal of Money, Credit and Banking*, 41, 673–709. 42

Stiglitz, J. E. and A. Weiss (1981): "Credit Rationing in Markets with Imperfect Information," *American Economic Review*, 71, 393–410. 1, 4

Wang, L. (2023): "Regulating Competing Payment Networks," *Working Paper*. 7