Data and Welfare in Credit Markets*

Mark Jansen[1] Fabian Nagel[2]

Constantine Yannelis[3,4] Anthony Lee Zhang[3]

July 1, 2025

Abstract

We show how to measure the welfare effects arising from increased data availability. When lenders have more data on prospective borrower costs, they can charge prices that are more aligned with these costs. This increases total social welfare and transfers surplus across borrower types. We show that under certain assumptions the magnitudes of these welfare changes can be estimated using only quantity and price data. Applying our methodology to bankruptcy flag removal, we find that in a counterfactual world where bankruptcy flags are never removed from credit reports, previously-bankrupt borrowers' surplus decreases substantially, whereas efficiency increases only modestly.

*We are grateful to Jacelly Cespedes, Tony Cookson, Eduardo Davila, Anthony DeFusco, Andreas Fuster, Andra Ghent, Zhiguo He, Kyle Herkenhoff, Ray Kluender, Christian Leuz, Alberto Manconi, Sanjog Misra, Matt Notowidigdo, Nicola Pavanini, Nathan Seegert, Alp Simsek, Huan Tang, Ansgar Walther, Jialan Wang, Wenyu Wang, Laura Veldkamp, Alexander Zentefis, and Wanran Zhao for thoughtful discussions and comments, as well as seminar participants at University of Chicago, Yale School of Management, the NYU Stern School of Business, London Business School, Cambridge University Judge School of Business, the Duke Fuqua School of Business, the Rice Jones School of Business, HEC Paris, INSEAD, Washington University in St. Louis, the University of Kentucky Gatton College of Business, the Swiss Finance Institute, the Swiss Federal Institute of Technology Lausanne, the Stockholm School of Economics, the Berkeley Conference on Consumer Law, Durham University Business School, SAIF, HEC Lausanne, UW Seattle, FIRS, the London Business School Summer Finance Symposium, the Swiss Society for Financial Market Research Conference, the Frankfurt Conference on Regulating Financial Markets, the Young Scholars in Finance Consortium, the Midwest Finance Association Conference, the Minnesota Corporate Finance Conference, and the NBER Summer Institute. Greg Tracey provided superb research assistance. We are especially grateful for the comments of an anonymous referee, which greatly improved the paper. Calculated (or derived) based on credit data provided by TransUnion, a global information solutions company, through a relationship with the Kilts Center for Marketing at the University of Chicago Booth School of Business. TransUnion (the data provider) has the right to review the research before dissemination to ensure that it accurately describes TransUnion data, does not disclose confidential information, and does not contain material it deems to be misleading or false regarding TransUnion, TransUnion's partners, affiliates or customer base, or the consumer lending industry.

[1]University of Utah, Eccles School of Business
[2]Stanford University, Stanford Graduate School of Business
[3]University of Chicago, Booth School of Business
[4]National Bureau of Economic Research

1 Introduction

The past half century saw an explosion of data that lenders can use to screen and score potential borrowers. In principle, increasing data availability should allow lenders to charge interest rates that are more closely aligned with borrowers' true risks, thereby improving the efficiency of credit allocation and raising social welfare. But data may also lead to changes in interest rates, shifting the distribution of social surplus between lenders and borrowers of different risk levels. How can we quantify the effects of increased data availability on social welfare and the distribution of surplus between lenders and different kinds of borrowers?

This paper develops a tractable framework for measuring the welfare effects of increased data availability, viewing changes in data access as enabling a form of *third-degree price discrimination*. In an empirical application to consumer lending markets, we show that keeping prior bankruptcy information on consumer credit reports would substantially reduce the social surplus of previously-bankrupt borrowers while only slightly increasing allocative efficiency in credit markets.

Whether lenders should be permitted to leverage granular data has been central to policy debates.[1] Recent policy decisions, such as the 2022 removal of medical debt collections under $500 from U.S. credit reports, highlight the tension between personalized pricing and redistribution. Laws such as the General Data Protection Regulation (GDPR) in the EU and the Fair Credit Reporting Act (FCRA) in the US aim to balance these tradeoffs by limiting the use of certain borrower data. For example, the European Data Protection Supervisor (2021) recommended allowing lenders to use some data for personalized pricing but "*clearly delineating the categories and sources of personal data that may be used for the purpose of creditworthiness assessment*." Regulatory scrutiny of personalized loan pricing has intensified as banks have increased their spending on IT and technology: this spending rose to $74 billion in 2022, an increase of 37% from 2017.

We develop a model to assess how data shapes welfare in credit markets, where competitive lenders serve both high- and low-cost borrowers. In the absence of data, lenders cannot distin-

[1] Lenders have screened borrowers based on informal data since antiquity (Calomiris and Neal, 2013). Modern credit scoring systems emerged in 1958 with Fair Isaac and Company (FICO), and expanded over subsequent decades. The introduction of machine learning techniques and alternative credit data, such as VantageScore, likely had significant welfare implications that our framework can help quantify.

guish high- and low-cost borrowers. Suppose, counterfactually, that data was available, and lenders could always distinguish high-cost consumers from low-cost consumers: how would this change each borrower group's consumer surplus and total social welfare?

Since consumer loans generally involve upfront borrowing and staggered repayment, following Marshall (1920) and Vives (1987), we think of consumption over future periods as a "large composite good" and approximate future-period utility as linear. Marshallian *consumer surplus* —the "area under the demand curve"—is then exactly equal to the dollar value consumers derive from borrowing. This allows us to define surplus in consistent dollarized units across consumer groups with different default rates, providing a microfoundation for the classical "surplus trapezoid" diagrams in the consumer lending setting.

Data policy entails a tradeoff between efficiency and redistribution. Retaining data raises social welfare by improving credit allocation, but also transfers surplus from high-cost to low-cost borrowers. The welfare gains scale quadratically with data-induced price changes, whereas the redistribution effects scale linearly; thus, the transfer effects tend to dominate when data-induced price changes are small. Regulators may be willing to limit data availability in settings where the resulting cross-subsidies are deemed valuable enough to justify the efficiency losses.

If loan demand is approximately linear in payments, the welfare and transfer effects of data can be expressed as simple functions of observed loan prices and quantities, with and without data. We apply this finding in the context of consumer bankruptcy flag removal. Under the Fair Credit Reporting Act (FCRA), flags indicating the occurrence of consumer bankruptcy are removed after seven (ten) years after a Chapter 13 (Chapter 7) bankruptcy filing. We exploit within-consumer variation, comparing loan terms just before and after flag removal to estimate counterfactual outcomes for previously-bankrupt borrowers. For never-bankrupt borrowers, we back out counterfactual prices using the lenders' zero-profit condition, and infer quantity changes by assuming equal demand elasticities across groups. With these price and quantity changes in hand, we can then calculate the welfare and redistribution effects under a policy that never removes bankruptcy flags.

Our empirical application studies the social welfare losses and transfers induced by bankruptcy flag removal policy in the US auto lending market, using administrative data from TransUnion.

A difference-in-differences design shows that individuals whose flags are removed experience a 17-point increase in credit scores, a 22.6-basis-point reduction in interest rates, and an $18 increase in borrowing. Using lenders' zero-profit condition and the assumption of equal demand elasticities, we back out the smaller price decrease and quantity increase that never-bankrupt individuals would counterfactually face. Our estimates imply that, if flags were never removed, social welfare would increase modestly, by $598,000 a year in aggregate; however, roughly $19 million in borrower surplus each year would be transferred from the previously-bankrupt to the never-bankrupt. In other words, the counterfactual policy generates only about $0.03 in social surplus for every dollar redistributed between groups. Thus, while the current policy generates allocative inefficiencies, its welfare costs are fairly small relative to its redistributional effects.

Our baseline model is stylized and makes strong assumptions in some settings. In theoretical extensions, we discuss how our approach can be extended to incorporate forces such as imperfect competition, adverse selection, interest rate menus, and endogenous default. While we frame our results around the application to bankruptcy flag removal, our methodology can be applied in any setting where lenders acquire pricing data which is informative about consumers' default risks.

This paper joins together two literatures: work on data in financial markets and work on price-theoretic approaches to study markets. Theoretical work illustrating different channels through which data can affect lending market outcomes include Begenau et al. (2018), Acemoglu et al. (2019), Farboodi et al. (2019), Farboodi and Veldkamp (2020), Jones and Tonetti (2020), He, Huang and Zhou (2020), and Liu et al. (2023).[2] Empirical papers on data include Tang (2019), who uses data from a Fintech to value privacy; Nelson (2018) and Blattner and Nelson (2021), who use structural models to analyze the efficiency and distributional con-

[2]There is a large theoretical literature studying the welfare effects of data sharing policies in credit markets in particular. These papers largely emphasize different channels to our paper. Pagano and Jappelli (1993) analyze lenders' incentives to share information, and how information sharing can alleviate adverse selection. Vercammen (1995), Padilla and Pagano (2000), and Elul and Gottardi (2015) focus on how the threat of updating credit scores and thus interest rates dynamically affect borrowers' default incentives. Kovbasyuk and Spagnolo (2021) and Blattner et al. (2022) apply Bayesian persuasion methods to analyze optimal information structures in credit markets, taking into account adverse selection and dynamically evolving consumer types. Chatterjee et al. (2020) emphasize how credit scores are a substitute for dynamic reputation in incentivizing repayment. Relative to the literature, our paper emphasizes exogenous heterogeneity in the willingness to pay for credit; this is key to creating allocative efficiency gains from information revelation, which is a core focus of our analysis. In the baseline model, we abstract away from forces such as dynamic types, reputation, moral hazard and adverse selection, though we discuss the effects of many of these effects in extensions and appendices.

sequences of information removal in the credit card market; and Fuster et al. (2022), who analyze the distributional effects of machine learning models in credit markets.

Our contribution to the data literature is to develop a *price-theoretic* approach to quantify the welfare and transfer effects of data. In the baseline, we make a number of simplifying assumptions relative to purely theoretical and structural papers. The benefit of these simplifications is that we attain sharp qualitative results about the effects of data in credit markets, and a simple "sufficient-statistics" approach to empirically estimate the effects of data. We show theoretically and empirically that—somewhat surprisingly—when data is not very informative about default rates, the transfer effects of data availability are large relative to the effects on aggregate welfare. We view our contribution as complementary to richer structural approaches, and we discuss in extensions and appendices the extent to which our methodology is robust to other theoretical channels we disregard in the baseline.

Our price-theoretic approach builds classical ideas about consumer and producer surplus and deadweight loss triangles (Harberger, 1964), brought into the context of insurance markets by Einav et al. (2010), and extended to credit markets in DeFusco et al. (2022). Our paper applies this price-theoretic approach to data addition or removal in credit markets, showing how price and quantity information are sufficient to measure data-induced changes in surplus and welfare. DeFusco et al. (2022) focuses on adverse selection—costs depending on prices—but does not discuss data policy. In the baseline, we assume away adverse selection but address data policy with adverse selection briefly in an extension of the model. Beyond the different focus, our paper also uses a different microfoundation from DeFusco et al. (2022), allowing us to show that the price-theoretic approach can still be used when incorporating realistic model features such as multiple periods, costly default, arbitrary lender cost structures, and extensive and intensive margins of loan demand. Another closely related paper is Liberman et al. (2019), who analyze the effect of information removal in Chile. Applying a similar price-theoretic methodology based on Einav et al. (2010), Liberman et al. estimate the surplus effects of removing default data from credit reports. Our analysis is complementary to that of Liberman et al. We build a microfounded model and illustrate the assumptions under which demand and supply curves can be interpreted in terms of consumer utility in the setting of credit markets. We simplify by assuming away adverse selection in the baseline model. This delivers a sharper qualitative insight: data is always welfare-improving in our setting, since it eliminates deadweight loss triangles for both groups, though the welfare effects of data addition may be small relative to the transfer effects. This helps to rationalize the finding of Liberman et al. (2019) that, under a variety of assumptions, data removal seems to reduce social surplus. In terms of data, our setting has the advantage that we observe prices, whereas Liberman et al. (2019) infer prices from observed default rates.

This paper is also related to papers around the classic idea that third-degree price discrimination has ambiguous effects on social welfare (Schmalensee, 1981; Varian, 1985, 1989). A related recent paper is Chen and Schwartz (2015), who theoretically analyze price discrimination for a monopolist with information about costs, and also find that differential pricing tends to improve pricing more generally than in the classic case.

More broadly, this paper also contributes to an empirical literature on the consequences of bankruptcy, and specifically bankruptcy flag removal. Musto (2004) is the first paper to study the empirical effects of flag removal, studying the impact on credit scores and default. Several studies, including Herkenhoff et al. (2021), Bos et al. (2018), Dobbie et al. (2020), and Gross et al. (2020), study how bankruptcy flag removal affects credit access, employment, entrepreneurship, consumption, and other outcomes. A recent literature on household bankruptcy also links theory to empirics. Gross et al. (2021) study the economic consequences of bankruptcy, Indarte (2021) studies moral hazard and liquidity in bankruptcy, Argyle et al. (2022) study disparities, and Dávila (2020) provides a theoretical framework for optimal bankruptcy exemptions. Relative to this literature, we are the first to quantify the relative magnitudes of the welfare and transfer effects of flag removal, showing that the allocative efficiency losses from flag removal are small relative to the induced transfers, in the context of auto lending.

2 Model

This section formalizes how to quantify the efficiency and redistributive effects of data in credit markets. In our framework, data enables a form of *third-degree price discrimination* based on consumers' underlying costs. In the absence of data, lenders charge the same pooling rate across consumers with different costs. The introduction of data allows lenders to align prices more closely with costs. The improved alignment enhances allocative efficiency by reducing underprovision and overprovision of credit, but also redistributes surplus across lenders and borrowers with different costs.

There are two groups of consumers—high-cost and low-cost—who borrow to finance a one-time expenditure in period $t = 0$, paying back their loans over periods $t = 1 \ldots T$. High-cost consumers (H) default on loans at the higher rate δ_H. Low-cost consumers (L) default on loans at the lower rate δ_L. Consumers borrow from a set of competitive firms who post identical interest rates.

In a world without data, lenders cannot distinguish H and L consumers, and thus are constrained to set equal prices for these two groups. Our goal is to estimate how market outcomes would change in a counterfactual world where data is available, so lenders can distinguish H and L consumers and set separate prices for the two groups. Since most of our results apply symmetrically across groups, rather than introducing a group subscript, we will occasionally state results for the H group, understanding that they apply analogously to the L group.

2.1 Consumers

Consider first a single consumer i, who wishes to finance a one-time expenditure in period $t = 0$; in our setting, for example a large durable asset such as an auto purchase. The borrower's utility is:

$$\underbrace{u_{i,0}\left(c_{i,0}\right)}_{Purchase} + \underbrace{\sum_{t=1}^{T} \beta^t (1-\delta)^t u\left(c_{i,t}\right)}_{Payment} + \underbrace{\sum_{t=1}^{T} (1-\delta)^{t-1} \delta \sum_{\tilde{t}=t}^{T} \beta^{\tilde{t}} u(c_D)}_{Default} \tag{1}$$

In words, the consumer gets concave utility $u_{i,0}\left(c_{i,0}\right)$ from period-0 consumption. The consumer defaults at exogenous rate δ each future period $t = 1 \ldots T$, which is δ_H for high-cost consumers and $\delta_L < \delta_H$ for low-cost consumers. With probability $(1-\delta)^t$, the consumer reaches period t without defaulting, consuming c_t and for utility $\beta^t u\left(c_{i,t}\right)$. With probability $(1-\delta)^{t-1}\delta$, she defaults in period t. Default may be harmful for future-period consumption, for example due to collateral repossession, or because the consumer has more difficulty borrowing in the future; we incorporate this by assuming that, after period-t default, the borrower consumes c_D in each period t through T.

Loans are self-amortizing with fixed payments, so if the principal is L and the interest rate is r, the borrower pays $L\phi(r)$ each future period, where $\phi(r)$ follows the amortization formula:

$$\phi(r) \equiv \frac{r(1+r)^T}{(1+r)^T - 1} \tag{2}$$

The function $\phi(r)$, which is increasing in r, can thus be thought of as the dollar payment in each future period for each dollar borrowed. We will generally use $\phi(r)$ instead of r to measure prices; while the notation is slightly tedious, this facilitates interpretation of our estimates because dollars are a more intuitive measure of welfare than interest-rate points.

The borrower is endowed with wealth w_t in period t, and is completely liquidity-constrained, so she consumes:

$$c_{i,0} = w_{i,0} + L_i,\ c_{i,t} = w_{i,t} - L_i\phi(r) \tag{3}$$

We simplify the general problem by approximating consumer utility in payment period $t = 1 \dots T$ by a linear function:

$$u\left(c_{i,t}\right) \approx u\left(w_{i,t}\right) + u'\left(w_{i,t}\right)\left(c_{i,t} - w_{i,t}\right) = u\left(w_{i,t}\right) - u'\left(w_{i,t}\right) L\phi(r) \tag{4}$$

Combining (1), (3), and (4), we can thus write the borrower's optimization problem as:

$$V(r) = \max_{L_i} u_{i,0}\left(w_{i,0} + L_i\right) + \sum_{t=1}^{T} \beta^t (1-\delta)^t \left[u\left(w_{i,t}\right) - u'\left(w_{i,t}\right) L_i \phi(r)\right] + \sum_{t=1}^{T} (1-\delta)^{t-1} \delta \sum_{\tilde{t}=t}^{T} \beta^{\tilde{t}} u(c_D) \tag{5}$$

For technical convenience we assume there is an interest rate $\bar{r}$ high enough that all consumers stop borrowing.

Let $L_i^*(r)$ denote the solution to (5), as a function of r. Differentiating (5) using the envelope theorem, we have:

$$\frac{dV}{dr} = -\underbrace{\left[\sum_{t=1}^{T} \beta^t (1-\delta)^t u'\left(w_{i,t}\right)\right]}_{Utility\ weight} \underbrace{\left[L_i^*(r)\frac{d\phi}{dr}\right]}_{Payment\ change} \tag{6}$$

The "payment change" term in (6) captures how an increase in r increases payments in each future period $t = 1 \dots T$, and the "utility weight" term captures the sum of the discounted marginal utility in each future period t, multiplied by the probability $(1-\delta)^t$ of reaching period t without defaulting.

The linearity assumption in (4) makes the utility weight term in (6) constant, so changes in $\frac{dV}{dr}$ are driven entirely by loan demand $L_i^*(r)$. This makes the consumer's preferences quasilinear in future-period payments, eliminating income effects. Hicksian and Marshallian demand thus coincide, and compensating variation (CV), equivalent variation (EV), and Marshallian consumer surplus (CS) are always equal. In the general case, CS differs from CV or EV, but the differences tend to be small when price changes and income effects are small.[3]

The classical justification for quasilinear utility is that, if a good is "small" as a fraction of overall consumption, utility concavity for the remaining "composite good" can be approximately ignored (Marshall, 1920; Vives, 1987). In our setting, loans reallocate consumption between a single present period and multiple future periods; our linearity assumption similarly ignores concavity over the "composite" future-period good.

We can define Marshallian *consumer surplus* for a single consumer simply as the integral over loan demand, dividing out the utility weight term in (6):

$$CS_i(r) = \int_r^{\bar{r}} L_i^*(r)\frac{d\phi}{dr} dr \tag{7}$$

An interpretation of $CS_i(r)$ is that, from (6), the consumer is indifferent between borrowing at r, and borrowing nothing and receiving $CS_i(r)$ dollars in each future period conditional on not defaulting. Technically, this statement follows from the fact that CS is equal to compensating variation in quasilinear settings. To see this formally, note that we defined $\bar{r}$ as a rate

[3] See, for example, Willig (1976), Chipman and Moore (1980), Jehle and Reny (2011, pp. 179–183), and Mas-Colell et al. (1995, pp. 80–91).

high enough that the consumer borrows nothing; the utility of borrowing at r relative to not borrowing at all is thus:

$$V(r)-V(\bar{r})=\left[\sum_{t=1}^{T}\beta^{t}(1-\delta)^{t}u'\left(w_{i,t}\right)\right]\left[\int_{r}^{\bar{r}}L_i^*(r)\frac{d\phi}{dr}dr\right] \tag{8}$$

From (4), the RHS is exactly the borrower's total utility gain from receiving a lump-sum of $\int_r^{\bar{r}}L_i^*(r)\frac{d\phi}{dr}dr$ dollars in each period $t=1\dots T$, conditional on not defaulting.

Expression (7) naturally suggests we can define total surplus across consumers in the H group as:

$$CS_H(r)\equiv\sum_{i\in\mathscr{A}_H}CS_i(r)=\int_r^{\bar{r}}\Lambda_H(\hat{r})\frac{d\phi(\hat{r})}{d\hat{r}}d\hat{r} \tag{9}$$

Where $\mathscr{A}_H$ is the finite set of high-cost consumers, and we use $\Lambda_H(r)$ to denote aggregate H group loan demand:

$$\Lambda_H(r)=\sum_{i\in\mathscr{A}_H}L_i(r) \tag{10}$$

The definitions for the L group are analogous.

Expressions (7) and (9) measure consumers' utility in the unintuitive units of dollars in each non-default future period. We thus normalize by multiplying by the expected number of non-default periods, for example:

$$DCS_H(r)=\psi_H\int_r^{\bar{r}}\Lambda_H(\hat{r})\frac{d\phi(\hat{r})}{d\hat{r}}d\hat{r} \tag{11}$$

$$\psi_H=(1-\delta_H)\left(\frac{1-(1-\delta_H)^T}{\delta_H}\right) \tag{12}$$

and analogously for the L group. For example, if DCS_H is \$1 million, this implies that H consumers would approximately require payments summing to \$1 million to be willing to completely stop borrowing.[4] While the notation is slightly complex, (11) simply corresponds to the integral of loan quantity over changes in interest rates, multiplied by a constant adjustment factor ψ_H which depends on default rates.

In our empirical application, we will measure *changes* in dollarized consumer surplus aris-

[4]If we interpret the \$1 million as a lump-sum, the interpretation is approximate because (12) and (11) simply sum dollars across future periods, ignoring discounting.

ing from an increase in interest rates from r to $\tilde{r}$ as:

$$DCS_H(\tilde{r}) - DCS_H(r) = -\psi_H \int_r^{\tilde{r}} \Lambda_H(\hat{r}) \frac{d\phi(\hat{r})}{d\hat{r}} d\hat{r} \tag{13}$$

Expressions (11) and (13) have the simple "welfare trapezoid" interpretation, shown in Figure 1: consumer surplus is the trapezoid corresponding to the integral of demand quantity over prices. We thus fit into a recent literature applying "price-theoretic" concepts of money-metric surplus to credit markets (Einav et al., 2010; DeFusco et al., 2022).

2.2 Producers

Loans are produced by a group of competitive lenders. We assume that lending has constant marginal costs: within each group, there is some interest rate at which lenders make zero profits regardless of the amount they lend. This break-even interest rate is the higher rate $r_{H,fair}$ for the H group, and the lower rate $r_{L,fair}$ for the L group. Our baseline exercise requires that lenders' marginal costs are constant; this assumes away any fixed costs of lending, as well as any forces that would make break-even rates dependent on loan amounts, such as adverse selection or moral hazard. However, we do not need to take a stance on what exactly the structure of lenders' marginal costs is, and how $r_{H,fair}$ and $r_{L,fair}$ are connected to the default rates δ_H, δ_L. We discuss how to relax this assumption and incorporate information asymmetry in Appendix B.6.[5]

[5] As a simple example, suppose that lenders must pay some origination cost c_H, c_L per dollar they lend to the H and L groups: we may have $c_H > c_L$, for example, if screening is more costly for riskier borrowers. Lenders may also have expected default-related losses of ζ_H, ζ_L, per dollar lent. Lenders' expected dollar profits at rate r, per dollar lent, are thus:

$$\underbrace{L\psi_H\phi(r)}_{Expected\ Loan\ Payment} - \underbrace{L}_{Upfront\ Amount\ Lent} - \underbrace{Lc_H}_{Origination\ Costs} - \underbrace{L\zeta_H}_{Default\ Losses}$$

The break-even interest rate $r_{H,fair}$ thus solves:

$$\psi_H\phi\left(r_{H,fair}\right) - 1 - c_H - \zeta_H = 0$$

Hence, at any other rate r, lenders earn the spread

$$\psi_H\phi(r) - 1 - c_H - \zeta_H = \psi_H\left(\phi(r) - \phi(r_{H,fair})\right)$$

on the loan amount. We define $r_{L,fair}$ correspondingly. Similarly, any other source of marginal costs—which scale linearly in the amount lent L—can be thought of incorporated within the definition of $r_{H,fair}$ and $r_{L,fair}$.

Lenders make or lose money at any interest rates other than $r_{H,fair}, r_{L,fair}$. Analogous to (11), we define *producer surplus* for group H at rate r_H as:

$$PS_H(r_H) = \psi_H \Lambda_H(r_H)\left(\phi(r_H) - \phi\left(r_{H,fair}\right)\right) \tag{14}$$

We define producer surplus for the L group analogously. In words, producers' dollar profits at rate r_H are the per-period difference between the actual price $\phi(r_H)$ and the break-even price $\phi\left(r_{H,fair}\right)$ per dollar of loans, multiplied by loan volume $\Lambda(r_H)$ and the expected number of non-default periods ψ_H. We then use (14) to define producers' *zero-profit condition*.

Assumption 1. *Suppose that no data is available, so producers cannot distinguish L consumers from H consumers, and thus set a pooled price r_{pool}. We assume that r_{pool} must make producer surplus sum to zero across the two groups:*

$$\psi_L \Lambda_L\left(r_{pool}\right)\left(\phi\left(r_{pool}\right) - \phi\left(r_{L,fair}\right)\right) = \psi_H \Lambda_H\left(r_{pool}\right)\left(\phi\left(r_{H,fair}\right) - \phi\left(r_{pool}\right)\right) \tag{15}$$

In words, the total dollar gains from lending to L consumers above their break-even rate, on the LHS, must equal the dollar losses from lending to H consumers below their break-even rate, on the RHS. In our empirical application, this has the effect of weighting per-period payments slightly lower for the high-cost group, since they default more often and thus make payments with lower probability; however, the ratio $\frac{\psi_L}{\psi_H}$ is not far from 1 in our application. As with consumer surplus, (15) is not totally innocuous: it assumes producers do not discount, and ignore idiosyncratic risk from individual defaults. We will ultimately find that the welfare effects for the L group are small overall; slightly different specifications of the zero-profit condition are thus unlikely to have large effects on our headline results.

2.3 Intuition

Our framework shows that changing data availability in credit markets can be thought of as a form of *third-degree price discrimination*. Figure 1 shows a stylized illustration of the H and L groups. In the absence of data, lenders cannot distinguish these groups, so they charge r_{pool} for both groups, determined by the zero-profit condition (15). This pricing scheme is socially

Figure 1: Price Discrimination in Credit Markets

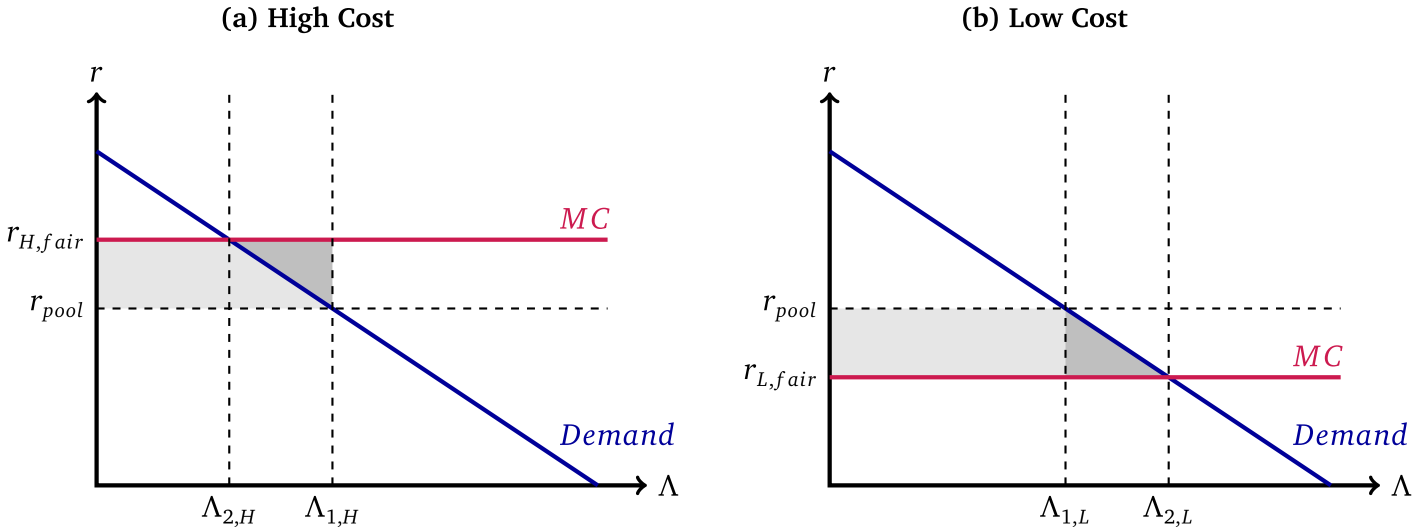

This figure illustrates how third-degree price discrimination affects welfare in credit markets. Suppose there are two groups of prospective borrowers, high-cost (panel a) and low-cost (panel b). The red lines show the cost of serving borrowers in each group, and the blue lines show borrowers' demand curves. Lenders are initially unable to distinguish between these borrowers, so they set the pooled price r_{pool}. Once lenders are able to distinguish between these prospective borrowers, they set $r_{H,fair}$ for the high-cost group (a) and $r_{L,fair}$ for the low-cost group (b). The dark gray shaded triangles illustrate the increase in social welfare for each group after the price change. In panel (a), the light gray shaded area shows the decrease in consumer welfare from the price change. In panel (b), the sum of the light gray shaded rectangle and the dark gray shaded triangle represents the increase in consumer welfare after the price change.

inefficient, generating two deadweight loss triangles (dark gray areas): credit is underprovided to low-cost borrowers, and overprovided to high-cost borrowers.

If lenders can distinguish between these consumer groups, prices increase to $r_{H,fair}$ for the high-cost group, and decrease to $r_{L,fair}$ for the low-cost group. Social surplus increases for both groups, since both DWL triangles are eliminated. The division of surplus also changes, since the pooling cross-subsidy between groups is eliminated. *H* borrowers' surplus decreases by the light gray area in the left panel of Figure 1, and *L* borrowers' surplus increases by the sum of the light gray and dark gray areas on the right panel of Figure 1.

An interesting feature of Figure 1 is that the DWL triangles scale quadratically with price changes, whereas the consumer surplus transfer trapezoids scale approximately linearly. Thus, when r_{pool} is close to the break-even rates $r_{H,fair}, r_{L,fair}$, the ratio of the DWL triangles and the transfer trapezoids decreases towards zero. Figure 1 thus has a nuanced conclusion: data is always welfare-increasing, but when data has a small effect on prices, the cross-subsidy effects of data tend to be large relative to the welfare effects.

Our conclusions differ from the classic literature on third-degree price discrimination in that we assume that lenders are competitive, and data is informative about consumers' costs. In the classic setting, sellers have *market power*, and data is informative only about consumers' *demand* (Schmalensee, 1981; Varian, 1985, 1989). It is well-known that the welfare effects of third-degree price discrimination are ambiguous in the classic setting. Appendix Figure E.1 shows the intuition: monopolists set prices above cost for both groups, and data may increase prices for one group and decrease prices for another, leading to partially offsetting changes in social welfare, so the sign of the net welfare effect is indeterminate.[6]

2.4 Estimation: General Data Policies

Our framework is particularly straightforward to implement empirically if we assume that loan demand in each group is linear in the payment function $\phi(r)$:

$$\Lambda(r) = a - b\phi(r) \tag{16}$$

Consumer, producer, and total surplus then have the following expressions.

Claim 1. *Suppose lenders receive data which allows them to distinguish between the H and L groups, so interest rates for the H group shift from r_{pool} to $r_{H,fair}$. The net increase in lenders' profits on group H is:*

$$\psi_H \Lambda_H\left(r_{pool}\right)\left(\phi\left(r_{H,fair}\right) - \phi\left(r_{pool}\right)\right) \tag{17}$$

The net decrease in H borrower surplus is:

$$\psi_H\left(\phi\left(r_{pool}\right) - \phi\left(r_{H,fair}\right)\right)\left(\frac{\Lambda_H\left(r_{pool}\right) + \Lambda_H\left(r_{H,fair}\right)}{2}\right) \tag{18}$$

[6] A related paper is Chen and Schwartz (2015), who analyze the case where a monopolist has information about costs, showing that welfare increases for a broader class of demand functions than the general case. We simplify further by assuming perfect competition in our baseline, implying that data is always welfare-improving in our setting.

Social welfare increases by:

$$\frac{1}{2}\psi_H\left(\phi\left(r_{H,fair}\right)-\phi\left(r_{pool}\right)\right)\left(\Lambda_H\left(r_{pool}\right)-\Lambda_H\left(r_{H,fair}\right)\right) \tag{19}$$

All expressions for the L group are analogous.

Intuitively, the expressions in Claim 1 are simply the geometric areas of the corresponding rectangular and triangular areas in Figure 1. These are straightforward to estimate, requiring only default rates to estimate the ψ term, and the prices $r_{H,fair}, r_{pool}$ and quantities $\Lambda_H\left(r_{H,fair}\right), \Lambda_H\left(r_{pool}\right)$ at the pooled and fair interest rates.

A convenient feature of Claim 1 is the break-even interest rates $r_{H,fair}, r_{L,fair}$ are directly inferred from the prices that customers face when lenders are able to distinguish the two groups, so we do not need to take a stance on how default rates affect lenders' break-even rates. Essentially, the core inputs into Claim 1 are loan rates and loan quantities with and without data; default rates for each group are needed only to compute the ψ term, which is used only to convert welfare into comparable dollarized units for each group.

2.5 Estimation: Bankruptcy Flag Removal

In our empirical application, we analyze the removal of bankruptcy flags from credit reports. In our framework, consumers who have never declared bankruptcy correspond to the low-cost group, who default at lower rates, while consumers with a prior bankruptcy correspond to the high-cost group, who default at higher rates. Among the previously-bankrupt, we distinguish two subgroups: "bankruptcy-with-flag" consumers, whose credit files still display the bankruptcy flag, and "bankrupt-no-flag" consumers, who have declared bankruptcy far enough in the past that, under current policy, bankruptcy flags are removed from their credit reports. Importantly, lenders cannot distinguish bankrupt-no-flag consumers from never-bankrupt consumers, and therefore apply pooled pricing to these groups in the status quo.

Our goal is to estimate how outcomes would differ in a counterfactual world where bankruptcy flags are never removed, allowing lenders to separate previously bankrupt consumers from others indefinitely. Although we do not observe an explicit change in flag policy, the staggered timing of flag removals allows us to observe a transition in which high-cost (previously-bankrupt) individuals move from being separated (with-flag) to being pooled (no-flag) with low-cost consumers. This variation maps into our model and enables a multi-step estimation strategy, described in Sections 2.5.1 and 2.5.2, to recover the components of Claim 1.

2.5.1 High-Cost/Previously-Bankrupt Borrowers

We can estimate the quantities in Claim 1 for the previously-bankrupt H group by analyzing how prices and loan quantities change around bankruptcy flag removal. Under current policy in the US, bankruptcy flags are removed from credit reports 7 to 10 years after bankruptcy, meaning that H customers, who lenders can distinguish from L customers, transition to becoming indistinguishable from L customers sharply after a cutoff in time.

Under our assumption that the H consumers have equal loan demand and default rates just before and after flag removal, customers should face the higher rate $\phi\left(r_{H,fair}\right)$ just before flag removal, and the lower rate $\phi\left(r_{pool}\right)$ just afterwards. In response, these customers should increase borrowing, from $\Lambda_H\left(r_{H,fair}\right)$ to $\Lambda_H\left(r_{pool}\right)$. We can thus estimate these four quantities, and thus all surplus expressions in Claim 1, simply by observing customers' interest rates and loan quantities just before and after bankruptcy flags are removed.

2.5.2 Low-Cost/Never-Bankrupt Borrowers

To identify counterfactual price changes for the never-bankrupt L group, we use lenders' zero-profit condition, (15). Rearranging:

$$\phi\left(r_{L,fair}\right)-\phi\left(r_{pool}\right)=-\frac{\psi_H\Lambda_H\left(r_{pool}\right)}{\psi_L\Lambda_L\left(r_{pool}\right)}\left(\phi\left(r_{H,fair}\right)-\phi\left(r_{pool}\right)\right) \tag{20}$$

The LHS of (20) is the price decrease that L consumers would face, if lenders could distinguish them permanently from H individuals. The RHS says that this can be estimated by multiplying the price change for H consumers by the ratio of loan volumes between the two groups, and a ψ-ratio term accounting for default rate differences between the groups. In our empirical application, the L group is much larger than the H group, so the ratio $\frac{\Lambda_H}{\Lambda_L}$ is small, and the counterfactual price change for the L group will be fairly small.

We then need to calculate the induced change in loan quantities for the never-bankrupt

L group. This could be calculated by multiplying the price change in (20) by an estimate of the demand elasticity for the L group.We assume that demand elasticities are equal across groups, although in principle these could be estimated separately given another source of price variation.

Assumption 2. *At the price r_{pool}, demand elasticities are equal across the L and H groups; that is, the slopes of demand b_L and b_H satisfy:*

$$\frac{b_L}{\Lambda_L\left(r_{pool}\right)} = \frac{b_H}{\Lambda_H\left(r_{pool}\right)} \tag{21}$$

Given this assumption, we can calculate the counterfactual quantity change facing the L group based on the price change; we can then calculate changes in consumer surplus and social welfare for this group using the expressions in Claim 1. While Assumption 2 is somewhat ad-hoc, our main quantitative results are not very sensitive to it. We show in Appendix B.1 that the social welfare effects of data policy are driven predominantly by the smaller H group, so errors in estimating outcomes for the larger L group actually have relatively little effect on our conclusions.

2.6 Discussion of Model Assumptions

Choice of Welfare Weights. Our approach builds on the tradition of Marshallian partial-equilibrium welfare analysis (Marshall, 1920). We take no explicit stance on welfare weights across groups: the implicit stance is that welfare is measured in dollars, and classic results imply that any Pareto-optimal outcome must maximize the sum of money-metric utility across agents. This approach has the benefit of being easily interpretable. We will find that bankruptcy flag removal burns around \$0.03 of social surplus per dollar redistributed across groups. The normative judgment of whether the policy's benefits outweigh its costs is left to policymakers.[7]

Demand Margins. Our model and empirical analysis ignores loan rejections. In practice, some prospective borrowers may be rejected by all lenders, and may thus be unable to borrow at any price. We cannot capture welfare for these hypothetical borrowers, since we cannot

[7]An alternative approach is to directly specify welfare weights on each agent, and then solve for optimal policy. In such approaches, welfare weights and optimal policy are entangled: there is little we can say about how efficient a policy is, independent of the particular welfare weights chosen.

estimate their willingness-to-pay. We do not expect this to be a substantial concern in the auto loan setting we analyze, since many lenders specialize in serving riskier borrowers, so we believe the vast majority of borrowers can get a loan offer at some price in our setting.

Interest Rate Menus. We assumed that borrowers can borrow any amount at a fixed rate r. In reality, borrowers may face a menu of interest rate–LTV (Loan-to-Value) combinations, typically subject to a maximum loan size. Data policy may affect both the structure of the menu and the loan size cap. In Appendix B.2, we show that our results still hold if data policy primarily influences outcomes through shifting interest rates. However, our analysis does not fully capture welfare effects if data policy primarily influences the size of binding loan size caps.

Default. Appendix B.3 explores extending the model to incorporate richer default behavior and post-default outcomes. Our baseline results hold even if borrowers default optimally, and regardless of how default affects consumption, as long as the interest rate does not influence post-default outcomes. However, if loan rates affect consumption *after* default, an additional term appears in the derivative of $V(r)$, meaning our measures will tend to understate the effects of interest rates on welfare. We are unable to quantify these effects in our setting, so we ignore this effect in our analysis.

Default Rate Heterogeneity. Appendix B.4 analyzes the case where default rates may be heterogeneous within groups. We show that Claim 1 remains approximately accurate as long as default rates are not too different within data groups, and demand elasticities are not too correlated with customers' default rates within groups.

Market Power. Appendix B.5 extends the framework to imperfect competition. In general, if data availability changes the markups that lenders can charge, outcomes are known to be complex (He et al., 2020; Huang, 2022). Nonetheless, we show that with additional data markups can be accommodated. Surprisingly, under a certain set of parameter restrictions—when the demand elasticities in the two borrower groups are equal, and when markups pre- and post-data availability are the same—the existence of markups does not, in fact, affect the welfare gains from data availability. While this set of assumptions may not always be satisfied exactly in reality, this finding suggests that the existence of markups per se does not dramatically affect our results.

Adverse Selection. In the main text, we assume that there is no selection—that is, that costs depend on borrowers' types but are not correlated with borrowers' willingness to pay. In Appendix B.6, we discuss how adverse selection could be accomdated in our framework.

3 Empirical Application

We now apply our framework to the specific setting we study empirically: bankruptcy flag removal. In the US, under the Fair Credit Reporting Act (FCRA), bankruptcy flags must be removed from credit records after ten years. There are two main types of consumer bankruptcy in the US: Chapter 7 (liquidation) and Chapter 13 (reorganization). Chapter 7 bankruptcy flags are typically removed 117 to 118 months after filing, while Chapter 13 bankruptcy flags are typically removed seven years after filing. Hence, previously-bankrupt individuals whose flags have been removed are pooled with never-bankrupt individuals. Flag removal can thus be thought of as a form of data policy: in a counterfactual world in which flags were never removed, social welfare would be higher, previously-bankrupt borrowers would be worse off, and never-bankrupt borrowers would be better off. This application illustrates the tradeoff inherent in our model. On the one hand, counterfactually keeping bankruptcy flags on consumer credit records increases efficiency, as lenders can use this information in pricing loans. On the other hand, keeping flags would make previously-bankrupt borrowers worse off, as they would face higher interest rates.

3.1 Data

To implement our analysis, we use the Booth TransUnion (TU) Consumer Credit Panel.[8] The data is an anonymized 10% sample of all TU consumer credit records from 2009 to 2020. We restrict the sample to the 2009-2018 period to allow at least two years for delinquency realizations after account openings. The sample is a panel. Individuals in the initial sample have their information updated monthly, and each month a new 10% sample of first time borrowers is added. At a monthly frequency, the data contain basic information about borrowers and

[8] These data, along with similar credit panel data, are described in more detail in Keys et al. (2020) and Yannelis and Zhang (2021).

loans, including the original balance, the current balance, the Vantage (credit) score, scheduled payments, the maturity of the loan, geography and importantly bankruptcy flags.

Our main outcomes of interest related to the welfare framework are credit scores, interest rates, new loan balances, and charge-offs. We do not directly observe interest rates, but we back them out from scheduled payments using the amortization formula.[9] To avoid selection concerns, we predict interest rates for all individual-month observations. We predict interest rates using a third-order polynomial of current and up to nine months lagged scores, and time and cohort dummies.[10] We measure balances as the sum of the balances of the new auto loan accounts that an individual opens in a given month, and as zero when the individual does not open an account in a given month. Charge-offs are measured as whether a loan becomes charged-off within two years of the account opening. We collapse the data to the borrower-month level, and restrict the sample to individuals who ever had a Chapter 7 or Chapter 13 bankruptcy flag and observations within six months around flag removal.

Table 1: Summary Statistics

This table displays basic summary statistics for the main analysis variables: the mean, median, and standard deviation. The first three columns show the statistics for the full sample, the next triplet for individuals pre-flag removal, and the final three columns show summary statistics post-flag removal. Source: TransUnion.

	Full			Pre			Post		
	Mean	Median	SD	Mean	Median	SD	Mean	Median	SD
Credit Score	630.67	638.00	83.76	620.63	631.00	76.68	639.28	643.00	88.47
Interest Rate	8.78	7.66	4.31	9.02	7.96	4.30	8.57	7.43	4.30
Quantity Opened Cond. on Opening	20783	19484	10773	20332	19031	10495	21143	19842	10976
Quantity Opened	326.38	0.00	2915.35	306.50	0.00	2792.55	343.41	0.00	3016.52
If Charged-off	0.02	0.00	0.15	0.02	0.00	0.15	0.02	0.00	0.15

Table 1 presents summary statistics for our main analysis variables. There are 865,499 individuals in our final sample over 13 months totaling 11,251,487 individual-month observations. In 1.57 percent of the individual-month observations at least one auto-debt account is opened, and the average opening balance conditional on opening is \$20,783. Hence, the average monthly account opening is \$326.38. The average rate is 8.78%. The average credit score in our sample is 631. Table 1 also shows summary statistics pre and post flag removal.

[9]Specifically, we let the monthly payment $A = \frac{P \times i}{1-(1+i)^{-n}}$, where P is the principal, n is the maturity, and i is the interest rate. We use a root-solving algorithm to solve for i. Note that we use scheduled, and not actual payments to construct interest rates.

[10]Appendix C shows the effect of flag removal on observed interest rates with fundamentally similar conclusions.

Following flag removal, mean credit scores increase, interest rates decrease and borrowing amounts increase. Of course, these changes may be due to both bankruptcy flag removals, and secular time trends, which motivates our empirical strategy.

Now, following the steps we have set up in Section 2.5, we show how we estimate counterfactual price and quantity changes for previously-bankrupt borrowers in Subsection 3.2 and for never-bankrupt borrowers in Subsection 3.3. We summarize our results in Subsection 3.4.

3.2 Previously-Bankrupt Borrowers

3.2.1 Design for Flag Removal

We argued in Subsection 2.5 that, in a counterfactual world where bankruptcy flags were never removed, previously-bankrupt borrowers whose flags are removed would see an increase in interest rates and decrease in loan quantities, reflecting their visible higher risk to lenders. Moreover, somewhat surprisingly, the magnitudes of the counterfactual rate increase and quantity decrease can be estimated simply by comparing rates and quantities between individuals observed in the present world, just before and after bankruptcy flags are removed from credit reports. To estimate these quantities empirically, we use variants of the following specification:

$$y_{it} = \gamma_c + \gamma_t + \delta^y \mathbb{1}[FlagRemoved] + \beta X_{it} + \varepsilon_{it} \tag{22}$$

where y_{it} are outcomes for individual i in month t, γ_c are cohort fixed effects, γ_t are calendar period fixed effects, X_{it} are individual controls, and ε_{it} is an error term which we assume is uncorrelated with $\mathbb{1}[FlagRemoved]$, conditional on observables. We cluster standard errors at the level of the month in which the bankruptcy flag is removed. $\mathbb{1}[FlagRemoved]$ is an indicator of whether bankruptcy flags have been removed. The main coefficients of interest are the δ^y terms, which identify the difference in the outcome y_{it} following the removal of information; as we argued in Subsection 2.5.1, under the assumptions of our framework, δ^y is directly informative about the difference between separating and pooling prices $\left(\phi(r_{H,fair}) - \phi(r_{pool})\right)$ and quantities $\left(\Lambda_H(r_{H,fair}) - \Lambda_H(r_{pool})\right)$ for the H group.

We explore three primary outcomes: Vantage credit scores, interest rates, and loan amounts. We observe credit scores and loan quantities in all time periods. We only observe interest rates conditional on contracting. Hence, we additionally assume that these reflect the lowest available offer rates for a contract.[11]

The coefficient δ^y captures the difference in the outcome y_{it} under the assumption that nothing changes other than the removal of information on previous bankruptcies. More precisely, we assume that flag removal is orthogonal to the error term ε_{it}. A potential concern is that, due to removal of individual flags at different points in time, the estimates for δ^y may be biased by individuals leaving the control group and the heterogeneity of the treatment effect over time (Goodman-Bacon, 2021; Barrios, 2021). To address this concern, we first provide sharp graphical evidence of breaks when flags are removed, and then implement modern dynamic difference-in-difference estimators in Appendix D.

To provide graphical evidence that the observed effects are indeed driven by flag removals, we further estimate an event-study regression to evaluate the identifying assumption using the following variant of equation (22):

$$y_{it} = \gamma_c + \gamma_{ts} + \sum_{t=-6}^{6} \delta_t \{e_{it} = t\} + \beta X_{it} + \varepsilon_{it} \tag{23}$$

where γ_c are cohort-month and γ_{ts} are year-month by score bucket fixed effects.[12]

We plot the coefficients δ_t, along with a 95% confidence interval. The coefficients capture the difference in an outcome in each month before and after flag removal. We exclude the relative time dummy for period -1 as well as relative time dummy for period -6 due to collinearity arising from the age-period-cohort problem common in similar specifications.[13]

[11] We focus on the case where the bankruptcy decision is already determined, and focus on the allocative effects. Appendix B.7 discusses the case where flag removal impacts the filing decision. We construct a simple model in which borrowers strategically decide whether to file for bankruptcy, and bankruptcy has social deadweight loss costs. In the model, borrowers default more than is socially optimal; removing the flag causes borrowers to default more, increasing the magnitude of this distortion. In a back-of-envelope calculation, we find that these costs can be large relative to the allocative welfare costs that we consider here. A policymaker interested in evaluating the full costs of bankruptcy flag removal should thus estimate these costs and take them into account. However, there is no evidence that we are aware of showing that flag removal impacts the decision to file. In fact, only a small fraction of bankrupt individuals seem to be aware of flag removal policies: just 9.2% of Chapter 7 filers correctly guess the number of years remaining for their flag (Gross et al., 2020, Table 6). In this paper, we thus focus only on quantifying the costs that bankruptcy flag removal imposes on allocative welfare.

[12] We measure score buckets by sorting individuals into one of 20 score buckets in the month before flag removal, and hold the sorting constant throughout the 13 months observed.

[13] To ensure that our results are not dependent on this design choice, we implement a stacked difference-in-difference comparing individuals with flag removal to individuals with bankruptcy flag removal 12 to 17 months later in Appendix D.

3.2.2 Effects of Flag Removal

We begin with graphical evidence showing point estimates of equation (23). The figures show specifications including cohort and score bucket by year-month fixed effects. Figure 2 shows estimates of the coefficients δ_t, where the outcomes are credit scores, interest rates, and loan quantities. The top panel shows the Vantage score, the middle panel shows interest rates at origination, and the bottom panel shows loan quantities. Consistent with prior work, we see a very sharp increase—almost 20 credit score points—for previously-bankrupt individuals following the flag removal. This translates into a reduction in borrowing costs. Interest rates for previously-bankrupt individuals show a clear drop following the flag removal, when these previously-bankrupt borrowers become indistinguishable from never-bankrupt individuals. Consistent with the decline in interest rates, we see a sharp rise in loan volumes. There is an approximately $20 increase in auto loan openings.

We next quantify the visual results in a regression framework. Table 2 presents variants of equation (22). Column (1) includes a linear time trend. Column (2) adds time period fixed effects. Column (3) adds cohort fixed effects, based on the bankruptcy filing date. Column (4) includes both cohort and year-month fixed effects. And, finally, column (5) includes cohort and score bucket by time period fixed effects. We measure score buckets by sorting individuals into one of 20 such score buckets in the month before flag removal and hold the sorting constant throughout the 13 months observed. In the top panel, the outcome is the Vantage credit score; in the middle panel, it is interest rates; and in the bottom panel, it is loan volumes.

The results are broadly in line with the graphical evidence. The top panel of Table 2 indicates that bankruptcy flag removals lead to a 17.1- to 17.2-point increase in credit scores, or an approximate 2.76% increase in credit scores. The middle panel indicates that this is associated with a 21- to 23-basis-point decrease in interest rates, or a 2.4%-2.5% decrease in interest rates. The bottom panel shows that average new auto loan balances increase by $17.7 to $18.4, or 5.8%-6%. In the majority of our specifications and in our preferred specification, the effect is significant at the 1% level. In columns (1) and (3), the effect size for loan quantities is significant at the 5% level.

Figure 2: Credit Scores, Interest Rates, and Loan Balances

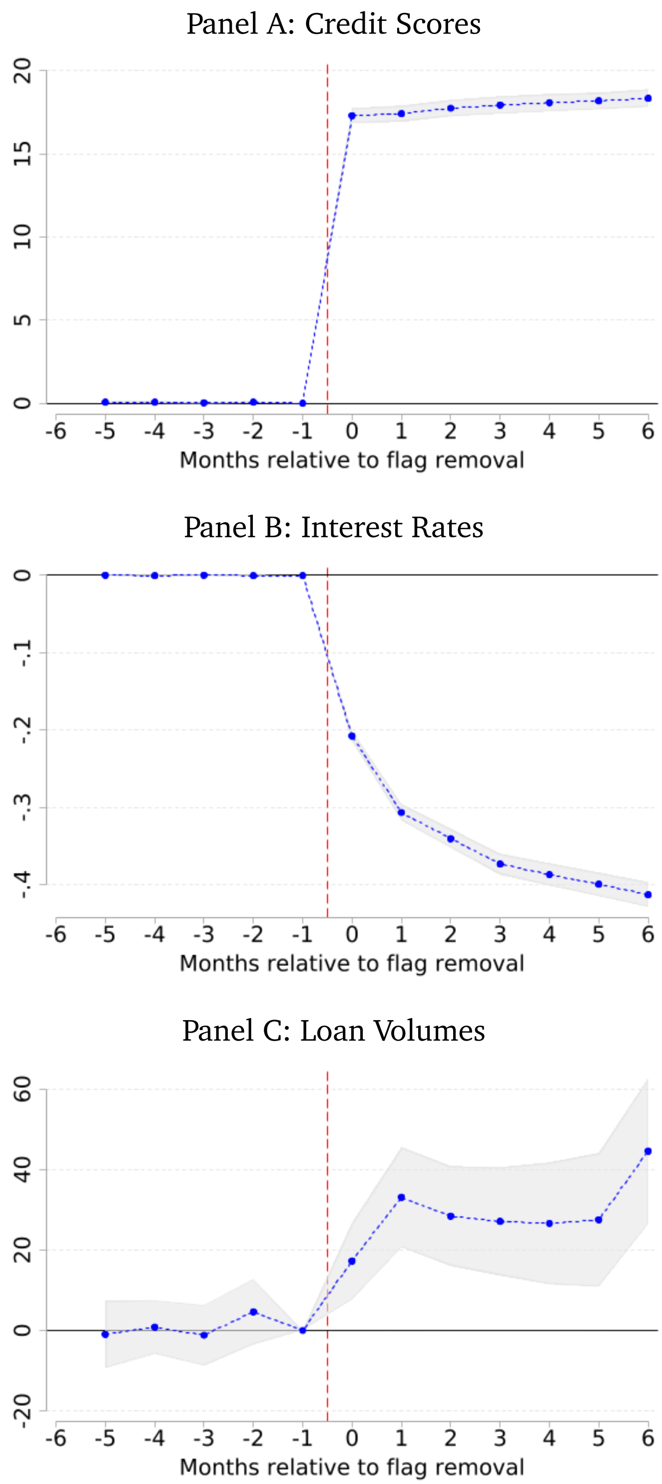

This figure shows estimates of the coefficients δ_t from the following specification $y_{it} = \gamma_c + \gamma_{ts} + \sum_{t=-6}^{6} \delta_t \{e_{it} = t\} + \beta X_{it} + \varepsilon_{it}$, along with a 95% confidence interval. In Panel A, the outcome y_{it} is credit scores; in Panel B, it is interest rates; and in Panel C it is loan volumes. γ_c are cohort fixed effects, and γ_{ts} are time period by score bucket fixed effects. Standard errors are clustered at the cohort level. Source: TransUnion.

Table 2: Credit Scores, Interest Rates, and Loan Volumes

This table shows estimates of the coefficients δ^y from the following specification $y_{it} = \gamma_c + \gamma_t + \delta^y \mathbb{1}[FlagRemoved] + \beta X_{it} + \varepsilon_{it}$. In the top panel, the outcome y_{it} is the Vantage Score, in the middle panel the outcome is interest rates, while in the bottom panel it is loan volumes. Interest rates are predicted with a polynomial of current and past credit scores, period, and cohort fixed effects. γ_c are cohort fixed effects, and γ_t are time period fixed effects. Standard errors are clustered at the cohort level. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Source: TransUnion.

	(1)	(2)	(3)	(4)	(5)
Panel A: Credit Scores					
$\mathbb{1}$[FlagRemoved]	17.216***	17.147***	17.216***	17.125***	17.118***
	(0.298)	(0.258)	(0.298)	(0.239)	(0.237)
Constant	620.075***	621.165***	620.075***	621.448***	621.452***
	(0.488)	(0.497)	(0.188)	(0.129)	(0.128)
Observations	11,251,487	11,251,487	11,251,487	11,251,487	11,251,487
Adjusted R^2	0.012	0.017	0.017	0.018	0.813
Panel B: Interest Rates					
$\mathbb{1}$[FlagRemoved]	-0.218***	-0.226***	-0.218***	-0.226***	-0.226***
	(0.018)	(0.006)	(0.018)	(0.004)	(0.004)
Constant	9.109***	8.979***	9.109***	8.898***	8.898***
	(0.096)	(0.024)	(0.018)	(0.002)	(0.002)
Observations	11,251,487	11,251,487	11,251,487	11,251,487	11,251,487
Adjusted R^2	0.003	0.044	0.044	0.046	0.878
Panel C: Loan Volumes					
$\mathbb{1}$[FlagRemoved]	18.344**	17.926***	18.344**	17.979***	17.793***
	(6.270)	(4.049)	(6.270)	(4.017)	(4.010)
Constant	299.363***	310.897***	299.363***	316.696***	316.796***
	(7.473)	(2.901)	(3.098)	(2.163)	(2.159)
Observations	11,251,487	11,251,487	11,251,487	11,251,487	11,251,487
Adjusted R^2	0.000	0.001	0.001	0.001	0.002
Linear Time Trend	Yes	Yes	Yes	No	No
Year-month FE	No	Yes	No	Yes	No
Cohort FE	No	No	Yes	Yes	Yes
Year-month by Score Bucket FE	No	No	No	No	Yes
Clustered SE	Cohort	Cohort	Cohort	Cohort	Cohort

3.2.3 Welfare Estimates

Using our regression estimates, we can now calculate the quantities:

$$\phi\left(r_{H,fair}\right), \phi\left(r_{pool}\right), \Lambda_H\left(r_{H,fair}\right), \Lambda_H\left(r_{pool}\right)$$

allowing us to calculate the consumer surplus decrease, and welfare increase, that would be generated for the H group in a counterfactual world where bankruptcy flags are never removed. We set $r_{H,fair}$ equal to 9.02%, the average interest rate for borrowers before their flags are removed in the data. After flag removal, previously-bankrupt individuals are pooled with never-bankrupt individuals; the pooled rate r_{pool} is equal to the pre-flag removal rate plus our regression coefficient, $9.02\% + \delta^{InterestRate}$. To simplify surplus accounting and to stay in accordance with our theory, we express these in terms of repayment fractions, ϕ. $\phi\left(r_{H,fair}\right)$ is 2.077% of the original loan balance per month, and the pooling payment, $\phi\left(r_{pool}\right)$, is 2.066%. Note, also, that we show both $\phi\left(r_{H,fair}\right)$ and $\phi\left(r_{pool}\right)$ for expositional purposes here and in Figure 3; however, from Claim 1, only the difference $\phi\left(r_{pool}\right) - \phi\left(r_{H,fair}\right)$ matters for welfare and surplus, meaning that our results are functionally driven mainly by the regression estimate $\delta^{InterestRate}$, with the levels of $r_{H,fair}$ and r_{pool} playing a small role.

We observe total loan quantity prior to flag removal: for previously-bankrupt individuals before flag removal, roughly 18.1% of individuals open a loan each year, and the average loan size conditional on opening a loan is \$20,332. We assume that previously-bankrupt individuals with or without flag have similar demand at the same price. Hence, previously-bankrupt consumers with no bankruptcy flag, when separated, have an annual average loan volume of $\Lambda_H\left(r_{H,fair}\right) = \$3,678$ per individual. From our regression estimates, flag removal causes annual loan quantity to increase to:

$$\Lambda_H(r_{pool}) = \Lambda_H(r_{H,fair}) + 12\delta^{LoanVolume}$$

where $\delta^{LoanVolume}$ is the coefficient on loan quantities from our difference-in-differences specification; we multiply $\delta^{LoanVolume}$ by 12 because we run the loan volume regression at a monthly level. We thus find that $\Lambda_H(r_{pool})$ is equal to $\$3,891.5$; that is, loan quantity increases by roughly \$213.5 per year after flag removal. Note, also, that from Claim 1, the total change in consumer surplus from flag removal depends on the level of $\Lambda_H(r_{pool})$, but the social surplus change depends only on the difference $\Lambda_H(r_{pool}) - \Lambda_H(r_{H,fair})$.

We then plug these quantities into Claim 1 to calculate counterfactual outcomes if bankruptcy flags are never removed. We depict these calculations graphically in the left panel of Figure 3. We report borrower surplus and social welfare in two ways: in dollars per non-default month, and in total expected dollars transferred to a borrower over the lifetime of a loan (taking into account borrowers' default probability). To calculate surplus in terms of expected dollars over the life of a loan as in Claim 1, we multiply per-period surplus by the default adjustment term ψ_H, as defined in (12). This only requires observing the default rate of previously-bankrupt borrowers. Plugging in 0.15% for the default rate δ_H, we get ψ_H equal to 57.29.[14]

We find that in a counterfactual world without flag removal, H group surplus decreases by \$0.41 per eligible borrower each month. Multiplying by ψ_H, surplus decreases by roughly \$23.75 in expectation over the lifetime of a five-year loan. When, in the counterfactual world, we restrict our attention to the 18.1% of borrowers that get a new auto loan each year, each borrower essentially loses an expected transfer of \$131.22 over the lifetime of a five-year loan. We find that, due to the elimination of credit overprovision to the H group in the absence of flag removals, there is a social welfare gain of \$0.012 per H borrower per month; multiplying by ψ_H, this is \$0.67 in expectation over a five-year loan. Restricting attention to the 18.1% of borrowers that get a new auto loan each year, there is a welfare gain of \$3.70 per borrower over a five-year loan due to the elimination of inefficiently high credit provision in the counterfactual.

Next, we aggregate these welfare estimates across borrowers. Approximately 800,000 individuals have their bankruptcy flags removed each year. Multiplying by average loan size, aggregate loan volume is approximately \$2.94 billion per year in a world without flag removal, and approximately \$3.11 billion per year in a world with flag removal. In the counterfactual world without flag removal, the borrower surplus decrease for the H group, illustrated as the

[14] For all auto loans ever opened by individuals who are ever bankrupt, we compute the ratio of loans that are charged-off within two years of loan opening as 3.6%. This implies a monthly default probability of 0.15% $= \delta_m = 1 - (1 - Pr(\text{charged off in first two years}))^{\frac{1}{24}}$. The monthly default rate corresponds to approximately 91% of loans reaching five-year maturity without being charged-off.

light gray trapezoid in the left panel of Figure 3, is \$0.33 million per month, or \$19 million in expectation over a five-year loan term. In other words, for loans originated in one year, previously-bankrupt borrowers lose \$19 million in expectation over a five-year loan term due to the higher monthly payments. For loans originated in a given year, the gain in social welfare due to the elimination of credit overprovision in the absence of flag removal, illustrated as the dark gray triangle in the left panel of Figure 3, is \$9,356 per month, or \$535,975 in expectation over a five-year loan term.

3.3 Never-Bankrupt Borrowers

Price Effects. We follow Subsection 2.5.2, estimating the counterfactual price change for never-bankrupt individuals in a world without bankruptcy flag removal. We use the rearranged zero-profit condition for lenders, equation (20). To apply (20), we need the relative loan quantities for previously-bankrupt and never-bankrupt borrowers, as well as the default rate adjustment terms, ψ. From the data, we calculate that approximately 10.6% of individuals go bankrupt at some point; thus, the ratio of loan quantity for never-bankrupt individuals to loan quantity for previously-bankrupt individuals with no flag is approximately:

$$\frac{\Lambda_H\left(r_{pool}\right)}{\Lambda_L\left(r_{pool}\right)} = \frac{0.106}{1-0.106} \approx 0.1186. \tag{24}$$

We thus can calculate total loan quantity, under the pooled rate r_{pool} for low-cost never-bankrupt individuals, simply as:

$$\Lambda_L\left(r_{pool}\right) = \frac{\Lambda_H\left(r_{pool}\right)}{0.1186} = \$26.26 \text{ billion}$$

The default rate adjustment terms, ψ_H and ψ_L, are simply functions of default rates. Monthly default rates are 0.15% among the previously-bankrupt, and 0.09% among the never-bankrupt. Thus, we have:

$$\psi_H = 57.29,\ \psi_L = 58.34$$

Note that, while we keep the ψ_H, ψ_L multipliers for internal consistency, ψ_H and ψ_L are very similar in practice, implying that the ratio $\frac{\psi_L}{\psi_H}$ is very close to 1; thus, the ψ values essentially have no quantitative role in our analysis besides scaling H and L group welfare by a constant.

We can then plug these estimates into (20), derived from the zero-profit condition of Assumption 1, to determine the hypothetical price change for never-bankrupt individuals. We find that, in a counterfactual world where bankruptcy flags are never removed, a never-bankrupt individual has to repay 0.001% of the principal less per month, reflecting the fact that in the counterfactual lenders can separate the safer never-bankrupt individuals from the riskier previously-bankrupt individuals. This price change is relatively small since the never-bankrupt group is approximately ten times larger than the previously-bankrupt group.

Quantity Effects. We then calculate the counterfactual quantity increase for never-bankrupt individuals simply by using Assumption 2, which states that demand elasticities are identical across groups. This implies that the never-bankrupt demand slope b_L is:

$$b_L = b_H * 1/(\Lambda_H/\Lambda_L)$$

where b_H, the demand slope for previously-bankrupt borrowers, is just the ratio of aggregate quantity changes and repayment rate changes implied by our regressions in Subsection 3.2 ($b_H = \frac{213.52*800{,}000}{0.011\%}$). Using the ratio of loan quantities from (24), we calculate a total counterfactual quantity increase, for never-bankrupt borrowers, of approximately \$170 million. The total loan quantity changes for never-bankrupt and previously-bankrupt borrowers are similar in magnitude. Essentially, this is because the never-bankrupt quantity change results from multiplying a smaller price change by a larger total borrower base, and these two effects offset under the assumption of equal elasticities.

Welfare Estimates. With counterfactual price and quantity changes for the never-bankrupt in hand, we can compute the changes in total welfare and borrower surplus generated by the never-bankrupt group under our counterfactual where flags are never removed. In the counterfactual, for each never-bankrupt individual over the five-year loan term, borrower surplus increases by \$2.9, but the welfare increase due to more efficient credit allocation is only \$0.01. Aggregating across never-bankrupt borrowers, we find that aggregate loan volume is approximately \$26.42 billion per year in a counterfactual world without flag removal, compared to approximately \$26.26 billion per year in the status quo where flag removal exists. In the counterfactual without flag removal, the increase in borrower surplus for the never-bankrupt, illustrated as the light gray rectangle and dark gray triangle in the right panel of Figure 3, is $335,984 per month, or $19.6 million in expectation over a five-year loan term. In other words, for loans originated in a year, never-bankrupt borrowers gain $19.6 million in expectation over a five-year loan term due to the lower monthly payments. This gain is slightly larger than the borrower surplus decrease for previously-bankrupt borrowers in the counterfactual, and the difference between the two constitutes the social welfare gain from keeping bankruptcy flags. For loans originated in a given year, the gain in social welfare due to the elimination of credit underprovision for never-bankrupt borrowers in the presence of data, illustrated as the dark gray triangle in the right panel of Figure 3, is $1,070 per month, or $62,410 in expectation over a five-year loan term. This gain is small relative to the welfare effect generated by the previously-bankrupt borrowers. As we discuss in Appendix B.1, this is because the never-bankrupt borrowers far outnumber the previously-bankrupt borrowers, implying that deadweight loss triangles are smaller for this group.

Note that the assumption that demand elasticities are equal across groups is not quantitatively important for our results, because the never-bankrupt L group contributes so little to total welfare. For the welfare effect in the L group to be quantitatively important, the group's demand elasticity would have to be unrealistically large; we show, at the end of Appendix B.1, that the L group would need a demand elasticity almost nine times larger than the H group's in order for the L group's welfare effect to equal that of the H group.

3.4 Summary: Counterfactual Welfare and Transfer Effects Across Groups

Combining our estimates, we can summarize the welfare and transfer effects of our counterfactual in which bankruptcy flags are never removed. We visualize these effects in Figure 3. The horizontal red lines show the break-even payments lending to previously-bankrupt $\phi(r_{H,fair})$ and never-bankrupt $\phi(r_{L,fair})$ types in the counterfactual world without flag removals. The dashed horizontal line shows the pooling payment in the presence of flag removal.

In the counterfactual, total welfare increases by $0.75 per previously-bankrupt individual with a flag removed. Table 3 also summarizes the relevant quantities expressed per individual with flag removal. Of the $0.75 increase, eliminating credit overprovision to previously-bankrupt contributes \$0.67 per individual, and eliminating credit underprovision to never-bankrupt borrowers contributes \$0.08 per previously-bankrupt individual.[15] Multiplying \$0.75 by the total number of borrowers affected by flag removals, we have a total welfare change of approximately \$598,385 for loans originated in the U.S. in a given year. This estimate reflects the elimination of two inefficiencies: one arising from credit overprovision to the previously-bankrupt, shown as the dark gray triangle in the left panel of Figure 3, and one arising from credit underprovision to the never-bankrupt (a much smaller contribution), shown as the dark gray triangle in the right panel of Figure 3.

The figure also depicts the surplus redistributed relative to the counterfactual. Per previously-bankrupt individual with no flag, borrower surplus decreases by \$23.75 over the loan term in the counterfactual. In aggregate, in the counterfactual, previously-bankrupt lose \$19 million in borrower surplus per year of loan originations, as we calculated in Subsection 3.2.3. This corresponds to the light gray trapezoid in the left panel of Figure 3. Per never-bankrupt individual, borrower surplus increases by \$2.9 over the loan term in the counterfactual. In aggregate, in the counterfactual, never-bankrupt gain of \$19.6 million in borrower surplus per year of loan originations; as we calculated in Subsection 3.3. This corresponds to the light gray rectangle and dark gray triangle in the right panel of Figure 3.

In our framework, bankruptcy flag removal, relative to the counterfactual where flags are never removed, can be thought of as a policy which redistributes surplus from the never- to the previously-bankrupt, at the cost of some social deadweight loss. We can now quantify the degree to which flag removal is an imperfect transfer tool: our estimates imply that for loan originations in a given year, flag removal transfers roughly \$19 million from never-bankrupt to previously-bankrupt borrowers, at the cost of destroying \$598,385 in social surplus. Thus, for each dollar of surplus transferred to previously-bankrupt borrowers, 3.15 cents of social surplus are destroyed. This substantiates the intuition, from Subsection 2.3, that when bankruptcy flags are removed from credit reports, the social deadweight losses from credit misallocation are small.[16]

[15] Note that this \$0.08 figure is per previously-bankrupt individual, while the \$0.01 reported in Section 3.3 is per never-bankrupt individual, and there are roughly 8.5 times as many never-bankrupt as previously-bankrupt individuals with no flag.

[16] A natural question is how our conclusions are affected by the informativeness of the removed data. In Appendix B.8, we show that the efficiency ratio is worse when the signal is more informative: that is, if data is more

Figure 3: Empirical Welfare Estimates

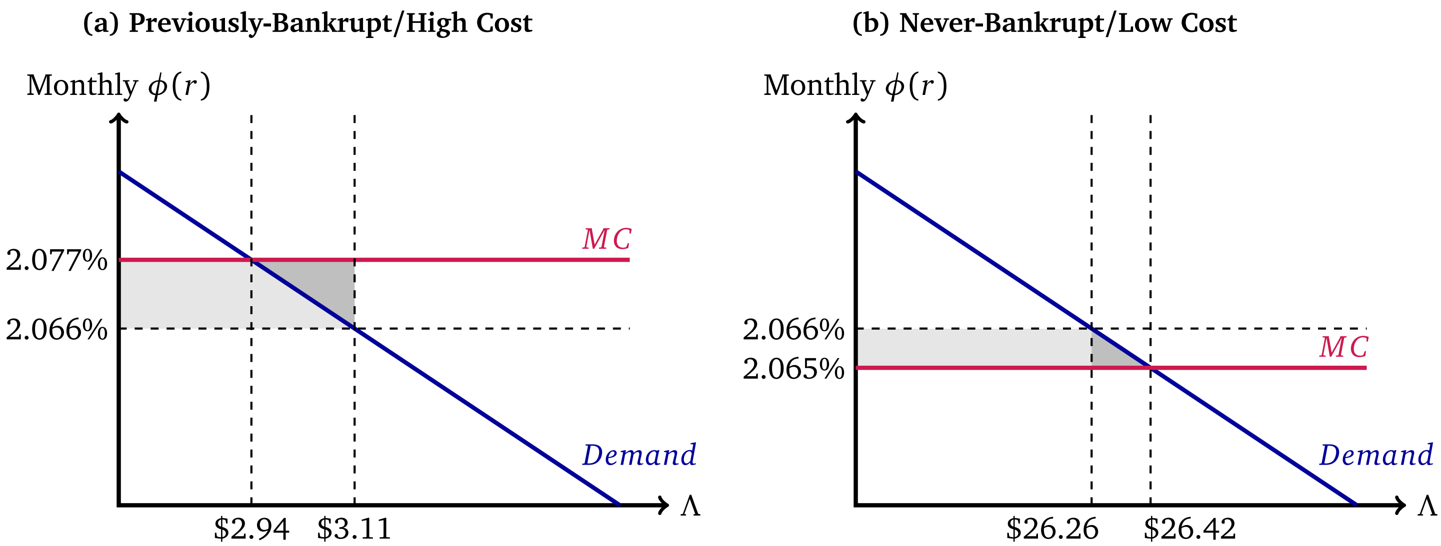

This figure illustrates the changes in borrower surplus and efficiency if bankruptcy flags are never removed. In each panel, the y-axis, $\phi(r)$, represents the nominal fraction of the loan amount repaid each month, which we calculate as the monthly payment divided by the principal balance. For example, a $\phi(r)$ value of 2.077% means that the borrower pays 2.077% of the principal amount each month over the course of a five-year loan. The x-axis shows the total loan amount per year, Λ, in billion dollars for previously-bankrupt/high-cost and never-bankrupt/low-cost borrowers, respectively. For example, $2.94 billion represents the loan amount that 800,000 previously-bankrupt individuals borrow each year when bankruptcy flags are never removed. In panel (a), the light gray trapezoid illustrates the borrower surplus lost by previously-bankrupt individuals in the counterfactual where flags are never removed. The dark gray triangle shows the efficiency gain from eliminating credit overprovision to previously-bankrupt individuals in the counterfactual. In panel (b), the light gray and dark gray area shows the borrower surplus gain for never-bankrupt individuals in the counterfactual. The dark gray triangle illustrates the efficiency gain due to eliminating credit underprovision to never-bankrupt individuals.

Table 3: Welfare Estimates

This table summarizes our main estimates implied by the specifications of Table 2. Panel A shows average interest rates in the six months before flag removal ($r_{H,fair}$), the interest rate effect of flag removal ($r_{pool} - r_{H,fair}$), and the effect of flag removal on the fraction of the principal repaid each month in a standardized five-year loan ($\phi(r_{pool}) - \phi(r_{H,fair})$). Panel B shows average loan quantities in the six months before flag removal and the quantity effect of flag removal. Panel C summarizes surplus changes relative to a counterfactual world in which bankruptcy flags are never removed. The first row shows the average change in consumer surplus for individuals with flag removal for the average five-year loan. It is the sum of monthly non-default period surpluses. The number of non-default periods is derived from the probability that loans to individuals who at some point have a bankruptcy flag will be charged-off within two years of opening. The second row shows the social surplus change in expected dollars over the term of a loan, for each individual whose flag is removed. The estimate combines the welfare loss due to credit overprovision to the previously-bankrupt H group and credit underprovision to the never-bankrupt L group. The third row provides the efficiency change per dollar redistributed to bankrupt individuals through the removal of bankruptcy flags. It is computed by dividing the second row by the first row. Source: TransUnion.

	(1)	(2)	(3)	(4)	(5)
Panel A: Prices					
Pre-flag-removal loan interest rate (%)	9.02%	9.02%	9.02%	9.02%	9.02%
Flag removal-induced change in interest rate (%)	-0.218%	-0.226%	-0.218%	-0.226%	-0.226%
Change in monthly repayment as share of loan principal (%)	-0.011%	-0.011%	-0.011%	-0.011%	-0.011%
Panel B: Quantities					
Pre-flag-removal loan quantity (Average $ per borrower per year)	$3,678.00	$3,678.00	$3,678.00	$3,678.00	$3,678.00
Flag removal-induced change in loan quantity (Average $ per borrower per year)	$220.13	$215.11	$220.13	$215.75	$213.52
Panel C: Average Surplus Changes					
Average consumer surplus redistributed to individuals with flag removal over 5 years ($ per eligible borrower with flag removal)	$22.93	$23.76	$22.93	$23.76	$23.75
Change in social surplus per individual over 5 years ($ per eligible borrower with flag removal)	-$0.74	-$0.75	-$0.74	-$0.76	-$0.75
Welfare change per dollar redistributed to bankrupt individuals	-0.0324	-0.0317	-0.0324	-0.0318	-0.0315

4 Conclusion

This paper presents a new framework for studying the role of data acquisition in consumer credit markets. In a simple price-theoretic model of lending markets, we show that the price and quantity changes resulting from new data are sufficient statistics for calculating the welfare and transfer effects of data policy. We apply our framework to the setting of bankruptcy flag removal; we show that keeping bankruptcy flags would increase social welfare relative to the status quo where bankruptcy flags are removed. However, the associated welfare gains would be fairly small relative to the induced transfers between previously-bankrupt and never-bankrupt borrowers.

While we present a specific application associated with the welfare benefits of data acquisition, the method is broadly applicable in financial markets. Future work could study data acquisition in other lending markets, and explore other contexts in which data acquisition leads to consumer benefits. For example, certain types of data acquisition may lead to large welfare gains relative to welfare transfers. Further, the welfare gains associated with data acquisition may be small in certain settings, suggesting that the privacy or equity gains from making data unavailable may outweigh the direct benefits of using the data to screen borrowers.

informative about default rates, then removing it has a larger negative impact on social welfare, for each dollar of surplus transferred between groups. Quantitatively, however, we show that data removal remains a low-cost way to transfer surplus between groups, costing less than 21 cents of deadweight loss for each dollar transferred, even for signals that induce price changes up to eight times as large as the price changes induced by bankruptcy flag removal.

References

Acemoglu, Daron, Ali Makhdoumi, Azarakhsh Malekian, and Asuman Ozdaglar, "Too Much Data: Prices and Inefficiencies in Data Markets," Technical Report, National Bureau of Economic Research 2019.

Adams, William, Liran Einav, and Jonathan Levin, "Liquidity Constraints and Imperfect Information in Subprime Lending," *American Economic Review*, 2009, *99* (1), 49–84.

Argyle, Bronson, Benjamin Iverson, Taylor D Nadauld, and Christopher Palmer, "Personal Bankruptcy and the Accumulation of Shadow Debt," Technical Report, National Bureau of Economic Research 2021.

_ , Sasha Indarte, Ben Iverson, and Christopher Palmer, "Explaining Racial Disparities in Personal Bankruptcy Outcomes," 2022. Working Paper.

_ , Taylor D Nadauld, and Christopher Palmer, "Real Effects of Search Frictions in Consumer Credit Markets," Technical Report, National Bureau of Economic Research 2020.

Barrios, John Manuel, "Staggeringly Problematic: A Primer on Staggered DiD for Accounting Researchers," *Working Paper*, 2021.

Begenau, Juliane, Maryam Farboodi, and Laura Veldkamp, "Big Data in Finance and the Growth of Large Firms," *Journal of Monetary Economics*, 2018, *97*, 71–87.

Bergemann, Dirk, Benjamin Brooks, and Stephen Morris, "The Limits of Price Discrimination," *American Economic Review*, 2015, *105* (3), 921–57.

Blattner, Laura and Scott Nelson, "How costly is noise? Data and disparities in consumer credit," *arXiv preprint arXiv:2105.07554*, 2021.

_ , Jacob Hartwig, and Scott Nelson, "Information Design in Consumer Credit Markets," Technical Report, Working paper 2022.

Bos, Marieke, Emily Breza, and Andres Liberman, "The Labor Market Effects of Credit Market Information," *The Review of Financial Studies*, 2018, *31* (6), 2005–2037.

Calomiris, CW and L Neal, "History of Financial Globalization, Overview," *Handbook of Key Global Financial Markets, Institutions and Infrastructure*, 2013, *1*, 3–14.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer, "The Effect of Minimum Wages on Low-wage Jobs," *The Quarterly Journal of Economics*, 2019, *134* (3), 1405–1454.

Chatterjee, Satyajit, Dean Corbae, Kyle P Dempsey, and José-Víctor Ríos-Rull, "A Quantitative Theory of the Credit Score," Technical Report, National Bureau of Economic Research 2020.

Chen, Yongmin and Marius Schwartz, "Differential Pricing when Costs Differ: A Welfare Analysis," *The RAND Journal of Economics*, 2015, *46* (2), 442–460.

Chiappori, Pierre-André and Bernard Salanie, "Testing for Asymmetric Information in Insurance Markets," *Journal of Political Economy*, 2000, *108* (1), 56–78.

Chipman, John S and James C Moore, "Compensating variation, consumer's surplus, and welfare," *The American Economic Review*, 1980, *70* (5), 933–949.

Dávila, Eduardo, "Using Elasticities to Derive Optimal Bankruptcy Exemptions," *The Review of Economic Studies*, 2020, *87* (2), 870–913.

DeFusco, Anthony A, Huan Tang, and Constantine Yannelis, "Measuring the Welfare Cost of Asymmetric Information in Consumer Credit Markets," *Journal of Financial Economics*, 2022, *146* (3), 821–840.

Dobbie, Will, Paul Goldsmith-Pinkham, and Crystal S Yang, "Consumer Bankruptcy and Financial Health," *Review of Economics and Statistics*, 2017, *99* (5), 853–869.

_ , _ , Neale Mahoney, and Jae Song, "Bad Credit, No Problem? Credit and Labor Market Consequences of Bad Credit Reports," *The Journal of Finance*, 2020, *75* (5), 2377–2419.

Einav, Liran, Amy Finkelstein, and Mark R. Cullen, "Estimating Welfare in Insurance Markets Using Variation in Prices," *Quarterly Journal of Economics*, 2010, *125* (3), 877–921.

Elul, Ronel and Piero Gottardi, "Bankruptcy: Is it enough to forgive or must we also forget?," *American Economic Journal: Microeconomics*, 2015, *7* (4), 294–338.

European Data Protection Supervisor, "Opinion on the Proposal for a Directive on Consumer Credits," 2021.

Farboodi, Maryam and Laura Veldkamp, "Long-run Growth of Financial Data Technology," *American Economic Review*, 2020, *110* (8), 2485–2523.

_ , Roxana Mihet, Thomas Philippon, and Laura Veldkamp, "Big Data and Firm Dynamics," in "AEA papers and proceedings," Vol. 109 2019, pp. 38–42.

Fuster, Andreas, Paul Goldsmith-Pinkham, Tarun Ramadorai, and Ansgar Walther, "Predictably Unequal? The Effects of Machine Learning on Credit Markets," *The Journal of Finance*, 2022, *77* (1), 5–47.

Goodman-Bacon, Andrew, "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 2021.

Gross, Tal, Matthew J Notowidigdo, and Jialan Wang, "The Marginal Propensity to Consume Over the Business Cycle," *American Economic Journal: Macroeconomics*, 2020, *12* (2), 351–84.

_ , Raymond Kluender, Feng Liu, Matthew J Notowidigdo, and Jialan Wang, "The Economic Consequences of Bankruptcy Reform," *American Economic Review*, 2021, *111* (7), 2309–41.

Harberger, Arnold C., "The Measurement of Waste," *American Economic Review*, 1964, *54* (3), 58–76.

He, Zhiguo, Jing Huang, and Jidong Zhou, "Open Banking: Credit Market Competition When Borrowers Own the Data," Technical Report, National Bureau of Economic Research 2020.

Herkenhoff, Kyle, Gordon M Phillips, and Ethan Cohen-Cole, "The Impact of Consumer Credit Access on Self-employment and Entrepreneurship," *Journal of Financial Economics*, 2021, *141* (1), 345–371.

Huang, Jing, "Fintech Expansion," Technical Report, Working Paper 2022.

Indarte, Sasha, "Moral Hazard versus Liquidity in Household Bankruptcy," 2021. Working Paper.

Jehle, Geoffrey A. and Philip J. Reny, *Advanced Microeconomic Theory*, 3rd ed., Pearson, 2011.

Jones, Charles I and Christopher Tonetti, "Nonrivalry and the Economics of Data," *American Economic Review*, 2020, *110* (9), 2819–58.

Keys, Benjamin J, Neale Mahoney, and Hanbin Yang, "What Determines Consumer Financial Distress? Place-and Person-Based Factors," Technical Report, National Bureau of Economic Research 2020.

Kovbasyuk, Sergey and Giancarlo Spagnolo, "Memory and Markets," *Available at SSRN 2756540*, 2021.

Liberman, Andres, Christopher A. Neilson, Luis Opazo, and Seth Zimmerman, "The Equilibrium Effects of Information Deletion: Evidence from Consumer Credit Markets," 2019. Working Paper.

Liu, Zhuang, Michael Sockin, and Wei Xiong, "Data Privacy and Algorithmic Inequality," Technical Report, National Bureau of Economic Research 2023.

Marshall, Alfred, *Principles of Economics*, 8 ed., Macmillan and Co., 1920.

Mas-Colell, Andreu, Michael Dennis Whinston, Jerry R Green et al., *Microeconomic theory*, Vol. 1, Oxford university press New York, 1995.

Mayer, Christopher, Edward Morrison, Tomasz Piskorski, and Arpit Gupta, "Mortgage Modification and Strategic Behavior: Evidence from a Legal Settlement with Countrywide," *American Economic Review*, 2014, *104* (9), 2830–57.

Musto, David K, "What Happens When Information Leaves a Market? Evidence from Post-bankruptcy Consumers," *The Journal of Business*, 2004, *77* (4), 725–748.

Nelson, Scott, "Private Information and Price Regulation in the US Credit Card Market," *Unpublished Working Paper*, 2018.

Oster, Emily, "Unobservable Selection and Coefficient Stability: Theory and Evidence," *Jour-*

nal of Business & Economic Statistics, 2019, *37* (2), 187–204.

Padilla, A Jorge and Marco Pagano, "Sharing default information as a borrower discipline device," *European Economic Review*, 2000, *44* (10), 1951–1980.

Pagano, Marco and Tullio Jappelli, "Information Sharing in Credit Markets," *The Journal of Finance*, 1993, *48* (5), 1693–1718.

Schmalensee, Richard, "Output and Welfare Implications of Monopolistic Third-degree Price Discrimination," *The American Economic Review*, 1981, *71* (1), 242–247.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

Tang, Huan, "The Value of Privacy: Evidence from Online Borrowers," 2019.

Varian, Hal R, "Price Discrimination and Social Welfare," *The American Economic Review*, 1985, *75* (4), 870–875.

_ , "Price Discrimination," *Handbook of Industrial Organization*, 1989, *1*, 597–654.

Vercammen, James A, "Credit Bureau Policy and Sustainable Reputation Effects in Credit Markets," *Economica*, 1995, pp. 461–478.

Vives, Xavier, "Small income effects: A Marshallian theory of consumer surplus and downward sloping demand," *The review of economic studies*, 1987, *54* (1), 87–103.

Willig, Robert D, "Consumer's surplus without apology," *The American Economic Review*, 1976, *66* (4), 589–597.

Yannelis, Constantine, "Strategic Default on Student Loans," Technical Report, Working paper 2016.

_ **and Anthony Lee Zhang**, "Competition and Selection in Credit Markets," Technical Report, National Bureau of Economic Research 2021.

Appendix

A Proofs

A.1 Proof of Claim 1

Lender profits. Since we have assumed $r_{j,fair}$ is equal to the marginal cost of providing credit, at rate $r_{j,fair}$, and payment $\phi\left(r_{j,fair}\right)$, lenders break even on borrowers. At rate r_{pool}, lenders' profit, per non-default period and per dollar lent, is:

$$\phi\left(r_{pool}\right)-\phi\left(r_{j,fair}\right)$$

Thus, lenders' total profit at rate r_{pool} is:

$$\psi_j\Lambda\left(r_{pool}\right)\left[\phi\left(r_{pool}\right)-\phi\left(r_{j,fair}\right)\right] \tag{25}$$

The change in profits on group j, when moving from rate r_{pool} to $r_{j,fair}$, is thus the negative of (25).

Borrower welfare. From (13) in the main text, we have:

$$DCS_j\left(r_{j,fair}\right)-DCS_j\left(r_{pool}\right)=-\psi_j\int_{r_{pool}}^{r_{j,fair}}\Lambda_j\left(\hat{r}\right)\frac{d\phi\left(\hat{r}\right)}{d\hat{r}}d\hat{r}$$

Changing variables to ϕ, we can alternatively write this as:

$$DCS_j\left(r_{j,fair}\right)-DCS_j\left(r_{pool}\right)=-\psi_j\int_{\phi\left(r_{pool}\right)}^{\phi\left(r_{j,fair}\right)}\Lambda_j\left(\hat{r}\right)d\phi \tag{26}$$

Since $\Lambda(\phi)$ is linear, (26) is equal to (18). Summing (17) and (18) and rearranging, we get (19).

A.2 Proof of Claim 2

Repeating the definitions in (70) and (71), we have:

$$s_L\equiv\frac{\psi_L\Lambda_L\left(r_{pool}\right)}{\psi_L\Lambda_L\left(r_{pool}\right)+\psi_H\Lambda_H\left(r_{pool}\right)} \tag{27}$$

$$\Delta\equiv\phi\left(r_{H,fair}\right)-\phi\left(r_{L,fair}\right) \tag{28}$$

Using (27) and (28), we can write (68) as:

$$\phi\left(r_{pool}\right)=\phi\left(r_{L,fair}\right)+(1-s_L)\,\Delta+m_{pool} \tag{29}$$

Or,

$$\phi\left(r_{pool}\right)=\phi\left(r_{H,fair}\right)-s_L\Delta+m_{pool} \tag{30}$$

To calculate welfare, we will calculate the welfare of each group, when data is available and is not, relative to the fully efficient case, using the result of Claim 1. For the low-cost L group, when data is available, the social welfare loss relative to the fully efficient case can be obtained using (19):

$$\frac{1}{2}\left(\phi\left(r_L\right)-\phi\left(r_{L,fair}\right)\right)\left(\Lambda_L\left(r_{L,fair}\right)-\Lambda_L\left(r_L\right)\right) \tag{31}$$

$$=\frac{1}{2}\left(\phi\left(r_{L,fair}\right)+m_L-\phi\left(r_{L,fair}\right)\right)\left(\left(a_L-b_L\phi\left(r_{L,fair}\right)\right)-\left(a_L-b_L\left(\phi\left(r_{L,fair}\right)+m_L\right)\right)\right) \tag{32}$$

$$Loss_{L,data}=\frac{b_L}{2}m_L^2 \tag{33}$$

Intuitively, there is a welfare loss from markups, which depends on the size of the markup m_L, and the slope of demand b_L. Similarly, for the high-cost H group, the welfare loss from markups is:

$$Loss_{H,data}=\frac{b_H}{2}m_H^2 \tag{34}$$

Using these expressions, we calculate the welfare loss, relative to the fully efficient benchmark, in the case of pooled pricing. Note that we can write (29) as:

$$\phi\left(r_{pool}\right)=\phi\left(r_{L,fair}\right)+(1-s_L)\,\Delta+m_{pool}$$

Plugging this into (31), expanding and simplifying, the welfare loss for the low-cost L group in the no-data case is:

$$Loss_{L,nodata}=\frac{b_L}{2}\left(m_{pool}+(1-s_L)\,\Delta\right)^2 \tag{35}$$

And for the high-cost H group, we have:

$$Loss_{H,nodata}=\frac{b_H}{2}\left(m_{pool}-s_L\Delta\right)^2 \tag{36}$$

Hence, the welfare change when going from the pooled case to the case with data is

$$\frac{b_L}{2}\left(\left(m_{pool}+(1-s_L)\Delta\right)^2-m_L^2\right) \tag{37}$$

for the low-cost L group, and

$$\frac{b_H}{2}\left(\left(m_{pool}-s_L\Delta\right)^2-m_H^2\right) \tag{38}$$

for the high-cost H group, in terms of dollars per non-default period. To convert these quantities into expected dollars over the term of a loan, we will multiply each by the expected number of non-default periods, ψ_L and ψ_H, defined in (12). The total welfare change from data availability is thus:

$$\Delta Welfare = \psi_H Loss_{H,nodata} + \psi_L Loss_{L,nodata} - \psi_H Loss_{H,data} - \psi_L Loss_{L,data}$$

$$\Delta Welfare = \underbrace{\psi_H \frac{b_H}{2}\left(\left(m_{pool}-s_L\Delta\right)^2-m_H^2\right)}_{High\ cost\ group} + \underbrace{\psi_L \frac{b_L}{2}\left(\left(m_{pool}+(1-s_L)\,\Delta\right)^2-m_L^2\right)}_{Low\ cost\ group} \tag{39}$$

This is (69).

B Extensions

B.1 Ratio of Welfare Changes between L and H groups

In this Appendix, we calculate the ratio of welfare changes between the high-cost H group and the low-cost L group when data is never removed from credit reports. We show that the L group welfare change tends to be a factor $\frac{\Lambda_H(r)}{\Lambda_L(r)}$ smaller than the H group welfare effect, multiplied by some constants; thus, in our setting where most individuals are not in the previously-bankrupt H group, the welfare effects will tend to be dominated by what is happening in the previously-bankrupt H group.

From (19) of Claim 1, the increase in social welfare for the L group from never removing data, in expected dollars over the lifetime of a loan, is:

$$\frac{1}{2}\psi_L\left(\phi\left(r_{L,fair}\right)-\phi\left(r_{pool}\right)\right)\left(\Lambda_L\left(r_{pool}\right)-\Lambda_L\left(r_{L,fair}\right)\right) \tag{40}$$

Writing (40) using the demand slope b_L, this becomes:

$$\frac{1}{2}\psi_L\left(\phi\left(r_{L,fair}\right)-\phi\left(r_{pool}\right)\right)\left(b_L\left(\phi\left(r_{L,fair}\right)-\phi\left(r_{pool}\right)\right)\right) \tag{41}$$

Now, we can express (41) in terms of the price change for the H group using the zero-profit condition. From (20):

$$\phi\left(r_{L,fair}\right)-\phi\left(r_{pool}\right)=-\frac{\psi_H\Lambda_H\left(r_{pool}\right)}{\psi_L\Lambda_L\left(r_{pool}\right)}\left(\phi\left(r_{H,fair}\right)-\phi\left(r_{pool}\right)\right) \tag{42}$$

In words, (42) states that the price reduction for L borrowers is the negative of the price increase to the H group, $\left(\phi\left(r_{H,fair}\right)-\phi\left(r_{pool}\right)\right)$, multiplied by the ratio of total loan volumes, $\frac{\Lambda_H(\phi(r))}{\Lambda_L(\phi(r))}$, and the ratio of expected non-default periods $\frac{\psi_H}{\psi_L}$. Applying this to (41), to get:

$$\begin{aligned}=\frac{1}{2}\psi_L&\left(-\frac{\psi_H\Lambda_H\left(r_{pool}\right)}{\psi_L\Lambda_L\left(r_{pool}\right)}\left(\phi\left(r_{H,fair}\right)-\phi\left(r_{pool}\right)\right)\right)\times\\ &\left(b_L\left(-\frac{\psi_H\Lambda_H\left(r_{pool}\right)}{\psi_L\Lambda_L\left(r_{pool}\right)}\left(\phi\left(r_{H,fair}\right)-\phi\left(r_{pool}\right)\right)\right)\right)\end{aligned} \tag{43}$$

If we further apply (21) of Assumption 2, stating that demand elasticities are equal across groups, then (43) rearranges to:

$$=\frac{1}{2}\psi_L\left(\frac{\psi_H}{\psi_L}\right)^2\frac{\Lambda_H\left(r_{pool}\right)}{\Lambda_L\left(r_{pool}\right)}\left(\phi\left(r_{H,fair}\right)-\phi\left(r_{pool}\right)\right)\left[b_H\left(\phi\left(r_{H,fair}\right)-\phi\left(r_{pool}\right)\right)\right] \tag{44}$$

Now, from (19) of Claim 1, and substituting for the Λ_H terms using demand linearity, the change in social welfare for the H group in a counterfactual world where is data is not removed is, in expected dollars over the lifetime of a loan:

$$\frac{1}{2}\psi_H b_H\left(\phi\left(r_{H,fair}\right)-\phi\left(r_{pool}\right)\right)\left(\phi\left(r_{H,fair}\right)-\phi\left(r_{pool}\right)\right) \tag{45}$$

Comparing (45) and (44), the change in welfare for the L group is a factor $\frac{\psi_H\Lambda_H\left(r_{pool}\right)}{\psi_L\Lambda_L\left(r_{pool}\right)}$ times the change in welfare for the H group. The factor $\frac{\psi_H}{\psi_L}$ will tend not to be very large or small, for default rates that are not unrealistically high. In our setting, $\Lambda_H\left(r_{pool}\right)$ is much smaller than $\Lambda_L\left(r_{pool}\right)$, so the change in welfare will tend to be much smaller for the L group than the H group, as we show empirically in Section 3. Quantitatively, note that if we assumed that the L group elasticity was k_L times larger than the H group elasticity, then the welfare change for the L group, (44), would become:

$$\frac{1}{2}k_L\frac{\psi_H^2}{\psi_L}\frac{\Lambda_H\left(r_{pool}\right)}{\Lambda_L\left(r_{pool}\right)}\left(\phi\left(r_{H,fair}\right)-\phi\left(r_{pool}\right)\right)\left[b_H\left(\phi\left(r_{pool}\right)-\phi\left(r_{H,fair}\right)\right)\right]$$

In order for the L group welfare change to be non-negligible, the factor $k_L$would have to be approximately equal to $\frac{\psi_L}{\psi_H}\frac{\Lambda_L(r)}{\Lambda_H(r)}$. Plugging in our Λ ratio ($\frac{\Lambda_H\left(r_{pool}\right)}{\Lambda_L\left(r_{pool}\right)}\approx 0.1186$) and ψ estimates ($\psi_H = 57.29$, $\psi_L = 58.34$) from the main text (e.g., compare Subsection 3.3), we get that the L group demand elasticity would have to be approximately 8.6 times the H group demand elasticity for the L group welfare change to be equal to that of the H group.

B.2 Interest Rate Menus and Loan Size Limits

In the baseline model, we assume the interest rate r is fixed and independent of loan size, and that a borrower can borrow arbitrarily much. In this appendix, extending the baseline model, we relax these assumptions, allowing the interest rate $r(L)$ to depend on loan size L, and allowing some upper bound $\bar{L}$ on loan size. This may capture, for example, hard PTI or LTV caps, or interest rates which are higher for higher LTV loans. We analyze the problem facing a single borrower. To model a shift in interest rates, we introduce a parameter ω which shifts the "menu" of interest rates vertically, so that the borrower's rate is $r(L)+\omega$. We also assume there may be a hard ceiling $\bar{L}$ on the size of the loan the borrower gets. We can thus write the borrower's optimization problem as:

$$V(r) = \max_L u_0(w_0 + L) + \sum_{t=1}^{T} \beta^t (1-\delta)^t \left[u(w_t) - u'(w_t) L\phi(r(L)+\omega)\right] + \sum_{t=1}^{T} (1-\delta)^{t-1} \delta \sum_{\tilde{t}=t}^{T} \beta^{\tilde{t}} u(c_D)$$
$$s.t.\ L \leq \bar{L}$$

We can thus analyze how a small change in ω, which shifts upwards or downwards the menu of interest rates facing the borrower, affects borrower utility.[17] Applying the envelope theorem, we have:

$$\frac{dV}{d\omega} = -\sum_{t=1}^{T} \beta^t (1-\delta)^t u'(w_t) L^*(\omega) \phi'(r(L^*(\omega)) + \omega) \tag{46}$$

Integrating, for a given change in ω, we have:

$$V(\tilde{\omega}) - V(\omega) = -\underbrace{\left[\sum_{t=1}^{T} \beta^t (1-\delta)^t u'(w_t)\right]}_{Utility\ weight} \underbrace{\left[\int_{\omega}^{\tilde{\omega}} L^*(\omega) \phi'(r(L^*(\omega)) + \omega) d\hat{\omega}\right]}_{Payment\ change} \tag{47}$$

Expressions (46) and (47) are unchanged from (6) and similar to (8). Thus, the analysis of the main text applies unchanged in the presence of credit constraints: if data policy affects market outcomes solely through shifting the menu of interest rates available to customers, expression (47) correctly captures the welfare effects of data policy.

It may seem somewhat surprising that our findings are completely unaffected by the presence of credit constraints. A simple intuition is that the envelope theorem states that, if a consumer is optimally borrowing some amount L^*, paying $\phi(r)L^*$ and interest rates shift by ϵ, the borrower has to pay $L\phi(r+\epsilon)$ instead of $L\phi(r)$, and this monetary difference fully captures the borrower's change in welfare. This logic applies whether L is chosen freely, at some cost to interest rates, or constrained to be below some upper bound $\bar{L}$.

An important caveat, however, is that this analysis only applies to policies which act through changing interest rates. Suppose for example the main channel through which bankruptcy flag removal affected outcomes was through changing the amount consumers could borrow, $\bar{L}$; it would then not be correct to summarize the effect of policy using $\frac{dV}{d\omega}$. We would instead want

[17] There is a more general version in which we allow $r(L,\omega)$ to depend flexibly on ω; the envelope theorem still holds, implying:

$$\frac{dV}{d\omega} = -\sum_{t=1}^{T} \beta^t (1-\delta)^t u'(w_t) L^*(\omega) \phi'(r(L^*,\omega)) \frac{\partial r}{\partial \omega}\Big|_{L=L^*(\omega)}$$

We focus on the linear-shift version for simplicity.

to calculate $\frac{dV}{d\bar{L}}$, a quantity which depends on the shadow price of the constraint $L \leq \bar{L}$, which we have not found any simple way to measure empirically.

B.3 Endogenous Default and Post-Default Value

In the baseline model, we assumed that default is exogenous, and that upon default borrowers receive the fixed amount c_D. In this appendix, we explore relaxing both these assumptions. We analyze the problem facing a single borrower; for notational simplicity, assume there are only two periods, so money is borrowed in period 0 and paid back in period 1. Assume the borrower can choose whether to default in period 1; hence, the borrower chooses loan size to solve:

$$V(r) = \max_L u_0(w_0 + L) + \int \beta \max\left[V_D(r, L, s), \left[u(w_t, s) - u'(w_t, s) L\phi(r)\right]\right] dF(s) \tag{48}$$

Where, in the second period, the borrower chooses between repaying the loan and attaining

$$u(w_t, s) - u'(w_t, s) L\phi(r)$$

or defaulting, and receiving some value function $V_D(r, L, s)$ from default. The variable s, distributed as $F(s)$, is a continuous random variable which affects borrowers' default decisions; we could think of s as representing default-relevant random factors, such as stochastic income shocks, or whether a borrower loses her job. s is allowed to affect $u(w_t, s), u'(w_t, s)$, and the value of default $V_D(r, L, s)$. We assume there is a threshold equilibrium, so borrowers default if $s > s^*$. We leave precisely what s represents unspecified, as it does not affect the conclusions from our analysis. Importantly, in this setting, the value borrowers receive upon default, $V_D(r, L, s)$, may depend on borrowers' loan size and interest rate. This represents, for example, that if lenders have partial recourse, borrowers may be worse off if they start a loan with a higher interest rate and then default on a loan with a larger balance. Accounting for such effects quantitatively is not straightforward, as it requires taking a stance on proceedings after loan default.

Differentiating (48) with respect to r using the envelope theorem, we have:

$$\frac{dV}{dr} = \beta \left[\underbrace{\int_{s<s*} -u'(w_t, s) L \frac{d\phi}{dr} dF(s)}_{A} + \underbrace{\int_{s \geq s*} \frac{\partial V_D(r, L, s)}{\partial r} dF(s)}_{B} \right] \tag{49}$$

Suppose that $\frac{\partial V_D(r,L,s)}{\partial r} = 0$; this is the case in the main text, model, where we assumed that consumption upon default, c_D, is exogenous. Expression (49) then reduces to term A, which is just (6) in the main text, with $1-F(s^*)$ playing the role of the exogenous default rate δ. Thus, in a model where borrowers optimally decide when to default, but outcomes upon default are unaffected by interest rates, our results continue to hold. However, in the general case where $\frac{\partial V_D(r,L,s)}{\partial r}$ may be nonzero, the derivative $\frac{dV}{dr}$ has the additional term B, reflecting how much changes in r on average change the value function upon default.

Expression (49) shows that endogenous default per se does not affect (6); the basic logic of the envelope theorem—that, when r increases, the total monetary amount paid in non-default states increases by an amount directly proportional to loan size L—continues to hold. What does matter is whether interest rates affect the value borrowers attain after default; if $\frac{\partial V_D(r,L,s)}{\partial r}$ is nonzero, there is an extra term in the derivative $\frac{dV}{dr}$. For example, higher interest rates could increase the balance at the time of default, which can affect post-default consumption when loans have recourse.

B.4 Partially Separating Prices

In our main model, as well as our empirical analysis, we consider shifting from a pooling equilibrium to a perfectly separating equilibrium, in which each borrower group has homogeneous default rates and faces a price which exactly matches their default rate. In most realistic settings, data will only be partially informative, and there will be some heterogeneity in default rates within borrower groups. In a more realistic model where data only leads to partial separation, so there are residual differences in borrowers' costs within groups separated through data, how accurate are the expressions in Claim 1? In this Appendix, we show that the simple expressions in Claim 1 are approximately equal to expressions for lender, borrower, and total surplus in a richer model with heterogeneity within groups, as long as default rates do not vary very much within groups, and borrowers' demand elasticities do not covary very much with default rates.

We assume there is a finite (but possibly large) set of possible default rates $\delta_1 \ldots \delta_N$. There are two groups, X and Y, with different distributions over default rates $p^X(\delta), p^Y(\delta)$. We write loan demand as $\Lambda_{X,i}(r)$ or $\Lambda_{Y,i}(r)$, so loan demand may vary across groups, and across types within a group. As in Subsection 2.4, we assume loan demand is linear in $\phi(r)$.

Suppose some data is introduced so that the market can distinguish between the two groups X and Y; the market thus shifts from the fully pooled price r_{pool} for all borrowers, to the partially pooled prices $r_{X,pool}$ and $r_{Y,pool}$ for groups X and Y. The prices $r_{X,pool}$ and $r_{Y,pool}$ are not perfectly separating prices, since we assumed there are differences in borrowers' default rates within groups X and Y. However, intuitively, if group identities are somewhat informative about default rates, then total welfare should tend to increase when shifting to $r_{X,pool}, r_{Y,pool}$.

Social welfare. To calculate the change in welfare, we can calculate total deadweight loss under each scenario, using (19), and then take the difference between these sums. Under r_{pool}, summing across all borrower types within groups X and Y, deadweight loss is:

$$\begin{aligned}&\frac{1}{2}\sum_{i=1}^{N}\psi(\delta_i)\left(\phi\left(r_{i,fair}\right)-\phi\left(r_{pool}\right)\right)\left(\Lambda_{X,i}\left(r_{pool}\right)-\Lambda_{X,i}\left(r_{i,fair}\right)\right)p_X^i\\ +&\frac{1}{2}\sum_{i=1}^{N}\psi(\delta_i)\left(\phi\left(r_{i,fair}\right)-\phi\left(r_{pool}\right)\right)\left(\Lambda_{Y,i}\left(r_{pool}\right)-\Lambda_{Y,i}\left(r_{i,fair}\right)\right)p_Y^i\end{aligned} \tag{50}$$

where we multiply each term by $\psi(\delta_i)$, from (12), so welfare is in terms of expected dollars over the lifetime of a loan. Since we assumed demand is linear in $\phi(r)$, we can write (50) instead as:

$$\frac{1}{2}\sum_{i=1}^{N}\psi(\delta_i)\left(\phi\left(r_{pool}\right)-\phi\left(r_{i,fair}\right)\right)^2 b_{X,i}p_X^i+\frac{1}{2}\sum_{i=1}^{N}\psi(\delta_i)\left(\phi\left(r_{pool}\right)-\phi\left(r_{i,fair}\right)\right)^2 b_{Y,i}p_Y^i \tag{51}$$

Correspondingly, total welfare loss under separated pricing r_X, r_Y can be written as:

$$\frac{1}{2}\sum_{i=1}^{N}\psi(\delta_i)\left(\phi\left(r_{X,pool}\right)-\phi\left(r_{i,fair}\right)\right)^2 b_{X,i}p_X^i+\frac{1}{2}\sum_{i=1}^{N}\psi(\delta_i)\left(\phi\left(r_{Y,pool}\right)-\phi\left(r_{i,fair}\right)\right)^2 b_{Y,i}p_Y^i \tag{52}$$

We wish to find the difference between (51) and (52). We will focus on the group X terms, since the group Y case is symmetric. Now, note that we can write the group X term from (51) as:

$$\frac{1}{2}\sum_{i=1}^{N}\psi(\delta_i)\left(\left(\phi\left(r_{pool}\right)-\phi\left(r_{X,pool}\right)\right)+\left(\phi\left(r_{X,pool}\right)-\phi\left(r_{i,fair}\right)\right)\right)^2 b_{X,i}p_X^i$$

$$\begin{aligned}=\frac{1}{2}\sum_{i=1}^{N}\psi(\delta_i)\Big[&\left(\phi\left(r_{pool}\right)-\phi\left(r_{X,pool}\right)\right)^2+\left(\phi\left(r_{X,pool}\right)-\phi\left(r_{i,fair}\right)\right)^2+\\ &2\left(\phi\left(r_{pool}\right)-\phi\left(r_{X,pool}\right)\right)\left(\phi\left(r_{X,pool}\right)-\phi\left(r_{i,fair}\right)\right)\Big]b_{X,i}p_X^i\end{aligned}$$

$$
\begin{aligned}
&= \underbrace{\frac{1}{2}\left(\phi\left(r_{pool}\right)-\phi\left(r_{X,pool}\right)\right)^{2}\sum_{i=1}^{N}\psi\left(\delta_{i}\right)b_{X,i}p_{X}^{i}+}_{1:\ Positive} \\
&\qquad\underbrace{\frac{1}{2}\sum_{i=1}^{N}\psi\left(\delta_{i}\right)\left(\phi\left(r_{X,pool}\right)-\phi\left(r_{i,fair}\right)\right)^{2}b_{X,i}p_{X}^{i}+}_{2:\ Separated\ welfare} \\
&\qquad\underbrace{\left(\phi\left(r_{pool}\right)-\phi\left(r_{X,pool}\right)\right)\sum_{i=1}^{N}\psi\left(\delta_{i}\right)\left(\phi\left(r_{X,pool}\right)-\phi\left(r_{i,fair}\right)\right)b_{X,i}p_{X}^{i}}_{3:\ Covariance} \qquad (53)
\end{aligned}
$$

Expression (53) can be thought of as an "approximate sum of squares" decomposition of DWL under r_{pool}. Term 1 in (53) is positive, since demand is downwards-sloping so $b_{X,i}$ is positive, the weights $\psi(\delta_i)$ and p_X^i are positive, and $\left(\phi\left(r_{pool}\right)-\phi\left(r_{X,pool}\right)\right)^2$ is always positive. Term 2 is exactly the left term of (52), that is, total DWL in group X under the separated price $\phi\left(r_{X,pool}\right)$. Thus, if term 3 is 0, (53) implies that DWL under r_{pool} is always greater than DWL under $r_{X,pool}, r_{Y,pool}$, analogous to the baseline model.

When will term 3 equal zero? Multiplying and dividing by $\Lambda_{X,i}\left(r_{X,pool}\right)$, term 3 in (53) can be written as:

$$
=\left(\phi\left(r_{pool}\right)-\phi\left(r_{X,pool}\right)\right)\sum_{i=1}^{N}\psi\left(\delta_{i}\right)\left(\phi\left(r_{X,pool}\right)-\phi\left(r_{i,fair}\right)\right)\Lambda_{X,i}\left(r_{X,pool}\right)\frac{b_{X,i}}{\Lambda_{X,i}\left(r_{X,pool}\right)}p_{X}^{i}
$$

The ratio $\frac{b_{X,i}}{\Lambda_{X,i}\left(r_{X,pool}\right)}$ is the demand elasticity of sub-group i; writing this as ε_i, this is:

$$
=\left(\phi\left(r_{pool}\right)-\phi\left(r_{X,pool}\right)\right)\sum_{i=1}^{N}\psi\left(\delta_{i}\right)\left(\phi\left(r_{X,pool}\right)-\phi\left(r_{i,fair}\right)\right)\Lambda_{X,i}\left(r_{X,pool}\right)\varepsilon_{i}p_{X}^{i} \qquad (54)
$$

Suppose demand elasticities are constant across borrower types, $\varepsilon_i=\varepsilon$ for all i. Factoring out ε, we have:

$$
=\left(\phi\left(r_{pool}\right)-\phi\left(r_{X,pool}\right)\right)\varepsilon\sum_{i=1}^{N}\psi\left(\delta_{i}\right)\left(\phi\left(r_{X,pool}\right)-\phi\left(r_{i,fair}\right)\right)\Lambda_{X,i}\left(r_{X,pool}\right)p_{X}^{i} \qquad (55)
$$

Now, generalizing Assumption 1 to more than two groups, if competitive lenders set prices so that they break even in expectation across types over the lifetime of a loan, lenders' zero-profit condition is that:

$$\sum_{i=1}^{n} \psi(\delta_i)\left(\phi\left(r_{X,pool}\right) - \phi\left(r_{i,fair}\right)\right) \Lambda_{X,i}\left(r_{X,pool}\right) p_X^i = 0 \tag{56}$$

Hence, if $\phi\left(r_{X,pool}\right)$ satisfies lenders' zero-profit condition and ε_i is constant, then (55) is equal to 0.

In this case, (53) shows that moving from pooled pricing to separated pricing always increases welfare. Moreover, the size of the welfare increase is exactly term 1 in (53). Note that we can write this term as:

$$\left(\phi\left(r_{pool}\right) - \phi\left(r_{X,pool}\right)\right)^2 \sum_{i=1}^{N} \psi(\delta_i)\, b_{X,i} p_X^i \tag{57}$$

If default rate δ_i's do not vary too much across borrowers, we can approximate $\psi(\delta_i)$ as a constant:

$$\psi(\delta_i) \approx \bar{\psi} \tag{58}$$

Then the welfare increase in (57) is approximately:

$$\approx \frac{1}{2}\bar{\psi}\left(\phi\left(r_{pool}\right) - \phi\left(r_{X,pool}\right)\right) \sum_{i=1}^{N} \left(\phi\left(r_{pool}\right) - \phi\left(r_{X,pool}\right)\right) b_{X,i} p_X^i$$

$$= \frac{1}{2}\bar{\psi}\left(\phi\left(r_{pool}\right) - \phi\left(r_{X,pool}\right)\right) \sum_{i=1}^{N} \left(\Lambda_{X,i}\left(r_{X,pool}\right) - \Lambda_{X,i}\left(r_{pool}\right)\right) p_X^i$$

$$= \frac{1}{2}\bar{\psi}\left(\phi\left(r_{X,pool}\right) - \phi\left(r_{pool}\right)\right)\left(\Lambda_X\left(r_{pool}\right) - \Lambda_X\left(r_{X,pool}\right)\right)$$

This is exactly (19), that is, the welfare increase assuming all borrowers in group X have the same default rate.

What happens when elasticities are not equal? For ease of interpretation, define the weighted average elasticity:

$$\bar{\varepsilon} = \sum_{i=1}^{N} p_X^i \varepsilon$$

Since (55) is 0 for any constant ε, for ease of interpretation, we can choose $\varepsilon = \bar{\varepsilon}$, and then subtract the sum (55), which is equal to 0, from (54), to write (54) as:

$$\left(\phi\left(r_{pool}\right) - \phi\left(r_{X,pool}\right)\right) \sum_{i=1}^{N} \psi(\delta_i) \Lambda_{X,i}\left(r_{X,pool}\right) p_X^i \left(\phi\left(r_{X,pool}\right) - \phi\left(r_{i,fair}\right)\right)(\varepsilon_i - \bar{\varepsilon}) \tag{59}$$

Expression (59) shows that term 3 in (53) can be thought of as a *weighted covariance* between the fair price $\phi\left(r_{i,fair}\right)$ and the elasticity ε_i within group X, weighted by $\psi\left(\delta_i\right)\Lambda_{X,i}\left(r_{X,pool}\right)p_X^i$. Suppose $\phi\left(r_{pool}\right) > \phi\left(r_{X,pool}\right)$, so pooled pricing tends to increase prices for group-X borrowers. This tends to decrease welfare (increase DWL) for low-risk agents with $\phi\left(r_{X,pool}\right) > \phi\left(r_{i,fair}\right)$, for whom even the pooled price $\phi\left(r_{X,pool}\right)$ is too high, and increase welfare (decrease DWL) for high-risk agents with $\phi\left(r_{X,pool}\right) < \phi\left(r_{i,fair}\right)$. When demand elasticities are constant across types, these effects exactly cancel, and the covariance term will be ignored. However, if for example the low-risk individuals tend to have higher elasticities than the high-risk individuals, then the covariance term will tend to be positive, increasing the DWL effect of pooled pricing, relative to the baseline. While demand elasticities may not be exactly the same, note that if the variance of demand elasticities is low, the covariance also cannot be very high; thus, if elasticities are roughly constant across borrowers, the covariance term will be low and Claim 1 of the main text will be approximately correct.

Producer surplus. Under separated pricing $r_{X,pool}, r_{Y,pool}$, by assumption, producers break even for groups X and Y. Hence, prices satisfy (56). Under any other price r_{pool}, constant across groups, producers' profits/losses in group X without data are:

$$\sum_{i=1}^{N}\psi\left(\delta_i\right)\left(\phi\left(r_{pool}\right)-\phi\left(r_{i,fair}\right)\right)\Lambda_{X,i}\left(r_{pool}\right)p_X^i \tag{60}$$

and with data:

$$\sum_{i=1}^{N}\psi\left(\delta_i\right)\left(\phi\left(r_{X,pool}\right)-\phi\left(r_{i,fair}\right)\right)\Lambda_{X,i}\left(r_{X,pool}\right)p_X^i$$

Hence, the change in producer surplus is:

$$\begin{aligned}&\sum_{i=1}^{N}\psi\left(\delta_i\right)(\phi\left(r_{pool}\right)-\phi\left(r_{X,pool}\right))\Lambda_{X,i}\left(r_{pool}\right)p_X^i\\+&\sum_{i=1}^{N}\psi\left(\delta_i\right)(\phi\left(r_{X,pool}\right)-\phi\left(r_{i,fair}\right))*(\Lambda_{X,i}\left(r_{pool}\right)-\Lambda_{X,i}\left(r_{X,pool}\right))p_X^i\end{aligned} \tag{61}$$

Under constant elasticity of substitution and zero lender profit, we get that the change in producer surplus is:

$$\sum_{i=1}^{N}\psi\left(\delta_i\right)(\phi\left(r_{pool}\right)-\phi\left(r_{X,pool}\right))\Lambda_{X,i}\left(r_{pool}\right)p_X^i$$

Again, if we approximate $\psi\left(\delta_i\right)$ as a constant using (58), producers' losses are:

$$\approx \bar{\psi}\left(\phi\left(r_{pool}\right)-\phi\left(r_{X,pool}\right)\right)\underbrace{\sum_{i=1}^{N}\Lambda_{X,i}\left(r_{pool}\right)p_X^i} \tag{62}$$

Now, the term with the underbrace is simply average loan demand across types i, weighted by the fractions of group X which have type i, which is just total loan demand in X; hence, (62) is approximately (17) of Claim 1, for $\bar{\psi} \approx \psi$.

Consumer surplus. We have shown that, if $\psi(\delta_i) \approx \bar{\psi}$ and $\varepsilon_i \approx \varepsilon$, then the simple expression for producer surplus changes, (17) in the main text, is approximately equal to the sum of producer surplus changes across all types, (61); and that if $\psi(\delta_i) \approx \bar{\psi}$ and $\varepsilon_i \approx \varepsilon$, then the simple expression for total surplus changes, (19) in the main text, is approximately equal to the sum of total surplus changes across all types, the X term in (51). The consumer surplus change from data, across all types, is for group X borrowers:

$$\frac{1}{2}\sum_{i=1}^{N}\psi(\delta_i)\left(\phi\left(r_{pool}\right)-\phi\left(r_{X,pool}\right)\right)\left(\Lambda_{X,i}\left(r_{pool}\right)+\Lambda_{X,i}\left(r_{X,pool}\right)\right)p_X^i \tag{63}$$

The change in total surplus—the X term in (51)—is equal to the sum of producer surplus, (61), and consumer surplus, (63). Hence, since the sum of consumer surplus (18) and producer surplus (17) is equal to total surplus (19) in the main text, it follows that the multi-type expression for consumer surplus (63) is approximately equal to the simple expression (18) in the main text, under the assumptions that $\psi(\delta_i) \approx \bar{\psi}$ and $\varepsilon_i \approx \varepsilon$.

Any model is an approximation to reality; the goal of this analysis is to quantify the extent of errors in the naive model in which each group can be thought of as having a single homogeneous default rate. We have shown that the errors in the approximation that all borrowers within a group have the same default rate depends on two quantities: how much default rates vary within groups, and how much default rates covary with demand elasticities. In settings where fine-grained data is available, in principle (59) could also be estimated directly.

B.5 Imperfect Competition

In the baseline model, we assume that markets are perfectly competitive. We can extend our framework to imperfect competition, under which lenders can charge a markup above marginal cost. In general, if data availability changes the markups that lenders can charge, outcomes are known to be complex (He et al., 2020; Huang, 2022). By analogy to the classic literature on third-degree price discrimination, it is not generally possible to say whether data will increase or decrease markups and thus social welfare (Schmalensee, 1981; Varian, 1985, 1989). Nonetheless, while markups make our theoretical conclusions less sharp, we show that markups can be accommodated in the empirical application of our methodology, with some additional data. Then, we show that, under a certain set of parameter restrictions—when the demand elasticities in the two borrower groups are equal, and when markups pre- and post-data availability are the same—the existence of markups does not, in fact, affect the welfare gains from data availability. While this set of assumptions may not always be satisfied exactly in reality, this finding suggests that the existence of markups per se does not dramatically affect our results. The two groups would need to have quite different demand elasticities, or markups would have to change significantly due to data availability, for our results to be substantially inaccurate.

As in the main model, let H denote high-cost borrowers and L denote low-cost borrowers. In the empirical application previously-bankrupt will correspond to the high-cost H group and never-bankrupt will correspond to the low-cost L group. Let $r_{H,fair}$ and $r_{L,fair}$ represent the interest rates for H and L, respectively, if markets were fully competitive; these variables are also equal to the social cost of providing credit to H and L. We will have $r_{H,fair} > r_{L,fair}$, since previously-bankrupt customers tend to be more costly to lenders (which is why their rates drop when flags are removed). As in the baseline model, we assume that the demand in both groups is linear, with possibly different slopes and intercepts:

$$\Lambda_L(r) = a_L - b_L \phi(r) \tag{64}$$

$$\Lambda_H(r) = a_H - b_H \phi(r) \tag{65}$$

Unlike in the main text, we assume that lenders may charge markups over marginal cost both before and after data is made available. Rather than take a stance on the particular theoretical model generating markups, we will simply take markups as exogenous and express welfare in terms of markups over marginal costs before and after data is made available. Let m_H, m_L be the markups charged over the competitive prices $\phi\left(r_{H,fair}\right), \phi\left(r_{L,fair}\right)$ when data is available to distinguish the two groups, and m_{pool} be the markup when data is not available.

When lenders charge markups, interest rates for each group when data are available, $r_{H,data}$ and $r_{L,data}$, will be higher than lenders' break-even interest rates $r_{H,fair}$ and $r_{L,fair}$. We will write these as:

$$\phi\left(r_{H,data}\right) = \phi\left(r_{H,fair}\right) + m_H,\ \phi\left(r_{L,data}\right) = \phi\left(r_{L,fair}\right) + m_L \tag{66}$$

When data is not available, the average cost across both groups is:

$$\frac{\psi_L \Lambda_L\left(r_{pool}\right) \phi\left(r_{L,fair}\right) + \psi_H \Lambda_H\left(r_{pool}\right) \phi\left(r_{H,fair}\right)}{\psi_L \Lambda_L\left(r_{pool}\right) + \psi_H \Lambda_H\left(r_{pool}\right)} \tag{67}$$

That is, (67) is a weighted average of the cost of serving L and H type borrowers, with weights equal to the loan volumes Λ_L, Λ_H multiplied by the expected number of non-default periods, ψ_H, ψ_L, defined in (12). If lenders set a markup m_{pool} above average costs, the price that borrowers face without data is then:

$$\phi\left(r_{pool}\right) = \frac{\psi_L \Lambda_L\left(r_{pool}\right)\phi\left(r_{L,fair}\right) + \psi_H \Lambda_H\left(r_{pool}\right)\phi\left(r_{H,fair}\right)}{\psi_L \Lambda_L\left(r_{pool}\right) + \psi_H \Lambda_H\left(r_{pool}\right)} + m_{pool} \tag{68}$$

The following claim characterizes the welfare effects of data availability in this setting.

Claim 2. *The change in total welfare when data is made available, in expected dollars over the term of a loan, is:*

$$\Delta Welfare = \underbrace{\psi_H \frac{b_H}{2}\left(\left(m_{pool} - s_L \Delta\right)^2 - m_H^2\right)}_{H\ group} + \underbrace{\psi_L \frac{b_L}{2}\left(\left(m_{pool} + (1 - s_L)\Delta\right)^2 - m_L^2\right)}_{L\ group} \tag{69}$$

Where:

$$s_L \equiv \frac{\psi_L \Lambda_L\left(r_{pool}\right)}{\psi_L \Lambda_L\left(r_{pool}\right) + \psi_H \Lambda_H\left(r_{pool}\right)} \tag{70}$$

is the share of loans given to low-cost borrowers, at the pooled price r_{pool},

$$\Delta \equiv \phi\left(r_{H,fair}\right) - \phi\left(r_{L,fair}\right) \tag{71}$$

is the difference in costs between the two groups, and ψ_H, ψ_L, are the expected number of non-default periods per group, defined in (12).

Expression (69) is the most general expression for the change in welfare when lenders set markups above marginal costs. For us to say anything about how data availability impacts welfare in the fully general case, we must estimate markups m_H, m_L, m_{pool} in addition to prices, quantities, and the terms ψ_L, ψ_H.

However, if markups are constant and demand elasticities are identical across groups, then equation (69) collapses to our earlier result, and we can estimate welfare changes using price and quantity data alone. To see this, first, suppose that markups are constant, across groups, and before and after data is available:

$$m_{pool} = m_H = m_L = m \tag{72}$$

This assumption is likely to approximately hold in our empirical setting. We calculate state-level Herfindahl-Hirschman indices (HHI) in our empirical setting for previously-bankrupt and never-bankrupt borrowers with similar credit scores. The HHI in a market is defined as the sum of squared market shares ($\sum_{l=1}^{N} s_l^2$), where l indexes lenders. In most models of imperfect competition, markups depend on measures of market concentration, so two markets for similar products which have similar HHI values are likely to have similar markups. We find that the state-level Herfindahl index is 0.0376 for previously-bankrupt borrowers, and 0.0330 for never-bankrupt borrowers with similar credit scores. The HHIs in both cases are low and fairly similar, supporting the assumption that markups are similar across groups in our empirical setting.

Given (72), the welfare change in (69) then simplifies to:

$$\Delta Welfare = \psi_H \frac{b_H}{2}\left(s_L^2\Delta^2 - 2ms_L\Delta\right) + \psi_L \frac{b_L}{2}\left((1-s_L)^2\Delta^2 + 2m(1-s_L)\Delta\right) \tag{73}$$

For additional intuition, note that we can write (73) as:

$$\Delta Welfare = \underbrace{\psi_H \frac{b_H}{2} s_L^2\Delta^2 + \psi_L \frac{b_L}{2}(1-s_L)^2\Delta^2}_{A} \underbrace{-\psi_H b_H m s_L\Delta + \psi_L b_L m(1-s_L)\Delta}_{H} \tag{74}$$

Suppose we set markups to zero in Expression (74). The change in welfare is then term A in (74). This term thus represents the welfare gain from data availability in competitive markets. Term A is equivalent to the sum of expression (19) of Claim 1 for the L and H groups.[18]

Term B in (74) thus captures how markups change welfare gains, relative to the competitive case. To understand term B, first we consider a special case of the result, where the two groups' demand elasticities around the pooled-pricing rate are the same. That is, assume that the slopes of demand are proportional to the size of each group:

$$\frac{b_H}{\Lambda_H\left(r_{pool}\right)} = \frac{b_L}{\Lambda_L\left(r_{pool}\right)} \tag{75}$$

If (75) holds, given the definition of s_L in (70), we have:

$$\frac{\psi_H b_H}{\psi_L b_L} = \frac{1-s_L}{s_L} \tag{76}$$

Now, under (76), term B in (74) then becomes $m\psi_L\Delta(-b_L(1-s_L) + b_L(1-s_L)) = 0$. Thus, when the two groups' demand elasticities are equal and markups pre- and post-data availability are equal, the welfare change with markups is exactly the same as if markets were competitive.

The intuition behind this result is illustrated in Figure B.1, which graphically depicts the

[18] To see this, note that when markets are competitive and data becomes available, prices for the H group increase by $s_L\Delta$, so quantities decrease by $b_H s_L\Delta$. Taking the product of the price change and the quantity change and dividing by 2, according to (19), we get $\frac{b_H}{2}s_L^2\Delta^2$. The same calculation for the L group gives $\frac{b_L}{2}(1-s_L)^2\Delta^2$.

Figure B.1: Price Discrimination with Imperfect Competition

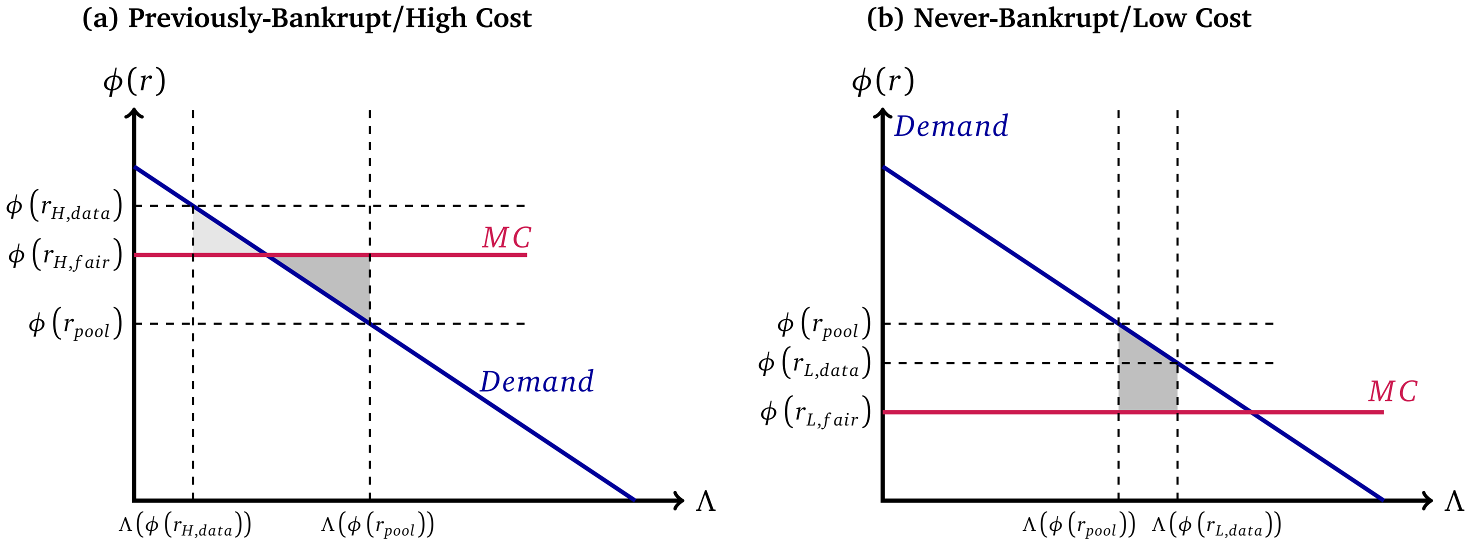

This figure illustrates how third-degree price discrimination affects welfare in credit markets when there is imperfect competition and prices may be higher than costs. Suppose that there are two groups of prospective borrowers, high-cost previously-bankrupt (panel a) and low-cost never-bankrupt (panel b). The red lines show the cost of serving each group and the blue lines show borrowers' demand curve. Lenders are initially unable to distinguish between these borrowers, so they set the price $\phi(r_{pool})$. After lenders become able to distinguish the two groups of borrowers, suppose they set $\phi(r_{H,data})$ for the high-cost group (panel a) and $\phi(r_{L,data})$ for the low-cost group (panel b). The dark gray (light gray) triangle in panel (a) shows the welfare gains (losses) for the high-cost previously-bankrupt group. The total welfare effect on the high-cost previously-bankrupt group is the difference between the size of the dark gray and light gray triangles. The dark gray area in panel (b) shows the welfare gain for the low-cost never-bankrupt group, whose prices decrease.

welfare effects of data availability in the presence of markups. The left panel of Figure B.1 shows that when there are markups, the welfare gains from *raising* prices a given amount for H borrowers are smaller, since prices are already closer to their marginal costs. This is reflected by the negative $-b_H\psi_H m s_L\Delta$ term in (74). However, the left panel shows that the welfare gains from *lowering* prices for L borrowers are larger, since prices are further above marginal costs. This is captured by the positive $b_L\psi_L m(1-s_L)\Delta$ term in (74). When markups and demand elasticities are the same across groups, these two effects exactly offset each other, so welfare gains in the case with imperfect competition are exactly the same as in the competitive case: releasing data will generally increase social welfare.

A simple intuition behind the result that data tends to improve welfare, even when markups are present, is as follows. Suppose that the prices for the H and L groups are the same, though prices may be much higher than costs. Then the willingness to pay of the marginal borrower in group H and group L are the same. Suppose we remove a small number of marginal H borrowers from the borrowing pool, and add an equal number of marginal L borrowers, so that the total loan amount across the two groups is unchanged. Since the marginal WTP is the same, total borrower utility across the two groups is unchanged. However, reallocating from H to L borrowers decreases the average social cost of serving these borrowers. Thus, social welfare must increase. This argument holds regardless of whether markets are competitive or not.

We can think of the general case in terms of how it deviates from the special case of constant markups and elasticities across the two groups. First, suppose we hold markups fixed, but relax the elasticity assumption in (75). The sign of term B in (74) depends on the relative elasticities in the two groups. When we have:

$$\frac{b_H}{\Lambda_H\left(r_{pool}\right)} > \frac{b_L}{\Lambda_L\left(r_{pool}\right)}$$

so that the elasticity in the H group is greater (smaller) than the elasticity in the L group, then the welfare gain from making data available is smaller (greater) than in the competitive case. The intuition is that when the high group has a higher demand elasticity, the decreased welfare gains from raising prices for H group borrowers tend to dominate, and vice versa.

Second, suppose we allow markups to vary before and after data availability. Note that (69) is strictly decreasing in m_H and m_L, the size of the post-data-availability markups. The intuition is simply that higher average markups are worse for social welfare. Thus, if data availability tends to increase (decrease) the level of overall markups, then social welfare will tend to decrease (increase).[19]

In summary, our results imply that if data availability does not substantially affect markups and if the demand elasticities in the two groups are similar, then data availability tends to increase welfare even when there is market power; this welfare increase arises through a similar mechanism of reallocating to lower-cost borrowers. If there were sufficient data available, one could quantify the most general expression for welfare changes, (69), by measuring pre- and post-change prices, quantities, ψ_L and ψ_H, and markups for all borrower groups.

Next, we apply the results from this extension to imperfect competition and empirically show how the efficiency ratio—the welfare cost per dollar redistributed—depends on market power in credit markets. The thought experiment considers the redistributional and welfare consequences if the observed data had been generated in a market with market power. Following the derivations in section B.5 and equation (74), we decompose the effects of data removal

[19] In the theoretical literature, third-degree price discrimination is known to have ambiguous effects on the effect of overall markups: with different demand functions in the two groups, essentially any pattern of markup increases or decreases is possible (Bergemann et al., 2015).

into the effect in the absence of market power and the consequences of data removal due to market power. Table B.1 shows that the efficiency ratio is in the range of a few cents per dollar redistributed. We compute the HHI-implied markup as 34 basis points (=9.02%$*$0.0376). For robustness, we also consider markups twice the size of HHI-implied markups. As argued above, Table B.1 illustrates that under similar markups across markets and with similar demand elasticity across H and L groups, the consequences of data removal are independent of the markup (compare columns (2) and (4)). When demand elasticities vary across groups, market power can either improve ($\frac{b_L}{\Lambda_L}/\frac{b_H}{\Lambda_H} = 0.5$) or deteriorate ($\frac{b_L}{\Lambda_L}/\frac{b_H}{\Lambda_H} = 1.5$) the efficiency ratio by a few cents. In cases with market power, data removal can even slightly increase efficiency in some cases.

Table B.1: Efficiency Ratio of Data Removal with Market Power

This table summarizes the efficiency ratio (welfare loss per dollar redistributed) of removing data under various elasticity and markup assumptions. The first three columns summarize the consequences of removing data when the observed data was generated in a market with a 34-basis-point markup. The markup is calibrated to the HHI. Columns (4) to (6) double the HHI-implied markup. The markup is assumed to be similar for H, L, and pooling market. Low Cost to High Cost Elasticity shows varying assumptions about the L group's demand elasticity in relation to the H group's. The efficiency ratio with market power includes losses due to both data removal and credit rationing – it combines the effects of A and B in equation (74). The efficiency ratio without market power only shows the effect of A in equation (74), which is the effect of data removal in the absence of market power. The efficiency ratio due to market power only shows the effect of B in equation (74), which is the additional losses (gains) from data removal due to prevailing market power.

	34bp markup			68bp markup		
Low Cost to High Cost Elasticity	0.5	1	1.5	0.5	1	1.5
Efficiency Ratio						
With Market Power	\$0.01	-\$0.03	-\$0.08	\$0.06	-\$0.03	-\$0.12
Without Market Power	-\$0.03	-\$0.03	-\$0.03	-\$0.03	-\$0.03	-\$0.03
Due to Market Power	\$0.04	\$0.00	-\$0.04	\$0.09	\$0.00	-\$0.09

B.6 Adverse Selection

In the main text, we assume that there is no selection—that is, that costs depend on borrowers' types but are not correlated with borrowers' willingness to pay. In this section, we relax this assumption and allow prices to be correlated with costs. The top panel of Figure B.2 depicts market outcomes in a case with adverse selection and competitive markets. There is a deadweight loss triangle (i.e., the light gray region in the figure) because markets reach the point where average costs are equal to the marginal borrower's willingness to pay. However, at this point, marginal costs are *below* willingness to pay.

Figure B.2: Price Discrimination with Adverse Selection

(a) Adverse selection

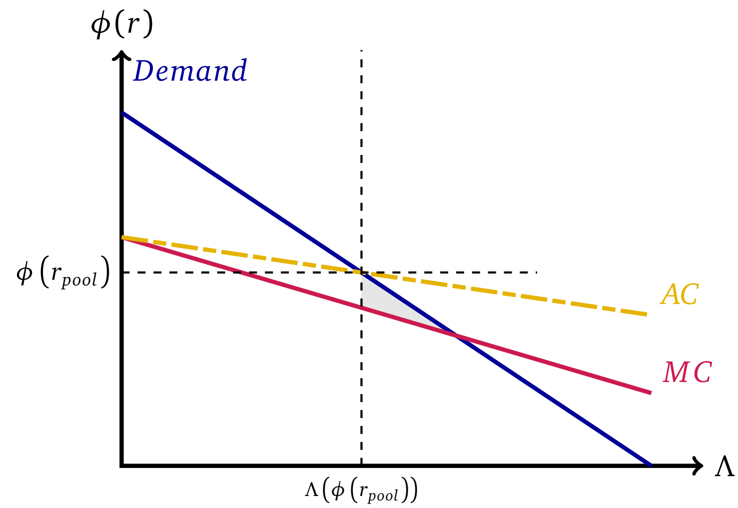

(b) Never-Bankrupt/Low Cost

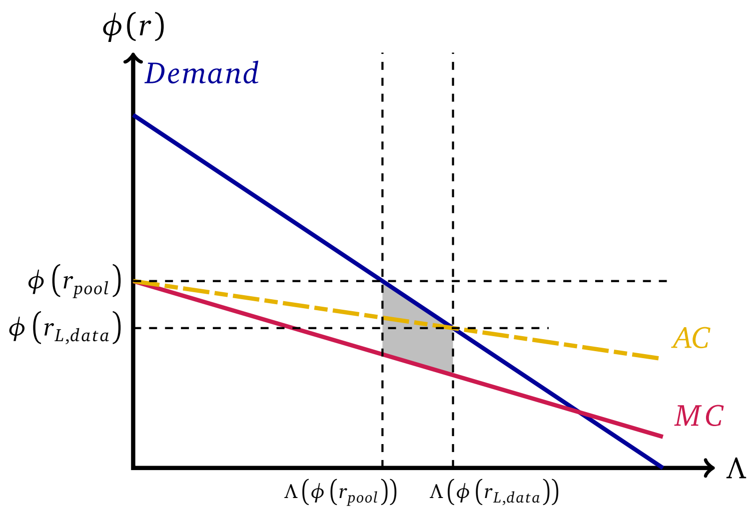

(c) Previously-Bankrupt/High Cost

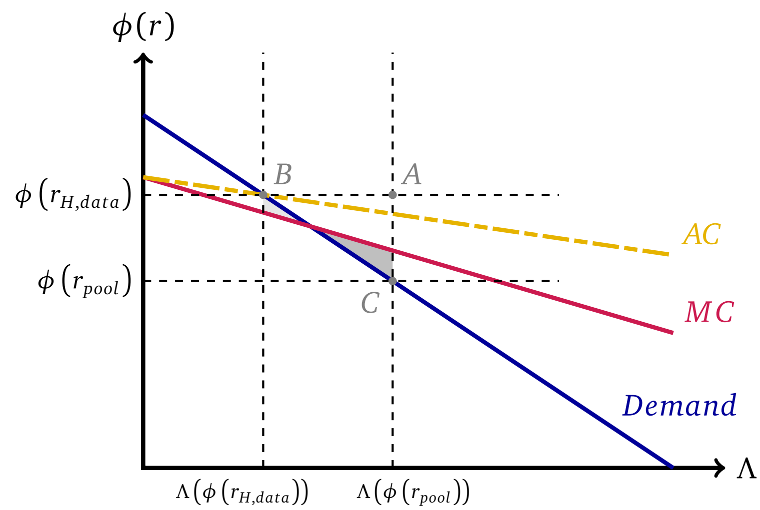

This figure illustrates how third-degree price discrimination affects welfare in credit markets in the presence of adverse selection or moral hazard. With these frictions, costs and prices vary. Panel (a) illustrates that under adverse selection, prices in competitive equilibrium will be equal to average costs. The light gray triangle shows the welfare loss, relative to the constrained optimum of setting prices equal to average costs. If data becomes available on low-cost L borrowers and high-cost H borrowers as illustrated in panels (b) and (c), prices will be set equal to average costs separately for each group in competitive equilibrium. Prices will tend to fall for the low-cost L group, as shown in panel (b), and rise for the high-cost H group, as shown in panel (c). After lenders become able to distinguish the two groups, suppose they set $\phi\left(r_{L,data}\right)$ for the low-cost L group and $\phi\left(r_{H,data}\right)$ for the high-cost H group. The dark gray area in panel (b) shows the welfare gain for the low-cost L group, whose prices decrease. The dark gray (light gray) triangle in panel (c) shows welfare gains (losses) for the high-cost H group.

Suppose that data becomes available but there is also adverse selection in each of the two submarkets. For the low-cost L group, prices decrease. When data is not available, prices are too high for these borrowers, relative to the socially efficient point, for two reasons. First, they are pooled with the high-cost H borrowers. Second, there is adverse selection. When lenders have data on these borrowers, they lower prices to the point where the average cost is equal to marginal WTP; this is the point $\Lambda(r_{L,data})$. This price reduction increases consumer surplus and social welfare, though not to the socially optimal point, as distortions from adverse selection remain. The bottom-left panel of Figure B.2 illustrates the welfare gain as the dark gray region.

For the high-cost H group, prices increase. When data is not available, prices may be too high or too low for these borrowers relative to the socially efficient price, as there are two counteracting forces. Adverse selection tends to cause prices to be too high relative to the social optimum. However, pooling with low-cost borrowers tends to make prices too low.

The bottom-right panel of Figure B.2 illustrates a case where the price without data is below the social optimum. After data is available, lenders set prices where the average cost curve crosses the demand curve. This is always above the socially optimal point, where the marginal cost curve crosses the demand curve. As a result, the effects on consumer welfare are ambiguous: there is a welfare gain as prices increase to the social optimum, represented by the dark gray triangle, but there is a welfare loss due to prices increasing further, represented by the light gray triangle.

We can also use this framework to examine how assuming no adverse selection might skew our estimates of changes in total surplus. Fixing prices $\phi\left(r_{pool}\right)$ and $\phi\left(r_{H,data}\right)$, and loan amounts $\Lambda\left(r_{H,data}\right)-\Lambda\left(r_{pool}\right)$, the calculated welfare gain from making data available is always larger if we assume there is no adverse selection, so the marginal cost curve is flat and equal to $\phi\left(r_{H,data}\right)$. To see this, note that the welfare gain that we calculate in the main text, (19) of Claim 1, corresponds to the triangular area enclosed by points A, B, and C in the bottom-right panel of Figure B.2. In contrast, the welfare gain under adverse selection is the dark gray area, which is weakly smaller than the ABC triangle, minus the light gray area.[20]

Thus, when there is adverse selection, the actual welfare gains from data availability for the H group must be even smaller than we find in the main text, whereas the change in consumer surplus is identical to expression (18). Conversely, the total welfare effects of *removing* data such as bankruptcy flags are *smaller* if there is adverse selection. Thus, our conclusion that flag removals are a quantitatively efficient way to redistribute surplus would not change significantly if there were adverse selection.[21]

[20] If adverse selection is sufficiently severe, so r_{pool} is higher than the efficient price, it is possible that the dark gray area is empty, and the light gray area is a trapezoid; it is then the case that data availability lowers aggregate welfare.

[21] We disregard the low-cost L group in this discussion; based on arguments in Appendix B.1, the welfare effects

If data were available on lenders' costs before and after lenders are able to use new data for pricing, the triangles in Figure B.2 could be quantified to calculate the welfare gains from making data available. DeFusco, Tang and Yannelis (2022) demonstrate this in the case of a single set of borrowers; to quantify the effects of data availability, the methodology in DeFusco, Tang and Yannelis (2022) could be applied to the high-cost H and low-cost L borrower groups separately. In our empirical application, however, adverse selection does not appear to be present in our setting. We conduct a test similar in spirit to Chiappori and Salanie (2000). Asymmetric information would suggest a positive correlation between rates and charge-offs: adverse selection would mean that riskier borrowers select higher interest loans, and moral hazard would mean that higher rates induce borrowers into default. This test is shown in Figure B.3. Figure B.3 is comparable to Figure 2 in the main text. However, it replaces the outcome variable with a dummy variable equal to one if a loan gets charged-off within two years of loan opening and zero if the opened loan does not get charged-off within two years of opening. In the graphical evidence, we cannot reject the null of no effect of flag removal on charge-offs. As charge-offs do not appear to decrease when prices decrease, we do not find evidence of adverse selection. We also quantify the visual result in a regression framework and present the results in Table B.2. The table shows that flag removal is associated with an insignificant increase in charge-offs.[22]

for the L group will tend to be much smaller, since the group is larger, and thus the change in prices is smaller.

[22]We do not mean to argue that adverse selection and moral hazard are not concerns in auto lending settings. Indeed, some studies, including Adams et al. (2009), have found that adverse selection and moral hazard are important and were particularly so in the period prior to the Great Financial Crisis. More recent papers, including Argyle et al. (2020), find no effects, which is consistent with auto loans being highly collateralized. In our setting, we do not find any effects of flag removal in particular on default rates, which suggests that there is no asymmetric information.

Figure B.3: Charge-offs

This figure shows estimates of the coefficients δ_t from the following specification $y_{it} = \gamma_c + \gamma_{ts} + \sum_{t=-6}^{6} \delta_t \{e_{it} = t\} + \beta X_{it} + \varepsilon_{it}$, along with a 95% confidence interval. The outcome y_{it} is charge-offs. γ_c are cohort fixed effects, and γ_{ts} are time period by score bucket fixed effects. Standard errors are clustered at the cohort level. Source: TransUnion.

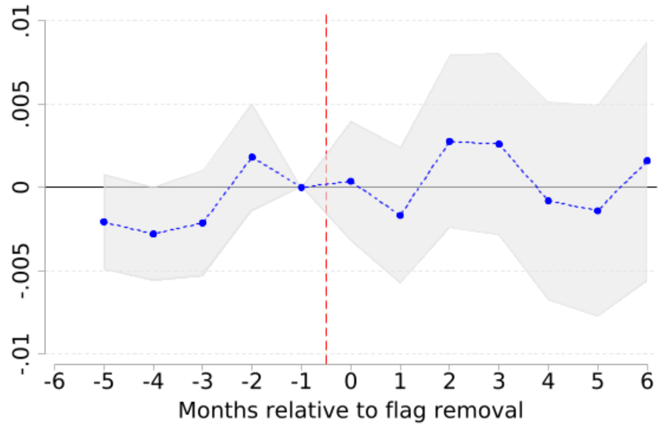

Table B.2: Charge-offs Around Flag Removal

This table shows estimates of the coefficients δ_y from the following specification $y_{it} = \gamma_c + \gamma_t + \delta^y \mathbb{1}[FlagRemoved] + \beta X_{it} + \varepsilon_{it}$. The outcome y_{it} is charge-offs. γ_c are cohort fixed effects, and γ_t are time period fixed effects. Standard errors are clustered at the cohort level. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Source: TransUnion.

	(1)	(2)	(3)	(4)	(5)
$\mathbb{1}$[FlagRemoved]	0.001	0.001	0.001	0.001	0.000
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Constant	0.024^{***}	0.024^{***}	0.024^{***}	0.023^{***}	0.023^{***}
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	176690	176690	176690	176690	176686
Adjusted R^2	0.000	0.001	0.001	0.001	0.030
Linear Time Trend	Yes	Yes	Yes	No	No
Year-month FE	No	Yes	No	Yes	No
Cohort FE	No	No	Yes	Yes	Yes
Year-month by Score Bucket FE	No	No	No	No	Yes
Clustered SE	Cohort	Cohort	Cohort	Cohort	Cohort

B.7 Strategic Bankruptcy and Incentive Effects of Flag Removals

Bankruptcy flags on credit reports provide information to the market about borrowers' default risks, increasing the efficiency of credit allocation. While only 39.6% of bankrupt individuals know the duration of bankruptcy flags on their credit file,[23] bankruptcy flags may still have an incentive effect: borrowers who declare bankruptcy face higher interest rates in the future, creating a disincentive to declare bankruptcy. If policymakers force credit reporting agencies to remove the bankruptcy flag from the credit report, then this would also affect the bankruptcy incentives facing borrowers, which in turn increases bankruptcy rates. If bankruptcies decrease social welfare on the margin, the incentive effect is important to account for in a full welfare accounting of the effects of flag removals. While the main focus of the paper is on the allocative effects of flag removal, in this appendix, we construct a simple model to illustrate how to evaluate the effects of consumer bankruptcy on social welfare, when flag removal affects the bankruptcy filing decision. Using the model, we then calculate the incentive effects for borrowers of the flag removal, and show that these incentive effects can be much larger than the allocative effects.

B.7.1 Model of Bankruptcy Decisions and Welfare

We consider a two-stage game. The second stage, which we call the "downstream" market, is identical to the model in the main text. Previously-bankrupt borrowers face some cost r of getting credit, which may be affected by policies such as bankruptcy flag removal. We add a first stage, in which borrowers have some heterogeneous cost $c \sim F(\cdot)$ of declaring bankruptcy. Costs may differ because borrowers have different subjective valuations of bankruptcy. In the first stage, borrowers can choose whether to declare bankruptcy, or not declare bankruptcy and receive value V_{NB}. There are three kinds of agents: borrowers, "downstream" lenders who lend to borrowers in the second stage, and "upstream" lenders who have outstanding loans to borrowers at the point where they can declare bankruptcy. We will separately characterize the surplus of each kind of agent, then add these terms to analyze social welfare.

First, we analyze the downstream market. For simplicity, suppose all previously-bankrupt borrowers have the same default rate, and thus the cost of serving previously-bankrupt borrowers is some constant $\phi\left(r_{fair}\right)$, where r_{fair} is the break-even interest rate for these borrowers. We focus on high-cost borrowers; by arguments analogous to the main text, the welfare effects for never-bankrupt borrowers will be small, since the never-bankrupt group is much larger than the previously-bankrupt group. Similar to the main text, let $L^*(r;c)$ denote the loan demand of type c of the prospective borrower, at interest rate r. We normalize the marginal value of a

[23] See Table 6 of Gross et al. (2020)

dollar in each future period to 1. Now, the total amount of loans made at price r is:

$$\int_{c\leq\bar{c}(r)} L^*(r;c)\,dF(c)$$

where $\bar{c}(r)$ reflects the fact that the opt-out condition depends on r.

Let V_{NB} represent the value from not declaring bankruptcy; for simplicity, assume this is a constant for all borrowers, though this can be relaxed without affecting the results. In the second stage, years after a borrower has declared bankruptcy, she faces some price r for loans. To calculate surplus in lending markets, note that the surplus of a borrower with type c is:

$$V_B(c)=\underbrace{\int_{\hat{r}=r}^{\bar{r}(c)} L^*(\hat{r};c)\,d\phi(\hat{r})-c}_{Downstream} \tag{77}$$

where $\bar{r}(c)$ is the maximum rate at which a consumer of type c borrows positive amounts. Expression (77) is just the surplus from the main text adjusted by the fixed cost of bankruptcy c. Taking into account the fixed cost c of bankruptcy, a borrower with cost c has a value of declaring bankruptcy:

$$\int_{\hat{r}=r}^{\bar{r}(c)} L^*(\hat{r};c)\,d\phi(\hat{r})-c \tag{78}$$

A borrower with bankruptcy cost c optimally declares bankruptcy if (78) is greater than the value of not declaring bankruptcy, V_{NB}. We can thus define a function $\bar{c}(r)$ as the marginal bankrupt borrower, given the downstream rate r:

$$\bar{c}(r)=\left\{c:\ V_{NB}=\int_{\hat{r}=r}^{\bar{r}(c)} L^*(\hat{r};c)\,d\phi(\hat{r})-c\right\} \tag{79}$$

All borrowers with $c\leq\bar{c}(r)$ declare bankruptcy, and all borrowers with $c>\bar{c}(r)$ do not. From (79), $\bar{c}(r)$ is always decreasing in r: the higher the rate post-bankruptcy, the lower consumer surplus in the post-bankruptcy market, and thus the less types c will declare bankruptcy. The function $\bar{c}(r)$ thus captures the elasticity of the bankruptcy decision to the post-bankruptcy interest rate r.

Consumer surplus. Integrating over all consumers with different bankruptcy costs c, consumers' surplus is thus:

$$CS=\underbrace{\int_{c>\bar{c}(r)} V_{NB}\,dF(c)}_{No\ bankruptcy}+\underbrace{\int_{c\leq\bar{c}(r)}\left[\int_{\hat{r}=r}^{\bar{r}(c)} L^*(\hat{r};c)\,d\phi(\hat{r})-c\right]dF(c)}_{Bankruptcy} \tag{80}$$

We wish to characterize how consumer surplus changes as we shift r, the interest rate facing previously-bankrupt borrowers in lending markets. Differentiating (80) with respect to r, we have:

$$\frac{\partial CS}{\partial r} = \underbrace{-\int_{c \le \bar{c}(r)} L^*(r;c)\,\phi'(r)\,dF(c) +}_{Term\ 1} \underbrace{\bar{c}'(r)\,f(\bar{c})\left(\int_{\hat{r}=r}^{\bar{r}(c)} L^*(\hat{r};\bar{c}(r))\,d\phi(\hat{r}) - \bar{c}(r) - V_{NB}\right)}_{Term\ 2} \tag{81}$$

Now, from the definition of $\bar{c}(r)$ in (79), the rightmost piece of term 2 is 0; thus, we have:

$$\frac{\partial CS}{\partial r} = -\int_{c \le \bar{c}(r)} L^*(r;c)\,\phi'(r)\,dF(c) \tag{82}$$

In words, (82) states that the derivative of total consumer surplus with respect to r is the standard envelope formula: it is the change in payments, $\phi'(r)$, multiplied by loan size $L^*(r;c)$, integrated over all consumers. Changing r also changes the set of consumers that declare bankruptcy. However, the marginal consumers are indifferent between declaring bankruptcy and not doing so, hence there is no first-order welfare effect of moving these consumers into or out of bankruptcy.

Downstream lender profits. As in (14), profits of downstream lenders, who lend to previously-bankrupt borrowers, are simply demand minus costs:

$$\Pi_D = \int_{c \le \bar{c}(r)} L^*(r;c)\left(\phi(r) - \phi\left(r_{fair}\right)\right) dF(c) \tag{83}$$

Differentiating (83) with respect to r, we have:

$$\frac{\partial \Pi_D}{\partial r} = \underbrace{\int_{c \le \bar{c}(r)} L^*(r;c)\,\phi'(r)\,dF(c) +}_{Term\ 1} \underbrace{\int_{c \le \bar{c}(r)} \frac{\partial L^*(r;c)}{\partial r}\left(\phi(r) - \phi\left(r_{fair}\right)\right) dF(c)}_{Term\ 2} + \underbrace{\bar{c}'(r)\,f(\bar{c}(r))\,L^*(r;c)\left(\phi(r) - \phi\left(r_{fair}\right)\right)}_{Term\ 3} \tag{84}$$

In (84), term 1, which is exactly the negative of $\frac{\partial CS}{\partial r}$ in (82), reflects the fact that, when rates increase, welfare is transferred from borrowers to downstream lenders. Term 2 is the marginal change in the deadweight loss triangle, as r increases: it is the height of the deadweight loss triangle, $\left(\phi(r)-\phi\left(r_{fair}\right)\right)$, multiplied by the change in loan amount, $\frac{\partial L^*(r;c)}{\partial r}$. Both these terms are also present in the baseline model, where there is no bankruptcy margin. Term 3 is novel to the setting where bankruptcy decisions are elastic. When $r \neq r_{fair}$ in downstream markets, the marginal consumer's decision to declare bankruptcy imposes an externality on downstream lenders, of size:

$$L^*(r;c)\left(\phi(r)-\phi\left(r_{fair}\right)\right) \tag{85}$$

For example, if downstream lenders lose money on previously-bankrupt consumers, so $\phi(r) < \phi\left(r_{fair}\right)$, then the marginal consumer who declares bankruptcy imposes a negative externality on lenders. The size of the effect depends on the size of the negative externality, (85), multiplied by the measure of marginal consumers, $\bar{c}'(r) f(\bar{c}(r))$

Upstream lender profits. Suppose that type c consumers, at the time that they declare bankruptcy, have some outstanding debt $D(c)$ with upstream lenders. Suppose that their decision to declare bankruptcy causes lenders to lose a fraction ψ of the debt. Upstream lenders' welfare, as a function of r, is thus:

$$\Pi_U = \int_{c \leq \bar{c}(r)} -\psi D(c)\, dF(c) \tag{86}$$

That is, upstream lenders lose $\psi D(c)$ on all consumer types that default, $c \leq \bar{c}(r)$. Differentiating with respect to r, we have:

$$\frac{\partial \Pi_U}{\partial r} = -\bar{c}'(r) f(\bar{c}(r))\, \psi D(\bar{c}(r)) \tag{87}$$

In words, decreasing r slightly causes a measure $-\bar{c}'(r) f(\bar{c}(r))$ of marginal consumers with type $\bar{c}(r)$ to declare bankruptcy (note that $\bar{c}'(r)$ is negative). This decreases upstream lenders' profits by the losses on their loans for these consumers, which is $\psi D(\bar{c}(r))$.

Now, total social welfare is just the sum of the welfare of consumers, upstream producer profits, and downstream producer profits. To find the effect of a small change in r on total welfare, we sum (82), (84), and (87). Consumer surplus (82) cancels with term 1 in downstream lenders' profits (84), so we get:

$$\begin{aligned}\frac{\partial CS}{\partial r} + \frac{\partial \Pi_D}{\partial r} + \frac{\partial \Pi_U}{\partial r} = & \underbrace{\int_{c \leq \bar{c}(r)} \frac{\partial L^*(r;c)}{\partial r}\left(\phi(r)-\phi\left(r_{fair}\right)\right) dF(c)}_{Deadweight\ Loss} \\ & + \underbrace{\bar{c}'(r) f(\bar{c}(r)) L^*(r;c)\left(\phi(r)-\phi\left(r_{fair}\right)\right)}_{Downstream\ externality} - \underbrace{\bar{c}'(r) f(\bar{c}(r))\, \psi D(\bar{c}(r))}_{Upstream\ externality}\end{aligned} \tag{88}$$

In (88), the first term is the marginal change in the size of the deadweight loss triangle, which is analogous to the depiction in Figure 1. When there is no incentive effect of bankruptcy, so $\bar{c}'(r) = 0$ and bankruptcy decisions are perfectly inelastic, then the change in welfare is simply the change in the deadweight loss triangle. When bankruptcy is elastic, there are two additional terms: the externality on downstream lenders, which is term 3 in (84), and the externality on upstream lenders, which is (87). The upstream externality term will always be positive (that is, decreasing rates will tend to lower welfare), since bankruptcies create negative externalities on upstream lenders. When $\phi(r) < \phi\left(r_{fair}\right)$, so prices are lower than marginal costs for previously-bankrupt borrowers, and the downstream externality term is also positive, so decreasing rates will tend to lower social welfare. Thus, if bankruptcy decisions are sensitive to rates in downstream markets, there are two additional forces causing lower rates to tend to decrease social welfare.

B.7.2 Estimate of Welfare Costs with Elastic Bankruptcy

Next, we do a back-of-the-envelope calculation of how large the incentive effects of bankruptcy on welfare could be in the data. There are a variety of estimates in the literature on how strategic borrowers are in their decisions to default on loans and declare bankruptcy. Yannelis (2016) provides evidence for strategic default on student loans, showing that introducing bankruptcy protection for student loans would increase loan default by 18%, and increasing garnishable income by $10,000 would lead to a 15% decrease in defaults. Mayer et al. (2014) argue that a legal settlement offering modifications to delinquent borrowers increased delinquency rates by 10%. Argyle et al. (2021) find that borrowers with increased cash flows tend to delay filing for bankruptcy.

It is difficult to extrapolate the effect of a particular policy, i.e., bankruptcy flag removal, on strategic bankruptcies. The lower bound of the incentive effect in the literature—that borrowers are completely non-strategic in their default and bankruptcy decisions—would imply that there is no incentive effect of flag removal on bankruptcy. Indeed, only 9.2% of Chapter 7 filers correctly guess the number of years remaining for their flag (Gross et al., 2020, Table 6), suggesting limited strategic implications of flag removal. Nevertheless, to gauge the size of the incentive effect, we do a quick back-of-the-envelope calculation: suppose that flag removal, relative to keeping bankruptcy flags on borrowers' credit records indefinitely, would increase bankruptcy filing rates by 1%. With approximately 800,000 bankruptcy filings annually, this implies 8,000 additional bankruptcies.

According to our model, we must estimate two numbers. For upstream lenders, we must estimate the effect of bankruptcy filing on lenders' losses. Since our main analysis focuses on auto loans, we also consider losses to auto lenders. At the time of bankruptcy, the average borrower has $19,865 in auto loan debt.[24] In most states, borrowers lose their cars in bankruptcy; thus, we will assume a loss rate of 40%.[25] Hence, the loss per loan is $7,946 (=0.4*19,865) or $3,575.7 per consumer.[26] With approximately 800,000 filers per year, the aggregate loss to lenders would be $28.6 million (=0.01*800,000*3,575).

For downstream lenders, we must estimate their losses on each customer, multiplied by the set of marginal consumers. From our baseline estimates in Table 2, the average loss per customer of lenders in downstream markets is $24.42 over a five year loan term or $19.5 million for 800,000 bankruptcy filers every year. Multiplying by the 1% increase in bankruptcy filings, this would create an additional loss of $0.195 million to social surplus.

Adding the effects on upstream and downstream lenders, we get a total effect of $28.8 million. This quantity is large relative to the allocative welfare quantities we calculated in the main text. Accounting for a 1% increase in bankruptcies due to incentive effects, flag removal transfers $19 million to previously-bankrupt consumers, at the cost of $29.4 million (=$28.8 million+$0.6 million) in social welfare.

Why are the welfare effects through the incentive channel so large, relative to the allocative effects? Intuitively, we showed in the main text that, because competitive lending markets lead to efficient credit allocations, the removal of small amounts of data has only a second-order effect on the allocative efficiency of lending. This does not apply to the effects of data removal on bankruptcy incentives. Borrowers do not internalize the costs to lenders of their bankruptcy decisions, so they default more than the socially optimal level. Data removal can affect borrowers' incentives to declare bankruptcy, and this generally has a first-order effect on social welfare. From a policy perspective, however, note that increased incentives from flag removal could be offset by increasing the cost of bankruptcy, e.g., through more stringent repayment plans or lower thresholds for asset protection in liquidation.

B.8 Varying Signal Informativeness

Different kinds of data may be differentially informative about customers' default rates, and thus the costs of lending to these customers. In this appendix, we show that, when data is more informative about default rates, the social welfare losses from data removal tend to be larger relative to the surplus transfers, so data removal is less efficient as a tool for transferring surplus. However, plugging in our demand elasticity estimate from the main text, we show that flag removal would remain a quantitatively efficient way to transfer surplus to previously-

[24] See Experian's auto loan debt study.

[25] See American Banker.

[26] Table 1 of Dobbie et al. (2017) indicates that approximately 45% of consumers have an auto loan when filing for bankruptcy.

bankrupt individuals, even if bankruptcy flags were much more informative about default rates than we find in our analysis. As in the main text, we assume demand to be linear in the payment $\phi(r)$ for each group:

$$\Lambda(r) = a - b\phi(r) \tag{89}$$

Let the data removal-induced change in interest rates be:

$$t = \phi(r) - \phi\left(r_{H,fair}\right) \tag{90}$$

Adjusting Claim 1 to a data *removal*, H group consumer surplus increases by

$$\psi_H\left(\phi\left(r_{H,fair}\right) - \phi(r)\right)\left(\frac{\Lambda_H(r) + \Lambda_H\left(r_{H,fair}\right)}{2}\right) \tag{91}$$

when removing the data. The high-cost H group gains as they are charged the lower pooling price. However, removing the data also induces an efficiency loss due to the over credit provision to the H group. Adjusting Claim 1, removing data social welfare decreases by:

$$-\frac{1}{2}\psi_H\left(\phi\left(r_{H,fair}\right) - \phi(r)\right)\left(\Lambda_H(r) - \Lambda_H\left(r_{H,fair}\right)\right) \tag{92}$$

In addition, removing data increases prices for low-cost L consumers leading to under credit provision for L consumers. The efficiency loss of removing data due to under credit provision to L consumers is:

$$-\frac{1}{2}\psi_L\left(\phi\left(r_{L,fair}\right) - \phi(r)\right)\left(\Lambda_L(r) - \Lambda_L\left(r_{L,fair}\right)\right) \tag{93}$$

Dividing efficiency changes by redistributive consequences, we obtain the efficiency ratio, that is, the welfare cost per dollar redistributed to high-cost individuals:

$$\text{Efficiency Ratio} = \frac{-\frac{1}{2}\psi_H\left(\phi\left(r_{H,fair}\right)-\phi(r)\right)\left(\Lambda_H(r)-\Lambda_H\left(r_{H,fair}\right)\right)-\frac{1}{2}\psi_L\left(\phi\left(r_{L,fair}\right)-\phi(r)\right)\left(\Lambda_L(r)-\Lambda_L\left(r_{L,fair}\right)\right)}{\psi_H\left(\phi\left(r_{H,fair}\right)-\phi(r)\right)\left(\frac{\Lambda_H(r)+\Lambda_H\left(r_{H,fair}\right)}{2}\right)} \tag{94}$$

Plugging linear demand (89) into the efficiency ratio (94), exploiting the zero-profit condition of the competitive equilibrium, and writing the pooling price as sum of the fair price and price distortion, we obtain:

$$\text{Efficiency Ratio} = \frac{1}{2}\left(\epsilon_H + \frac{\Lambda_H(r)}{\Lambda_L(r)}\frac{\psi_H}{\psi_L}\epsilon_L\right)\frac{t}{\frac{\Lambda_H\left(r_{H,fair}\right)}{\Lambda_H(r)} - \frac{1}{2}\epsilon_H t} \tag{95}$$

where, ϵ_H and ϵ_L are respectively the demand elasticities in the H and L groups at r, that is:

$$\epsilon_H \equiv \frac{b_H}{\Lambda_H(r)}, \epsilon_L \equiv \frac{b_L}{\Lambda_L(r)}$$

Taking the derivative with respect to the data induced price distortion:

$$\frac{\partial \text{Efficiency Ratio}}{\partial t} = \frac{1}{2}\left(\epsilon_H + \frac{\Lambda_H(r)}{\Lambda_L(r)}\frac{\psi_H}{\psi_L}\epsilon_L\right)\frac{\frac{\Lambda_H(r_{H,fair})}{\Lambda_H(r)}}{\left(\frac{\Lambda_H(r_{H,fair})}{\Lambda_H(r)} - \frac{1}{2}\epsilon_H t\right)^2} \tag{96}$$

When t is small, expression (96) is approximately constant. Thus, the efficiency ratio will change roughly linearly in the price change t induced by flag removal, for relatively small values of t. Intuitively, this is because the social welfare loss is a triangle, which is quadratic in the size of the deviation of prices from their efficient level, whereas the transfer is a trapezoid. The ratio of the two is therefore larger when data is more informative, leading to larger price changes upon its removal.

We can bring this theory to our data, to evaluate how efficient flag removals would be as a tool for transferring surplus, in a counterfactual scenario where bankruptcy flags were more informative about default rates, so their removal decreased interest rates more. Under our baseline estimates, flag removal decreases interest rates by 0.226%. We consider counterfactual scenarios in which flag removal induces a rate change two, four, eight, and sixteen times larger. We can then plug these changes into our expressions for welfare changes, surplus transferred, and the efficiency ratio, holding fixed the demand elasticity at our estimate in the main text, and evaluate how efficiency ratios would change if bankruptcy flags were more informative about default rates.

The results of this exercise are shown in Table B.3. As expected, consumer surplus changes approximately linearly with the induced change in interest rates. Welfare changes vary approximately quadratically with price distortions. Thus, the efficiency ratio changes approximately linearly with the price distortion. As we double and quadruple the price effect of bankruptcy flag removals from 22.6bps reductions to 45.2bps and 90.4bps reductions, we can see that the efficiency cost per dollar redistributed to previously-bankrupt H individuals doubles and approximately quadruples from 3 cents to 6 cents and 12 cents, respectively. Thus, even when flag removals are fairly strong signals of default rates—4 or 8 times more informative than we find in the main text—flag removal remains a relatively low-cost way to redistribute surplus, costing less than \$0.21 in social surplus per dollar transferred between groups. However, the efficiency ratio deteriorates for larger rate changes.

Table B.3: Welfare cost by price effect size

This table summarizes consumer surplus changes, welfare consequences, and efficiency ratios for the average five-year loan, under counterfactual scenarios in which bankruptcy flag removal is increasingly informative about costs. We consider scenarios in which flag removal induces a rate change equal to our baseline estimate of 0.226%, and then two, four, eight, and sixteen times higher. "Multiple of Effect" shows the x fold of the true price change. "Induced Rate Change" is the counterfactual interest rate variation induced by the flag removal (in percentage points). The true effect size is shown in the first row. Counterfactual price changes are depicted in rows two to five. "Consumer Surplus Redistribution" is the $ change in consumer surplus per individual, and "Welfare Change" depicts the change in social welfare per individual. Both variables are in units of expected dollars per individual, over the course of a five-year loan. "Efficiency Ratio" is the ratio of the welfare change to consumer surplus redistributed; that is, the dollars of social surplus lost, per dollar redistributed to bankrupt individuals.

Multiple of Effect	Induced Rate Change	Consumer Surplus Redistribution	Welfare Change	Efficiency Ratio
1x	-0.226%	23.75	-0.75	-0.0315
2x	-0.452%	48.76	-2.98	-0.0612
4x	-0.904%	102.51	-11.86	-0.1156
8x	-1.808%	224.56	-46.82	-0.2085
16x	-3.616%	524.04	-182.52	-0.3483

B.9 Efficiency Ratios with Quantity Data Only

The baseline model only requires price variation and quantity data to estimate the efficiency and redistributional consequences of data availability. However, researchers cannot always observe price data. In this section, we show that under our baseline assumptions a key quantity of our model—the efficiency loss per USD redistributed—can be computed with quantity data and variation in data availability only.

First, from Assumption 1 and equation (20), we can express the price increase due to data deletion for the low-cost types as:

$$\phi(r) - \phi\left(r_{L,fair}\right) = \frac{\psi_H}{\psi_L}\frac{\Lambda_H(r)}{\Lambda_L(r)}\left(\phi\left(r_{H,fair}\right) - \phi(r)\right) \tag{97}$$

Second, note that we can write the price change underlying the quantity changes for the high-cost types as a function of those quantity changes and the demand elasticity:

$$\Lambda_H\left(r_{H,fair}\right) - \Lambda_H(r) = a_H - b_H\phi\left(r_{H,fair}\right) - a_H + b_H\phi(r) = b_H(\phi(r) - \phi\left(r_{H,fair}\right))$$

Divide by the pooling quantity for high-cost types:

$$\frac{\Lambda_H\left(r_{H,fair}\right) - \Lambda_H(r)}{\Lambda_H(r)} = \frac{b_H}{\Lambda_H(r)}(\phi(r) - \phi\left(r_{H,fair}\right)) = \epsilon_H * (\phi(r) - \phi\left(r_{H,fair}\right))$$

Rearrange for the quantity implied price change:

$$\frac{\Lambda_H\left(r_{H,fair}\right) - \Lambda_H(r)}{\Lambda_H(r)} * \frac{1}{\epsilon_H} = \phi(r) - \phi\left(r_{H,fair}\right)$$

Third, we can write the quantity change for the low-cost types as a function of high-cost type quantities and price elasticities:

$$\Lambda_L\left(r_{L,fair}\right)-\Lambda_L(r)=a_L-b_L\phi\left(r_{L,fair}\right)-a_L+b_L\phi(r)=b_L(\phi(r)-\phi\left(r_{L,fair}\right))$$

Plug in for the loss implied price change and use the price change calculated for the high type as a function of elasticity:

$$\begin{aligned}\Lambda_L\left(r_{L,fair}\right)-\Lambda_L(r)&=b_L(\frac{\psi_H}{\psi_L}\frac{\Lambda_H(r)}{\Lambda_L(r)}\left(\phi\left(r_{H,fair}\right)-\phi(r)\right))\\&=\epsilon_L\frac{\psi_H}{\psi_L}\Lambda_H(r)\left(\phi\left(r_{H,fair}\right)-\phi(r)\right))\\&=\epsilon_L\frac{\psi_H}{\psi_L}\Lambda_H(r)\frac{\Lambda_H(r)-\Lambda_H\left(r_{H,fair}\right)}{\Lambda_H(r)}*\frac{1}{\epsilon_H}\end{aligned}$$

With those three expressions:

$$\phi(r)-\phi\left(r_{L,fair}\right)=\frac{\psi_H}{\psi_L}\frac{\Lambda_H(r)}{\Lambda_L(r)}\left(\phi\left(r_{H,fair}\right)-\phi(r)\right)$$

$$\frac{\Lambda_H\left(r_{H,fair}\right)-\Lambda_H(r)}{\Lambda_H\left(r_{H,fair}\right)}*\frac{1}{\epsilon_H}=\phi(r)-\phi\left(r_{H,fair}\right)$$

$$\Lambda_L\left(r_{L,fair}\right)-\Lambda_L(r)=\frac{\epsilon_L}{\epsilon_H}\frac{\psi_H}{\psi_L}\frac{\Lambda_H(r)}{\Lambda_H(r)}(\Lambda_H(r)-\Lambda_H\left(r_{H,fair}\right))$$

We can rewrite the efficiency ratio as a function of quantities only.

$$\text{Efficiency Ratio}=\frac{-\frac{1}{2}\psi_H\left(\phi\left(r_{H,fair}\right)-\phi(r)\right)\left(\Lambda_H(r)-\Lambda_H\left(r_{H,fair}\right)\right)-\frac{1}{2}\psi_L\left(\phi\left(r_{L,fair}\right)-\phi(r)\right)\left(\Lambda_L(r)-\Lambda_L\left(r_{L,fair}\right)\right)}{\psi_H\left(\phi\left(r_{H,fair}\right)-\phi(r)\right)\left(\frac{\Lambda_H(r)+\Lambda_H\left(r_{H,fair}\right)}{2}\right)}\tag{98}$$

Replace the low-cost price change:

$$=\frac{-\frac{1}{2}\psi_H\left(\phi\left(r_{H,fair}\right)-\phi(r)\right)\left(\Lambda_H(r)-\Lambda_H\left(r_{H,fair}\right)\right)+\frac{1}{2}\psi_L\left(\frac{\psi_H}{\psi_L}\frac{\Lambda_H(r)}{\Lambda_L(r)}\left(\phi\left(r_{H,fair}\right)-\phi(r)\right)\right)\left(\Lambda_L(r)-\Lambda_L\left(r_{L,fair}\right)\right)}{\psi_H\left(\phi\left(r_{H,fair}\right)-\phi(r)\right)\left(\frac{\Lambda_H(r)+\Lambda_H\left(r_{H,fair}\right)}{2}\right)}$$

Cancel the price change:

$$=\frac{-\psi_H\left(\Lambda_H(r)-\Lambda_H\left(r_{H,fair}\right)\right)+\psi_L\left(\frac{\psi_H}{\psi_L}\frac{\Lambda_H(r)}{\Lambda_L(r)}\right)\left(\Lambda_L(r)-\Lambda_L\left(r_{L,fair}\right)\right)}{\psi_H\left(\Lambda_H(r)+\Lambda_H\left(r_{H,fair}\right)\right)}$$

Plugging in the low-cost type quantity change:

$$= \frac{-\psi_H\left(\Lambda_H(r) - \Lambda_H\left(r_{H,fair}\right)\right) - \psi_L\left(\frac{\psi_H}{\psi_L}\frac{\Lambda_H(r)}{\Lambda_L(r)}\right)\frac{\epsilon_L}{\epsilon_H}\frac{\psi_H}{\psi_L}\frac{\Lambda_H(r)}{\Lambda_H(r)}\left(\Lambda_H(r) - \Lambda_H\left(r_{H,fair}\right)\right)}{\psi_H\left(\Lambda_H(r) + \Lambda_H\left(r_{H,fair}\right)\right)}$$

$$= -\frac{\left(\Lambda_H(r) - \Lambda_H\left(r_{H,fair}\right)\right) + \left(\frac{\Lambda_H(r)}{\Lambda_L(r)}\right)\frac{\epsilon_L}{\epsilon_H}\frac{\psi_H}{\psi_L}\left(\Lambda_H(r) - \Lambda_H\left(r_{H,fair}\right)\right)}{\left(\Lambda_H(r) + \Lambda_H\left(r_{H,fair}\right)\right)}$$

and rearranging:

$$\text{Efficiency ratio}_{\text{High and Low Cost}} = -\left(1 + \frac{\epsilon_L}{\epsilon_H}\frac{\psi_H}{\psi_L}\frac{\Lambda_H(r)}{\Lambda_L(r)}\right)\underbrace{\frac{\Lambda_H(r) - \Lambda_H\left(r_{H,fair}\right)}{\Lambda_H(r) + \Lambda_H(r_{H,fair})}}_{\text{Efficiency ratio}_{\text{High Cost}}} \tag{99}$$

The last expression shows that if high- and low-cost demand elasticities are the same, we can write the efficiency ratio as a function of quantities only (assuming $\epsilon_L = \epsilon_H$).

$$\text{Efficiency ratio}_{\text{High and Low Cost}} = -\left(1 + \frac{\psi_H}{\psi_L}\frac{\Lambda_H(r)}{\Lambda_L(r)}\right)\underbrace{\frac{\Lambda_H(r) - \Lambda_H\left(r_{H,fair}\right)}{\Lambda_H(r) + \Lambda_H(r_{H,fair})}}_{\text{Efficiency ratio}_{\text{High Cost}}} \tag{100}$$

Furthermore, we can bound the efficiency ratio with quantity data only if we have some theory-guided intuition for the maximum low-cost type demand elasticity and the minimum demand elasticity for the high-cost type.

C Flag Removals and Observed Interest Rates

In the main text, we show the effect of flag removals on interest rates. This section shows that the effect of bankruptcy flag removals is qualitatively and quantitatively similar using observed interest rates only. Table C.1 shows variations of the main specification with observed interest rates as the outcome variable. Note that the effect size is of comparable magnitude to the effect on predicted interest rates in Table 2 in the main text. One potential concern with our estimates of the effect of flag removal on interest rates is that we only observe interest rates for loan offers that were actually taken up by customers, whereas our model is about offered rates. This could introduce downwards bias in our estimates of effects on interest rates, if customers with different characteristics are offered different rates, but customers who receive higher rate offers are less likely to accept. However, if unobserved heterogeneity between customers affects loan rates in this manner, then controlling for *observable* heterogeneity between customers should also affect our coefficient estimates (Oster, 2019). Our estimates of effects on interest rates in columns 1 to 5 of Table C.1 are very stable, suggesting that this bias is not likely to be quantitatively important.

The graphical evidence in Figure C.1 confirms the findings of Table C.1. The removal of bankruptcy flags reduces interest rates by approximately 20bps.

Table C.1: Interest Rates Around Flag Removal

This table shows estimates of the coefficients δ_y from the following specification $y_{it} = \gamma_c + \gamma_t + \delta^y \mathbb{1}[FlagRemoved] + \beta X_{it} + \varepsilon_{it}$. The outcome y_{it} is observed interest rates. γ_c are cohort fixed effects, and γ_t are time period fixed effects. Standard errors are clustered at the cohort level. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Source: TransUnion.

	(1)	(2)	(3)	(4)	(5)
$\mathbb{1}$[FlagRemoved]	-0.192***	-0.202***	-0.183***	-0.198***	-0.169***
	(0.066)	(0.060)	(0.064)	(0.059)	(0.050)
Constant	8.160***	8.131***	8.176***	7.799***	7.782***
	(0.053)	(0.033)	(0.029)	(0.033)	(0.028)
Observations	176690	176690	176690	176690	176686
Adjusted R^2	0.003	0.009	0.009	0.010	0.328
Linear Time Trend	Yes	Yes	Yes	No	No
Year-month FE	No	Yes	No	Yes	No
Cohort FE	No	No	Yes	Yes	Yes
Year-month by Score Bucket FE	No	No	No	No	Yes
Clustered SE	Cohort	Cohort	Cohort	Cohort	Cohort

Figure C.1: Interest rates

This figure shows estimates of the coefficients δ_t from the following specification $y_{it} = \gamma_c + \gamma_{ts} + \sum_{t=-6}^{6} \delta_t \{e_{it} = t\} + \beta X_{it} + \varepsilon_{it}$, along with a 95% confidence interval. The outcome y_{it} is observed interest rates. γ_c are cohort fixed effects, and γ_{ts} are time period by score bucket fixed effects. Standard errors are clustered at the cohort level. Source: TransUnion.

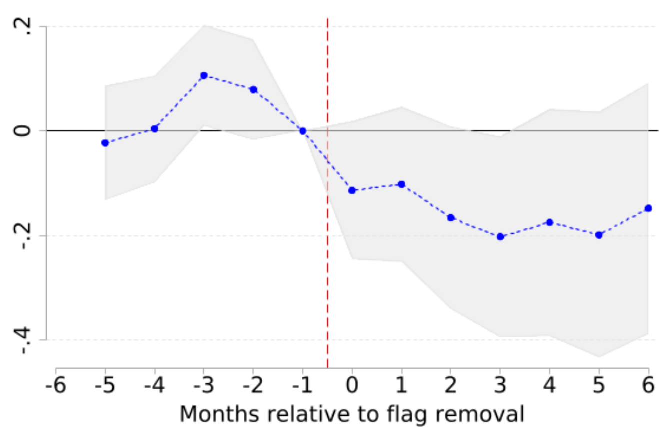

D Stacked Dynamic Difference in Difference Estimation

Our main estimator is a two-way fixed effect estimator with heterogeneous treatment timing. If fully saturated and under homogeneous treatment effects, our estimator provides an unbiased estimate of a treatment effect. In the presence of heterogeneous treatment effects, the estimator may suffer from negatively weighting contrasts, and leads may reflect lags (Sun and Abraham (2021)). To ensure that the choice of the estimator does not drive our results, we follow best practices in Barrios (2021) and Cengiz et al. (2019) (Appendix D) in implementing a Stacked Difference in Differences Estimation. The results remain qualitatively unchanged.

In particular, we implement the stacked difference in differences as follows: For each treated cohort, which is defined by the month in which the individuals have their bankruptcy flags removed, we construct a separate control group. The control group consists of individuals with their bankruptcy flags removed 12 to 17 months after the treated cohort. We then restrict the dataset of treated and control individuals to the six months surrounding the flag removal of the treated cohort. That is, for individuals who have their bankruptcy flag removed in July of 2009, we compare credit scores, interest rates, and auto loan quantities during 2009 to outcomes for individuals who will have their bankruptcy flags removed from July to December 2010. The identifying assumption is that the change in outcomes for individuals with flag removal in July 2009 would have been the same as the change in outcomes for the control cohorts, in the absence of the flag removal for the treated cohort (parallel trends assumption). We repeat this dataset construction for all treated cohorts from July 2009 to June 2017 and stack the separate datasets together. We call each dataset a group and run variants of the following regression:

$$y_{itg} = \gamma_{cg} + \gamma_{tsg} + \sum_{t=-6}^{6} \delta_t \{e_{itg} = t\} + \varepsilon_{itg} \tag{101}$$

e_{itg} indicates time relative to the treatment of the treated cohort. We plot the coefficients δ_t, along with a 95% confidence interval. The coefficients capture the difference in an outcome in each month before and after flag removal relative to the months prior to flag removal.[27] We include cohort-month fixed effects, as well as year-month by score bucket fixed effects. We allow those to differ by the respective dataset. Standard errors are clustered at the cohort-month by group level. Figure D.1 plots estimation results and validates our findings from the main specifications.

We further validate the graphical evidence by implementing a regression framework and showing the results in Table D.1. To be computationally able to account for group-specific linear

[27] We exclude the relative time dummy for period -1.

time trends, we run the stacked regression at the month-cohort month-score bucket-group level and weight by the number of observations. Following standard practice, standard errors are clustered at the cohort-month by group level. When not prohibited by multi-collinearity, the specifications are chosen to match the main specifications in Table 2. Overall, the results in Table D.1 confirm the visual results in Figure D.1. To ensure that the estimates from the stacked specification do not substantially change our welfare computations, we repeat the exercise illustrated in Table 3 and replace the regression estimates with the estimates from the stacked specifications. The resulting welfare estimates are shown in Table D.2. We find an efficiency ratio of approximately 0.059 (=1.83/31.04): for each dollar transferred, 5.9 cents of social welfare is lost, not changing our conclusions from the main text.

To address concerns that one particularly influential group drives the results, we also aggregate treatment and control outcomes for each generated dataset and, subsequently, average relative time means across datasets. Hence, each generated dataset has the same weight in the plotted mean scores, quantities, and interest rates. Figure D.2 illustrates that mean score, interest, and quantity outcomes move in parallel in the pre-period. Besides, the estimated treatment effects appear to be driven by trend breaks in the treated group at the time of flag removal. While it is strictly speaking not a necessary condition for identification, we find this observation comforting. To further address treatment effect heterogeneity, we sort and plot the treat × post coefficients for each of the stacked datasets in Figure D.3. The majority of point estimates are in line with our overall conclusions.

Table D.1: Credit Scores, Interest Rates, and Loan Volumes

This table shows estimates of the coefficients δ_y from the following specification $y_{itg} = \gamma_{cg} + \gamma_{tsg} + \delta^y \mathbb{1}[FlagRemoved] + \beta X_{itg} + \varepsilon_{itg}$. In the top panel, the outcome y_{itg} is the Vantage Score, in the middle panel the outcome is observed interest rates, while in the bottom panel it is loan volumes. γ_{cg} are cohort fixed effects, and γ_{tg} are time period fixed effects that can vary by group. Standard errors are clustered at the cohort by group level. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$ Source: TransUnion.

	(1)	(2)	(3)	(4)	(5)
Panel A: Credit Scores					
Post=1 × Treat=1	16.916***	17.186***	17.198***	17.186***	17.407***
	(0.381)	(0.208)	(0.219)	(0.208)	(0.220)
Observations	174,720	174,720	174,720	174,720	174,720
Panel B: Interest Rates					
Post=1 × Treat=1	-0.269**	-0.303***	-0.302***	-0.303***	-0.288***
	(0.085)	(0.006)	(0.009)	(0.006)	(0.007)
Observations	174,720	174,720	174,720	174,720	174,720
Panel C: Loan Volumes					
Post=1 × Treat=1	28.547***	34.309***	33.715***	34.309***	34.193***
	(7.744)	(2.489)	(3.130)	(2.489)	(2.656)
Observations	174,720	174,720	174,720	174,720	174,720
Group Specific Linear Trend	Yes	No	Yes	No	No
Year-month by Group FE	No	Yes	No	Yes	No
Cohort by Group FE	No	No	Yes	Yes	Yes
Year-month by Score Bucket by Group FE	No	No	No	No	Yes
Clustered SE Cohort by Group	Yes	Yes	Yes	Yes	Yes

Table D.2: Summarizing Estimates Implied by Stacked Specifications

This table is comparable to Table 3 in the main text. We replace the regression coefficients obtained from the main specifications in Table 2 with coefficients of the stacked specifications in Table D.1. This table then summarizes our estimates implied by the stacked specifications of Table D.1. Panel A shows average interest rates in the six months before flag removal ($r_{H,fair}$), the interest rate effect of flag removal ($r_{pool}-r_{B,fair}$), and the effect of flag removal on the fraction of the principal repaid each month in a standardized five-year loan ($\phi(r_{pool})-\phi(r_{H,fair})$). Panel B shows average loan quantities in the six months before flag removal and the quantity effect of flag removal. Panel C shows the market demand elasticity implied by our estimates, the inverse demand slope in terms of the interest rate ($\frac{r_{pool}-r_{H,fair}}{\Lambda_{pool}-\Lambda_{H,fair}}$), and the inverse demand slope in terms of the repayment fraction ($\frac{\phi(r_{pool})-\phi(r_{H,fair})}{\Lambda_{pool}-\Lambda_{H,fair}}$). Panel D summarizes surplus changes implied by the estimates in Table 2. The first row shows the average change in consumer surplus for individuals with flag removal for the average five-year loan. It is the sum of monthly non-default period surpluses. The number of non-default periods is derived from the probability of loans to individuals who ever have a bankruptcy flag to be charged off within two years of loan opening. The second row shows the aggregate change in consumer surplus for individuals with bankruptcy flags when flags are removed for 800,000 individuals. The third row of Panel D shows the implied consumer surplus loss for never-bankrupt individuals for the average 5 year loan scaled by the number of flag removals. It is the sum of non-default period surpluses. The number of non-default periods is derived from the probability of loans to individuals who never have a bankruptcy flag to be charged off within two years of loan opening. The fourth row scales the implied consumer surplus loss for never-bankrupt individuals over 5 years by the occurrence of never-bankrupt individuals and is, consequently, showing the average burden carried by individuals in the never-bankrupt group. The fifth row calculates the consumer surplus loss for never-bankrupt people when 800,000 bankruptcy flags are removed. The sixth row shows the social surplus change over five years scaled by the number of people with flag removal. It is the sum of first and third row. The sixth row shows the total change in social surplus when flags are removed for 800,000 individuals. It is the sum of the second and fifth row. The seventh row provides the efficiency change per dollar redistributed to bankrupt individuals by removing the bankruptcy flag. Source: TransUnion.

	(1)	(2)	(3)	(4)	(5)
Panel A: Prices					
Pre-flag-removal loan interest rate (%)	9.02%	9.02%	9.02%	9.02%	9.02%
Flag removal-induced change in interest rate (%)	-0.269%	-0.303%	-0.302%	-0.303%	-0.288%
Change in monthly payments (%)	-0.013%	-0.015%	-0.015%	-0.015%	-0.014%
Panel B: Quantities					
Pre-flag-removal loan quantity (Average $ per borrower per year)	$3,678.00	$3,678.00	$3,678.00	$3,678.00	$3,678.00
Flag removal-induced change in loan quantity (Average $ per borrower per year)	$342.56	$411.71	$404.58	$411.71	$410.32
Panel C: Elasticity and Slope					
Market Demand Elasticity	-3.12	-3.33	-3.29	-3.33	-3.49
Inverse demand slope (Interest rate % per $100)	-0.0785	-0.0736	-0.0746	-0.0736	-0.0702
Inverse demand slope (Repayment fraction % per $100)	-0.0038	-0.0036	-0.0036	-0.0036	-0.0034
Panel D: Surplus Changes					
Average consumer surplus redistributed to individuals with flag removal over 5 years ($ per eligible borrower with flag removal)	$28.74	$32.66	$32.52	$32.66	$31.04
Total consumer surplus redistributed to individuals with flag removal over 5 years ($)	$22,995,245	$26,128,165	$26,018,217	$26,128,165	$24,832,830
Average consumer surplus taken from never-bankrupt individuals over 5 years ($ per eligible borrower with flag removal)	-$30.17	-$34.59	-$34.42	-$34.59	-$32.87
Average consumer surplus taken from never-bankrupt individuals over 5 years ($ per eligible never bankrupt borrower)	-$3.58	-$4.10	-$4.08	-$4.10	-$3.90
Total consumer surplus taken from never-bankrupt individuals over 5 years ($)	-$24,137,612	-$27,674,278	-$27,532,558	-$27,674,278	-$26,297,587
Change in social surplus per individual over 5 years ($ per eligible borrower with flag removal)	-$1.50	-$1.43	-$1.93	-$1.93	-$1.83
Total change in social surplus over 5 years ($)	-$1,142,368	-$1,546,113	-$1,514,341	-$1,546,113	-$1,464,757
Welfare change per dollar redistributed to bankrupt individuals	-0.0497	-0.0592	-0.0582	-0.0592	-0.0590

Figure D.1: Credit Scores, Interest Rates, and Loan Balances

This figure shows estimates of the coefficients δ_t from the following specification $y_{itg} = \gamma_{cg} + \gamma_{tsg} + \sum_{t=-6}^{6} \delta_t \{e_{itg} = t\} + \varepsilon_{itg}$, along with a 95% confidence interval. In the first panel, the outcome y_{itg} is credit scores, while in the second panel it is interest rates. In the third panel, the outcome is loan balances. γ_{cg} are cohort by group fixed effects, and γ_{tsg} are time period by score bucket by group fixed effects. Standard errors are clustered at the cohort by group level. Source: TransUnion.

Panel A: Credit Scores

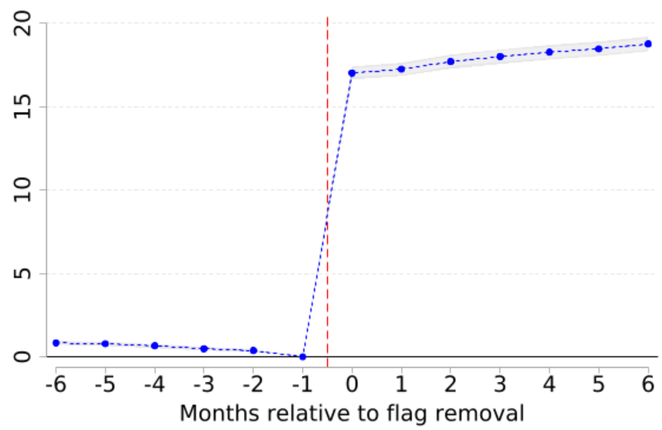

Panel B: Interest Rates

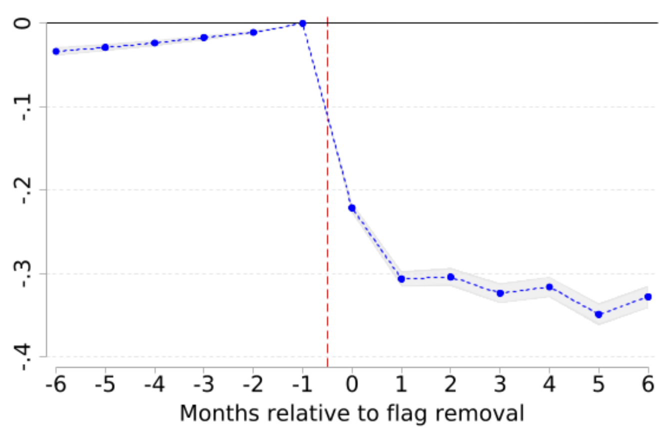

Panel C: Loan Balances

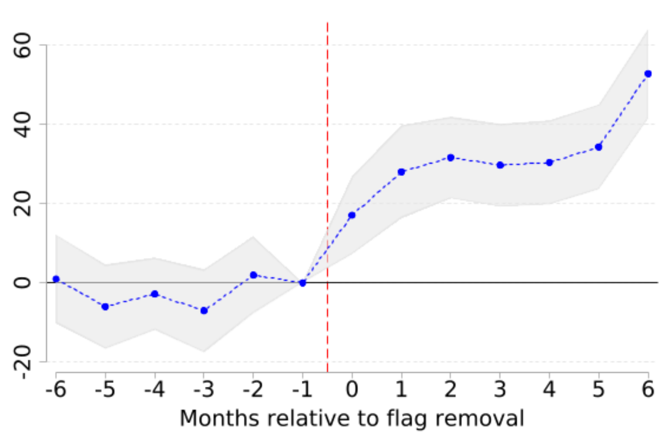

Figure D.2: Mean Outcomes: Credit Scores, Interest Rates, and Loan Volumes

This figure shows average treatment and control outcomes in relative time. We aggregate treatment and control outcomes for each generated dataset and, subsequently, average relative time means across datasets. Each generated dataset has the same weight in the plotted mean scores, predicted interest rates, and quantities. Source: TransUnion.

Panel A: Credit Scores

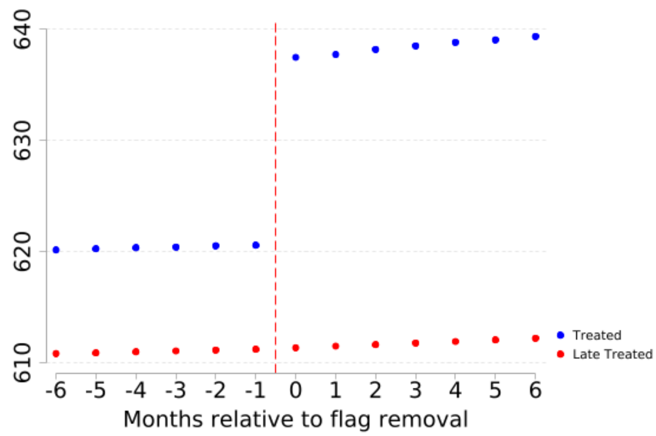

Panel B: Interest Rates

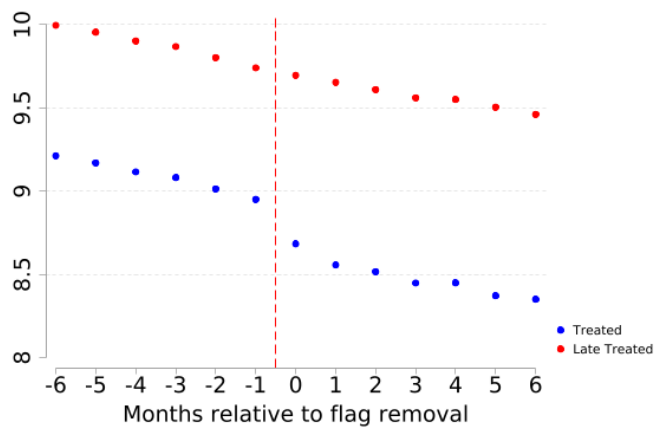

Panel C: Loan Volumes

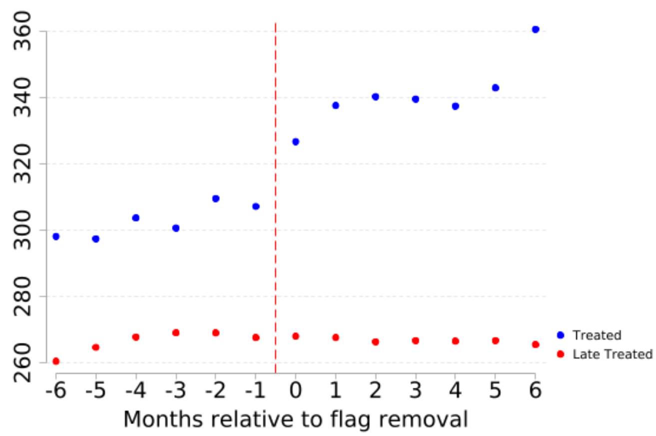

Figure D.3: Individual Events: Credit Scores, Interest Rates, and Loan Volumes

This figure shows event-specific point estimates. It plots the coefficients δ_g from the following specification $y_{itg} = \gamma_{cg} + \gamma_{tsg} + \delta_g \text{Treat} \times \text{Post} + \varepsilon_{itg}$. Point estimates across Panels do not correspond to each other as the coefficient sorting is Panel specific. Source: TransUnion.

Panel A: Credit scores

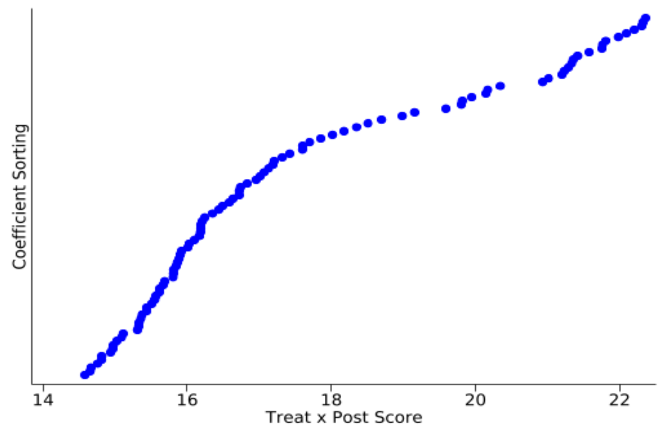

Panel B: Interest rates

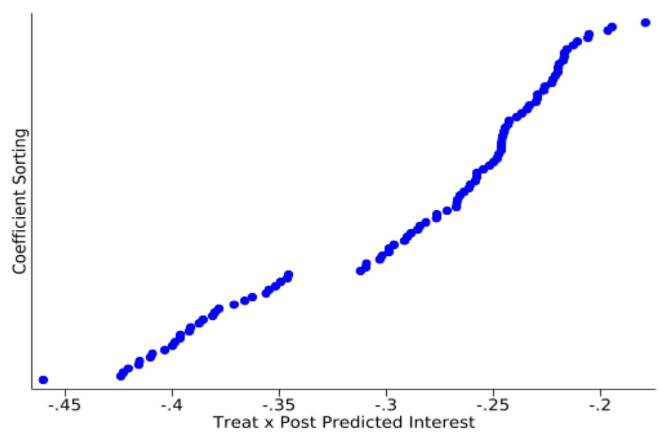

Panel C: Loan volumes

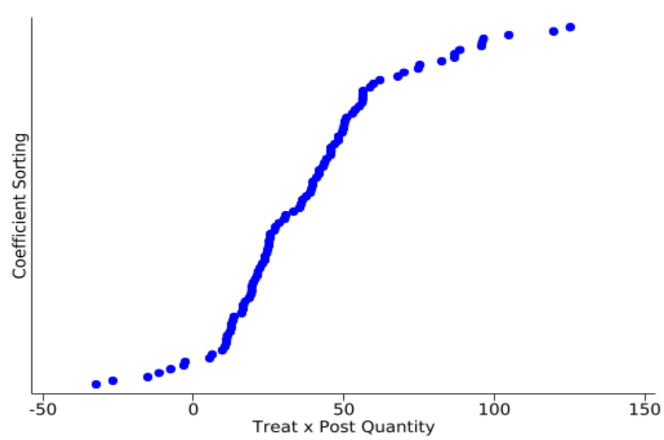

E Illustration of Classic Third Degree Price Discrimination

This section shows an illustration of classic third degree price discrimination contrasting with Figure 1 showing price discrimination on individuals' costs. As Section 2.3 explains, there are two core differences between our setting and the classic literature on third-degree price discrimination: market power and data being informative about demand. Figure E.1 shows an illustration of classic third degree price discrimination.

Figure E.1: Price Discrimination in Classic Markets

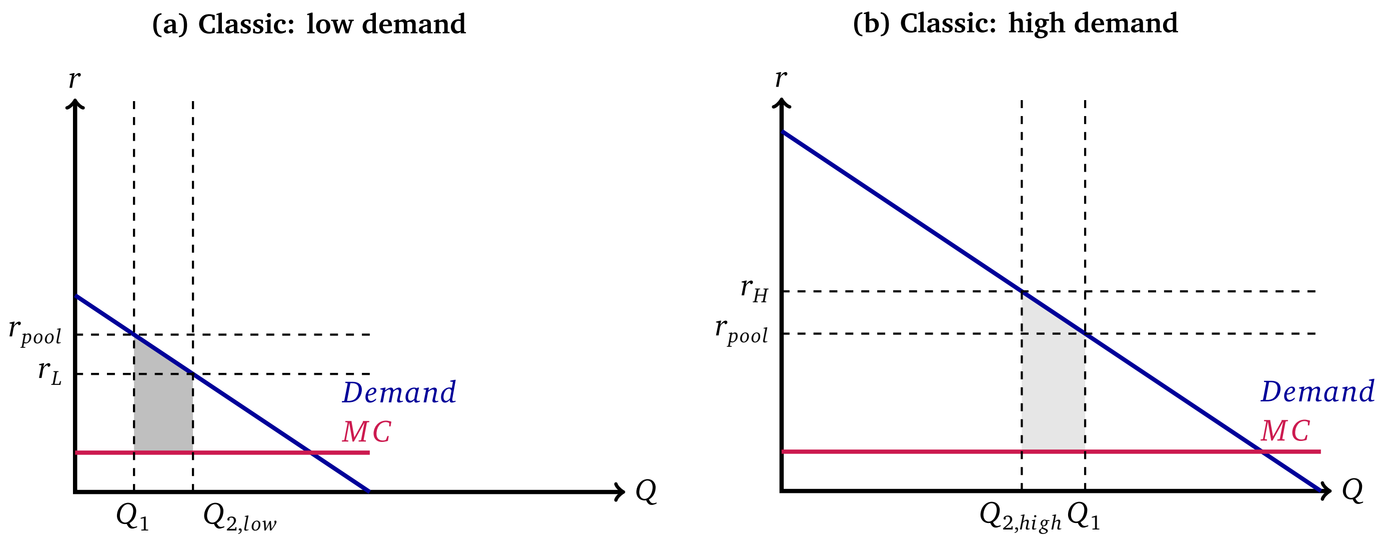

This figure illustrates how third-degree price discrimination affects welfare in classic markets. Suppose there are two groups of prospective borrowers—low demand (panel a) and high demand (panel b). The red lines show the cost of serving these borrowers, and the blue lines show borrowers' demand curve. Lenders are initially unable to distinguish between these prospective borrowers, so set price r_{pool}. After lenders are able to distinguish the two groups of borrowers, they set r_L for the low-demand group (as shown in panel a) and r_H for the high-demand group (panel b). The dark gray shaded area in panel (a) shows the welfare gain for the low-demand group, where prices decrease, and the light gray shaded area in panel (b) shows the welfare loss for the high-demand group, where prices increase.

F Variable Definitions

Table F.1: Variable Description

This table denotes the construction of the main analysis variables. The source for all variables is TransUnion.

Variable	Description
Credit Score	VantageScore 3.0
Quantity Opened	Sum of balances on new auto accounts opened by an individual in a given month; zero when no account opened by the individual in the given month
Quantity Opened Cond. on Opening	Sum of balances on new auto accounts opened by an individual in a given month conditional on an opening being reported
Auto Interest Rate	Credit amount weighted interest of auto accounts at opening. Missing when no auto account opened by individual in a given month
Charged-off	1 if one of the auto loans opened by an individual in a given month is charged-off within the 2 years after opening and zero otherwise
Score Bucket	One of 20 score buckets assigned in the month before flag removal and held constant throughout