Rate Caps on Revolving Credit Lines*

Gajendran Raveendranathan[†] **Georgios Stefanidis**[‡]
Guillaume Sublet[§]

February 7, 2025

Abstract

We show how the revolving nature of credit lines—long-term contracts for access to short-term debt—matters for the regulation of credit card interest rates. First, revolving credit lines introduce history dependence to credit access. As a result, the transition path is key for the magnitude of gains. Second, the competition for credit offers to existing accounts dilutes the profitability of revolving contracts, which reduces credit offers to consumers without existing credit access. In a model calibrated to target salient features of the U.S. credit card market, we find that rate caps tailored to credit access are twice as effective in generating efficiency gains compared to the commonly studied uniform cap. The gains from tailoring rate caps beyond credit access are small.

Keywords: credit cards, revolving credit lines, rate cap, market power, efficiency, welfare
JEL classification: E21, E44, G28, G50

*We thank the editor and the anonymous referees at the Journal of Political Economy Macroeconomics, Igor Livshits, Barış Kaymak, Kyle Dempsey, Jim MacGee, Joaquin Saldain, and the audience of talks at Bank of Canada, SAET, MEA, Annual Workshop of Southern Ontario Macroeconomists, Université de Montréal, and Dalhousie University, where this work was presented. Gajendran Raveendranathan and Georgios Stefanidis thank the SSHRC for the Insight Development Grant 430-2021-00075.

[†]McMaster University. E-mail: raveeg1@mcmaster.ca
[‡]York University. E-mail: stefa107@yorku.ca
[§]Université de Montréal. E-mail: guillaume.sublet@umontreal.ca

1 Introduction

In the United States, interest rates charged to credit card debt are not subject to a cap. Studies have repeatedly documented that credit card interest rates are high not solely because of default risk or the cost of lending but also because of high markups.[1] While some policy makers are strong proponents of capping credit card rates, others raise concerns about credit access. Recent quantitative work on consumer credit builds on early models to emphasize the trade-off between markups and credit access via search and matching between consumers and lenders.[2] The existing literature on the regulation of the credit card market, however, has abstracted from the revolving nature of credit card contracts—i.e., the lender commits to a long-term contract of short-term debt at a fixed interest rate—as well as the ensuing on-the-credit-search for better credit lines.

We show that the *revolving* nature of credit card contracts matters for the design of rate regulation. First, modeling long-term contracts introduces *history dependence to credit access*. As a result, consumers with identical credit worthiness differ in their assessment of the trade-off between markups and credit access because some consumers have pre-existing access to credit lines. Second, the consumers' lack of commitment to not search for better credit lines leads to a *contract dilution problem*. Potential lenders anticipate that the competition will target profitable consumers for balance transfers, which dilutes the profitability of the long-term contract. As a result, the competition for balance transfers ex-post undermines the competition for credit ex-ante, which lowers the likelihood of credit access for consumers who do not have access to credit lines.

We derive three implications for interest rate regulation of the U.S. credit card market. First, credit card interest rates are inefficiently high. The optimal uniform cap, despite being a blunt regulatory tool, reduces the average interest rate by 4.9 percentage points. Second, an access-dependent cap on interest rates—a simple alternative in which the stringency of the cap depends on whether the consumer has a credit card—doubles the efficiency gains that a uniform cap can achieve. The additional gains are specific to the regulation of revolving credit lines because of the history dependence in credit access and the contract dilution problem. Finally, the additional gains from tailoring the cap to consumer characteristics beyond credit access are small.

To quantify the benefits of rate regulation in a framework with revolving credit lines as well as high markups and frequent customer solicitations, we extend the framework of Raveendranathan (2020) and Raveendranathan and Stefanidis (2024) to allow lenders and consumers to bargain over the terms of the long-term credit card contract (which are a borrowing premium and a credit limit).

[1] See Herkenhoff and Raveendranathan (2024) and Dempsey and Ionescu (2024) for recent empirical evidence.

[2] For the early quantitative papers of the credit card market, see Athreya (2002), Chatterjee, Corbae, Nakajima, and Rios-Rull (2007), and Livshits, MacGee, and Tertilt (2007). For the recent strand of the quantitative literature, see Drozd and Nosal (2008), Mateos-Planas and Ríos-Rull (2013), Galenianos and Gavazza (2022), Chatterjee and Eyigungor (2023), Galenianos, Law, and Nosal (2023), Bethune, Saldain, and Young (2024), and Braxton, Herkenhoff, and Phillips (2024).

The extent to which the high markups and frequent customer solicitations observed in the data are inefficient in our model depends on the nature of competition in the credit card industry. In our model, the competition results from the lenders' targeted search for consumers, bargaining between consumers and lenders, the consumers' outside option, and the anticipation by lenders of contract dilution due to the future competition for balance transfers. We exploit the block-recursivity of our environment to show that, despite the complications associated with modeling revolving credit lines, our analysis of optimal rate regulation remains tractable.

We begin our quantitative analysis by calibrating the model to salient moments of the 2019 U.S. credit card market. We target key moments related to the average markup, charge-offs, and the extensive and intensive margins of credit access (that is, the population with credit cards and total credit limit per card holder). We then evaluate the model's fit to non-targeted moments relevant to the policy exercises. The calibrated model replicates the distribution of interest rates and the default rate in the data reasonably well. Furthermore, the model performs well in replicating high levels of credit card utilization rates, which we measure in the data as the share of card holders that use more than 75 percent of their total credit limit.[3] Additionally, our model rationalizes qualitative patterns observed in credit statistics in the cross-section by income and by age (e.g., credit, credit limits, rates, and the population with credit cards) as well as borrowing premiums that are rigid with respect to one-year-ahead default risk (see Dempsey and Ionescu (2024)).

After validating the model calibrated to the U.S. economy, we turn to our main exercise of analyzing the regulation of rates, namely, (1) uniform caps, (2) access-dependent caps, and (3) submarket-specific caps and submarket-and-time-specific caps. The regulator maximizes efficiency gains, which are computed as the sum of onetime compensation transfers that make agents indifferent between the status quo and the counterfactual economy (following the Kaldor-Hicks criterion). Importantly, we compute transitions from the steady state equilibrium of the status quo economy to the steady state equilibrium of the counterfactual economy.

First, we quantify the efficiency gains from the optimal uniform rate cap. We find that a uniform rate cap generates efficiency gains, which is evidence that credit card interest rates are inefficiently high on average. However, a regulatory tool as blunt as a uniform rate cap is rather ineffective at regulating credit card interest rates: the gains amount to a onetime transfer worth 0.32 percent of initial steady state annual income. Although a rate cap reduces the average rate, it also lowers credit access, and due to the heterogeneity among consumers, a uniform cap has to trade off inefficiently high rates for some consumers with lack of credit access for other consumers. For instance, while a stringent cap would significantly benefit consumers with a pre-existing credit card, it would hurt

[3]The model understates lower levels of credit card utilization rates, and, therefore, total credit usage as well. By targeting credit limits instead of total credit, our model produces an empirically plausible estimate for acquisition and marketing costs relative to income. We provide a sensitivity analysis in which we target total credit and show that our baseline calibration produces conservative estimates for the welfare gains.

both those without a credit card and future generations. The trade-off between lowering rates and the incentive for lenders to send credit card offers is nonlinear: the fall in credit access from a marginal tightening of the policy is small when the policy is loose and big when the policy is tight. This nonlinear trade-off is a robust feature of models with competitive entry, which implies a hyperbolic relationship between the incentive for lenders to send a credit card offer and expected profits.[4]

Second, we quantify the gains from access-dependent caps on interest rates and show how the revolving nature of credit lines matters for rate regulation. In this exercise, there is one rate cap for customers with pre-existing accounts in the period of the transition and for balance transfers, and another rate cap for customers without a credit card. We find that it is optimal to set a stringent cap on pre-existing accounts and on balance transfers and a lax cap for those without a credit card. The optimal access-dependent cap leads to gains that amount to a onetime transfer worth 0.62 percent of annual income, which are twice as much as the gains from the optimal uniform rate cap. Most of the efficiency gains are from regulating pre-existing credit card accounts. A stringent cap on pre-existing accounts can reduce the deadweight loss due to inefficiently high rates without compromising the likelihood of credit access for consumers who do not have a credit card at the time of the reform. The stringent cap on balance transfers also benefits consumers who do not have a credit card and future generations by alleviating the contract dilution problem.

Third, we quantitatively evaluate the gains from fine-tuning caps to consumers' cross-sectional characteristics—age, assets, earnings, and the terms of current credit access—and over time. We find that the efficiency gains from fine-tuning the rate cap regulation beyond conditioning the stringency of the cap on credit access are relatively small. We first focus on regulating pre-existing accounts. We find that a single cap on pre-existing accounts captures 80 percent of the gains that can be achieved by fine-tuning the cap on such accounts to all other consumers' characteristics. This is because consumers' characteristics are rather persistent and, hence, having a credit card is an informative measure of a consumer's credit worthiness. While there is no simple way to attribute the efficiency gains to different frictions, we use counterfactual exercises to attribute the efficiency gains to reducing the following three frictions: (1) the lender market power, (2) rate rigidity and default deadweight loss, and (3) the partially uninsurable earnings risk. We find that most of the efficiency gains from capping rates on pre-existing accounts—62 to 87 percent—can be attributed to addressing inefficiencies due to the lenders' market power.

Lastly, to evaluate the regulatory authority's temptation to revise the policy over time, we consider sequentially optimal rate cap regulations. We show that the problem of a regulator who

[4] The trade-off between lower rates and the extensive margin of credit access is corroborated by empirical evidence from Cuesta and Sepúlveda (2021). They document that in Chile, following the implementation of a rate cap, the number of unsecured non-revolving consumer loans fell.

cannot commit to a time-varying policy remains tractable because sequential optimality preserves the block-recursivity of our environment. Our findings caution against the temptation of the regulatory authority to revise the policy over time. The long-term nature of revolving contracts makes the optimal rate cap policy time-inconsistent. The source of the time inconsistency is the contract dilution problem. The ex-post benefits of competition through balance transfers undermines the ex-ante benefits of competition for credit offers to consumers without access to credit.

Related literature Recently, several studies have highlighted the potential benefits of uniform interest rate caps in the credit card market. Hatchondo and Martinez (2017) illustrate how a rate cap can lead to gains even in a perfectly competitive credit card market, as in the models of Chatterjee, Corbae, Nakajima, and Rios-Rull (2007) and Livshits, MacGee, and Tertilt (2007). The authors find that an interest rate cap improves welfare in a model with debt dilution due to non-exclusive borrowing contracts.[5] The remaining studies highlight welfare gains from a rate cap in structural models in which the lender has some degree of market power. Galenianos and Gavazza (2022) propose a static framework with informational frictions, search intensity, and product differentiation to account for credit card interest rates that are high and dispersed. They consider a mechanism in which a rate cap lowers search intensity, thereby leading to an increase in the average interest rate. They find that quantitatively this mechanism does not overturn the direct effect of a rate cap on the average interest rate, and that the cap generates gains. Galenianos, Law, and Nosal (2023) consider an environment with non-revolving credit contracts in which consumers search for competitive lenders in every period, with the outside option of borrowing from a monopolist in the event of not matching with a competitive lender. This means that the borrower might not be able to roll over debt if they do not match with a new lender, whereas in our model, the contracts are revolving. In Galenianos, Law, and Nosal (2023), a cap leads to an increase in the extensive margin of borrowing (number of loans), whereas in our paper a (uniform) cap leads to a decrease in the extensive margin (both number of credit cards and number of loans). This is because in their paper with non-revolving contracts, reducing the market power of the monopoly lender via a cap incentivizes more households to search in the competitive market. Consequently, Galenianos, Law, and Nosal (2023) find larger gains from uniform caps, whereas in our paper there is a standard trade-off between rates and credit access. We view our results as complementary because we highlight how the revolving nature of contracts affects the outcomes of rate regulation. Bethune, Saldain, and Young (2024) study how rate caps can deliver gains in the presence of lenders with market power and households with some degree of financial illiteracy. We depart from these papers by analyzing rate caps in a model with revolving credit lines—a unique feature of our model. This departure

[5] Nakajima (2017) finds that without debt dilution, consumers' shortsightedness does not justify imposing an interest rate cap on credit cards. In complementary work, Saldain (2023) finds that a cap reduces welfare for consumers with temptation in a model calibrated to payday loans instead of credit cards.

allows us to show that, due to the revolving nature of credit lines, rate caps generate larger gains along the transition path in comparison to gains across steady states and by alleviating the contract dilution problem.

Our paper also contributes to the structural literature that extends the standard consumer credit model of Chatterjee, Corbae, Nakajima, and Rios-Rull (2007), and Livshits, MacGee, and Tertilt (2007)) to allow for revolving credit lines, which tend to be computationally demanding (e.g., Drozd and Nosal (2008), Mateos-Planas and Ríos-Rull (2013), Braxton, Herkenhoff, and Phillips (2024), Raveendranathan (2020), Raveendranathan and Stefanidis (2024), and Chatterjee and Eyigungor (2023)). Our analysis of rate caps contributes to this recent class of models by eliciting the efficiency gains along the transition path resulting from the rate regulation of revolving credit lines. However, we note that across steady states, the rate caps in our framework lead to qualitatively similar conclusions to those from the standard consumer credit model.

The contract dilution problem in our model with revolving credit lines and balance transfers complements other sources of time-inconsistency in credit markets. While contract dilution is specific to on-the-credit-search for a better credit line, debt dilution is specific to the lack of commitment to repay debt (as in Bizer and DeMarzo (1992) and Hatchondo and Martinez (2017)).

Lastly, our work complements the structural literature that analyzes other regulatory interventions in the credit card market. Chatterjee, Corbae, Nakajima, and Rios-Rull (2007) and Nakajima (2017) analyze the Bankruptcy Abuse Prevention and Consumer Protection Act of 2005. Exler, Livshits, MacGee, and Tertilt (2024) analyze various policies related to the credit card market (e.g., financial literacy education, reducing default costs, increasing borrowing costs, and debt limits) in the presence of consumers who are over-optimistic about their earnings. Nelson (2023) studies the restrictions on interest rate hikes imposed by the Credit CARD Act of 2009. Raveendranathan and Stefanidis (2024) study the "Ability to Pay" provision of the CARD Act, which imposed quantity restrictions on credit card limits.

2 A simple model of revolving credit lines

In this section, we study the economics of revolving credit lines in a simplified version of our quantitative model. The simple model corresponds to a submarket of the credit card market from the quantitative model in which we abstract from the consumers' risk aversion, their life-cycle profile of earnings, earnings and survival risk, consumer default, and endogenous borrowing limits. We elicit two features that are inherent to revolving credit lines. First, despite the short-term nature of the debt, the long-term nature of revolving credit contracts introduces history dependence in credit access. Second, we elicit a source of time inconsistency for the consumer that is specific to revolving credit lines facing competition for balance transfers.

A unit mass of consumers live for three periods, $t = 0, 1, 2$. Consumers have linear utility from borrowing b with a premium τ; that is, the indirect utility function for two consecutive periods is $u(\tau) = b - \beta(1 + \tau)b$, where $\beta \in (0, 1)$ denotes the discount factor. A consumer with access to credit at premium $\tau \leq \frac{1}{\beta} - 1$ borrows up to the exogenous borrowing limit $\bar{b}$. If the consumer does not have access to credit, the consumer does not borrow, $b = 0$. Lenders have unlimited access to credit at an exogenous interest rate, which is normalized to zero in this section. Profits from lending b, as a function of the premium τ, are $\pi(\tau) = \tau b$.

The timing is as follows. In period 0, the consumer receives an offer of a revolving credit line with probability p_0, which is determined by a free entry condition. The borrowing premium τ_0 on a revolving credit line is determined through Nash bargaining between the consumer and the lender, and is fixed for the duration of the contract. The revolving contract lasts until period 2 unless the consumer accepts an offer of a balance transfer in period 1. The probability p to receive an offer of a balance transfer is exogenous in this section, and Nash bargaining determines the premium on the balance transfer $\tau_1(\tau)$. A consumer without credit access in period 1 receives a credit offer with exogenous probability p_1^*. The premium on this credit offer τ_1^* is also determined through Nash bargaining. Debt is short-term and there is no default.

Importantly, despite the short-term debt of revolving credit lines, the long-term contract introduces *history dependence in credit access*. In period 1, a fraction p_0 of consumers already have credit access and expect a balance transfer with probability p, and a fraction $1 - p_0$ of consumers do not have credit access and expect a credit offer with probability p_1^*.

We solve for the Stackelberg equilibrium $(p_0, \tau_0, \tau_1(\tau), \tau_1^*)$ by backward induction. In period 2, debt is repaid. In period 1, a consumer with access to credit at premium τ rejects an offer of a balance transfer at premium τ_1 if $\tau_1 > \tau$. The premium offered in period 1 for a balance transfer solves

$$\max_{\tau_1 \geq 0} \pi(\tau_1)^{\theta^b} (u(\tau_1, b) - u(\tau, b))^{1-\theta^b},$$

where θ^b denotes the lender's bargaining power. The premium for a balance transfer equates a weighted average of the elasticities of the per period profit and of the utility function,

$$\theta^b \frac{\pi'(\tau_1)}{\pi(\tau_1)} = -(1 - \theta^b) \frac{u'(\tau_1)}{u(\tau_1) - u(\tau)}.$$

The best response of the competition for balance transfers in period 1 is to undercut the premium τ on the revolving credit line, $\tau_1(\tau) = \theta^b \tau$, which implies that in equilibrium, the consumer accepts a balance transfer.

If the consumer does not have a credit card in period 1, the premium offered solves

$$\max_{\tau \geq 0} \pi(\tau)^{\theta^b} u(\tau)^{1-\theta^b},$$

which again implies that the premium τ_1^* equates a weighted average of the elasticity of the per period profit and of the utility function; that is, $\tau_1^* = \theta^b \frac{1-\beta}{\beta}$.

A lender in period 0 anticipates the competition in period 1. That is, it anticipates that the consumer will receive an offer of a balance transfer with probability p and premium $\tau_1(\tau_0)$ (i.e., revolving credit lines with on-the-credit search for balance transfers subject the lender to a stochastic participation constraint). The expected profits from offering premium τ for a lender who matches with the consumer in period 0 are

$$\Pi(\tau) = \pi(\tau)(1 + (1 - p)).$$

The lender also anticipates that the consumer's outside option is to get a credit card with probability p_1^* and premium τ_1^*. The premium offered in period 0 solves

$$\max_{\tau \geq 0} \Pi(\tau)^{\theta^b} (V(\tau) - V^*)^{1-\theta^b},$$

where the payoff of the consumer from accepting the offer is

$$V(\tau) = u(\tau) + \beta(1 - p)u(\tau) + \beta p u(\tau_1(\tau)),$$

and the consumer's outside option is $V^* = \beta p_1^* u(\tau_1^*)$.

The borrowing premium offered in period 0 equates the elasticity of the lender and the consumer surpluses from the revolving credit line, weighted by the lender's and the consumer's respective bargaining powers. After some algebra,

$$\tau_0 = \theta^b \frac{1}{\beta} \frac{1 - \beta^2 - p_1^*(1 - \theta^b)\beta(1 - \beta)}{1 + \beta - p(1 - \theta^b)\beta}.$$

Unless the lender has all the bargaining power, i.e., $\theta^b = 1$, future competition for credit offers matters for the premium offered on a revolving credit line. In particular, the competition for new offers in period 1, governed by $p_1^* > 0$, induces the lender to offer a lower premium than it would if there was no competition for offers of credit access in period 1 (i.e., $p_1^* = 0$). In contrast, the competition for balance transfers, parameterized by the probability of offers $p > 0$, induces a lender to front-load profits by offering a higher premium than it would if there was no competition for balance transfers (i.e., $p = 0$). We interpret the borrowing premium as a measure of the *market* *power* of lenders, as a catch-all for the bargaining power of the lender, the outside option of the consumer, and the ability of the lender to front-load profits in anticipation of future competition for balance transfers.

We show that the competition for balance transfers lowers the probability of credit access in period 0. This probability results from the lenders' targeted search for consumers and the free entry of lenders (after paying a fixed cost κ). That is, the probability of credit access $p_0 = M(1, \mu)$ in period 0 is determined by market tightness μ and a matching function $M(1, \mu)$ that exhibits constant returns to scale and is strictly increasing in both its arguments (as is the case for most standard matching functions such as Cobb-Douglas, den Haan, Ramey, and Watson (2000), and urn-ball matching functions). Because there is a unit mass of consumers, μ corresponds to the mass of offers in this section and $p_0 = M(1, \mu)$. For consumers with expected profitability Π_0, market tightness satisfies the following free entry condition

$$M(1, \mu)\Pi_0 = \kappa\mu,$$

where $\Pi_0 = \pi(\tau_0)(2 - p)$.

We obtain the following comparative statics for the equilibrium probability of offers of revolving credit lines p_0 as a function of the probability of offers of balance transfers p, and show that competition for balance transfers dilutes the profitability of the contract.

Proposition 1 (Contract dilution and likelihood of access)**.** *The competition for balance transfers dilutes the profitability of the contract and undermines the competition for offers of revolving credit lines, in that* Π_0 *and* p_0 *are strictly decreasing functions of* p*.*

The proof shows that with more competition for balance transfers, i.e., higher p, the increase in lender profitability resulting from a higher premium outweighs the reduction in profitability resulting from the shortened contract duration.

Proof. Because p_0 is a strictly increasing function of Π_0, it suffices to show that Π_0 is strictly decreasing as a function of p. Comparative statics with respect to p gives

$$\Pi_0'(p) = \pi'(\tau_0(p))\tau_0'(p)(2 - p) - \pi(\tau_0(p)),$$

where the equilibrium profits and premium are expressed as functions of p. Substituting in the equilibrium values for τ_0 gives

$$\Pi_0'(p) = b\tau_0(p)\left[\frac{2(1-\theta^b)\beta - p(1-\theta^b)\beta}{1+\beta - p(1-\theta^b)\beta} - 1\right],$$

which implies that $\Pi_0'(p) < 0$ because $2(1 - \theta^b)\beta < 1 + \beta$. □

Importantly, as shown in Proposition 1, the long-term nature of revolving credit contracts adds inter-temporal considerations to the trade-off between borrowing premium and credit access. In fact, as the next corollary shows, contract dilution leads to a time-inconsistency problem for the consumer.

Corollary 1. *The consumer would like to commit to not search for balance transfers.*

Proof. The proof contains two parts. First, Proposition 1 shows that p_0 is decreasing in p. Second, the value function of the consumer $V(\tau_0(p))$ does not depend on p. This result follows simply by introducing $\tau_0(p)$ to the value function of the consumer. The resulting expression is

$$V(\tau_0(p)) = (1-\theta^b)\big(1-\beta^2+\theta^b\beta(1-\beta)p_1^*\big)b,$$

More competition for balance transfers exactly offsets the benefit of the increased likelihood of a balance transfer with the cost of the increased premium, leaving the matched consumer indifferent ex-post (that is, $V(\tau_0(0)) = V(\tau_0(p))$). As a result, $p_0(0)V(\tau_0(0)) > p_0(p)V(\tau_0(p))$ for $p > 0$. $\square$

In addition to the mechanisms specific to revolving credit lines we have presented, the quantitative model features risk-averse consumers who are heterogeneous in their life-cycle profile of earnings and face earnings and survival risk, and endogenous borrowing limits. Also, in the quantitative model, the probabilities of offers of a balance transfer are endogenously determined by targeted search and free entry. The asset market is incomplete so that consumers cannot perfectly insure the risk they face, which is a source of inefficiency. In addition, consumers may default, which is associated with a stigma cost. The rigidity of the rate on a credit line implies inefficiencies associated with the stigma cost of default, which could be avoided if rates were not rigid.

3 Quantitative model

In this section, we introduce a structural model that can capture three salient features of the credit market. First, credit lines, characterized by a limit and a borrowing premium, are revolving. Second, markups are endogenous; they are determined based on the bargaining power of the lender and the outside option of the consumer. And third, the frequency of customer solicitations results from search and matching and free entry. We show that our analysis remains tractable because targeted search and free entry allow for a block-recursive equilibrium characterization.

We introduce the model in Section 3.1 and discuss the details of search and matching in Section 3.2, and the determination of credit card limits in Section 3.3. In Section 3.4 we present the consumer and lender value functions. In Section 3.5, we present a definition of a block-recursive quasi-equilibrium that takes the policy function for the borrowing premium as given. In Section

3.6 we introduce three types of equilibria that determine the policy function for the borrowing premium: (1) unregulated equilibrium, (2) rate cap equilibrium, and (3) optimal regulated equilibrium.

3.1 Overview

Consumers live up to J periods and j denotes their age. A consumer of age j survives to $j+1$ with probability ψ_j. Before retirement, a consumer supplies labor inelastically with a productivity that evolves over the consumer's life-cycle. The idiosyncratic productivity consists of a permanent component θ, a persistent $AR(1)$ component η, and a transitory iid component γ. Earnings before retirement, as a function of the idiosyncratic state of the consumer are $y_j = (1-\tau_{SS})\theta\eta\gamma\nu_j$, where ν_j is the deterministic life-cycle productivity and τ_{SS} is the Social Security tax rate. In retirement, the consumer does not work, and earns Social Security benefits, denoted by SS_θ, which are a function of the consumer's permanent earnings component.

A lender commits to a borrowing premium τ and a "lower bound" on the credit card limit $\bar{b}$ for the duration of the credit contract.[6] The limit is a lower bound because in each period the limit can be relaxed via Nash bargaining. The credit line is revolving because it provides ongoing access to credit. Consumers, however, do not commit to the credit contract.[7] A credit contract may end because of balance transfers or because the consumer defaulted or died.

Consumers are heterogeneous in their age j and the vector of productivity components (θ, η, γ), debt or asset level b, and the credit card terms $(\bar{b}, \tau)$. Additionally, the consumer faces extreme value type 1 taste shocks ζ^R and ζ^D over repayment and default, which are drawn from a Gumbel distribution centered at zero with scale parameter ξ.

The credit card market is segmented into submarkets indexed by the characteristics observed prior to the determination of new credit offers and terms. These characteristics are age, debt or asset level, existing credit card terms, and the permanent and persistent earnings components. Let $s = (j, b, \bar{b}, \tau, \theta, \eta)$ index a submarket. We assume that the transitory earnings component γ and the taste shocks ζ^R and ζ^D are realized after new credit card offers and terms are determined.[8] The distribution of submarkets Γ captures the rich heterogeneity in age, access to credit, and earnings

[6]The borrowing premium captures the interest rate margin on credit cards over the index rate. Our assumptions about the borrowing premium are motivated by features of the U.S. credit card market in 2019, the year to which we will calibrate the baseline model. The Credit CARD Act of 2009, which is still in effect in 2019, restricts rate hikes on existing balances. We model the restriction on rate hikes by having the lender commit to the borrowing premium. In Appendix B.4, we show that our model produces a weak but still positive relationship between the borrowing premium and the likelihood of one-year-ahead default, as documented by Dempsey and Ionescu (2024).

[7]While we focus on a model with on-the-credit-search in the main text, in Appendix C.4 we extend the analysis to give the consumer the choice whether to commit to search-on-the-credit or not.

[8]This timing assumption facilitates computation by adding smoothness to the lender's ex-post profit function, as in Mateos-Planas and Ríos-Rull (2013), where the lender does not observe the consumer's default and switching costs.

across consumers. Although the distribution of submarkets, which is the aggregate state of the economy, evolves endogenously, our analysis remains tractable because, as we will show, our model admits a block-recursive equilibrium. We consider a consumer without a credit card as having a borrowing limit $\bar{b} = 0$ and, for notational convenience, as facing a prohibitively high borrowing premium $\tau = \infty$.

The timing of events within each period is as follows: realization of persistent component of earnings, search and matching, new credit card terms, realization of transitory earnings and default/repayment taste shocks, repayment or default decision, consumption and saving/borrowing if possible, and realization of profits.

3.2 Search and matching

Lenders target their search by choosing the submarket in which they solicit consumers. The probability that a consumer in a given submarket receives a credit offer is the fraction of successful matches $p(s) = M(1, \mu(s))$. Here, M denotes a matching function with constant returns to scale that is strictly increasing in both its arguments. Additionally, $\mu(s)$ denotes the market tightness defined as the ratio of the mass of lenders soliciting consumers in submarket s to the mass of consumers in that submarket. In other words, market tightness refers to the frequency of customer solicitations in a submarket.

The frequency of customer solicitations results from the free entry of lenders in a submarket. For $\mu(s) > 0$,

$$M(1, \mu(s))\Pi(s) = \kappa\mu(s), \tag{1}$$

where κ denotes the cost of sending an offer, and $\Pi(s)$ denotes the expected net present value of profits of the long-term contract, conditional on matching with a consumer in submarket s. Note that we restrict (1) to submarkets with strictly positive entry. Additionally, we assume that consumers expect a submarket to not have entry if and only if the submarket is not profitable; that is, $\mu(s) = 0$ if and only if $\Pi(s) \leq \kappa$. This assumption rules out the possibility of self-fulfilling consumers' expectations of submarket closure.[9]

3.3 Credit card limits

The evolution of access to credit and of the lending terms depends on whether the consumer receives a credit offer this period. First, consider a consumer in submarket $s = (j, b, \bar{b}, \tau, \theta, \eta)$ who does not receive a credit offer from an entrant lender this period. The borrowing premium remains

[9]Otherwise, a submarket could be closed because consumers avoid that submarket because they expect that firms will not enter that submarket, and, in turn, firms do not enter because consumers avoided that submarket.

τ, which captures that the incumbent lender committed to a rate for the duration of the contract. The incumbent lender negotiates with the consumer a potential increase in the credit limit. We assume that the negotiation takes the form of Nash bargaining in which the increase in surplus is divided between the incumbent and the consumer according to the Nash bargaining parameter $\theta_b \in [0,1]$,

$$\max_{\bar{b}_I \geq \bar{b}} \Big\{ \big(E_\gamma \left[\pi \big(j, b, \bar{b}_I, \tau, \theta, \eta, \gamma \big) - \pi \big(s, \gamma \big) \right] \big)^{\theta_b} \big(E_\gamma \left[W \big(j, b, \bar{b}_I, \tau, \theta, \eta, \gamma \big) - W \big(s, \gamma \big) \right] \big)^{1-\theta_b} \Big\}. \quad (2)$$

Here $\pi\big(s,\gamma\big)$ denotes the expected profits of an incumbent lending to a consumer with characteristics (s,γ), and $W(s,\gamma)$ denotes the expected utility of a consumer with characteristics (s,γ). The solution to (2) is a submarket-specific borrowing limit $\bar{b}_I(s)$. While the lender cannot commit to relaxing the credit limit in the future, the lender anticipates that the credit limit this period sets a lower bound for future credit limits.

Second, consider a consumer in submarket $s = (j, b, \bar{b}, \tau, \theta, \eta)$ who receives a credit offer from an entrant lender this period. The credit offer specifies a borrowing premium $\tau(s)$, to which the lender commits for the duration of the relationship, and a borrowing limit, which can be relaxed in the future. While we postpone describing the choice of the borrowing premium until Section 3.6.2, the credit limit for a new match, denoted by $\bar{b}_E(s)$, is the solution to the following Nash bargaining problem in which the borrowing premium $\tau(s)$ is taken as given,[10]

$$\begin{aligned} \max_{\bar{b}_E \geq \bar{b}} \Big\{ & \big(E_\gamma \left[\pi \big(j, b, \bar{b}_E, \tau(s), \theta, \eta, \gamma \big) \right] - \max(b,0) \big)^{\theta_b} \\ & \big(E_\gamma \left[W \big(j, b, \bar{b}_E, \tau(s), \theta, \eta, \gamma \big) - W \big(s_I, \gamma \big) \right] \big)^{1-\theta_b} \Big\}. \end{aligned} \quad (3)$$

The expected profits of an entrant, which govern entry in (1), are the profits that result from the updated borrowing premium and limit (net of the balance transfer if the entrant poaches a consumer from an incumbent lender),

$$\Pi\big(s\big) = E_\gamma \pi \big(j, b, \bar{b}_E, \tau(s), \theta, \eta, \gamma \big) - \max(b,0). \quad (4)$$

For a submarket to be active, ex-ante profits must exceed the cost of entry, κ, as suggested by the free entry condition in (1). To capture the revolving nature of credit lines, we assume that the incumbent lender does not make a competing offer if the consumer receives an offer of balance transfer. As we will see in Section 3.6.1, however, lenders anticipate the competition for balance transfers. For instance, a lender can tailor the premium or the limit so that it is less lucrative for

[10] To keep track of the evolution of the credit terms within a period, we express the submarket as a function of the updated term $s_I(s) = (j, b, \bar{b}_I(s), \tau, \theta, \eta)$, and, similarly $s_E(s) = (j, b, \bar{b}_E(s), \tau(s), \theta, \eta)$, and where there is no ambiguity, suppress the argument, i.e., s_I and s_E.

the competition to target the first lender's consumers with offers of balance transfer.

3.4 Consumer problem and lender profits

Repayment or default. A consumer can either repay or default. The value function to the consumer before the realization of the extreme value shocks is

$$W(s,\gamma) = E_{\zeta^R,\zeta^D}\left[\max_{D\in\{0,1\}}(1-D)\Big(V\big(s,\gamma\big)+\zeta^R\Big)+D\Big(V^D\big(s,\gamma\big)+\zeta^D\Big)\right]. \tag{5}$$

Once the taste shocks (ζ^R, ζ^D) have realized, the consumer chooses between repaying, $D = 0$, or defaulting, $D = 1$. The value of either choice reflects the "modelled payoff," specified next, and the (ζ^R, ζ^D) terms reflect the non-modeled payoffs affecting the default decision.[11] The probability of default, prior to realization of (ζ^R, ζ^D), is

$$d(s,\gamma) = \frac{\exp\Big(\xi V^D(s,\gamma)\Big)}{\exp\Big(\xi V^D(s,\gamma)\Big)+\exp\Big(\xi V(s,\gamma)\Big)}, \tag{6}$$

where ξ is a scaling parameter that determines the variance of the extreme value shocks.

Consumption and saving or borrowing. We augment a standard earnings fluctuation problem with revolving credit lines and evolving credit access over the life-cycle. Consumers solve an earnings fluctuation problem by choosing to save or borrow in response to the evolution of the stochastic earnings over their life-cycle and the evolution of credit access,

$$V\big(j,b,\bar b,\tau,\theta,\eta,\gamma\big) = \max_{c\geq 0,\ b'\leq \bar b}\ \Big\{U(c)+\beta\psi_j E_{\eta',\gamma'|\eta}E_{o'}W\big(j+1,b',\bar b',\tau',\theta,\eta',\gamma'\big)\Big\} \tag{7}$$

subject to $c + b = y_{j,\theta,\eta,\gamma} + q_j\big(\tau\big)b'_+ + q_j\big(0\big)b'_-$

and

$$(\bar b',\tau') = \begin{cases}\Big(\bar b_I\big(j+1,b',\bar b,\tau,\theta,\eta'\big),\ \tau\Big) & \text{if } o'=0\\ \Big(\bar b_E\big(j+1,b',\bar b,\tau',\theta,\eta'\big),\ \tau\big(j+1,b',\bar b,\tau,\theta,\eta'\big)\Big) & \text{if } o'=1,\end{cases}$$

$p_{(o'=1)} = p\big(j+1,b',\bar b,\tau,\theta,\eta'\big)$, and $p_{(o'=0)} = 1 - p_{(o'=1)}$.

Here $U(c)$ denotes the utility of consumption, U is CRRA, b' is next period debt (assets if negative), o' is an indicator variable taking the value of one if a new match is formed in the next period and

[11] Consumers without a credit card cannot hold debt because revolving credit is the only modeled debt. Hence, default for a consumer without a credit card occurs only due to extreme value taste shocks capturing expense shocks such as health care bills, lawsuits, and divorce.

zero otherwise, and $E_{\eta',\gamma'|\eta}$ is the expectation over next period earnings conditional on the current period earnings. Current credit card holders who do not default can borrow up to the limit $\bar{b}$ at prices determined by the borrowing premium τ,

$$q_j(\tau) = \frac{\psi_j}{1+r+\tau}. \tag{8}$$

There is no premium on savings. In the subsequent period, with probability $p_{(o'=1)}$, the consumer receives a credit limit $\bar{b}'$ and borrowing premium τ'. The next period's credit card limit, credit card borrowing premium, and probability of receiving an offer $p_{(o'=1)} = p\big(j+1, \theta, \eta', b', \bar{b}, \tau\big)$ depend on the submarket in which the consumer will be next period. While the consumer takes the probabilities of matching in each submarket as given, the choice of the consumer partially determines the submarket in which the consumer will be in the beginning of next period. With probability $p_{(o'=0)}$, the consumer will not receive a credit offer, in which case, the incumbent lender—if there is one—may loosen the credit limit. If there is no incumbent lender, the borrowing premium is prohibitively large, i.e., $\tau = \infty$, and the credit limit does not matter. Note that consumers rationally expect that Nash bargaining and the policy function for the borrowing premium determine the evolution of the limit and the premium, respectively.

Consumers may also default. The costs of defaulting are stigma (denoted by χ) and exclusion from financial markets for a period. The defaulting consumers' payoff is

$$V^D\big(s,\gamma\big) = U(y_{j,\theta,\eta,\gamma}) - \chi + \beta\psi_j E_{\eta',\gamma'|\eta} E_{o'} W\big(j+1, 0, \bar{b}', \tau', \theta, \eta', \gamma'\big), \tag{9}$$

where

$$(\bar{b}', \tau') = \begin{cases} (0,\ \tau = \infty) & \text{if } o' = 0 \\ \big(\bar{b}_E\big(j+1, 0, \bar{b} = 0, \tau = \infty, \theta, \eta'\big),\ \tau\big(j+1, 0, \bar{b} = 0, \tau = \infty, \theta, \eta'\big)\big) & \text{if } o' = 1. \end{cases}$$

Lender profits. The expected profits from committing to a premium τ on borrowing depends on the submarket. Consider a submarket $s = (j, b, \bar{b}, \tau, \theta, \eta)$. Given policy functions $\bar{b}(s)$, $b'(s,\gamma)$, and $d(s,\gamma)$, and probabilities $p(s)$, the net present value of expected profits, conditional on offering τ and $\bar{b}$, is

$$\begin{aligned} \pi\big(s,\gamma\big) = \Big(1 - d(s,\gamma)\Big)\Big(&\max\{b, 0\} - \max\{b', 0\}(q_j(\tau) + \tau_c) + \\ &\frac{\psi_j}{1+r} E_{\gamma'\eta'|\eta}\big[(1-p)\pi\big(j+1, b', \bar{b}', \tau, \theta, \eta', \gamma'\big) + p\max\{b', 0\}\big]\Big), \end{aligned} \tag{10}$$

where

$$b' = b'(s,\gamma),\ p = p(j+1, b', \bar{b}', \tau, \theta, \eta'),\ \bar{b}' = \bar{b}_I(j+1, b', \bar{b}, \tau, \theta, \eta').$$

3.5 Block-recursive quasi-equilibrium

We define an equilibrium in two steps. First, we introduce quasi-equilibria, which is the part of the equilibrium definition that is common across our positive and counterfactual scenarios. Second, in the next section, we introduce equilibria by augmenting the quasi-equilibria with an endogenous determination of the policy function for the borrowing premium.

A *block-recursive quasi-equilibrium given* $\tau(s)$ *and an initial distribution of submarkets* consists of consumer value functions $V\big(s,\gamma\big)$, $V^D\big(s,\gamma\big)$, and $W(s,\gamma)$; consumer policy functions $d\big(s,\gamma\big)$, $c\big(s,\gamma\big)$, and $b'\big(s,\gamma\big)$; credit card firm profit functions $\pi(s,\gamma)$ and $\Pi(s)$; policy functions $\bar{b}_I(s)$ and $\bar{b}_E(s)$ for borrowing limits; market tightness function $\mu(s)$; and probabilities of credit access $p(s)$ such that

- in each submarket s, the competitive entry condition (1) holds for $\Pi(s)$ and $\mu(s)$;
- targeted search determines the probabilities of receiving a credit offer, $p(s) = M\big(1, \mu(s)\big)$;
- the policy functions $\bar{b}_I(s)$ and $\bar{b}_E(s)$ solve (2) and (3), respectively, and the profit functions $\Pi(s)$ and $\pi(s,\gamma)$ are obtained, respectively, from (4) and (10) given the policy functions $d(s,\gamma)$, $b'(s,\gamma)$, $\tau(s)$, $\bar{b}_I(s)$;
- the policy function $d\big(s,\gamma\big)$ satisfies (6), and $c\big(s,\gamma\big)$ and $b'(s,\gamma)$ solve the consumer's problem in (7) given the policy functions $\tau(s)$, $\bar{b}_I(s)$, $\bar{b}_E(s)$ and probabilities $p(s)$, with resulting value functions $W(s,\gamma)$, $V(s,\gamma)$, and $V^D(s,\gamma)$ as defined in (5), (7), and (9), respectively;
- the tax rate τ_{SS} balances the budget of the Social Security.

A block-recursive quasi-equilibrium is a recursive quasi-equilibrium in which the policy functions, value functions, and market tightness are independent of the distribution of submarkets.

Lemma 1. *If the policy function for the borrowing premium τ is independent of the distribution of submarkets, then a block-recursive quasi-equilibrium exists.*

The proof is in Appendix A.1. The argument is by backward induction as in Menzio, Telyukova, and Visschers (2016). A first part shows that the value and policy functions for a submarket with consumers at the end of their life-cycle does not depend on the distribution of submarkets. A second part shows that iterating backward preserves the block-recursivity. Intuitively, because lenders can target their search to consumers in a given submarket, they need to account only for the evolution of the conditional distribution of submarkets next period (i.e., conditionally on being in the submarket in which the lender searches this period). The evolution of the conditional distribution of submarkets need not depend on the unconditional distribution, so long as future policy functions do not depend on the unconditional distribution of submarkets. Hence, iterating backward from the block-recursive value functions at the end of the life-cycle delivers the block-recursivity of policy and value functions at all stages of the life-cycle.

3.6 The borrowing premium

In this section, we introduce the three types of equilibria that inform our policy counterfactuals and show that they are block-recursive, which makes them tractable to compute. The three equilibria differ in the way in which the borrowing premiums are determined. First, the unregulated equilibrium serves as a positive benchmark, in which Nash bargaining over the surplus from a match between entrant lender and borrower determines the credit terms. Second, we consider uniform and access-dependent caps in which regulation imposes constraints on the outcome of Nash bargaining. Third, the regulator chooses the premium for each submarket to maximize the surplus under the constraint that the resulting allocation is a quasi-equilibrium.

3.6.1 Positive benchmark

Our starting point is the unregulated equilibrium in which the borrowing premium is determined by Nash bargaining between matched consumers and entrant lenders. The borrowing premium $\tau(s)$ in credit offers sent to submarket $s = (j, b, \bar{b}, \tau, \theta, \eta)$ splits the expected surplus from the credit relationship according to the lender's bargaining power parameter θ_b,

$$\begin{aligned} \max_{\tau(s)} \Big\{ & \big(E_\gamma \left[\pi\big(j, b, \bar{b}_E, \tau(s), \theta, \eta, \gamma\big)\right] - \max(b, 0)\big)^{\theta_b} \\ & \big(E_\gamma \left[W\big(j, b, \bar{b}_E, \tau(s), \theta, \eta, \gamma\big) - W\big(s, \gamma\big)\right] \big)^{1-\theta_b} \Big\}, \end{aligned} \tag{11}$$

where the borrowing limit, $\bar{b}_E = \bar{b}_E(s; \tau(s))$ determined in (3), depends on the premium set in the bargaining process, $\tau(s)$.

A *block-recursive equilibrium (BRE)* given an initial distribution of submarkets is a block-recursive quasi-equilibrium $(V, V^D, W, d, c, b', \pi, \Pi, \bar{b}, \mu, p)$ given a policy function $\tau(s)$ and distribution Γ such that $\tau(s)$ solves (11) for π and W.

Proposition 2 (Existence of a BRE)**.** *A block-recursive unregulated equilibrium exists.*

The proof is in Appendix A.1.

3.6.2 Rate caps

In the second type of equilibrium that we study—the regulated equilibrium—a rate cap policy curtails the market power of lenders by imposing an upper bound on the borrowing premium. In particular, when bargaining over the borrowing premium in submarket s, entrant lenders and consumers are constrained to a borrowing premium less than the cap $\bar{\tau}(s)$. A rate cap is said to be uniform if the cap is the same in all submarkets, i.e., $\bar{\tau}(s) = \bar{\tau}$ for all s. We also consider access-dependent rate caps in which the cap may differ between submarkets in which consumers have access to credit and ones in which they do not, i.e., where $\bar{\tau}(s) = \bar{\tau}_a$ for submarkets in which consumers have access to credit $\tau < \infty$ and $\bar{b} \geq 0$, and $\bar{\tau}(s) = \bar{\tau}_n$ for submarkets in which consumers do not have access to credit $\tau = \infty$ and $\bar{b} = 0$.

The borrowing premium $\tau(s)$ in credit offers sent to submarket $s = (j, b, \bar{b}, \tau, \theta, \eta)$ splits the expected surplus from the credit relationship according to the lender's bargaining power parameter $\theta_b \in [0, 1]$, as described in equation (11), with the constraint that the resulting premium $\tau(s)$ does not exceed the cap $\bar{\tau}(s)$.

A *block-recursive equilibrium with a rate cap* $\bar{\tau}(s)$ given an initial distribution of submarkets Γ is a block-recursive quasi-equilibrium $(V, V^D, W, d, c, b', \pi, \Pi, \bar{b}, \mu, p)$ given a policy function $\tau(s)$ and distribution Γ such that $\tau(s)$ solves equation (11) with the constraint that the resulting premium $\tau(s)$ does not exceed the cap $\bar{\tau}(s)$ for Π and W.

Proposition 3 (Existence of BRE with rate cap)**.** *If the rate cap policy plan $\bar{\tau}(s)$ does not depend on the distribution of submarkets, then a block-recursive equilibrium with a rate cap $\bar{\tau}(s)$ exists.*

The proof, which is in Appendix A.1, follows the proof of existence of a BRE (Proposition 2) with an additional subtlety regarding the rate cap and the distribution of submarkets. Although the optimal rate cap may well depend on the distribution of submarkets at the time of the reform and on its evolution, given that the regulator commits to a time-invariant rate cap, the policy function for premiums—which solves (11) subject to the cap—is invariant to the evolution of the distribution of submarkets.

3.6.3 Optimal regulated equilibria

Since our welfare measure aggregates dissimilar agents' preferences, i.e., risk-averse consumers and risk-neutral lenders, we first convert all agents' preferences into comparable units using compensated variation.

Compensated variation. The welfare measure compares the surplus of heterogeneous consumers and lenders in a counterfactual scenario—due to a policy or the plan of the regulator—to the surplus in the status quo stationary equilibrium. To aggregate the surplus across heterogeneous consumers, we first translate welfare changes into dollar compensation. We use the onetime compensation transfer that makes the consumer indifferent between the counterfactual scenario and the status quo. Aggregating the onetime compensation transfer across consumers and lenders gives a welfare measure that focuses on efficiency while placing equal value to a marginal dollar in the budget constraint of the different consumers, as in the Kaldor-Hicks criterion. This criterion consists in maximizing the surplus as if lump-sum redistribution was available to address distributional considerations. Since the policy benefits some customers while making others worse off, in our quantitative analysis we also depict the distribution of welfare changes across heterogeneous groups of consumers and lenders.

To define the onetime compensation transfer, we augment the value function of the consumer in (7) to capture transfers, denoted by T. Let $V\big(s,\gamma;T\big)$ denote the value function with the following budget constraint instead of the one in (7):

$$c+b=y_{j,\theta,\eta,\gamma}-T+q_j\big(\tau\big)b'_{+}+q_j\big(0\big)b'_{-}. \tag{12}$$

Similarly, define $V^D\big(s,\gamma;T\big)$ by subtracting T from earnings, and $W(s,\gamma;T)$ based on the augmented value functions. Note that the transfer is not recurring because the continuation value W in (7) is without the transfer.

A counterfactual borrowing premium function affects consumer welfare through offers in the current period as well as through expectations of offers in the future. While our choice of notation suppresses the dependence of the value functions on expectations of offers throughout the text, we make this dependence explicit to define the total compensated variation (TCV). The TCV of the counterfactual borrowing premium function $\hat{\tau}(s)$, instead of the status quo borrowing premium function $\tau(s)$, is the transfer that makes the consumer indifferent between the counterfactual scenario and the status quo

$$E_\gamma[W\big(s|_{\tau=\hat{\tau}},\gamma;T=TCV\big)]=E_\gamma[W\big(s,\gamma;T=0\big)]. \tag{13}$$

A positive $TCV(s,\hat{\tau})$ indicates gains in terms of consumer surplus. Given two borrowing premium functions, the compensation is specific to the submarket at the time of the reform, and the determinants of the value function W. These determinants include the consumer policy functions, the anticipations of borrowing limit functions, and anticipations of probabilities of receiving a credit offer.

Optimal cap. As discussed in Section 3.6.2, we consider two types of rate caps, a uniform rate cap and an access-dependent one. In both cases, the rate caps are time-invariant and unanticipated before the reform is implemented.

We quantify the efficiency gains of the reform by aggregating the TCV for each consumer and the change in the net present value of incumbent lenders' profits (entrant lenders make zero profits by the free entry condition (1)).[12] We say that a reform is optimal if it maximizes the efficiency gains.

[12] We calculate the net present value of the TCV of future generations using the risk-free interest rate.

Submarket-specific rate caps. In this counterfactual exercise, the regulator chooses the borrowing premium for consumers with a credit card at the time of the reform only (we consider the regulation of new credit offers in the next counterfactual exercise). For each submarket with pre-existing accounts, the regulator sets the premium maximizing the objective function

$$\max_{\hat{\tau}}\Big\{TCV(s,\hat{\tau})+E_{\gamma}\big[\pi(s|_{\tau=\hat{\tau}},\gamma;T=TCV(s,\hat{\tau}))-\pi(s,\gamma;T=0)\big]\Big\}, \tag{14}$$

where $\pi(s,\gamma;\tau)$ denotes the incumbents' profits in (10) with policy functions $d, b', \bar{b}$ and probabilities of consumer poaching p that are part of a quasi-equilibrium given $\tau(s)$ and the initial stationary distribution of submarkets.

Submarket-and-time-specific rate caps. Next, we solve the problem of a regulator setting borrowing premiums on credit offers by entrant lenders each period so that it is sequentially optimal. In the first period of the reform, the regulator sets the borrowing premium as in (14). After the reform is implemented, however, there is perfect foresight. Therefore, the regulator trades off lower borrowing premiums with the associated lower likelihood of credit access. Formally, the regulator sets the policy sequentially to maximize the following efficiency gains from the reform

$$\begin{aligned}\max_{\hat{\tau}}\Big\{p(s|_{\tau=\hat{\tau}})\Big(TCV(s,\hat{\tau})+E_{\gamma}\big[\pi(s|_{\tau=\hat{\tau}},\gamma;T=TCV(s,\hat{\tau}))-\pi(s,\gamma;T=0)\big]\Big)-\\ \kappa\mu(s|_{\tau=\hat{\tau}})\Big\},\end{aligned} \tag{15}$$

where the frequency of customer solicitations μ and the likelihood of access p satisfy the free entry condition in (1).

The next proposition shows that the problem of a regulator that lacks the ability to commit to a policy plan remains tractable.

Proposition 4 (Block-recursive regulator's problem)**.** *The problem of the regulator setting submarket-and-time-specifc rate caps sequentially has a solution that is independent of the distribution of submarkets.*

The proof is in Appendix A.1. Intuitively, sequential optimality and targeted search jointly imply that when setting the premium, the regulator needs to account only for the evolution of the conditional distribution of submarkets, which is independent of the (unconditional) distribution of submarkets.

4 Functional forms, parameterization, and model validation

We first provide the functional forms in Section 4.1. To parameterize the quantitative model, we calibrate some parameters externally and others internally. In particular, we externally assign values to some of the parameters related to the preferences and life-cycle of consumers and of the credit card market. Internally calibrated parameters primarily target moments of the U.S. earnings distribution and moments related to the credit card market in 2019. We choose 2019 because it is ten years after the Great Recession and the Credit CARD Act of 2009, and prior to the 2020 COVID-19 pandemic. The externally calibrated parameters and the internally calibrated parameters are discussed in Sections 4.2 and 4.3, respectively. Finally, we validate the model against non-targeted moments in Section 4.4.

4.1 Functional forms

We assume a CRRA utility function given by

$$U(c) = \frac{\left(\dfrac{c}{1+\alpha_j}\right)^{1-\sigma}}{1-\sigma}, \tag{16}$$

where α_j is an adult equivalence scale for children and σ is the constant relative risk aversion. We assume a CRS matching function given by

$$M(u, v) = \frac{uv}{[u^{\zeta} + v^{\zeta}]^{1/\zeta}}, \tag{17}$$

where ζ determines the matching elasticity, u is the mass of consumers, and v is the mass of credit offers.

Social Security replaces a fraction λ of the average earnings from ages 30–64 given the permanent component θ and the unconditional expected values of η and γ ($E(\eta)$ and $E(\gamma)$). The Social Security function is given by

$$SS_{\theta} = \lambda\theta E(\eta)E(\gamma)\sum_{j=11}^{45} \nu_j/35. \tag{18}$$

4.2 Externally calibrated parameters

Table 1 presents the externally calibrated parameters. Panel A reports values for the risk aversion and risk-free interest rate. The risk aversion parameter is set to 2, a standard choice in the literature. A period is assumed to be one year, and the risk-free interest rate is set to 0.0344 (Gourinchas and Parker, 2002).

Panel B reports age-specific parameters. Consumers of age 1 in our model correspond to 25-year-olds in the data. The maximum life span, denoted by J, is set to 55 years, which corresponds to age 79 in the data. Retirement, denoted by j_r, is at age 41, which corresponds to age 65 in the data. The conditional survival probability ψ_j and the deterministic life-cycle productivity ν_j are estimated using data from the U.S. Census Bureau and the Global Repository of Income Dynamics (GRID), respectively. The adult equivalence scale is set to 0.3 using an estimate from the Organisation for Economic Co-operation and Development (OECD).

Table 1: Externally calibrated parameters

Parameter		**Notes**	**Value**
Panel A: General			
σ	Risk aversion	Standard in the literature	2
r	Risk-free interest rate	Gourinchas and Parker (2002)	0.0344
Panel B: Age-specific			
J	Maximum age	Life span, 25–79	55
j_r	Retirement age	Retire at 65	41
ψ_j	Conditional survival probability	U.S. Census Bureau	
ν_j	Life-cycle productivity	Global Repository of Income Dynamics	
$\alpha_{j=6,\dots,23}$	Adult equivalence scale	Organisation for Economic Co-operation and Development	0.3
Panel C: Credit card market			
ζ	Matching elasticity parameter	Herkenhoff (2019)	0.370
τ_c	Transaction cost	Agarwal et al. (2015)	0.029

Notes: The table presents parameters calibrated externally for the baseline model.

Finally, Panel C presents parameters related to the credit card market. The parameter that determines the matching elasticity in the credit card market ζ is set to 0.37, following Herkenhoff (2019). Herkenhoff (2019) estimates the matching elasticity parameter based on data on direct mail credit card offers from Synovate and credit applications from the Survey of Consumer Finances (SCF).[13] The transaction cost τ_c is set to 0.029 based on a direct estimate from data on operational costs in Agarwal et al. (2015).[14]

[13] In Appendix C.5, we perform sensitivity analyses for this parameter.

[14] See Table III and Appendix A.1.2 of the cited paper for more details. The operational costs include collection expenses and costs for servicing, card holder billing, processing interchange, processing payments, card issuing, authorizations, card administration, and outside services or outsourcing. We exclude marketing and acquisition costs from our estimate for the transaction cost because they are captured by the cost of posting credit offers in our model.

4.3 Internally calibrated parameters

We calibrate the remaining parameters in two steps. In the first step, we calibrate the productivity process for earnings, which does not require us to solve the whole model because the earnings distribution is determined exogenously based on the parameters of the earnings process. The idiosyncratic productivity process includes permanent, persistent, and transitory components. The idiosyncratic persistent productivity η is assumed to follow a $\log\ AR(1)$ process given by $\log\eta' = \rho_\eta \log\eta + \upsilon'$. The innovations for all three processes are assumed to be iid normal with mean zero and variances given by σ^2_θ, σ^2_η, and σ^2_γ, respectively. This is discretized using Tauchen (1986). We use 3, 11, and 5 grid points for the permanent, persistent, and transitory components, respectively. We jointly calibrate the four parameters of the earnings process, σ^2_θ, σ^2_η, σ^2_γ, and ρ_η, to target various moments of the U.S. earnings distribution in 2019. The parameter values and the target moments are presented in Table 2. For the target moments, we employ the Global Repository of Income Dynamics (GRID), which uses longitudinal administrative data to compute various statistics related to earnings. Our target moments are the standard deviations of 1- and 5-year changes in age-free log earnings, standard deviation of age-free log earnings over 3 years, and the autocorrelation of age-free log earnings from lags 1 to 5.

Table 2: Internally calibrated parameters: Step 1

Parameter	**Description**	**Value**
σ^2_γ	Variance transitory component	0.081
σ^2_η	Variance persistent component	0.108
σ^2_θ	Variance permanent component	0.337
ρ_η	Persistence	0.732
Target (year = 2019)	**Data**	**Model**
Standard deviation log earnings change: 1 year	*0.540*	*0.551*
Standard deviation log earnings change: 5 year	*0.764*	*0.749*
Standard deviation log permanent earnings	*0.799*	*0.804*
Auto-correlation log earnings lag-1	*0.788*	*0.800*
Auto-correlation log earnings lag-2	*0.728*	*0.736*
Auto-correlation log earnings lag-3	*0.692*	*0.689*
Auto-correlation log earnings lag-4	*0.664*	*0.655*
Auto-correlation log earnings lag-5	*0.639*	*0.630*

Notes: The table presents parameters of the stochastic earnings process as well as the target moments for the baseline model. The numbers in italics are the target moments. Data source: The Global Repository of Income Dynamics (GRID).

In the second step, we calibrate the six remaining parameters, which requires us to solve the whole model. The six parameters are the lender's bargaining power θ^b, the cost of credit offer κ, the discount factor β, the stigma χ, the scaling parameter for extreme value shocks ξ, and the Social Security replacement rate λ. They are calibrated jointly to match six targets in the year

2019 (Table 3). Although we describe a relation between each parameter and a target moment, we emphasize that we estimate them jointly because they affect other target moments as well.

Table 3: Internally calibrated parameters: Step 2

Parameter		Value	Target moment (year = 2019)	Data	Model
θ^b	Lender bargaining power	0.832	Average markup per card holder	*6.44*	*6.50*
κ	Cost of offer	0.001	Population w. CC | Lowest income quintile	*41.50*	*41.56*
β	Discount factor	0.926	Avg. total limit per card holder to avg. income	*27.13*	*27.09*
χ	Stigma	0.784	Charge-off rate	*3.70*	*3.70*
ξ	Scaling parameter	17.692	Defaults: health care, divorce, lawsuits	*44.81*	*44.99*
λ	Social Security replacement rate	0.280	Social Security to income per capita	*29.57*	*29.57*

Notes: The table presents the parameters calibrated internally as well as the target moments for the baseline model. The numbers in italics are the target moments. Data sources: The average "markup" is approximated as the average credit card interest rate per card holder (FRB G.19) net of the risk-free rate (annual returns on three-month Treasury bills from Federal Reserve Economic Data), charge-off rate (Federal Reserve Economic Data), and an estimate for the transaction cost (Table 1). The data on population with credit cards (CC) in the lowest income quintile and the average total limit per card holder to average income per household are from the Survey of Consumer Finances (SCF). The share of defaults related to health care, divorce, and lawsuits is estimated by Chakravarty and Rhee (1999) using the Panel Study of Income Dynamics (PSID). Social Security data are from the Bureau of Economic Analysis.

The lender's bargaining power is calibrated to match an approximate measure of an average markup, which is computed as the difference between the average credit card interest rate per card holder net of the risk-free rate, charge-off rate, and the transaction cost (6.44 percentage points). Figure 8 in Appendix B.1 illustrates, through numerical simulations, the sensitivity of the average markup to the lender's bargaining power. The cost of a credit offer is calibrated to match the population with credit cards in the lowest income quintile (41.50 percent).[15] The discount factor is calibrated to match the average total credit card limit per card holder to average income per household in the SCF for households who are 25 years and older (27.13 percent). The stigma associated with default is calibrated to match the charge-off rate (3.70 percent of outstanding credit). The scaling parameter for taste shocks is calibrated so that the fraction of "taste shock" defaults matches the fraction of defaults resulting from health care bills, divorce, and lawsuits (44.81 percent). Taste shock defaults are ones for which the Bellman value of repayment is greater than the

[15] We target the population with credit cards in the lowest income quintile (41.50 percent) rather than in the whole population (75.07 percent) for a cleaner calibration of the model. While the data reflect credit card use for both revolving credit and transactional benefits, our model abstracts from transactional benefits. Transactional benefits, however, are relevant mostly for high-income consumers, who tend to have a higher purchase volume and generate higher interchange income. Hence, targeting the population with credit cards in the lowest income quintile mitigates the concern that some credit card offers, mainly to high-income consumers, are for transactional benefits. In addition, our calibration improves the measurement of the lender's market power. This is because an alternative calibration that would target the whole population instead of the lowest income quintile could overstate the lender's bargaining power, which is a key input to our analysis. This is because the consumer's likelihood of receiving a credit offer would be higher in the alternative calibration, which would increase the competition that lenders face in the model. Consequently, to match the markup observed in the data, the lender's bargaining power in the alternative calibration would increase. In fact, the calibrated value for the lender's market power is 0.832 in our baseline calibration, whereas it would increase to 0.860 in the alternative calibration.

Bellman value of default (false positives). Finally, the Social Security replacement rate is calibrated to match average Social Security to disposable income per capita (29.57 percent).

4.4 Model validation

The first two columns of Table 4 compare non-targeted moments from U.S. credit card market data in 2019 to their counterparts in the baseline model.[16] The credit card moments are divided into five panels related to statistics on rates, issuance cost, credit usage, default, and access. The model does well in explaining non-targeted moments related to rates, issuance cost to income, default, and access but understates credit usage.[17]

Panel A of Table 4 shows that the model does reasonably well in explaining dispersion in rates: the model accounts for nearly 80 percent of the standard deviation observed in the data (6.32 percent in the data and 4.97 percent in the model). The last four rows of Panel A report the 10th, 20th, 80th, and 90th percentiles of the rate distribution. The model does well for all of them except the 10th percentile. Table 9 in Appendix B.2 shows that on-the-credit-search is an important feature that allows the model to generate high dispersion in rates. Without on-the-credit-search, the standard deviation of interest rates decreases from 4.97 to 2.60–2.98 percent depending on whether the model without on-the-credit-search is re-calibrated.[18]

Panel B of Table 4 shows that the acquisition cost to income in the model is close to the range estimated in the data. In the model, the acquisition cost is computed as the mass of credit offers in a period multiplied by the cost of one credit offer. The estimation for the data counterpart is described in footnote *a* of Table 4.[19]

Panel C of Table 4 reports statistics related to credit usage. Credit to income in the data is 5.60 percent, whereas it is 2.00 percent in the model. While the model does well in accounting for utilization rates of credit cards above 75 percent of the total credit card limit that a household has

[16]Some targeted moments are also included in the table because the second and third columns will be used later to compare the baseline model to the model with the optimal access-dependent cap. These variables are marked with a "†" sign, and the corresponding target estimates are italicized.

[17]In Panel C of Table 10 in Appendix B.5, we show that if we were to target credit usage, then the model would overstate the issuance cost to income, a key moment in our study. The appendix also shows that this result holds across various model specifications. Furthermore, in Appendix C.5, we evaluate the welfare implications of regulation in a calibration that targets credit usage. In comparison to this alternative, our baseline calibration leads to conservative estimates for the welfare gains.

[18]Galenianos and Gavazza (2022) and Nelson (2023) document high dispersion in credit card interest rates even within credit score bins, which suggests that dispersion in rates is not driven entirely by differences in the likelihood of delinquency, which is what is captured by a credit score. Our model result is complementary to their empirical finding.

[19]While the model does reasonably well in regard to acquisition costs relative to income, it overstates the acquisition cost relative to outstanding credit (0.5–0.8 percent in the data versus 2.96 percent in the model). This result is robust to several alternative specifications we have tried in which we target the same set of moments as in Table 3. See Appendix B.5 for more details.

Table 4: Model fit and steady state comparison to the optimal access-dependent cap

Credit statistic	(1) Data (2019)	(2) Baseline model	(3) Optimal cap | Access
	(unit = percentage or percentage points)		
Panel A: Rates			
Average "markup" per card holder†	*6.44*	*6.50*	4.80
Standard deviation	6.32	4.97	2.95
Rate: 10th percentile	7.50	10.72	8.64
Rate: 20th percentile	11.00	12.76	8.64
Rate: 80th percentile	22.00	20.92	15.64
Rate: 90th percentile	24.24	22.97	15.64
Panel B: Issuance cost			
Acquisition cost to income	0.03-0.05	0.06	0.03
Panel C: Credit usage			
Revolving credit to income	5.60	2.00	2.03
Population with CC debt	29.26	8.05	8.53
Utilization rate: above 50 percent	16.62	7.90	7.63
Utilization rate: above 75 percent	7.82	5.78	5.94
Panel D: Default			
Default rate	0.14	0.14	0.10
Charge-off rate†	*3.70*	*3.70*	2.49
Panel E: Credit access			
Population with CC	75.07	62.00	70.55
Avg. total limit per card holder to avg. income†	*27.13*	*27.09*	23.34

Notes: The table reports credit card market moments in the data and the steady states of the baseline model and the model with the optimal access-dependent cap. In the table "†" indicates moments that were targeted in the calibration of the baseline model (see Table 3 for details of the calibration). The numbers in italics are the target moments. Data sources: The average markup is approximated as the average credit card interest rate per card holder (FRB G.19) net of the risk-free rate (annual returns on three-month Treasury bills from Federal Reserve Economic Data), charge-off rate (Federal Reserve Economic Data), and an estimate for the transaction cost (Table 1). The data for rate dispersion and percentiles of the rate distribution, population with credit card (CC) debt, utilization rates, population with CC, and average total limit per card holder to average income per household are from the Survey of Consumer Finances (SCF). Revolving credit data are from FRB G.19. Disposable income data are from the Bureau of Economic Analysis. The acquisition cost estimates are based on Agarwal et al. (2015) and Consumer Financial Protection Bureau (2013).[a] Chapter 7 bankruptcy filings data are from the American Bankruptcy Institute.

[a] Agarwal et al. (2015) report that acquisition and marketing costs account for 0.5 percent of outstanding credit. Revolving credit to income in 2019 was 5.60. These two estimates imply the 0.03 estimate. Consumer Financial Protection Bureau (2013) found that the financial services industry spends approximately 17.5 billion dollars (2012) annually in marketing consumer financial products and services. This estimate is the sum of expenditure on awareness advertising equal to 5.5 billion dollars (e.g., spending to reach consumers via television, radio, and newspapers) and direct marketing equal to 12 billion dollars (e.g., internet advertising, direct mail). For awareness advertising, spending on credit cards accounts for 38.18 percent. We apply the same share for direct marketing, which is not reported. This leads to the 0.05 estimate.

(7.82 and 5.78 percent of card holders in the model and data, respectively), the model accounts for a significantly smaller share of households with credit cards who use more than 50 percent of their credit card limit and the share of households with positive credit card balances. These statistics imply that the model is not able to generate enough households who carry small positive balances.

Finally, Panels D and E of Table 4 report moments related to default and access. Panel D shows that although we targeted the charge-off rate in our calibration, the model is also able to account for the default rate (0.14 percent of households file for Chapter 7 bankruptcy annually in the data and model). The share of the population with credit cards is 75.07 percent in the data and 62.00 percent in the model. The model estimate is lower because in the calibration we targeted the population with credit cards in the lowest income quintile to mitigate the concern that our model abstracts from transactional benefits of credit card use (see footnote 15 for more details).

Thus far, we have discussed how our model does well in explaining several moments related to the credit card market in the aggregate. In Appendix B.3, we show that our model also rationalizes patterns in the cross-section by income and by age. In Appendix B.4, we show that our model accounts for borrowing premiums that are rigid with respect to one-year-ahead default risk, as documented by Dempsey and Ionescu (2024).

5 Quantitative evaluation of rate caps

We use the calibrated model of revolving credit lines to quantify efficiency gains from rate regulation. Specifically, in Section 5.1, we analyze the positive and the welfare implications of a uniform rate cap, in which all credit card premiums are subject to one common maximum. A uniform rate cap can generate welfare gains, that is, markups are inefficiently high. In Section 5.2, we analyze access-dependent caps and highlight how the revolving nature of credit lines matters for rate regulation. Finally, in Section 5.3, we analyze submarket-specific and submarket-and-time-specific caps to compute the gains from fine-tuning caps to characteristics beyond credit access and over time.

5.1 Uniform rate cap

In Section 5.1.1, we quantify the key trade-offs when choosing a uniform rate cap. In Section 5.1.2, we show that a uniform cap can generate welfare gains and that credit access plays a key role in determining the optimal uniform rate cap.

5.1.1 Positive implications of a uniform rate cap

To quantify the trade-off between rates and credit access—a central tension in the design of optimal rate cap policies—we depict in Figure 1 the average borrowing premium, the share of card holders for whom the cap is binding, the extensive margin of credit access (i.e., the share of the population with credit cards), and the intensive margin of credit access (i.e., borrowing limit per card holder)—for different stringencies of the uniform interest rate cap.

For each of the panels in Figure 1, a value of 10 on the x-axis refers to a maximum borrowing premium of 10 percentage points ($\bar{\tau} = 0.1$). From right to left, a panel depicts the comparative statics of tighter rate caps. The red-solid lines depict averages computed in the stationary distribution in the equilibrium with a uniform rate cap, whereas the blue-dashed lines depict averages computed from a partial equilibrium exercise in which the probabilities of credit offers are fixed at their values from the baseline economy without the rate cap.

While a uniform rate cap lowers the borrowing premium, it also lowers credit access, and this trade-off is nonlinear.[20] The slope of the average borrowing premium depicted by the red-solid line in Figure 1a is indicative of the effectiveness of the cap in lowering the average rate. The cap is increasingly effective, simply because the share of contracts for which the cap is binding is increasing (see Figure 1b for the share of binding accounts). Figure 1c shows that the extensive margin of credit access is highly nonlinear in the tightness of the cap. This nonlinearity comes from the competitive entry condition, which implies a hyperbolic relationship between the probability of credit access and profits. Intuitively, a rate cap lowers expected profits conditional on a match, which lowers the probability of credit offers at an increasing rate.

The intensive margin of credit access—the average credit limit per card holder—increases with tighter caps (see the red-solid line in Figure 1d). This is because the rate cap changes the distribution of consumer characteristics that have a credit card towards consumers that are less likely to default. To see this, notice that the blue-dashed line in Figure 1d depicts the average credit card limit keeping the extensive margin of credit access constant. With lower interest rates, lenders curb their exposure to default risk by tightening borrowing limits. In equilibrium, however, the rate cap induces lenders to target their credit offers to consumers that are less likely to default. As a result, the average credit limit per card holder is looser. In contrast, the average borrowing premium is not affected by the same selection channel, as can be seen by comparing the red-solid and blue-dashed lines in Figure 1a.

[20] The feature that the uniform cap reduces the population with credit cards, that is, the extensive margin of credit access, is consistent with recent empirical evidence on non-revolving consumer loans. For example, Cuesta and Sepúlveda (2021) analyze a policy change in Chile that lowered interest rate caps by 20 percentage points on unsecured non-revolving consumer loans and find that the policy decreased the number of loans.

Figure 1: Positive implications of a uniform cap on rates

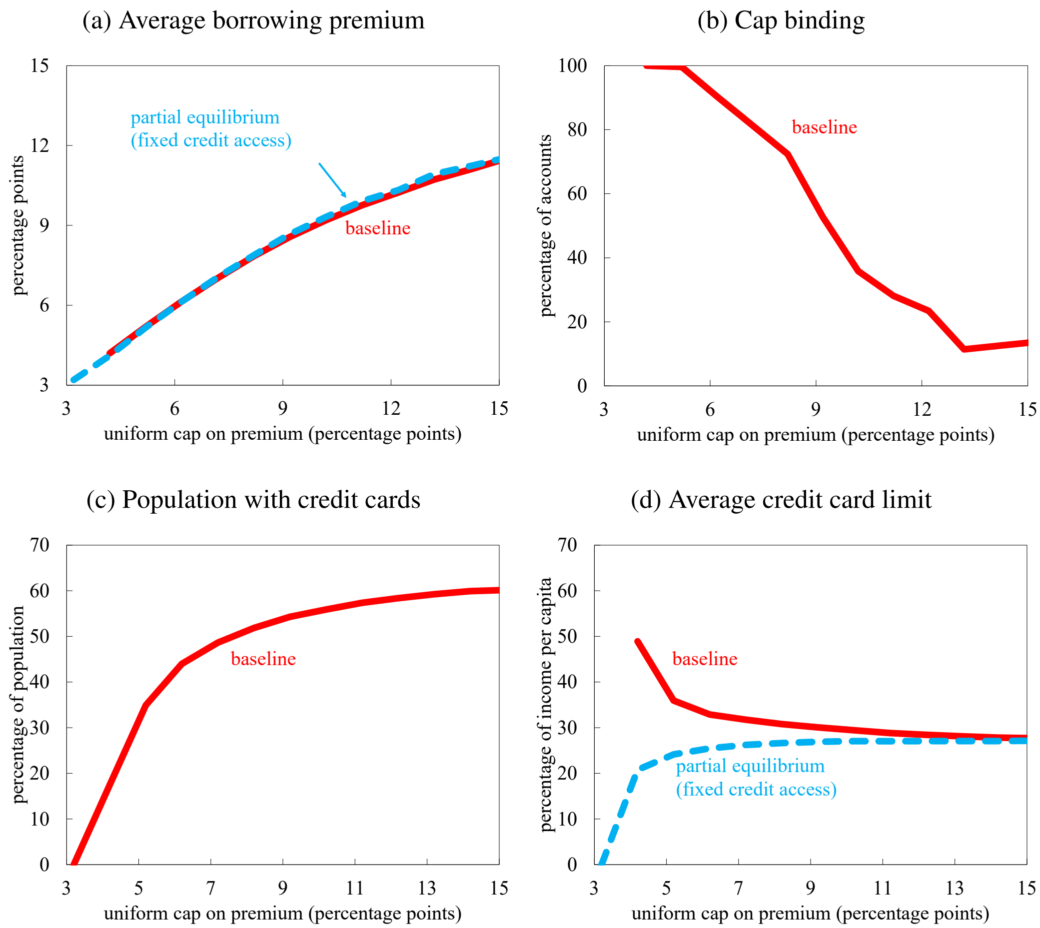

Notes: The figure depicts comparative statics with respect to a uniform rate cap computed in the stationary distribution in the equilibrium with a uniform rate cap. The x-axis refers to the cap on the borrowing premium $\bar{\tau}$. For example, a cap equal to 10 implies that the maximum premium over the risk-free rate is 10 percentage points. Figure 1a depicts the average borrowing premium per card holder, and Figure 1b depicts the share of card holders for whom the cap is binding. Figure 1c depicts the percentage of the population with credit cards. Figure 1d depicts the average borrowing limit for the population with credit cards. In Figures 1a and 1d, the red-solid lines depict equilibrium outcomes (the competitive entry condition (1) holds and determines credit access in equilibrium). The blue-dashed lines depict the outcome of an economy in which the probabilities of credit offers are fixed at their values from the baseline economy without the rate cap (i.e., the competitive entry condition (1) is relaxed).

5.1.2 Optimal uniform rate cap

We find that on average, markups are inefficiently high. Figure 2a, which depicts the total efficiency gains at different stringencies of the rate cap, illustrates the trade-off between consumer protection (via lowering rates) and credit access. As can be seen in the figure, the optimal uniform rate cap restricts borrowing premiums to be below 8.20 percentage points, which is to be compared to the average premium at 13.1 percentage points in the baseline equilibrium without the cap. Tightening the cap beyond 8.20 percentage points, however, reduces economic efficiency and may in fact bring welfare losses relative to the status quo. This is because, as shown in Figure 1c, credit access decreases sharply for premiums capped below 6 percentage points.[21]

Figure 2: Welfare implications of a uniform cap on the borrowing premium

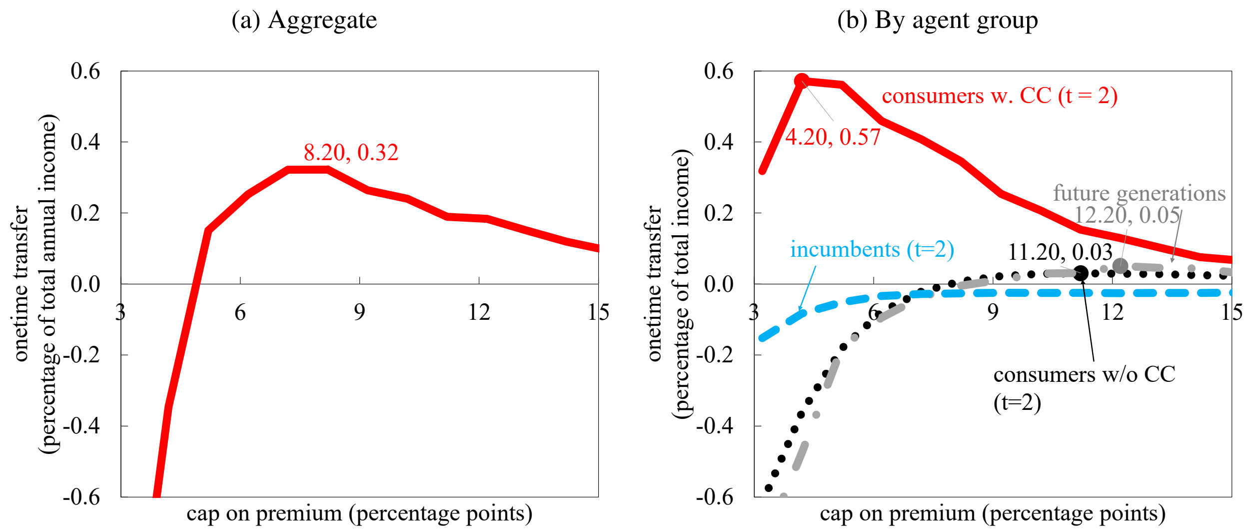

Notes: Figures 2a and 2b depict total onetime transfers as a percentage of total annual income for all agents and by agent group, respectively. The groups of agents are (1) consumers with credit cards in the first period of the transition (red-solid line), (2) consumers without a credit card in the first period of the transition (black-dotted line), (3) future generations whose welfare is discounted at the rate $1/(1+r)$ (gray-dash-dotted line), and (4) incumbent lenders in the first period of the transition (blue-dashed line). The x-axis refers to the uniform rate cap (maximum borrowing premium). For example, a cap equal to 10 implies that the maximum borrowing premium is 10 percentage points. The cap is applied to both pre-existing accounts and new issuances. In the figure "t = 2" refers to the first period of the transition.

The welfare benefits from a uniform rate cap are, however, strongly heterogeneous across consumers based on their credit access at the time of the reform. Figure 2b decomposes the aggregate welfare changes for (1) consumers with a credit card at the time of the reform, (2) future generations and consumers without a credit card at the time of the reform, and (3) incumbent lenders. Importantly, the optimal uniform cap for consumers who have a credit card at the time of the reform is tight (red marker in Figure 2b), and the optimal uniform cap for consumers who do not have a credit card at the time of the reform is lax (black and gray markers in Figure 2b). Furthermore, while the consumers with a credit card benefit significantly more from a tight cap, at those levels of stringency, the consumers without a credit card are significantly worse off. Consequently, the optimal uniform rate cap achieves moderate welfare gains, which amount to a onetime transfer worth 0.32 percent of initial steady state annual income.

[21]Figure 2a shows that a rate cap can lead to gains in efficiency. That is, there exists a set of lump sum transfers that make all agents weakly better off. In Appendix C.2, we show that our findings are robust to evaluating counterfactuals based on policy functions of consumers who receive the compensating transfers, instead of policy functions of uncompensated consumers—as we report in the main text.

5.2 Access-dependent rate cap

In this section, we turn our attention to access-dependent caps, whereby there are two caps. One cap is for rates charged on pre-existing accounts and for rates offered for balance transfers. The other cap is for rates on credit offers to consumers without a credit card. In Section 5.2.1, we evaluate the gains from the optimal access-dependent cap, and discuss how it generates gains by tailoring the stringency of the cap to credit access. In Sections 5.2.2 and 5.2.3, we compare the magnitude of the gains to the value of the credit card market and analyze the distributional implications. In Appendix C.3, we elaborate on the positive implications of the optimal access-dependent cap across steady states.

5.2.1 Optimal access-dependent rate cap

The first row of Table 5 shows that it is optimal to set a tight cap for those with a credit card and a lax cap for those without a credit card. The optimal cap for those with a credit card (that is, for pre-existing accounts and for balance transfers) is 5.2 percentage points, whereas the optimal cap for those without a credit card is 12.2 percentage points. Furthermore, the optimal access-dependent cap leads to gains that amount to a onetime transfer worth 0.62 percent of annual income. These gains are twice as much as the gains from the optimal uniform rate cap.

Table 5: Optimal access-dependent caps and welfare implications

Optimal cap | access for:	**Optimal cap | access** (unit = percentage points)		**Welfare gains for:** (unit = percentage of total annual income)				
	No CC	**CC**	**Consumers (t = 2)**		**Future generations**	**Incumbents (t = 2)**	**Aggregate**
			w. CC	**w/o CC**			
aggregate	12.2	5.2	0.54	0.09	0.09	-0.10	0.62
consumers w. CC (t=2)	10.2	4.2	0.58	0.08	0.07	-0.16	0.57
consumers w/o CC (t=2)	10.2	5.2	0.52	0.09	0.07	-0.10	0.59
future generations	12.2	4.2	0.57	0.08	0.09	-0.15	0.58

Notes: In the table, each row reports results when the access-dependent cap is chosen to maximize welfare for the group of agents specified in the given row. The table reports the optimal access-dependent cap on the premium as well as onetime transfers as a percentage of initial steady state total annual income by agent group and for all agents under each policy. The groups of agents are (1) consumers with credit cards in the first period of the transition, (2) consumers without a credit card in the first period of the transition, (3) future generations whose welfare is discounted at the rate $1/(1+r)$, and (4) incumbent lenders in the first period of the transition. The cap is applied to both pre-existing accounts and new issuances. In the table "t = 2" refers to the first period of the transition; "No CC" refers to the cap for those without a credit card, and "CC" refers to the cap for pre-existing accounts and for balance transfers.

The additional gains accrue from setting a stringent cap on rates of pre-existing accounts at the time of the reform, as well as on rates offered for balance transfers. These additional gains can be seen by comparing the access-dependent cap gains decomposed by agent groups in the first row of Table 5 to the uniform cap gains decomposed by agent groups in Figure 2b. Consumers with credit cards at the time of the reform are the ones who benefit the most from the policy. Their gains amount to a onetime transfer worth 0.54 percent of initial steady state annual income, which is nearly the gains they would realize from the uniform rate cap that is optimal for them (0.57 percent, red marker in Figure 2b). However, most of these gains are not due to a redistribution from incumbent lenders to these consumers because the losses to the incumbent lenders amount to only 0.1 percent.[22] Future generations and consumers without a credit card at the time of the reform also experience larger gains from the optimal access-dependent cap in comparison to the maximum gains they could realize from a uniform rate cap. For each of these groups, the gains from the access-dependent cap amount to 0.09 percent, which is approximately two-fold of the maximum gains from the uniform rate cap (0.03–0.05 percent, black and gray markers in Figure 2b).

Figure 3: Transition path to optimal access-dependent cap

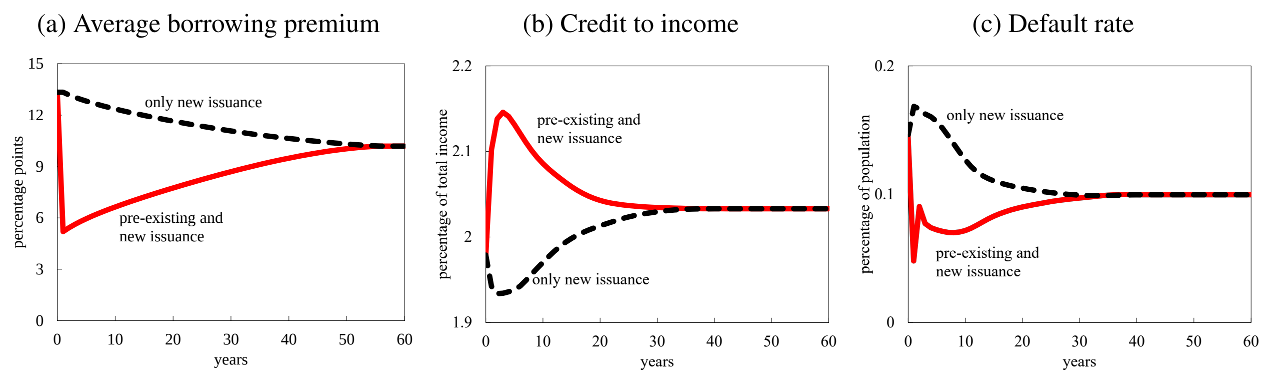

Notes: The figure depicts the transition path for the following variables after the implementation of the optimal access-dependent cap that is chosen to maximize welfare for all agents (see row 1 of Table 5): the average borrowing premium, total revolving credit as a percentage of total income, and the total number of defaults as a percentage of the total population. The red solid line refers to the case in which the access-dependent cap is applied to both pre-existing accounts and new issuances, whereas the black-dashed line refers to the case in which the access-dependent cap is applied to only new issuances.

Gains from targeting pre-existing access. The optimal access-dependent cap lowers the rate on pre-existing accounts for consumers with credit cards (see red-solid line of Figure 3a). This reduction in rates curbs the lenders' market power, which reduces the deadweight loss by allowing consumers to tilt consumption to the present, smooth consumption given increasing life-cycle earnings, and better insure themselves against earnings risk. This leads to increased borrowing, as depicted by the red-solid line in Figure 3b.[23]

[22] In reality, incumbent lenders making losses on a pre-existing account due to the policy change may want to close that account, which is not an option in our model. However, in our analysis, the net positive gains indicate that the incumbent lenders can be compensated for their losses so that they do not close the pre-existing account.

[23] In the quantitative model, lowering rates is also beneficial because it alleviates inefficiencies due to rigid rates (e.g., deadweight loss of default). The number of defaults decreases along the transition path, as depicted by the red-solid line in Figure 3c. In the next section, we argue that the primary driver of the efficiency gains from regulating

Our model of revolving credit reveals larger efficiency gains from regulating pre-existing accounts relative to those from regulating new credit card issuances. In Table 6, we compare the gains from the model with the optimal access-dependent cap applied to both pre-existing accounts and new issuances to an alternative in which the same optimal access-dependent cap is applied only to new issuances (the takeaway is not sensitive to re-optimizing the access-dependent cap on new issuances). We find that the gains are dampened from 0.62 percent to 0.16 percent when the cap is not applied to pre-existing accounts. The intuition is based on Figure 3, which plots the transition path for the borrowing premium, credit, and default rate when the cap is implemented under the two scenarios. If the cap is applied only on new issuances, the average borrowing premium on credit cards is slow to update (black-dashed line in Figure 3a). This inertia in the average borrowing premium is due to the long-term nature of revolving credit contracts. In turn, we no longer observe a rise in credit or a fall in default during the transition (black-dashed lines in Figures 3b and 3c).

Table 6: Role of enforcing the cap on pre-existing accounts

Regulation type	**Cap | Access** (unit = percentage points)		**Welfare gains for:** (unit = percentage of total annual income)				
	No CC	**CC**	**Consumers (t=2) w. CC**	**w/o CC**	**Future generations**	**Incumbents (t=2)**	**Aggregate**
Optimal cap | access on pre-existing and new issuances (Table 5, row 1)	12.2	5.2	0.54	0.09	0.09	-0.10	0.62
Optimal cap | access on only new issuances	12.2	5.2	-0.03	0.09	0.09	0.02	0.16

Notes: The table reports the access-dependent cap on the borrowing premium as well as the onetime transfers as a percentage of initial steady state total annual income by agent group and for all agents under two regulations. The groups of agents are (1) consumers with credit cards in the first period of the transition, (2) consumers without a credit card in the first period of the transition, (3) future generations whose welfare is discounted at the rate $1/(1+r)$, and (4) incumbent lenders in the first period of the transition. The two regulations are (1) optimal access-dependent cap on both pre-existing accounts and new issuances chosen to maximize welfare for all agents (see Table 5, row 1) and (2) optimal access-dependent cap on only new issuances chosen to maximize welfare for all agents. In the table, "No CC" refers to the cap for consumers without a credit card; in row 1, "CC" refers to the cap for pre-existing accounts and for balance transfers, and in row 2, "CC" refers to the cap on balance transfers only.

Gains from alleviating contract dilution due to balance transfers. In addition to lowering rates, the contingency of the optimal access-dependent cap is valuable to both future generations and consumers without a credit card at the time of the reform because it alleviates the contract dilution problem. That is, that the stringent cap on rates offered to consumers when they have a credit card reduces the incentive for new lenders to offer balance transfers. From the perspective of a consumer without a credit card, the access-dependent cap reduces excess offers to their future selves when they have a credit card, and alleviates the time inconsistency that arises due to their

rates is to limit the market power of lenders (see the discussion surrounding Table 7).

lack of commitment to a credit relationship.[24] As a result, a consumer without a credit card receives more offers with the optimal access-dependent cap than with the optimal uniform rate cap (see the red-solid and black-dotted lines in Figure 4). This increase in credit offers arises even though the optimal uniform rate cap is less stringent than the optimal access-dependent cap. Young consumers without credit cards, in particular, receive more offers with the access-dependent cap than they do even in the unregulated equilibrium (see the red-solid and blue-dashed lines in Figure 4a).

Figure 4: Probability of credit offer in the baseline, optimal uniform cap, and optimal access-dependent cap

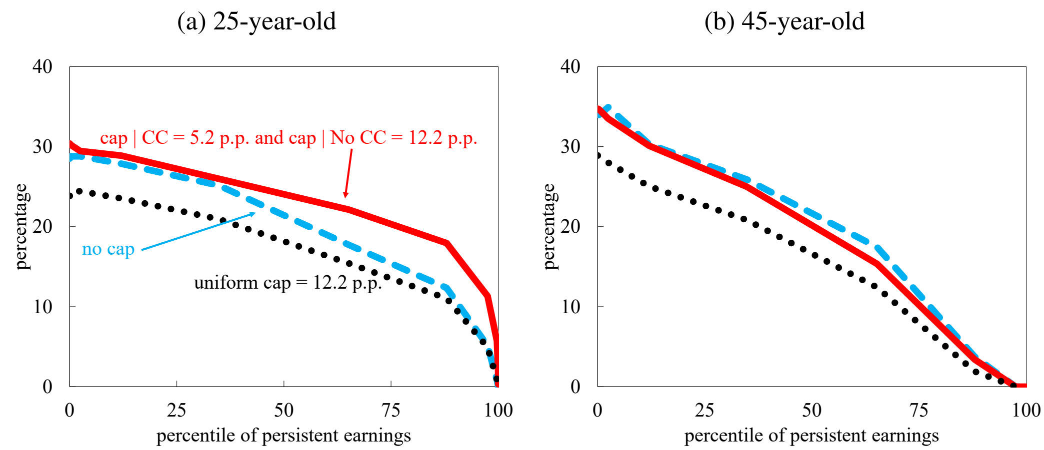

Notes: The figure depicts the probability of receiving a credit card for a 25-year-old and a 45-year-old who has no credit card, who has no assets, and has medium permanent earnings productivity by their persistent component of earnings on the x-axis under three scenarios: (1) baseline (blue-dashed line labeled as "no cap"), (2), uniform cap (black-dotted line), and (3) optimal access-dependent cap (red-solid line).

To summarize, an access-dependent cap generates larger efficiency gains across the consumer groups that we consider. Consumers with pre-existing accounts, as well as those without a credit card at the time of the reform and future generations, benefit from a stringent cap when they have a credit card. As a result, while these two consumer groups disagree on the stringency of a uniform cap, they agree on the stringencies of an access-dependent cap. This result is verified in rows 2 to 4 of Table 5, which report the optimal access-dependent cap chosen for the respective consumer groups.

5.2.2 Gains relative to the credit card market

Although the efficiency gains from the optimal access-dependent cap are arguably small relative to the size of the economy (i.e., 0.62 percent of annual income), the gains are large relative

[24] The time-inconsistency element is quantified in Appendix C.4. Table 11 shows that consumers without credit cards and future generations would benefit from committing to not search on-the-credit. The gains sum to a onetime transfer worth 0.09 percent of annual income for these two consumer groups.

to the size of the credit market. To gauge the size of the welfare gains relative to the size of the credit card market, we compute the gains from a transition from an economy without credit cards to the baseline with credit cards. The gains amount to a onetime transfer worth 2.62 percent of initial steady state annual income. Therefore, the gains from the optimal access-dependent cap are worth roughly 25 percent ($\approx 0.62/2.62$) of the value of the credit card market. The small gains relative to the size of the economy are to be somewhat expected because savers, who represent a significant fraction of the population, are not significantly affected by the rate cap. We expand on this in the next section and show that some consumers experience large gains from the optimal access-dependent cap.[25]

5.2.3 Distributional implications

In what follows, we go beyond the broad groups of agents in determining who benefits from the policy. To do so, we plot the distribution of welfare changes due to the optimal access-dependent rate cap and the associated consumer characteristics for the consumers in the first period of the transition.

While the optimal access-dependent rate cap brings in aggregate efficiency gains, it also has strong distributional implications. Figure 5a plots the distribution of welfare gains and losses for the consumers in the first period of the transition. The x-axis is an individual's onetime compensation transfer normalized by their annual income. A majority of consumers (roughly 50 percent) are barely affected by the policy. However, some consumers experience large gains: roughly 4 percent of consumers experience gains worth 5 percent of their income (rightmost bin). In contrast, only 0.003 percent of consumers experience large losses worth 5 percent of their income (leftmost bin). Furthermore, among consumers who experience losses (14 percent), nearly 85 percent experience losses that are smaller than 0.4 percent of their annual income.

We dive deeper by looking into the distribution of individual characteristics across welfare changes. Figures 5b, 5c, 5d, and 5e depict the average age, average net assets, average income, and the share of credit card holders across welfare changes due to the reform. The consumers that are barely affected tend to be older than 50 and asset rich. The consumers that benefit the most from the policy tend to have the lowest net assets in comparison to everyone else. Furthermore, in comparison to those barely affected by the policy, the biggest beneficiaries tend to be younger, have lower income, and are more likely to have a credit card.

[25] The gains relative to the size of the economy that we find are commensurate with estimates in other papers. For example, Herkenhoff (2019) finds gains equivalent to 0.12 percent of consumption to move from the level of credit access in 1977 to the level of credit access in 2010 (the share of households with access to unsecured revolving credit increased from 37.6 to 65.0 percent). Athreya, Tam, and Young (2012) show that improvements in information available to lenders accounts for a significant share of the rise in unsecured consumer credit and bankruptcies since the 1980s, and find that the welfare gains amount to less than 0.2 percent.

Figure 5: Distribution of welfare gains and corresponding consumer characteristics

(a) Population share

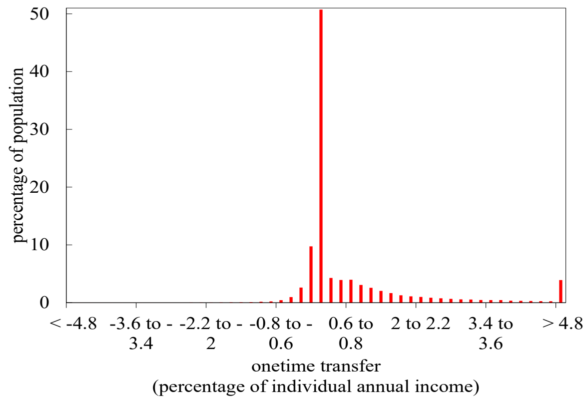

(b) Average age

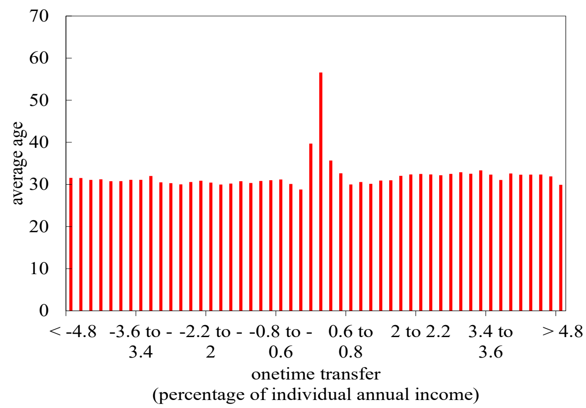

(c) Average net assets

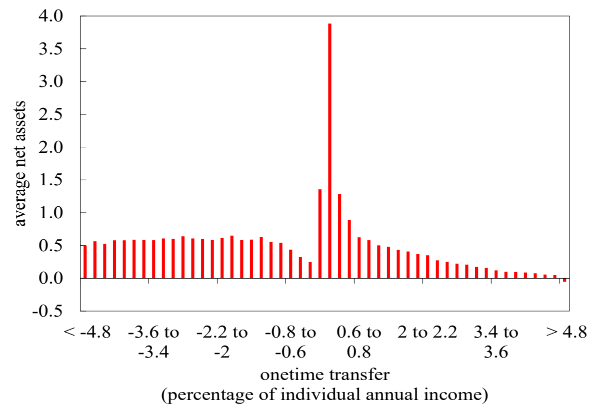

(d) Average income

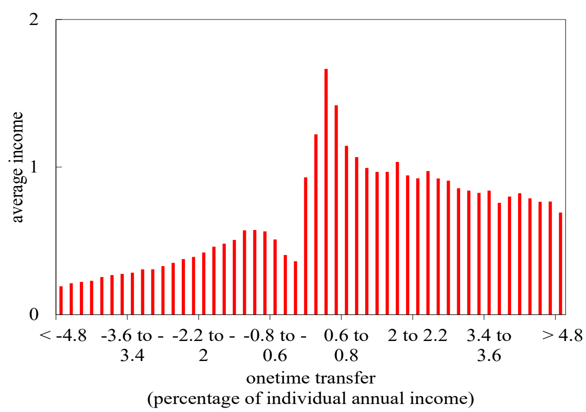

(e) Population with credit cards

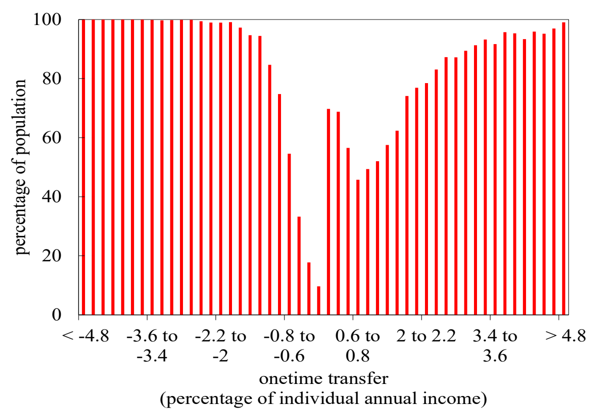

Notes: Figure 5a plots the histogram of welfare gains and losses following the implementation of the optimal access-dependent cap that is chosen to maximize welfare for all agents (see row 1 of Table 5). Welfare is computed for the consumers in the period of the transition. The x-axis is the onetime equivalent transfer normalized by individual annual income. The y-axis is the share of population for a given bin on the x-axis. Figures 5b, 5c, 5d, and 5e depict key characteristics associated with consumers in each bin: average age, average net assets, average income, and population with credit cards, respectively.

5.3 Submarket-specific and time-specific rate caps

In this section, we study the efficiency gains from fine-tuning caps to cross-sectional characteristics and over time. In Section 5.3.1, we study the welfare gains from tailoring the stringency of rate caps to all observable characteristics—i.e., age, assets, earnings, and the terms of current credit access. We find that the gains from fine-tuning the policy to consumer characteristics beyond credit access are relatively small. Furthermore, we explain why fine-tuning the regulation beyond access leads to relatively small gains, by analyzing the key inefficiencies that determine the optimal rate. We also quantify the extent to which the gains are attributable to these inefficiencies. In Section 5.3.2, we show that a regulator who re-optimizes over time is subject to the same time-inconsistency problem as the consumer. We find that this added flexibility is welfare dampening in comparison to committing to a time-invariant cap.

5.3.1 Optimal submarket-specific rate caps

We argue that fine-tuning rates to observable characteristics beyond access generates relatively small gains. In this section, we focus our analysis on gains from regulating pre-existing accounts only. We extend the analysis beyond pre-existing accounts in the next subsection. The first two rows of Table 7 show that an optimal uniform cap on only pre-existing accounts achieves 80 percent ($\approx 0.522/0.647$) of the maximum gains achieved by submarket-specific caps on pre-existing accounts. That is, while the gains from submarket-specific caps on pre-existing accounts are necessarily greater than those from the optimal uniform cap applied only on pre-existing accounts, the additional gains are relatively small.[26]

The additional gains from fine-tuning beyond credit access are small because, as we show in Figure 6, the distribution of optimal submarket-specific rates can be well approximated by a uniform rate cap on pre-existing accounts. This is because the uniform cap forces the bunching of rates at the cap, and our quantitative model implies a highly concentrated distribution of optimal submarket-specific rates. In Figure 6, although the distribution of rates in the unregulated equilibrium is dispersed (blue bars), the regulator sets the submarket-specific rate caps (red bars) to the zero profit risk-free premium of lending for the majority of accounts. Our welfare computations indicate that the bunching of rates on pre-existing accounts at the optimal uniform cap (imagine stacking up blue bars at the black-dashed line) is close, in terms of welfare, to the bunching of rates at the zero profit risk-free premium (the tallest red bar).[27]

[26] The optimal uniform cap on pre-existing accounts leads to a cap of 4.20 percentage points. This cap on pre-existing accounts is tighter than the one implied by the optimal uniform cap imposed on pre-existing accounts and new issuances (see the peak of Figure 2a).

[27] The zero profit risk-free premium of lending refers to the zero profit premium implied by the transaction cost of lending with zero default risk. When there is no default risk, the zero profit premium is equal to $(1 + r)/(1 - \tau_c(1 + r)) - 1 - r$.

Table 7: Welfare implications of submarket-specific caps: pre-existing accounts only

Regulation type	**Welfare gains for:** (unit = percentage of total annual income)				
	Consumers (t = 2)		**Future**	**Incumbents**	**Aggregate**
	w. CC	**w/o CC**	**generations**	**(t = 2)**	
Optimal uniform cap	0.664	0	0	-0.142	0.522
Submarket-specific caps	0.784	0	0	-0.137	0.647
Incumbent NPV ≥ 0	0.664	0	0	-0.098	0.566
One-off renegotiation	0.101	0	0	0.076	0.177
Incumbent NPV ≥ 0 or surplus ≥ 0	0.671	0	0	-0.096	0.575

Notes: The table reports the onetime transfers as a percentage of initial steady state total annual income by agent group and for all agents. The regulation is applied to only pre-existing accounts, and results are reported for five regulations: (1) optimal uniform cap on only pre-existing accounts chosen to maximize welfare for all agents, (2) optimal submarket-specific caps chosen to maximize welfare for all agents, (3) optimal submarket-specific caps subject to weakly positive net present value (NPV) of profits to the incumbent lender, (4) consumers and lenders are allowed to renegotiate on pre-existing accounts one-off, and (5) optimal submarket-specific caps subject to weakly positive net present value of profits to the incumbent lender or the incumbent lender is weakly better off with the regulation. The groups of agents are (1) consumers with credit cards (CC) in the first period of the transition, (2) consumers without a credit card in the first period of the transition, (3) future generations whose welfare is discounted at the rate $1/(1+r)$, and (4) incumbent lenders in the first period of the transition. By construction, the welfare changes for consumers without a credit card and for future generations are zero.

Figure 6: Distribution of borrowing premiums: pre-existing accounts

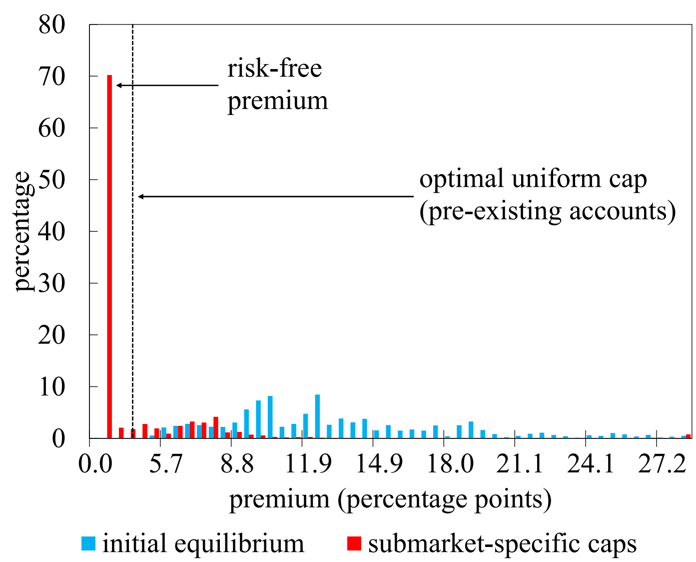

Notes: Figure 6 depicts the distribution of borrowing premiums for consumers with a credit card in the initial stationary equilibrium in the following cases: initial stationary equilibrium (no regulation) and submarket-specific rate caps that yield maximum efficiency gains.

Three sources of inefficiencies make it optimal to set a stringent cap for most submarkets. First, the lenders have market power determined by their bargaining power, which is calibrated to match markups observed in the data. Second, credit card rates are rigid. For example, the rates are not contingent on ex-post default risk. And third, consumers face partially uninsurable labor earnings risk. We base our discussion of the three sources of inefficiencies on the regulator's objective functions in the quantitative model, which we depict in Figure 7. Additionally, we analytically corroborate the quantitative results using a stylized model in Appendix A.2.

First, the regulator has an incentive to lower rates to the extent that the lenders' market power is inefficiently high. To isolate the role of the lenders' market power, we compute the optimal rate cap in the absence of earnings risk or default risk. As expected, in this case, the optimal premium is equal to the zero profit premium. That is, rates are priced competitively at a price equal to marginal cost (see the red-solid line labeled with no default and no uncertainty in Figure 7a, which depicts the total surplus as a function of the premium for a 20-year-old consumer with zero assets and a credit card with a prohibitively high borrowing premium). Therefore, when the lender has market power, unregulated rates are too high, which the regulator has an incentive to reduce.

Figure 7: Submarket-specific caps: regulator's objective functions

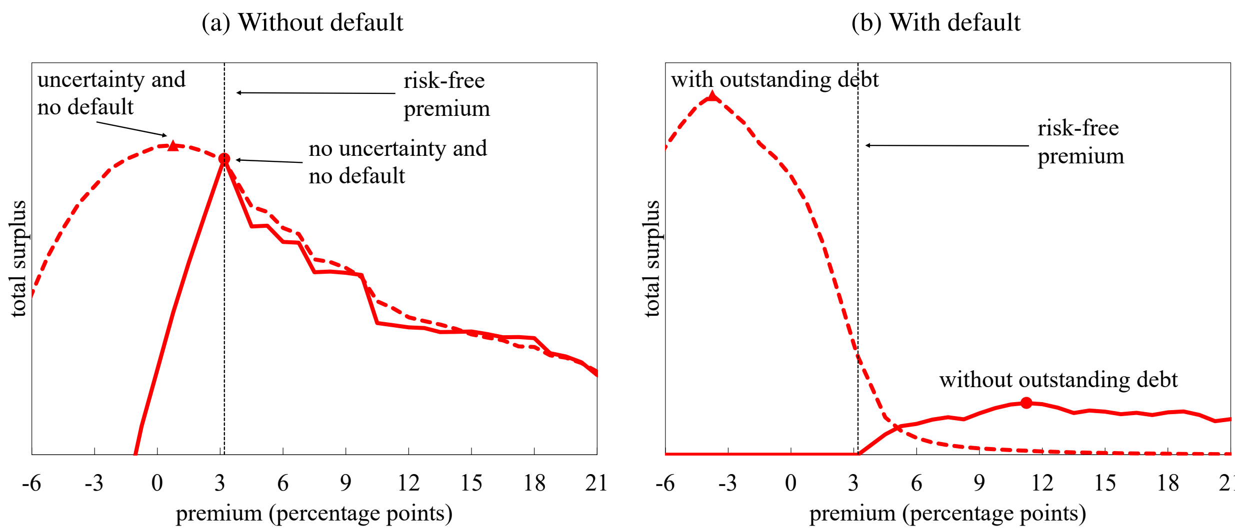

Notes: Figures 7a and 7b depict objective functions of the regulator under various cases. In Figure 7a, there is essentially no default (the stigma cost of default is equal to $1e38$ but the taste shocks are still active), and objective functions are depicted for a 20-year-old who has zero assets and a credit card with a prohibitively high premium. In the figure, in addition to the lack of default, the labels represent the following cases: (1) the line labeled with no uncertainty and no default refers to the case without earnings uncertainty and a large enough limit, and (2) the line labeled with with uncertainty and no default refers to the case with earnings uncertainty incorporated into (1). In Figure 7b, we allow for default (the stigma cost of default is equal to the baseline calibration value) and earnings uncertainty, and objective functions are depicted for a 20-year-old who has a credit card with a prohibitively high premium. In the figure, the line labeled without outstanding debt refers to a consumer who has zero net assets, and the line labeled with outstanding debt refers to a consumer who has a positive amount of outstanding debt.

Second, to address the incompleteness of credit card contracts due to rigid rates, the regulator has an incentive to tailor the rate to the credit worthiness of the consumer. A lower rate for indebted consumers acts as a haircut averting a costly default, whereas a higher rate prior to borrowing compensates lenders for the higher marginal cost of lending. We illustrate this incentive using the quantitative model in Figure 7b, which plots the total surplus as a function of borrowing premium for a 20-year-old consumer with a credit card with a prohibitively high borrowing premium in a framework with earnings uncertainty and default risk. The dashed-red line labeled with outstanding debt shows that the total surplus peaks below the risk-free premium when the consumer has an outstanding debt balance. In contrast, default risk adds a premium on new debt (see the red-solid line labeled without outstanding debt in Figure 7b). We note that in the quantitative welfare analysis of submarket-specific rates in Table 7, the optimal premiums are restricted to be weakly above the risk-free premium of borrowing.

Third, uninsurable earnings risk puts downward pressure on the efficient lending rate (in Figure 7a the red-dashed line labeled with uncertainty and no default peaks at a borrowing premium below the risk-free premium). We also establish this result in the stylized model without default risk in Appendix A.2. Through the lens of the stylized model, we highlight how the regulator has an incentive to set rates below the risk-free premium to facilitate insurance in response to low earnings realizations.

We argue that a large share of the efficiency gains from capping rates can be attributed to addressing inefficiencies due to the lenders' market power. While there is no simple way to attribute efficiency gains to different frictions due to nonlinearities in our model, we use counterfactual exercises on regulating pre-existing accounts—reported in the last four rows of Table 7—to attribute the welfare gains from regulation to the different sources of inefficiencies. The counterfactual with submarket-specific rates in which the net present value of profits to the incumbent lender cannot be reduced below zero leads to efficiency gains equal to 0.566 percent, which amounts to 87 percent of the maximum gains ($\approx 0.566/0.647$). We view this as an upper bound on the efficiency gains that can be attributed to curbing the lenders' market power. That is, as we have discussed, eliminating market power is achieved by ensuring lenders make no profits. Gains achieved by reducing profits below zero must be associated with another inefficiency (e.g., uninsurable earnings risk). Next, based on the counterfactual exercise in which consumers and incumbent lenders renegotiate one-off on the existing terms, we attribute at most 27 percent of the efficiency gains to overcoming the rigidity of rates for the duration of the relationship ($\approx 0.177/0.647$). Using the counterfactual in the last row, we attribute 11 percent of the gains to alleviating inefficiencies solely due to partially uninsurable labor earnings risk.[28] Together, the gains attributed to inefficiencies due to the rigidity in rates and partially uninsurable earnings risk amount to at most 38 percent ($\approx 27 + 11$). We view the remaining 62 percent as a lower bound on the inefficiencies that can be attributed to market power. Therefore, roughly 62 to 87 percent of the efficiency gains are attributable to the

[28] In the counterfactual, the regulator chooses submarket-specific rates that either lead to weakly positive net present value of profits to the incumbent lender or make the incumbent lender weakly better off. The residual gains of this counterfactual accrue in cases in which the incumbent lender makes a loss and is worse off, which echoes the case in which the regulator chooses the premium below marginal cost to provide partial insurance against earnings risk (red-dashed line labeled with uncertainty and no uncertainty in Figure 7a). As a result, we attribute the residual gains to partially uninsurable labor earnings risk ($\approx 1 - 0.575/0.647$)

reduction in lender market power.

5.3.2 Optimal submarket-and-time-specific rate caps

In this section, we study the fine-tuning of rate regulation over time. We study the submarket-specific rate cap policy that is optimal on pre-existing accounts in the first period of the transition and sequentially optimal for new issuances. The problem is computationally tractable because sequential optimality makes the regulator's problem block-recursive.[29] Our choice to allow the cap to be re-optimized every period is motivated by two considerations. First, in practice, the regulatory authority may be tempted to revise its policy so that it is sequentially optimal. Second, such an exercise quantitatively evaluates the time-inconsistency problem of the regulator (i.e., there are benefits to committing to limit the frequency of customer solicitations for balance transfers).

Table 8: Welfare implications of submarket-and-time-specific caps

	Consumers (t = 2)		**Future**	**Incumbents**	**Aggregate**
	w. CC	**w/o CC**	**generations**	**(t = 2)**	
		(unit = percentage of total annual income)			
Optimal cap | access (Table 5, row 1)	0.54	0.09	0.09	-0.10	0.62
Submarket-and-time-specific caps	0.78	0.05	0.06	-0.16	0.73

Notes: The table reports the onetime transfers as a percentage of initial steady state total annual income by agent group and for all agents. The regulation is applied to both pre-existing accounts and new issuances, and results are reported for two regulations: (1) optimal access-dependent cap chosen to maximize welfare for all agents (see Table 5, row 1), and (2) submarket-and-time-specific caps. The groups of agents are (1) consumers with credit cards (CC) in the first period of the transition, (2) consumers without a credit card in the first period of the transition, (3) future generations whose welfare is discounted at the rate $1/(1+r)$, and (4) incumbent lenders in the first period of the transition.

Overall, our findings caution against the temptation to revise the policy. In Table 8, we report the efficiency gains from the submarket-specific caps set by a regulator who re-optimizes every period ("Submarket-and-time-specific caps"). While there are additional efficiency gains from regulating future accounts (i.e., $0.05 + 0.06$ percent of total annual income), the additional gains from submarket-and-time-specific caps are smaller than the gains achieved by a time-invariant access-dependent rate cap (i.e., $0.09 + 0.09$ percent of total annual income). This is a telltale sign of the time inconsistency of the optimal rate cap policy. In fact, for consumers without a credit card and for future generations, the cost of not committing to future rate caps outweighs the benefits from fine-tuning the future rate caps to characteristics other than credit access. Commitment to access-dependent rate caps limits the balance transfers of consumers who already have a credit card. A regulator who re-optimizes ex-post, instead, treats the pre-existing revolving credit relationship as a sunken investment by the incumbent lender. In turn, the attempt by the re-optimizing

[29] Solving the problem of a regulator who can commit to a submarket-specific rate cap plan would require enriching the state space with promised utility as a state variable. We leave this exercise as an interesting avenue for future research.

regulator to further improve upon the allocation ex-post makes it harder to attract credit offers ex-ante. This time inconsistency is also revealed in the sensitivity analyses provided in Table 12 in Appendix C.5. In the cases with on-the-credit-search (panels A to D and F), the optimal access-dependent cap leads to higher gains for consumers without credit cards and for future generations than those from submarket-and-time-specific caps, and the opposite is true in the cases without on-the-credit-search (panels E and G).

6 Conclusion

A distinct feature of the U.S. credit card market that has gained attention among both academics and policy makers for several decades is high markups. While the Credit CARD Act of 2009 placed restrictions on prices by limiting interest rate hikes and fees charged, it did not go as far as imposing a cap on interest rates. Since the enactment of the CARD Act, the consensus in the structural literature on credit card rate regulation is that a rate cap can be beneficial. However, this literature has abstracted from the *revolving* nature of credit lines and the ensuing on-the-credit-search for better credit lines. We show how these features matter for the design of rate regulation.

First, modeling long-term contracts introduces history dependence to credit access. As a result, our model reveals that the transition path is key for the magnitude of gains. Second, the consumers' lack of commitment to not search for better credit lines leads to a contract dilution problem. Potential lenders anticipate that the competition will target profitable consumers for balance transfers, which dilutes the profitability of the long-term contract. Therefore, the competition for balance transfers lowers the likelihood of credit offers for consumers without a credit card.

In a model calibrated to the U.S. credit card market, we find that a rate cap tailored to credit access is twice as effective in generating efficiency gains compared to the commonly studied uniform cap. Furthermore, the additional gains from this simple alternative are specific to the regulation of revolving credit lines because of the history dependence in credit access and the contract dilution problem. We also find that the gains from tailoring the cap to consumer characteristics beyond credit access are small.

In our study of rate regulation, we abstract from analyzing optimal time-dependent policies with commitment. We believe this is a promising and interesting direction for future research.

References

Albrecht, J. W., and B. Jovanovic (1986), "The Efficiency of Search under Competition and Monopsony," *Journal of Political Economy,* 94(6), 1246-1257.

Agarwal, S., S. Chomsisengphet, N. Mahoney, and J. Stroebel (2015), "Regulating consumer financial products: evidence from credit cards," *The Quarterly Journal of Economics,* 130(1), 111-164.

Athreya, K. B. (2002), "Welfare implications of the bankruptcy reform act of 1999," *Journal of Monetary Economics,* 49(8), 1567-1595.

Athreya, K., X. S. Tam, and E. R. Young, "A quantitative theory of information and unsecured credit," *American Economic Journal: Macroeconomics* 4(3) (2012), 153-183.

Bethune, Z., J. Saldain, and E. R. Young (202), "Consumer credit regulation and lender market power," Working paper.

Bizer, D. S. and P. M. DeMarzo (1992), "Sequential banking," *Journal of Political Economy,* 100(1), 41-61.

Braxton, J. C., K. Herkenhoff, and G. Phillips (2019), "Can the unemployed borrow? Implications for public insurance," *Journal of Political Economy,* 132(9), 3025-3076.

Consumer Financial Protection Bureau (2013), "Navigating the market: A comparison of spending on financial education and financial marketing."

Chakravarty, S. and E. Rhee (1999), "Factors affecting an individual's bankruptcy filing decision," *Available at SSRN.*

Chatterjee, S., D. Corbae, M. Nakajima, and J. Ríos-Rull (2007), "A quantitative theory of unsecured consumer credit with a risk of default," *Econometrica,* 75(6), 1525-1589.

Chatterjee, S. and B. Eyigungor (2023), "Heterogeneity in the credit card market," Working paper.

Chatterjee, S. and B. Eyigungor (2024), "Credit and welfare implications of credit card interchange fees," Working paper.

Cuesta, J. I. and A. Sepúlveda (2021), "Price Regulation in Credit Markets: A Trade-off between Consumer Protection and Credit Access", Working paper.

Dempsey, K. and F. Ionescu (2024), "Borrowing premia in unsecured credit markets," Working paper.

den Haan, W. J., G. Ramey, and J. Watson (2000), "Job descruction and propagation of shocks," *American Economic Review,* 90(3), 482-498.

Drozd, L. A. and M. Kowalik (2019), "Credit cards and the Great Recession: The collapse of teasers," Working paper.

Drozd, L. A. and J. B. Nosal (2008), "Competing for customers: A search model of the market for unsecured credit," Working paper.

Exler, F., I. Livshits, J. MacGee, and M. Tertilt (2024), "Consumer credit with over-optimistic borrowers," *Journal of the European Economic Association,* forthcoming.

Exler, F. and M. Tertilt (2020), "Consumer debt and default: A Macroeconomic perspective," *Oxford Research Encyclopedia of Economics and Finance.*

Galenianos, M. and A. Gavazza, "Regulatory interventions in consumer financial markets: The case of credit cards," *Journal of the European Economic Association,* 20(5), 1897-1932.

Galenianos, M., T. H. Law, and J. B. Nosal (2023), "Market power in credit markets," Working paper.

Gourinchas, P. and J. A. Parker (2002), "Consumption over the Life Cycle," *Econometrica,* 70(1), 47-89.

Hatchondo, J., and L. Martinez (2017), "Credit risk without commitment," Working paper.

Herkenhoff, K. (2019), "The impact of consumer credit access on unemployment," *Review of Economic Studies,* 86(6), 2605-2642.

Herkenhoff, K. and G. Raveendranathan (2024), "Who bears the welfare cost of monopoly? The case of the credit card industry," *Review of Economic Studies,* Accepted.

Hosios, A., J. (1990), "On the Efficiency of Matching and Related Models of Search and Unemployment," *Review of Economic Studies,* 57(2), 279-298.

Jagtiani, J. A. and W. Li (2015), "Credit access and credit performance after consumer bankruptcy filing: New evidence," *American Bankruptcy Law Journal,* 89(2), 327-362.

Krueger, D. and A. Ludwig (2016), "On the optimal provision of social insurance: Progressive taxation versus education subsidies in general equilibrium," *Journal of Monetary Economics,* 77, 72-98.

Livshits, I., J. MacGee, and M. Tertilt (2007), "Consumer bankruptcy: A fresh Start," *The American Economic Review,* 97(1), 402-418.

Livshits, I., J. MacGee, and M. Tertilt (2010), "Accounting for the rise in consumer bankruptcies," *American Economic Journal: Macroeconomics,* 2(2), 165-193.

Livshits, I., J. MacGee, and M. Tertilt (2016), "The democratization of credit and the rise in consumer bankruptcies," *Review of Economics Studies,* 83(4), 1673-1710.

Madeira, C. (2019), "The impact of interest rate ceilings on households' credit access: Evidence from a 2013 Chilean legislation," *Journal of Banking and Finance,* 106, 166-179.

Mateos-Planas, X. and J. Ríos-Rull (2013), "Credit lines," Working paper.

McKinley, V. (1997), "Ballooning bankruptcies: Issuing blame for the explosive growth," *Regulation,* 20(4), 33–40.

Menzio, G., I. A. Telyukova, and L. Visschers (2016), "Directed search over the life cycle," *Review of Economic Dynamics,* 19, 38-62.

Moschini, E. G., G. Raveendranathan, and M. Xu (2023), "Over-optimism about graduation and college financial aid," Working paper.

Nakajima, M. (2017), "Assessing bankruptcy reform in a model with temptation and equilibrium default," *Journal of Public Economics,* 145, 42-64.

Nelson, S. (2023), "Private information and price regulation in the US credit card market," *Econometrica,* Condtionally Accepted.

Raveendranathan, G. (2020), "Revolving credit lines and targeted search," *Journal of Economic Dynamics and Control,* 118.

Raveendranathan, G., and G. Stefanidis (2024), "The unprecedented fall in U.S. revolving credit," *International Economic Review.*

Rogerson, R., R. Shimer, and R. Wright (2005), "Search-Theoretic Models of the Labor Market: A Survey," *Journal of Economic Literature,* 43(4), 959-988.

Saldain, J. (2023), "High-cost consumer credit: desperation, temptation and default," Working paper.

Storesletten, K., C. Telmer, and A. Yaron (2004), "Consumption and risk sharing over the life cycle," *Journal of Monetary Economics,* 51(3), 609-633.

Tauchen, G. (1986), "Finite state markov-chain approximations to univariate and vector autoregressions," *Economics Letters,* 20(2), 177-181.

Zinman, J. (2009), "Where is the missing credit card debt? Clues and implications," *The Review of Income and Wealth,* 55(2), 249-265.

A Model appendix

A.1 Proofs

Proof of Lemma 1. Consider a submarket $s = (J-1, b, \bar{b}, \tau, \theta, \eta)$. Because a consumer cannot borrow in the last period of their life-cycle, the value of not defaulting is $V(J, b', \bar{b}', \tau', \theta, \eta', \gamma') = U(y_J - b')$, which is independent of the distribution of submarkets. Similarly, the value of defaulting $V(J, b', \bar{b}', \tau', \theta, \eta', \gamma') = U(y_J) - \xi$ is also independent of the distribution of submarkets. Hence the continuation value $W(J, b', \bar{b}', \tau', \theta, \eta', \gamma')$ in (5) exists, and it is independent of the distribution of submarkets. The probability of receiving a credit offer in the last period is zero, independently of the distribution of submarkets. As a result, the value functions $V(s, \gamma)$ in (7), and $V^D(s, \gamma)$ in (9), and the policy functions $c(s, \gamma)$ and $b'(s, \gamma)$ exist and are independent of the the distribution of submarkets for any submarket s with consumers of age $J-1$. As a result, the policy function $d(s, \gamma)$ and, in turn, the value function $W(s, \gamma)$ exist and are also independent of the distribution of submarkets for any submarket s with consumers of age $J-1$.

The profit function $\pi(s, \gamma)$, defined in (10) exists and is also independent of the distribution of submarkets for any submarket s with consumers of age $J-1$ because, as we have shown, the probability of the consumer matching next period is zero and the policy functions of the consumer for that submarket are independent of the distribution of submarkets. This implies that the credit limit $\bar{b}_I(s)$ which solves (2) is also well-defined and independent of the distribution of submarkets for any submarket s with consumers of age $J-1$. Because the policy function τ is independent of the distribution of submarkets, the expected profit function $\Pi(s)$ in (4) exists and is independent of the distribution of submarkets for any submarket s with consumers of age $J-1$. In turn, because search is targeted, incentives for entry are submarket-specific, and market tightness $\mu(s)$ in (1) exists, and it is also independent of the distribution of submarkets for any submarket s with consumers of age $J-1$.

Consider a submarket $s = (J-2, b, \bar{b}, \tau, \theta, \eta)$. The argument that value and policy functions exist is the same as that for a submarket with $s = (J-1, b, \bar{b}, \tau, \theta, \eta)$. Because, as we have shown, the continuation value function $W(s', \gamma)$, the policy functions $\tau(s')$ and $\bar{b}(s')$, and market tightness $\mu(s')$ are independent of the distribution of submarkets for any submarket s' with consumers of age $J-1$, the probability of matching is also independent of the distribution of submarkets for any submarket s' with consumers of age $J-1$. As a result, the value functions $V^D(s, \gamma)$, $V(s, \gamma)$, and the policy functions $c(s, \gamma)$ and $b'(s, \gamma)$, and, in turn, the policy function $d(s, \gamma)$ and the value function $W(s, \gamma)$, are independent of the the distribution of submarkets for any submarket s with consumers of age $J-2$.

Because, as shown above, the probability of matching and the policy functions of the consumer in any submarket s' with consumers of age $J-1$ or $J-2$ are independent of the distribution of submarkets, the profit function $\pi(s,\gamma)$, defined in (10), is also independent of the distribution of submarkets for any submarket s with consumers of age $J-2$. This implies that the credit limit $\bar{b}(s)$ which solves (2) is also independent of the distribution of submarkets for any submarket s with consumers of age $J-2$. Because the policy function τ is independent of the distribution of submarkets, the expected profit function $\Pi\big(s\big)$ is also independent of the distribution of submarkets for any submarket s with consumers of age $J-2$. In turn, because search is targeted, incentives for entry are submarket-specific in (1), and market tightness $\mu(s)$ is also independent of the distribution of submarkets for any submarket s with consumers of age $J-2$.

The rest of the proof follows by induction by replacing $J-2$ with j and $J-1$ with $j+1$ for $j \leq J-2$. □

Proof of Proposition 2. First, the guess is that there is a policy function τ that is independent of the distribution of submarkets and τ solves (11) given Π and W from a block-recursive quasi-equilibrium given τ. To verify that the result holds, note that Lemma 1 implies that given that the guess τ is independent of the distribution of submarkets, there exists a block-recursive quasi-equilibrium given τ and, hence, Π and W are independent of the distribution of submarkets. This, in turn, verifies that there is a solution τ to (11) that is independent of the distribution of submarkets. □

Proof of Proposition 3. The proof consists in guessing and verifying the result. First, the guess is that there is a policy function τ that is independent of the distribution of submarkets, and τ solves (11) subject to the cap (given Π and W from a block-recursive quasi-equilibrium given τ). To verify that the result holds, note that Lemma 1 implies that given that the guess τ is independent of the distribution of submarkets, there exists a block-recursive quasi-equilibrium given τ and, hence, Π and W are independent of the distribution of submarkets. This, in turn, verifies that there is a solution τ to (11) subject to the cap that is independent of the distribution of submarkets. □

Proof of Proposition 4. The proof consists in guessing that the solution to the regulator's problem is independent of the distribution of submarkets, and using Lemma 1 to verify that it is the case. Let $\tau^*(s)$ denote the solution to (14) for submarket s at the time of the reform and the solution to (15) for submarket s after the reform. If $\tau^*(\cdot)$ is independent of the distribution of submarkets, Lemma 1 implies that there is a block-recursive quasi-equilibrium given $\tau^*(\cdot)$. To verify that (14) and (15) admit solutions for each submarket that are independent of the distribution of submarkets, note that the objectives in (14) and (15) are both independent of the distribution of submarkets because the quasi-equilibrium is independent of the distribution of submarkets. □

A.2 Stylized model on the insurance motive

In this section, we elicit the regulator's motive to reduce the borrowing premium τ below marginal cost to facilitate precautionary savings on the part of the consumer in the face of uninsurable idiosyncratic risk. Consider a simplified two-period version of our quantitative model that isolates the role of uninsurable idiosyncratic risk. This simplified model abstracts from the consumers' life-cycle profile of earnings, survival risk, consumer default, endogenous borrowing limits, and search and matching. The earnings process is simplified so that permanent earnings in both periods is θ, and transitory earnings in period 0 is γ. The transitory earnings is the only source of uninsurable risk. The consumer is risk averse, and the lender does not discount future payoffs, i.e., $r = 0$. There is no default.

We present the consumer's problem including the onetime compensation transfer T. The consumer's problem in period 0 after the realization of the transitory earnings shock is

$$\begin{aligned} V(\gamma, \tau; T) &= \max_{c_0, c_1, b'} \{U(c_0) + \beta U(c_1)\} \\ \text{subject to: } c_0 &= \theta + \gamma - T + q(\tau)b', \\ c_1 &= \theta - b'. \end{aligned}$$

The lender's profit function is $\pi(\gamma, \tau; T) = -q(\tau)b'(\gamma, \tau; T) + b'(\gamma, \tau; T)$, where $b'(\gamma, \tau; T)$ is the consumer's borrowing policy function. The total compensated variation TCV translates the efficiency gains from the change in premium from τ to $\hat{\tau}$ for the consumer into present-value dollars so that

$$E_\gamma \left[V(\gamma, \hat{\tau}; TCV(\tau, \hat{\tau})) \right] = E_\gamma \left[V(\gamma, \tau; 0) \right].$$

Finally, the regulator solves

$$\max_{\hat{\tau}} E_\gamma \left[TCV(\tau, \hat{\tau}) + \pi(\gamma, \hat{\tau}; TCV(\tau, \hat{\tau})) - \pi(\gamma, \tau; 0) \right]. \tag{19}$$

The main insight from our analysis is that the solution to the regulator's problem (19) sets the premium below zero. The argument shows that the slope of the objective function in (19) is strictly negative for any premium $\hat{\tau} \geq 0$. Formally, the slope is

$$-\left(-\frac{\partial q(\hat{\tau})}{\partial \hat{\tau}}\right) \left(E_\gamma \left[\frac{U'(c_0)}{-q^2 U''(c_0) - \beta U''(c_1)} \right] (1 - q(\hat{\tau})) + \frac{E_\gamma \left[\frac{qU''(c_0) + \beta U''(c_1)}{q^2 U''(c_0) + \beta U''(c_1)} \right]}{E_\gamma [U'(c_0)]} \text{Cov}_\gamma \left[b', U'(c_0) \right] \right),$$

where c_0, c_1, and b' denote the policy functions of the consumer. The utility function being concave and increasing, and the price being decreasing in the premium means that all terms in parentheses other than $(1-q(\hat{\tau}))$ and $\mathrm{Cov}_{\gamma}\left[b', U'(c_0)\right]$ are positive. The latter two terms govern the two motives of the regulator. The first motive, which is present even in the absence of uncertainty, is to set the price equal to the marginal cost of borrowing, which is normalized to 1 in this Appendix. A positive premium results in a price that is less than 1 making this term positive and as a result contributing to a negative slope of the objective in (19).

The second motive, governed by the covariance between borrowing and the marginal utility of period 0 consumption, reflects the insurance motive. Lowering the premium effectively reallocates consumption from higher realizations of γ to lower ones, thereby providing insurance to the consumer. The covariance term is positive because consumers with a lower realization of γ borrow more and value consumption more in period 0, i.e., the two variables co-move. When the premium is lowered, a consumer who borrows more stands to benefit disproportionately. The covariance term captures that the consumers who benefit disproportionately are the ones who value consumption the most. The two terms in the formula for the slope imply that the slope of the regulator's objective is strictly negative for any premium $\hat{\tau} \geq 0$. As a result, the solution to the regulator's problem in (19) is a negative premium.

B Calibration and model validation appendix

B.1 Average markup and the lender's bargaining power

This section numerically shows the sensitivity of the average markup approximated as the average credit card interest rate per card holder net of the risk-free rate, charge-off rate, and the estimate for the transaction cost (Table 1) to the lender's bargaining power. Figure 8 plots the average markup as a function of the lender's bargaining power by varying only the lender's bargaining power in the baseline calibration (that is, other parameters are held fixed at their calibrated values from the baseline calibration). As expected, the markup increases with the lender's bargaining power.

Figure 8: Average "markup" as a function of the lender's bargaining power

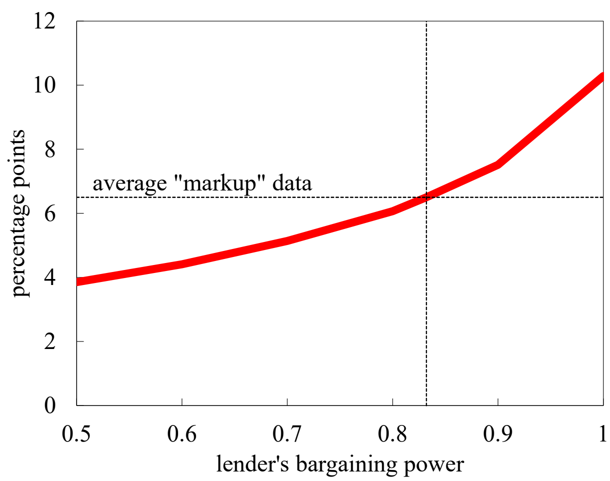

Notes: The figure depicts the average "markup" approximated as the average credit card interest rate per card holder net of the risk-free rate, charge-off rate, and the estimate for the transaction cost (Table 1) as a function of the lender's bargaining power.

B.2 Rate dispersion and on-the-credit-search

This section shows that on-the-credit-search is an important feature of our model to generate a large share of the dispersion in credit card interest rates in the data. Table 9 reports the standard deviation and various percentiles of the credit card interest rate distribution in our baseline model that features on-the-credit-search and in alternative specification without on-the-credit-search (both re-calibrated and not re-calibrated). The first row of the table shows that on-the-credit-search contributes significantly to the dispersion in rates. The last four rows of the table illustrate how the dispersion decreases. In comparison to the baseline model, rates increase at the 10th and

20th percentiles and decrease at the 80th and 90th percentiles in the model without on-the-credit-search. The lack of on-the-credit-search can rationalize both patterns. The higher rates at the 10th and 20th percentiles in the model without on-the-credit-search can be explained by the inability of consumers with credit cards to switch to credit cards with lower rates. The lower rates at the 80th and 90th percentiles in the model without on-the-credit-search can be explained by results from the simple model in Section 2. With on-the-credit-search, the lender has an incentive to charge a higher rate to front load profits.

Table 9: Rate dispersion with and without on-the-credit-search (OTCS)

Variable	**Baseline model**	**Model without OTCS**	
		Re-calibrated	**Not re-calibrated**
	(unit = percentage points)		
Standard deviation	4.97	2.98	2.60
Rate: 10th percentile	10.72	14.80	13.78
Rate: 20th percentile	12.76	15.31	15.31
Rate: 80th percentile	20.92	18.88	18.37
Rate: 90th percentile	22.97	20.41	18.88

Notes: The table compares credit card interest rate statistics between the baseline model that features on-the-credit-search and an alternative model without on-the-credit-search (both re-calibrated and not re-calibrated).

B.3 Model validation by income and by age

In this section, we show that our baseline model can account for patterns observed in cross-sectional credit statistics by income and by age, using data from the SCF. Figure 9 plots the average credit per household, average total credit card limit per household with a credit card (that is, the average of the sum of credit limits across all credit cards in the data), average spread (average credit card interest rate per household with a credit card net of the risk-free rate), and the share of the population with credit cards by income quintile in the data and model. The model accounts for the following patterns qualitatively: average credit and average limits increase with income, average interest rates decrease with income, and the share of the population with credit cards increases with income.[1]

[1] While Table 4 shows that the model understates credit usage in the aggregate, Figure 9 shows that the model can account for the levels of credit usage by income. This is because outstanding revolving credit is underreported in the SCF. See Zinman (2009) for a detailed discussion on the underreporting issue. As a result, we focus on qualitative patterns rather than on the levels.

Figure 9: Credit statistics by income

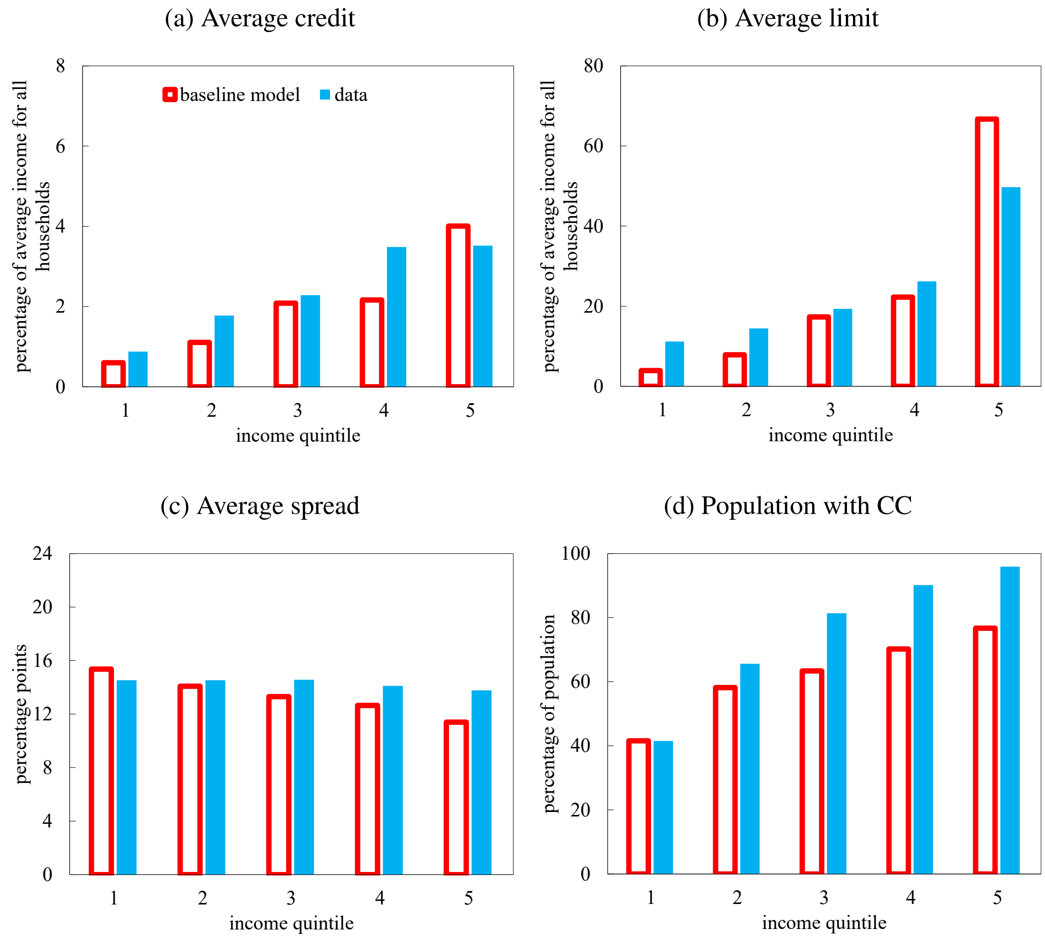

Notes: The figure depicts the following credit statistics by income quintile in the baseline model and data (2019): average credit per household normalized by average income per household, average total credit card limit per card holder normalized by average income per household, average credit card interest rate per card holder net of the risk-free rate (average spread), and the share of the population with credit cards (CC). Data source: Survey of Consumer Finances (SCF).

Figure 10 plots analogous statistics to Figure 9 by age instead of by income. Again, the model does reasonably well given that these cross-sectional patterns were not targeted in the calibration. In particular, the model qualitatively accounts for the hump-shaped profile of average credit and the increase in the average credit limits in the early stages of the life-cycle.[2] It is worth noting that this model calibration performs better in regard to the profile of credit over the life-cycle in comparison to Raveendranathan and Stefanidis (2024). In Raveendranathan and Stefanidis (2024), borrowing by young households in the model is drastically higher in comparison to the data. There are two reasons for the improvement: first, we start the model at age 25 rather than at age 20; and second,

[2]In the model, credit limits plummet post age 65 because lenders can decrease credit card limits to current debt levels when the consumer retires. Otherwise, the model generates too many defaults towards the end of the life-cycle.

we incorporate a household size shifter into the utility function for the part of the life-cycle in which households tend to have children. The former dampens the extent to which the agent wants to and can borrow against an increasing life-cycle earnings profile.[3] The latter incentivizes the household to save (or not borrow as much).[4]

Figure 10: Credit statistics by age

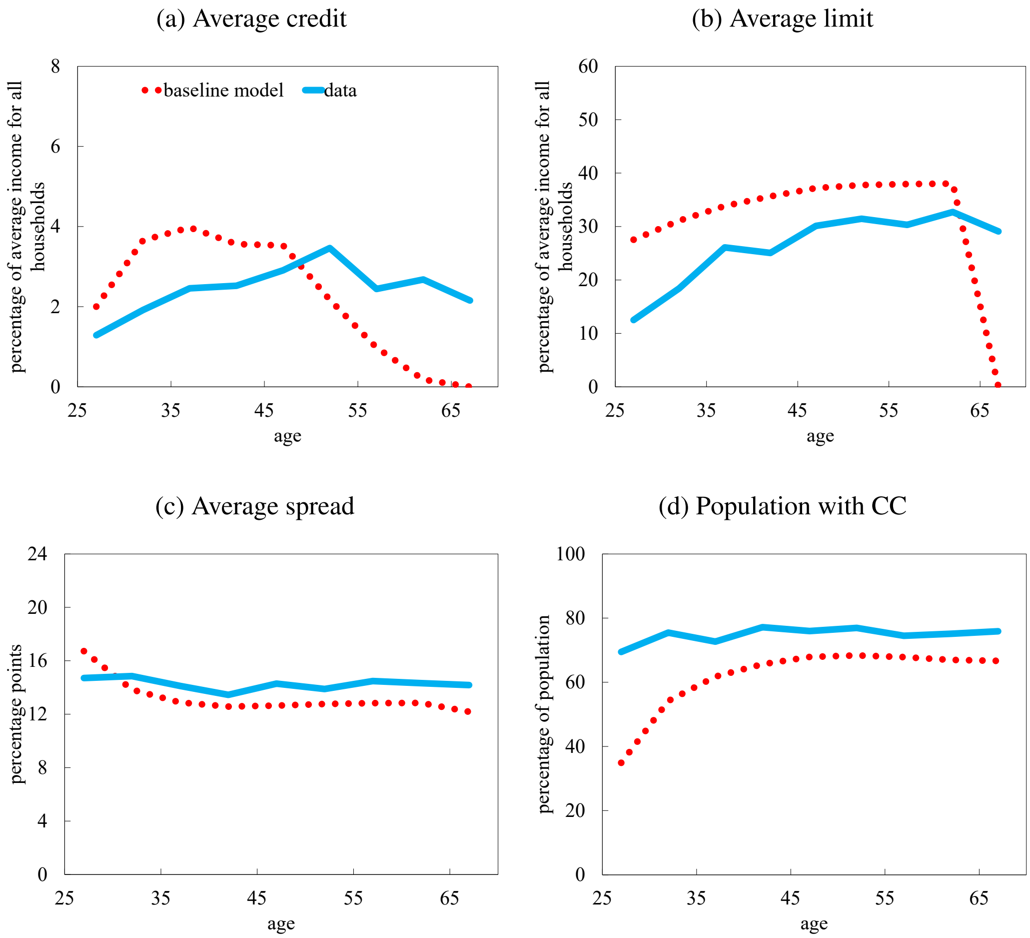

Notes: The figure depicts the following credit statistics by age in the baseline model and data (2019): average credit per household normalized by average income per household, average total credit card limit per card holder normalized by average income per household, average credit card interest rate per card holder net of the risk-free rate (average spread), and the share of the population with credit cards (CC). Data source: Survey of Consumer Finances (SCF).

[3]In the data, credit among households of age between 20 and 24 is a very small share of total credit (less than 1 percent). Therefore, starting the model at age 25 is not a major concern.

[4]Chatterjee, Corbae, Nakajima, and Rios-Rull (2007) and Livshits, MacGee, and Tertilt (2007) are the canonical quantitative models of unsecured consumer credit. The former features infinitely lived agents, whereas the latter features overlapping generations. The latter also includes a shifter for household size over the life-cycle.

B.4 Borrowing premium and default risk

Although there is a positive relationship between credit card borrowing premiums and the likelihood of default in the data, Dempsey and Ionescu (2024) document that this relationship is relatively flat in comparison to the linear relationship predicted by the standard consumer credit models. Figure 11 shows that our model produces a relationship between borrowing premiums and default likelihood that is consistent with the pattern documented in Dempsey and Ionescu (2024).

Figure 11: Borrowing premium and default risk

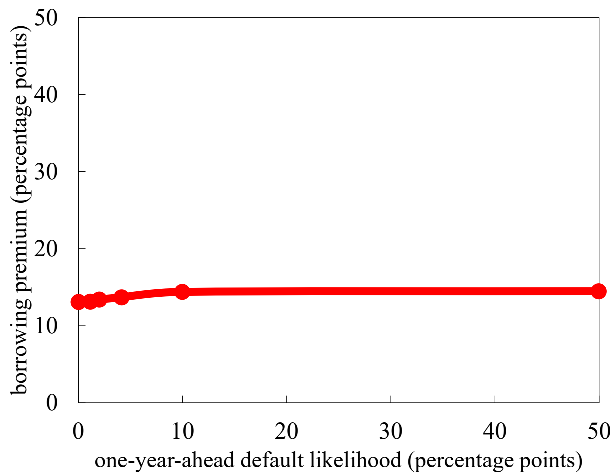

Notes: The figure depicts the average borrowing premium as a function of the average likelihood of default for the current period after sorting consumers with credit cards in ascending order based on default risk. The masses in each bin from left to right are 97.77, 0.64, 0.51, 0.53, 0.23, and 0.32 percent, respectively.

B.5 Acquisition cost

Table 10 reports the acquisition cost to income and acquisition cost to credit in the data (Panel A) and in various model specifications. Recall that in the model the acquisition cost is computed as the mass of credit offers in a period multiplied by the cost of one credit offer. Panel B reports results from model specifications which target the same set of moments as the baseline calibration in Table 3. The results show that these specifications produce a reasonable estimate for the acquisition cost to income but overstate the acquisition cost to credit. Panel C reports results from model specifications which target total revolving credit instead of the average credit card limit that was targeted in the baseline calibration (all other moments are the same as the baseline). In all these specifications, both acquisition cost to income and acquisition cost to credit are overstated. While our results indicate that it is challenging to match acquisition cost relative to credit given our set of target moments, it is possible to generate a reasonable magnitude for the acquisition cost to income.[5]

Table 10: Acquisition cost relative to income and relative to credit

Data and model specification	**Acquisition cost to income**	**Acquisition cost to credit**
	(unit = percentage)	
Panel A: Data		
(1) Data	0.03–0.05	0.50–0.83
Panel B: Target baseline moments		
(1) Baseline	0.06	2.96
(2) Lower ζ	0.05	2.74
(3) Higher ζ	0.05	2.55
(4) Urn-ball matching function	0.05	2.24
(5) Leontief matching function	0.06	2.96
(6) Wage garnishment	0.05	4.79
Panel C: Target total credit		
(1) Baseline variation	0.17	3.11
(2) Lower ζ	0.16	2.82
(3) Higher ζ	0.16	2.90
(4) Urn-ball matching function	0.14	2.45
(5) Leontief matching function	0.16	2.69
(6) Wage garnishment	0.29	5.27

Notes: The table reports the acquisition cost to income and acquisition cost to credit in the data (Panel A) as well as in various model specifications (Panels B and C). See footnote *a* of Table 4 in the main text for more details regarding the data. In Panel B, the target moments are the ones used for the baseline calibration in Table 3. In Panel C, the target moment is total revolving credit to income instead of the average total credit card limit per card holder (all other target moments are the same as in the baseline). The following model specifications are presented in each row of Panels B and C: (1) baseline or same target moments as in the baseline calibration except target total revolving credit to income instead of average credit card limit per card holder; (2) and (3) set the parameter that determines the matching elasticity ζ to 50 percent lower and higher values to those used in the baseline calibration, respectively, and re-calibrate the model as in (1); (4) and (5) use an urn-ball matching function and Leontief matching function instead of the baseline calibration matching function, respectively, and re-calibrate the model as in (1); (6) use wage garnishment as a default cost, where the garnishment amount is collected by the lender, instead of stigma as in the baseline, and re-calibrate the model as in (1). Data sources: Agarwal et al. (2015), Consumer Financial Protection Bureau (2013), and Bureau of Economic Analysis (BEA).

[5] It is possible to generate a lower acquisition cost to credit by reducing the lender's bargaining power, but with such a change the model will not match the markups observed in the data (a key moment for our analysis).

C Results appendix

C.1 How a uniform cap affects surpluses across steady states

In this section, we analyze how a uniform cap affects the total surplus, the number of credit contracts, the surplus per credit contract, and the consumer's share of the surplus from credit contracts. Because we perform this analysis across steady states, and to analyze welfare across steady states we measure welfare for 25-year-olds behind the veil of ignorance, we focus on 25-year-olds entering the economy.

Figure 12: Trade-offs of a uniform cap on interest rates across steady states for 25-year-olds

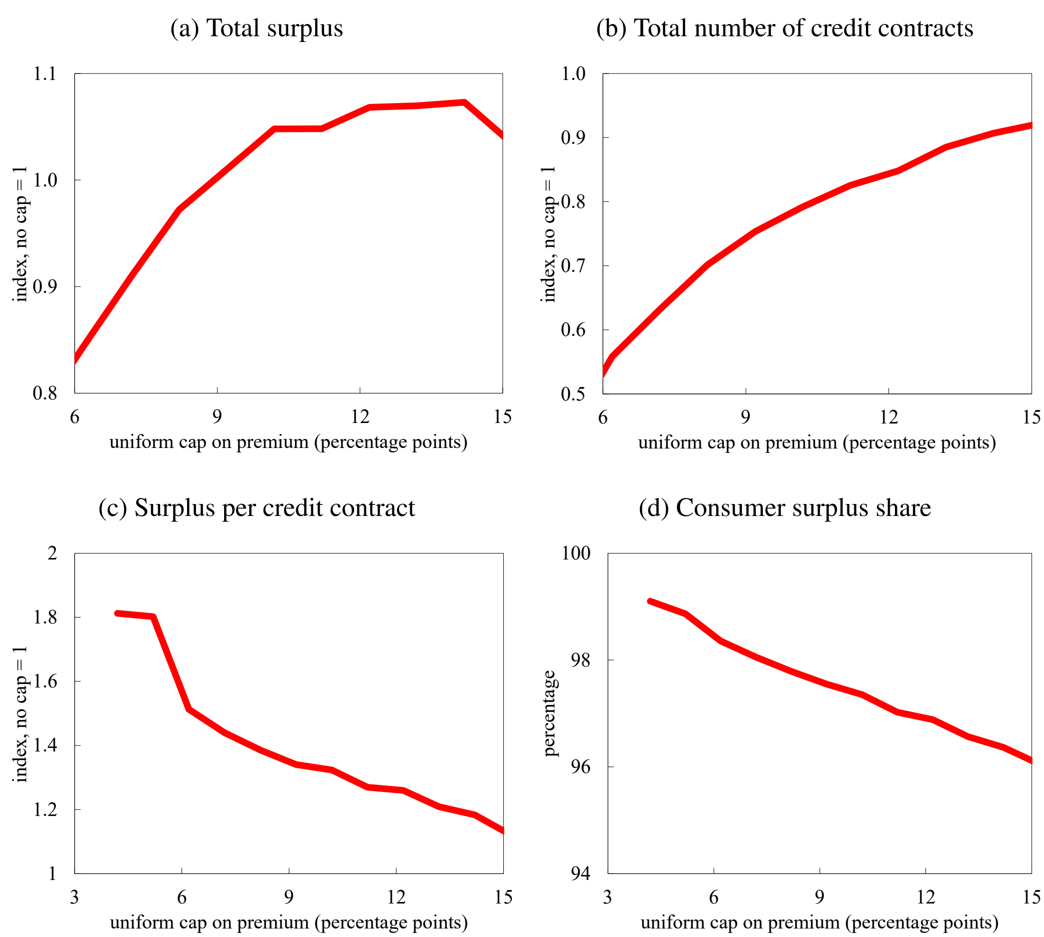

Notes: The figure depicts comparative statics with respect to a uniform rate cap computed in the stationary distribution in the equilibrium with a uniform rate cap. The statistics are computed for credit contracts between 25-year-olds entering the economy and lenders. The x-axis refers to the cap on the borrowing premium $\bar{r}_q$. For example, a cap equal to 10 implies that the maximum premium over the risk-free rate is 10 percentage points. Figure 12a depicts the total surplus created due to credit contracts between 25-year-olds entering the economy and lenders. Similarly, Figures 12b, 12c, and 12d depict the number of credit contracts, the surplus per credit contract, and the share of consumer surplus from credit contracts, respectively.

Figure 12a shows that a uniform cap can increase the total surplus as long as the rate cap is not too tight. The intuition can be gained from Figures 12b and 12c, which decompose the total surplus into total number of credit contracts and surplus per credit contract. The uniform cap reduces the number of credit contracts, which is consistent with the discussion regarding the population with credit cards in Figure 1c. However, the surplus per match increases with a rate cap. The latter effect dominates for a less stringent cap, leading to a higher total surplus. Figure 12d shows that the 25-year-old consumer's share of the surplus increases as the cap becomes tighter. Figures 12b and 12c highlight the trade-off for the regulator. The surplus per match increases with a rate cap, but it comes at the cost of fewer credit contracts. This result also explains why capping premiums on pre-existing accounts leads to larger gains than capping premiums on new issuances, as discussed in Section 5.2.1.

C.2 Aggregate welfare implications of a uniform cap with transfers

Our welfare measure adds the compensated variation and the change in profits to measure efficiency gains associated with a reform. In this section, we show that our welfare computations are robust to computing changes in profits based on the policy functions of consumers after consumers are compensated for the policy change. Figure 13 depicts the net gains when the effect of implementing lump sum transfers on incumbent lender profits is taken into account (blue-dashed line) as well as the gains depicted in Figure 2a (red-solid line) in which such an effect is not taken into account. The difference between the two lines is negligible. This figure implies that there exist a set of lump sum transfers that make all agents weakly better off for rate caps that generate positive welfare gains, as discussed in footnote 21 in the main text.

Figure 13: Aggregate welfare implications of a uniform cap on the borrowing premium

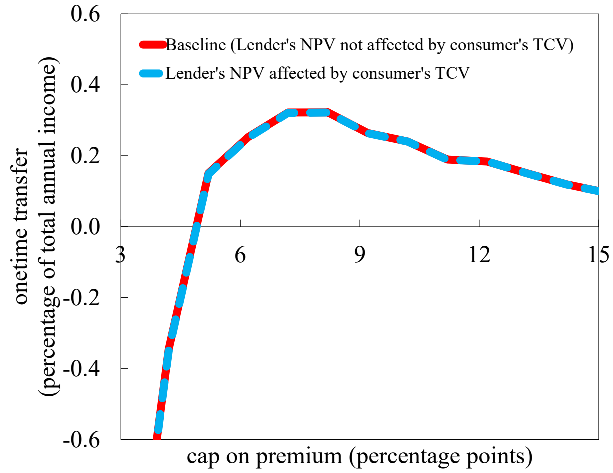

Notes: The figure depicts total onetime transfers as a percentage of total annual income for all agents under two cases: (1) onetime compensation transfers are computed for each agent type and aggregated across all agents (“Baseline ...") and (2) onetime compensation transfers are computed for each agent type taking into account that the compensation transfers affect consumer decision rules, and therefore, may affect incumbent lender profits, and then aggregated across all agents (“Lender’s NPV ..."). The x-axis refers to the uniform rate cap (maximum borrowing premium). For example, a cap of 10 implies that the maximum borrowing premium is 10 percentage points. The cap is applied to both pre-existing accounts and new issuances.

C.3 Positive steady state implications of the optimal access-dependent cap

In this section, we discuss the positive steady state implications of the optimal access-dependent cap. The second and third columns of Table 4 present various credit card market statistics by five panels on rates, issuance cost, credit usage, default, and access.

Panel A shows that the standard deviation of rates decreases, and that there is bunching at 8.64 percent (the lower end of the rate distribution) and at 15.64 percent (the higher end of the rate distribution). These are the interest rates implied by the optimal access-dependent caps on the borrowing premiums.

Panel B shows that the issuance cost to income decreases by half despite the increase in the population with credit cards (see Panel E). The optimal access-dependent cap is tight for consumers with a credit card. This reduces the incentive for new lenders to offer balance transfers. On the other hand, offers to young consumers without credit cards increase (Figure 4a) because the optimal access-dependent cap alleviates the contract dilution problem. The former effect dominates the change in the issuance cost.

Panel C shows that across steady states, revolving credit to income and the population with positive credit card balances increase, albeit slightly. As the third row of Panel C shows, the share of card holders who use more than 50 percent of their credit card limit decreases slightly. That is, the utilization rate above 50 percent of the credit card limit decreases among credit card holders. This is because the utilization rate is reported as a share of the population with credit cards, and the share of population with credit cards increases more than the share of the population that uses more than 50 percent of its credit card limit. The increase in credit usage is a result of lower rates and an increase in the population with credit cards, but the rise is mitigated by the fall in the intensive margin of access (i.e., credit limits), as shown in Panel E.[6] The lender has an incentive to lower the limit when rates are low because the lender cannot charge a higher premium for higher default risk.

Panel D shows that both the default rate and the charge-off rate decrease despite the rise in credit usage. There are two explanations for this pattern. First, the lower rate reduces the incentive to default, controlling for other characteristics such as the level of debt and income. Second, the average credit card limit is lower, which limits borrowing and, therefore, default risk.

C.4 Contract dilution due to offers of balance transfers

In this section, we quantify the contract dilution problem due to on-the-credit search for a balance transfer (in the baseline model). If the consumer can commit to not search on-the-credit, the model retains block-recursivity because submarkets will be additionally segmented by contracts with and without commitment. As a result, the welfare changes of committing to not search on-the-credit are equivalent to comparing the welfare changes from eliminating on-the-credit-search.

Table 11 reports the welfare changes for consumers who commit to not accept offers of a balance transfer. The table shows that future generations and consumers without a credit card benefit by 0.09 percent ($\approx$ 0.04 + 0.05) from this commitment.[7] It is because committing to not accept offers of a balance transfer increases the likelihood of a credit offer for a consumer searching for access to credit. For example, such a commitment for a 25-year-old entering the economy with low, medium, and high permanent earnings would increase the probability of credit access by 0.6, 8.9, and 15 percentage points, respectively. The incentive to commit is stronger for higher permanent earners because the targeted search for balance transfers is based on the profitability of the account. As a result, the severity of contract dilution increases with the profitability of the account (and the profitability of an account increases with the permanent earnings of a 25-year-old).

Table 11 also shows that incumbent lenders at the time of the reform benefit from the consumer committing not to accept balance transfers. Consumers with a credit card, however, would prefer

[6]Part of the decline in credit limits is driven by selection because the probabilities of receiving credit offers change. However, even if we simulate a partial equilibrium in which probabilities of credit offers do not change, we observe lower limits in comparison to the initial steady state. The tightening of the credit limit with a tighter cap is also shown in Figure 1d in Section 5.1.1.

[7]Among those without credit cards, 99 percent would choose to commit.

to renege on their commitment to not accept offers of balance transfers. Overall, commitment to not accept offers of balance transfers is beneficial across steady states but costly when one takes the transition into account.

Table 11: Welfare implications of eliminating on-the-credit-search

Welfare gains for: (unit = percentage of total annual income)				
Consumers (t = 2)		**Future**	**Incumbents**	**Aggregate**
w. CC	**w/o CC**	**generations**	**(t = 2)**	
-0.39	0.04	0.05	0.21	-0.09

Notes: The table reports the onetime transfers as a percentage of initial steady state total annual income by agent group and for all agents from eliminating on-the-credit-search in the baseline model. The groups of agents are (1) consumers with credit cards (CC) in the first period of the transition, (2) consumers without a credit card in the first period of the transition, (3) future generations whose welfare is discounted at the rate $1/(1+r)$, and (4) incumbent lenders in the first period of the transition.

C.5 Sensitivity analyses

In this section, we present several sensitivity analyses to further understand how revolving credit lines and the ensuing on-the-credit-search matter for rate regulation. The results are presented in Tables 12 and 13.

Lender's bargaining power. Panels B and C of Table 12 report welfare gains from regulation as we vary the lender's bargaining power. The magnitude of the gains increases with the lender's bargaining power. This result highlights the importance of the lender's bargaining power for our quantitative results. Recall that in our baseline calibration, we calibrate the lender's bargaining power to match a measure of the average markup (see Table 3).

Size of the credit card market. Panel D of Table 12 shows that the magnitude of the gains from regulation increases when the size of the credit card market is larger. This suggests that the gains from our baseline calibration, which understates credit usage, are conservative.

On-the-credit-search. The following sensitivity analysis reveals the importance of the contract dilution problem relative to the benefits of fine-tuning the rate regulation to consumer credit worthiness. Without on-the-credit-search, there is no contract dilution problem. As a result, Panels E and G of Table 12 show that the gains for those without a credit card and future generations are highest under submarket-and-time-specific caps. With on-the-credit-search, however, panels A to D and F show that the corresponding gains are highest under the optimal access-dependent cap.

Wage garnishment instead of stigma. Panel F of Table 12 reports the welfare gains from regulation when the main cost of default is wage garnishment instead of stigma. In this variation, the deadweight loss of default is reduced because garnished wages are transferred to the lender (Exler and Tertilt (2020)). In comparison to the baseline, there is only a moderate decline in gains, suggesting that a reduction in the deadweight loss of default is not the primary driver of gains from a rate cap.

Alternative matching elasticities and matching functions. The following comparisons reveal the sensitivity of the contract dilution problem to alternative matching functions. Table 13 reports welfare gains from regulation for alternative values of the matching elasticity parameter ζ and alternative constant returns to scale matching functions (urn-ball and Leontief). Panels D and E report specifications in which credit offers are more elastic with respect to expected profits. In these two cases, the contract dilution problem is more severe, and the optimal access-dependent cap generates larger gains than submarket-and-time-specific caps for those without a credit card in the period of the transition and for future generations. In contrast, Panels B, C, and G report specifications in which credit offers are less elastic with respect to expected profits. As a result, the contract dilution problem is dampened and the relative importance of the contract dilution problem relative to the benefits of fine-tuning the rate regulation to consumer credit worthiness flips. That is, for consumers without a credit card in the period of the transition and future generations, submarket-and-time-specific caps generate weakly larger gains than those from the optimal access-dependent cap.

Table 12: Sensitivity analyses: part 1

Regulation type	Optimal cap (unit = percentage points)		Welfare gains for: (unit = percentage of total annual income)				
			Consumers (t = 2) w. CC	Consumers (t = 2) w/o CC	Future generations	Incumbents (t = 2)	Aggregate
Panel A: Baseline							
Optimal uniform cap	8.2	8.2	0.35	0.01	0.00	-0.03	0.32
Optimal cap ∣ access	12.2	5.2	0.54	0.09	0.09	-0.10	0.62
Submarket-and-time-specific caps			0.78	0.05	0.06	-0.16	0.73
Panel B: Low bargaining power							
Optimal uniform cap	8.2	8.2	0.19	-0.03	-0.04	-0.01	0.10
Optimal cap ∣ access	12.2	5.2	0.31	0.05	0.04	-0.08	0.31
Submarket-and-time-specific caps			0.54	0.01	0.00	-0.14	0.42
Panel C: High bargaining power							
Optimal uniform cap	6.2	6.2	0.94	0.04	-0.01	-0.06	0.92
Optimal cap ∣ access	10.2	4.2	1.09	0.22	0.19	-0.15	1.35
Submarket-and-time-specific caps			1.27	0.17	0.14	-0.15	1.43
Panel D: Target total revolving credit							
Optimal uniform cap	7.2	7.2	1.07	-0.05	0.01	-0.15	0.88
Optimal cap ∣ access	11.2	5.2	1.50	0.18	0.14	-0.50	1.32
Submarket-and-time-specific caps			2.00	0.09	0.07	-0.60	1.55
Panel E: No on-the-credit-search							
Optimal uniform cap	7.2	7.2	0.72	-0.03	-0.01	-0.23	0.46
Optimal cap ∣ access	10.2	6.2	0.90	0.01	0.03	-0.35	0.59
Submarket-and-time-specific caps			1.13	0.03	0.06	-0.44	0.79
Panel F: Wage garnishment							
Optimal uniform cap	7.2	7.2	0.28	0.03	0.06	-0.05	0.32
Optimal cap ∣ access	9.2	4.2	0.51	0.14	0.34	-0.11	0.88
Submarket-and-time-specific caps			0.61	0.05	0.07	-0.14	0.59
Panel G: Wage garnishment and no on-the-credit-search							
Optimal uniform cap	6.2	6.2	0.68	0.05	0.11	-0.18	0.66
Optimal cap ∣ access	8.2	4.2	0.96	0.07	0.20	-0.31	0.92
Submarket-and-time-specific caps			1.00	0.09	0.22	-0.32	0.98

Notes: The table reports the onetime transfers as a percentage of initial steady state total annual income by agent group and for all agents under various model specifications in each panel. Within each panel, the gains are reported for three regulations: (1) optimal uniform cap chosen to maximize welfare for all agents, (2) optimal access-dependent cap chosen to maximize welfare for all agents, and (3) submarket-and-time-specific caps. The following model specifications are presented in each panel: (A): baseline, (B), lower bargaining power (baseline calibration's bargaining power is reduced by 20 percent), (C) lender has all the bargaining power (baseline calibration's bargaining power is increased by 20 percent), (D) target total revolving credit instead of average limit per card holder, (E) no on-the-credit-search, (F) the main cost of default is wage garnishment instead of the stigma cost of default, and (G) the main cost of default is wage garnishment instead of the stigma cost of default and no on-the-credit-search. Each model variation is re-calibrated to target the same set of moments as the baseline to the extent possible or unless specified otherwise. The groups of agents are (1) consumers with credit cards (CC) in the first period of the transition, (2) consumers without a credit card in the first period of the transition, (3) future generations whose welfare is discounted at the rate $1/(1+r)$, and (4) incumbent lenders in the first period of the transition.

Table 13: Sensitivity analyses: part 2

Regulation type	**Optimal cap** (unit = percentage points)		**Welfare gains for:** (unit = percentage of total annual income)				
			Consumers (t = 2)		**Future**	**Incumbents**	**Aggregate**
			w. CC	**w/o CC**	**generations**	**(t = 2)**	
Panel A: Baseline							
Optimal uniform cap	8.2	8.2	0.35	0.01	0.00	-0.03	0.32
Optimal cap ∣ access	12.2	5.2	0.54	0.09	0.09	-0.10	0.62
Submarket-and-time-specific caps			0.78	0.05	0.06	-0.16	0.73
Panel B: Lower ζ							
Optimal uniform cap	8.2	8.2	0.26	0.02	0.01	-0.07	0.21
Optimal cap ∣ access	10.2	5.2	0.48	0.03	0.02	-0.22	0.31
Submarket-and-time-specific caps			0.67	0.05	0.02	-0.25	0.48
Panel C: Low ζ							
Optimal uniform cap	7.2	7.2	0.40	-0.02	-0.03	-0.07	0.28
Optimal cap ∣ access	10.2	5.2	0.55	0.06	0.03	-0.17	0.46
Submarket-and-time-specific caps			0.79	0.06	0.03	-0.23	0.65
Panel D: High ζ							
Optimal uniform cap	7.2	7.2	0.31	-0.04	-0.08	-0.01	0.17
Optimal cap ∣ access	12.2	5.2	0.40	0.07	0.06	-0.08	0.46
Submarket-and-time-specific caps			0.65	0.02	0.03	-0.13	0.56
Panel E: Higher ζ							
Optimal uniform cap	8.2	8.2	0.23	-0.02	-0.08	-0.01	0.11
Optimal cap ∣ access	14.2	5.2	0.36	0.07	0.07	-0.07	0.43
Submarket-and-time-specific caps			0.59	0.02	0.01	-0.12	0.49
Panel F: Urn-ball matching							
Optimal uniform cap	6.2	6.2	0.38	-0.10	-0.08	0.00	0.20
Optimal cap ∣ access	11.2	4.2	0.49	0.09	0.01	-0.12	0.47
Submarket-and-time-specific caps			0.68	0.01	-0.02	-0.15	0.52
Panel G: Leontief matching							
Optimal uniform cap	6.2	6.2	0.53	-0.09	0.11	-0.05	0.50
Optimal cap ∣ access	14.2	4.2	0.71	0.00	0.14	-0.12	0.72
Submarket-and-time-specific caps			0.96	0.03	0.19	-0.20	0.98

Notes: The table reports the onetime transfers as a percentage of initial steady state total annual income by agent group and for all agents under various model specifications in each panel. Within each panel, the gains are reported for three regulations: (1) optimal uniform cap chosen to maximize welfare for all agents, (2) optimal access-dependent cap chosen to maximize welfare for all agents, and (3) submarket-and-time-specific caps. The following model specifications are presented in each panel: (A): baseline, (B), matching elasticity parameter ζ is lowered by 75 percent of the baseline value, (C) matching elasticity parameter ζ is lowered by 50 percent of the baseline value, (D) matching elasticity parameter ζ is increased by 50 percent of the baseline value, (E) matching elasticity parameter ζ is increased by 100 percent of the baseline value, (F) urn-ball matching function, and (G) Leontief matching function. Each model variation is re-calibrated to target the same set of moments as the baseline to the extent possible or unless specified otherwise. The groups of agents are (1) consumers with credit cards (CC) in the first period of the transition, (2) consumers without a credit card in the first period of the transition, (3) future generations whose welfare is discounted at the rate $1/(1+r)$, and (4) incumbent lenders in the first period of the transition.