Finance and Economics Discussion Series

Federal Reserve Board, Washington, D.C.
ISSN 1936-2854 (Print)
ISSN 2767-3898 (Online)

Automated Credit Limit Increases and Consumer Welfare

Vitaly M. Bord, Agnes Kovacs, and Patrick Moran

2025-088

Please cite this paper as:
Bord, Vitaly M., Agnes Kovacs, and Patrick Moran (2025). "Automated Credit Limit Increases and Consumer Welfare," Finance and Economics Discussion Series 2025-088. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2025.088.

Automated Credit Limit Increases and Consumer Welfare*

Vitaly M. Bord† Agnes Kovacs‡ Patrick Moran§

September 17, 2025

Click for most recent version

Abstract

In the United States, credit card companies frequently use machine learning algorithms to proactively raise credit limits for borrowers. In contrast, an increasing number of countries have begun to prohibit credit limit increases initiated by banks rather than consumers. In this paper, we exploit detailed regulatory micro data to examine the extent to which bank-initiated credit limit increases are directed towards individuals with revolving debt. We then develop a model that captures the costs and benefits of regulating proactive credit limit increases, which we use to quantify their importance and evaluate the implications for household well-being.

*This paper was prepared for the Carnegie-Rochester-NYU Conference Series on Public Policy. The views expressed in this paper are solely those of the authors and do not represent the views of the Federal Reserve Board or the Federal Reserve System. Cord Barnes provided excellent research assistance.

†Federal Reserve Board (vitaly.bord@frb.gov)

‡King's College London, University of Manchester, and IFS (agnes.kovacs@kcl.ac.uk)

§Federal Reserve Board, IFS, and CEBI (patrick.e.donnellymoran@frb.gov)

1 Introduction

As algorithmic decision-making reshapes consumer finance, a critical tension has emerged between the efficiency of automated credit decisions and the protection of vulnerable consumers. In the credit card market, limit increases are a particularly important but understudied source of credit, affecting more than 12% of accounts each year. Countries differ in their approach to regulating limit increases. For example, in the United States, the overwhelming majority of these increases are implemented automatically by lenders using proprietary models rather than requested by consumers. By contrast, reflecting concerns about indebtedness and consumer protection, several countries have restricted banks' ability to raise credit limits: for instance, the UK now prohibits limit increases for borrowers who have been in persistent revolving debt, while Canada prohibits bank-initiated credit limit increases without consumers' consent.

Policy makers face an important question: to what extent should they regulate algorithmic decision-making in credit markets, particularly when banks increase credit limits automatically? On the one hand, automatic credit limit increases can be beneficial, as they relax credit constraints and give households greater flexibility to smooth consumption over adverse shocks. On the other hand, such increases may also be detrimental. Most of a bank's credit card profits come from consumers who carry persistent debt (Adams et al., 2022), creating incentives for banks to direct limit increases toward these individuals. If some consumers struggle with self-control, additional credit may lead to greater indebtedness (Laibson, 1997; Gul and Pesendorfer, 2004). Indeed, empirical evidence shows that consumers borrow more after credit limit increases, even when they did not request the increase and are not observably constrained (Gross and Souleles, 2002).

In this paper, we exploit regulatory data on credit card lending to investigate who receives bank-initiated credit limit increases, then develop a quantitative model to evaluate the costs and benefits of allowing banks to proactively raise credit limits. We make three main sets of contributions. First, in the data, we present several new stylized facts about the importance of limit increases in credit card lending strategies, the role of proactive bank-initiated increases that are not requested by consumers, and their relationship to revolving debt. Second, we examine the extent to which bank-initiated credit limit increases are directed towards individuals with revolving debt, and consistent with existing literature, document that debt rises after limit increases. Third, we develop a model of household behavior that allows us to perform the first quantitative analysis of the positive and normative implications of restricting bank-initiated credit limits increases, examining both the UK approach of prohibiting increases for revolving borrowers and the Canadian approach of requiring explicit consumer consent. Overall, we believe our results have important implications for consumer protection and the emerging field of algorithmic regulation in consumer finance. Indeed, while many aspects of credit card markets are regulated, there is little oversight of the factors lenders can use to proactively raise credit limits.

Our empirical analysis utilizes regulatory data on credit card lending from Federal Reserve's Capital Assessments and Stress Testing Reports (Y-14) filed by large credit card issuers, covering more than 70% of the U.S. credit card market. This dataset is uniquely suited to our purpose because it allows us to differentiate between bank- and consumer-initiated credit limit increases, something which is impossible to observe in most other datasets, such as credit bureau data. The Y-14 data contain monthly, account-level information about granular aspects of credit usage including balances, purchase volumes, finance charges, and fees.

We begin by cataloging several new stylized facts about the role and prevalence of limit increases in the U.S. credit card market. We show that limit increases are an important source of consumer credit, with half as much additional available credit coming from limit increases as from new account originations. Limit increases are relatively more important for lower credit score borrowers, as lenders often follow a "low and grow" strategy on these accounts— with low limits at origination, followed by increases depending on borrower behavior. The share of revolving balances — the part of total balances that are carried from the previous month rather than paid off and thus accrue interest — made possible by limit increases is similarly higher for low-credit score borrowers.

Importantly, the overwhelming majority of limit increases in the U.S. are proactively implemented by banks rather than requested by consumers. Credit card lenders have extensive amounts of data on their existing customers, which they can mine with machine learning and artificial intelligence algorithms to ascertain the most profitable customers to whom to give limit increases. Consistent with their use of algorithms, we show that banks that more often refer to 'artificial intelligence' or 'machine learning' in their annual filings tend to support a larger share of revolving balances with limit increases, rather than the credit limit granted at origination.

Next, building on these stylized facts, we examine the extent to which credit card utilization is correlated with the likelihood of receiving a limit increase. We distinguish between two different types of credit card utilization: revolving utilization, which reflects debt that is carried from prior months, and transacting utilization, which reflects new purchases. While both revolving utilization and transacting utilization are positively correlated with a limit increase, we find that the likelihood of receiving a limit increase varies across the two types of utilization. More specifically, the correlation with revolving utilization follows an inverse-U shape with the highest probability of receiving a limit increase occurring for moderate levels of utilization. By contrast, the correlation with transacting utilization follows a logarithmic growth pattern: the probability of a limit increase rises with utilization until about a utilization of 0.3, and above that level, utilization does not appear to vary with the probability of an increase. We find that these patterns vary somewhat with credit score, likely reflecting underlying differences in risk.

We conclude the empirical results with a simple event study examining what happens to utilization after a limit increase. Our findings are consistent with the results established in the literature: in the months after a limit increase, utilization rebounds to pre-increase levels, as consumer increase their revolving debt. Notably, this effect occurs even in accounts that are not near their credit limit and so are unlikely to be liquidity constrained. This is consistent with the presence of self-control issues, which also predicts that higher credit limits may lead to greater borrowing.

After presenting the empirical results, we develop a model that allows us to perform the first quantitative analysis of novel real-world policies that restrict bank-initiated credit limit increases. In the model, households make consumption, saving, and borrowing decisions over the life-cycle while faced with uninsurable income and employment risk. We allow for heterogeneous preferences following Nakajima (2017), allowing for two types of households with and without self-control issues à la Gul and Pesendorfer (2001, 2004). We assume hidden information on household type, such that credit card companies cannot make lending decisions based on type, but only based on observed consumer behavior. In the model, the credit card company initially offers a similar product to all consumers, but with the option of tailoring credit limits later. We calibrate the credit limit increase function using Y-14 data, then internally estimate the preference parameters to match aggregate statistics on the U.S. credit card market.

The above model captures two opposing effects from allowing proactive limit increases based on machine learning algorithms. According to traditional theory, consumers benefit from additional credit, as it gives them greater flexibility to smooth consumption and self-insure against risk. At the same time, if bank algorithms implicitly target consumers with self-control issues, it may lead to greater temptation and increased borrowing, which may be detrimental to consumers even if it is profitable to banks. Taken together, our model captures the benefit of relaxed constraints, as well as the potential risk of giving too much credit to consumers with behavioral biases (see e.g. Livshits, 2020).[1]

We begin by analyzing the properties of the baseline model, where consumers receive credit limit increases based on the empirical evidence from the Y-14 data. Through the lens of our model, we find that most credit limit increases go towards consumers with self-control issues, since these borrowers are more likely to maintain revolving balances. As a result, many limit increases are detrimental from the perspective of the consumer.

[1] We focus on self-control issues due to their well-documented importance in credit card borrowing (Meier and Sprenger, 2010; Gathergood, 2012; Kuchler and Pagel, 2021).

That said, we find important heterogeneity regarding which limit increases are beneficial: while customers with zero or low utilization almost always benefit from greater access to credit, customers with moderate or high utilization are frequently harmed by proactive limit increases, as these consumers are more likely to suffer from self-control issues.

Using the above model, we analyze the positive and normative implications of alternative regulations that restrict banks from raising credit limits in certain situations. In our first counterfactual exercise, we evaluate a policy that prohibits banks from raising the credit limits of revolving borrowers, motivated by a recent policy in the UK. In our model, we find that this policy results in a large decrease in the share of credit limit increases going towards consumers with self-control issues. As a result of fewer limit increases, the debt-to-income ratio declines by roughly 2 percentage points, even as the average utilization rate increases slightly due to reduced credit. Turning towards welfare, we find that prohibiting limit increases to revolving borrowers improves welfare by 1.1% in terms of consumption equivalent variation. While agents without self-control issues are slightly harmed by reduced access to credit, these costs are small compared to the benefits of reduced temptation and lower interest expenditure by agents with self-control issues. Finally, we find that when the firm is able to re-optimize its credit limit increase function in response to the new policy restriction, the policy helps to shift credit limit increases from agents with self-control issues to agents without such issues, and thus the counterfactual policy continues to improve overall well-being.

The second counterfactual policy we analyze restricts banks from increasing credit limits without consumer consent, inspired by recent regulations in Canada, Singapore, and New Zealand, which will be implemented across EU member states in 2026. This policy shows similar positive and normative implications as prohibiting limit increases for revolvers, under the assumption that consumers are fully aware of their self-control issues. That said, we find meaningful differences if we alter our assumptions about the share of sophisticated versus naïve consumers. Requiring consumer consent improves welfare when consumers are sophisticated, whereas prohibiting limit increases for revolving borrowers remains effective even if all households are naïve.

Taking stock, our results have important policy implications for algorithmic regulation in credit card markets. While the U.S. has extensive regulations on credit rejections and limit decreases, there is a notable lack of regulation surrounding proactive bank-initiated limit increases.[2] Although we do not observe the exact algorithms used by banks to determine their optimal policy for proactively raising credit limits, our empirical analysis

[2]Lenders have to follow two main types of regulations around credit limit changes: when decreasing a credit limit or rejecting an application to increase a credit limit, lenders needs to provide a reason under the Equal Credit Opportunity Act; and when increasing limits, they need to abide by "ability-to-pay" rules and only make loans that borrowers can reasonably pay given their income and other obligations, although Fulford and Stavins (2025) find that these rules are generally non-binding.

demonstrates a revealed preference for giving additional credit to revolving borrowers. In our calibrated model, this implies that credit increases are disproportionately allocated to consumers with self-control issues, raising concerns about exploitative contracting in the spirit of Heidhues and Kőszegi (2010). Taken together, we conclude that if some households have self-control issues, there are strong consumer protection reasons to regulate the algorithms used by banks to proactively to raise credit limits.

Related Literature. Our paper contributes to several strands of literature. First, an extensive literature has examined how credit limits affect consumption and borrowing in the credit card market (Agarwal et al., 2023, 2017; Fulford and Schuh, 2023; Gross and Souleles, 2002). These papers generally find that borrowers alter their spending with changes in their credit limit even if they are unconstrained and have utilization below their credit limit. Most of these studies do not establish a causal interpretation to these findings, with the exception of Aydin (2022) who uses a field experiment to identify a causal effect and Chava et al. (2023) who examine the effects of credit limit changes arising from bank funding shocks. A causal effect is generally difficult to identify, as automated changes to credit limits may be responses to borrower actions. Kovrijnykh et al. (2023) find that lenders increase limits on borrowers' existing credit card accounts when borrowers open new credit cards, consistent with the new accounts revealing positive information about the borrowers' creditworthiness. [3]

We add to this literature by examining the account and borrower characteristics that help predict limit increases. To our knowledge, the only previous paper that has focused on account-level drivers is Fulford and Stavins (2025), who examine how the ability-to-pay rules implemented as part of the 2009 Credit Card Accountability Responsibility and Disclosure (CARD) Act affected lenders' granting of credit limit increases. They find that although most income updates are followed by limit increases, lenders often increase limits without receiving any updates about borrower income. We add to this literature by being the first to evaluate the importance of utilization and revolving behavior in driving bank-initiated limit increases, and to document the differential roles of bank- and consumer-initiated limit increases.

Second, our analysis of the welfare effects of restricting bank-initiated credit expansion contributes to the literature on consumer protection regulation for consumers with biased or nonstandard preferences. Heidhues and Kőszegi (2010) show that when some borrowers have present-bias preferences, it is welfare-improving to ban large penalties for deferring small repayments, and Heidhues and Köszegi (2015) argue that in the credit card market, the welfare costs of lenders taking advantage of consumers' misunderstandings may be large. Although regulation may help prevent over-borrowing, it may nevertheless be

[3] In addition, Yin (2022) shows that borrowers infer both higher growth in future personal income and and higher macroeconomic growth from lender-initiated credit limit increases.

welfare reducing for irrational borrowers (Exler et al., 2024). Along these lines, Jolls and Sunstein (2005) argue that the law can help de-bias consumers by nudging them towards the rational outcome, rather than by enacting regulation. On the empirical side, Bertrand and Morse (2011) perform a field experiment and find that additional information disclosures about payday loans reduce customers' borrowing. Stango and Zinman (2011) show that weaker enforcement of Truth in Lending Act disclosures widens the gap in prices between more- and less- biased consumers. Gathergood (2012) shows that self-control issues are strongly correlated with over-indebtedness in the UK credit card market. Relative to the existing literature, we are the first to document that bank-initiated credit limit increases are often targeted towards revolving borrowers, and the first to analyze the recently-implemented real-world policies that restrict this behavior.

Third, we contribute to an emergent literature on the regulation of algorithmic decision making (see e.g. Blattner et al., 2021). While banks release little public information on their approach to proactive credit limit increases, many credit card companies use large-scale experiments to evaluate and optimize the profitability of credit limit adjustment functions (Botella, 2022). This approach closely resembles reinforcement learning, where a bank adopts a policy rule for credit limit increases, which they optimize based on the observed rewards to following that policy versus occasionally deviating from that policy, as described by Sutton and Barto (2018). Indeed, recent research shows that a machine learning algorithms bank-initiated credit limit increases may outperform other approaches.[4] We contribute to this literature by examining how credit card companies give limit increases based on observed borrowing behavior and the potential implications for household well-being.

Finally, we contribute to the literature on the revolving and transacting, or convenience, functions of credit cards. Adams and Bord (2020) stress the importance of separating accounts by their credit card usage to understand credit market dynamics during the COVID-19 pandemic. Grodzicki and Koulayev (2021), Fulford and Schuh (2023), and Lee and Maxted (2023) show that revolver status is highly persistent. Adams et al. (2022) emphasize that revolvers are the primary driver of credit card lenders' income, with interest income comprising the vast majority of lenders' revenues and profits. While other papers such as Agarwal et al. (2017) have discussed profitability in the context of credit limit increases, we focus on disentangling the roles that the revolving and transacting utilizations play in driving limit increases.

The rest of the paper is organized as follows. Section 2 introduces the data and establishes several motivating stylized facts on the importance of proactive limit increases, the prevalence of "low-and-grow" strategies, and their role in making possible higher ag-

[4]See (Alfonso-Sánchez et al., 2024) for an analysis of the outperformance of the Double Q-Learning algorithm

gregate revolving debt. Section 3 evaluates the relationship between revolving utilization and limit increase at the loan level. Building on these findings, section 4 presents the quantitative model and section 5 discusses the main model results and counterfactuals. Section 6 concludes and discusses directions for further research.

2 Data and stylized facts

In this section, we first discuss the regulatory data we use in the empirical parts of the paper. Next, we present several new stylized facts on the prevalence and importance of credit card limit increases. Finally, we turn to an overview of the regulatory framework regarding limit increases.

2.1 Data

Our data come from Form FR Y-14M of the Federal Reserve's Capital Assessments and Stress Testing Report (henceforth, Y-14) for the January 2014 to December 2024 time period. These data provide monthly, account-level information on all credit cards issued by large stress-tested banks with more than $100B in assets. The 26 banks in the dataset during our sample period comprise, in aggregate, more than 70% of credit card balances in the U.S.. Because of the large size of the dataset, for our analysis, we use a 0.5% sample of accounts, consisting of more than 150 million observations corresponding to more than 3.6 unique active credit cards.

The data include detailed information on the characteristics and usage of each credit card account, including credit limit, balances, utilization, total purchases, finance charges, and delinquency. The data also include measures of the borrower's credit score, which are available both updated and at origination. Self-reported income is available at origination for most borrowers. Until 2020, lenders also reported updated income whenever it was refreshed by the borrower; however, this information is not available since 2020 (see Fulford and Stavins (2025) for more details).

The Y-14 dataset has two advantages over credit bureau data, which is often used to examine consumer financial health and behavior. First, because the data include information on balances and actual payments, it is possible to identify revolvers—accounts that utilize the credit feature of a credit card and carry balances month to month on which they pay interest. It is important to identify revolvers because interest comprises the vast majority of the profitability of credit card portfolios, and so revolvers are a credit card lender's most profitable customers (Adams et al., 2022). Second, the data allow us to identify credit limit changes and whether they were initiated by the lender or the borrower. To our knowledge, this is the only currently-available dataset that identifies the source of a change in credit limit, thus allowing us to study bank-initiated limit changes separately from consumer-initiated ones.

One difference between the Y-14 data and more-commonly-used credit bureau data is that the Y-14 dataset is at the account level. Although we are able to examine some borrower characteristics such as credit score, we are unable to link to the borrower's other loans or even other credit cards. A further discussion of the data and summary statistics is presented in section A.2 of the Appendix.

2.2 Motivating stylized facts about credit card limit increases

Credit card limit increases are an understudied aspect of the provision of credit in consumer credit markets. In this subsection, we present a variety of aggregate stylized facts that motivate the model we present in section 4.

Fact 1. Limit increases are an important source of credit.

Figure 1, Panel A, plots the total credit limit—that is, the amount of available credit—coming from newly originated cards and from credit card limit increases each quarter. Panel B plots the total number of credit cards that are issued each quarter and the total number of cards that undergo limit increases.

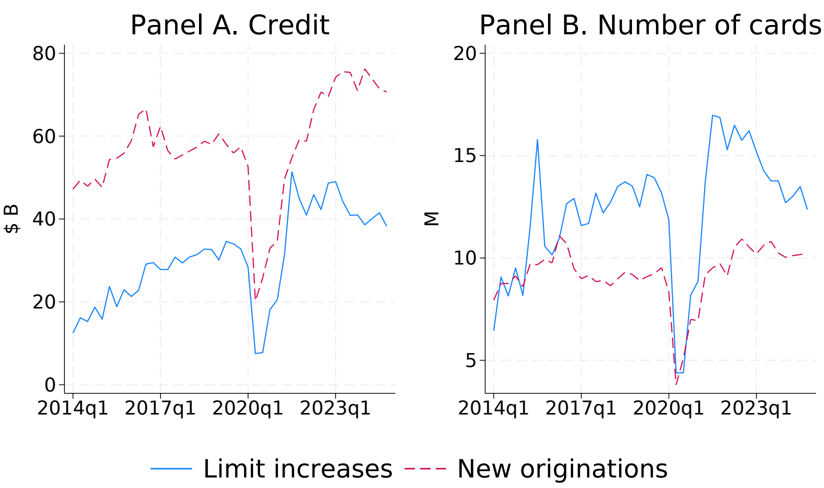

Figure 1: Credit card limit increases and new issuance

As the figures show, during the post-pandemic period, credit card limit increases result in more than \$40B of additional available credit each quarter, which is almost 60% of the approximately \$70B of available credit that arises from new card issuance. Prior to the pandemic, credit limit increases were about \$30B, or about half of new issuance. In addition, the total number of cards that undergo a credit limit increase each quarter is 30% higher, on average, than the total number of new cards issued.

Fact 2. Limit increases are a particularly important source of credit for higher-credit-risk borrowers.

Figure 2, Panel A, shows the average credit limit over time by credit score at origination for cards originated since 2014. We define as superprime all borrowers with credit scores above 760; as prime those with credit scores of 680-760; as near prime those with credit scores of 620-680; and as subprime those with credit scores below 600.[5]

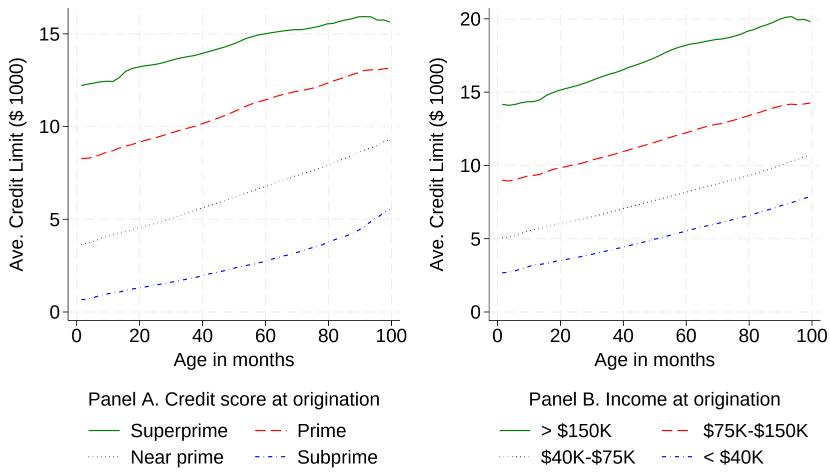

Figure 2: Average credit limit by account age

The panel shows a vast difference in the average credit limit at origination across the credit score groups. The average superprime credit card is granted a credit limit of more than $12,000 at origination, while the average subprime credit card limit is only $700. By five years after origination, the average superprime credit card limit is increased to $15,000 (a 25% increase), and it remains approximately unchanged by 8 years after origination. By contrast, the credit limit on the average subprime credit card increases to $2,700 (a 285% increase), and it continues to increase to almost $5,000 by 8 years after origination.[6] This is consistent with "low-and-grow" strategies of giving higher-risk borrowers low initial credit limits and then increasing them based on borrower behavior. The increases we document are driven to some extent by credit score improvements, although this cannot explain the full increase in credit.[7]

[5] A variety of credit scores and credit scoring methods are reported in the Y-14M data, which we standardize to these bins. Different reported methodologies use different formulas and even different scales. However, credit scores are essentially ordinal scales of risk, and so grouping them into these large bins, rather than more granular bins, avoids the bin definitions being influenced by particular credit scoring methodologies or scales.

[6] These magnitudes are smaller when limiting to accounts that are not closed or charged off. Credit cards that are observed through 2024 have an average credit limit of $1,300 at origination, which increases to $2,300 by five years after origination (a 75% increase) and $3,200 by 8 years after origination. Credit cards of near prime borrowers show similar trends, while those of prime and superprime borrowers are mostly unchanged since they are less likely to be closed or charged-off.

[7] Credit limits among borrowers who remain subprime increase to only $1600, on average, by 5 years

Panel B shows the average credit limit over time by four income groups: more than \$150,000, \$75,000-\$150,000, \$40,000-\$75,000, and less than \$40,000. Income is self-reported at origination. Differences in credit limits at origination by income group are somewhat smaller than by credit score—consistent with credit score playing a larger role in determining credit at origination—and the trends over time are similar.

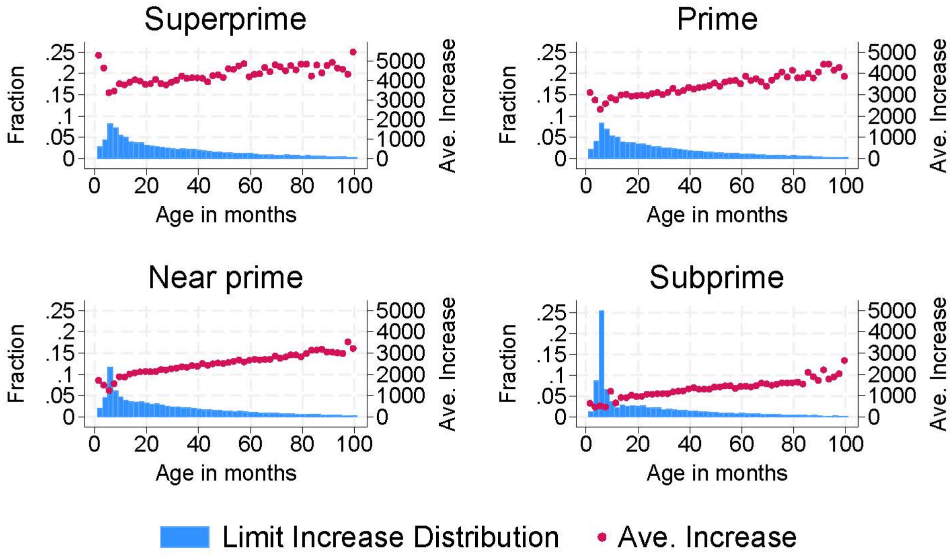

Figure 3: Distribution of limit increases by account age and credit score

Figure 3 shows the distribution of limit increases by account age, as well as the average amount of the increase, for the four credit score bins.[8] All credit score groups tend to experience more limit increases earlier in the lifecycle of their accounts, but the effect is particularly pronounced for subprime and near prime borrowers. Almost 33% (20%) of limit increases at subprime (near prime) accounts occur in the first 6 months after origination, but only 10% (14%) occur after 5 years. By contrast, about 15% of limit increases at prime and superprime accounts occur in the first 6 months after origination and a similar 15% occur after 5 years. Notably, about 55% of subprime accounts, 35% of near prime accounts, and 25% of prime and superprime accounts undergo at least one limit increase within the first 6 months of account origination. Section A.1 includes the distributions by the four income groups we use. Again, the patterns are similar.

Fact 3. Most credit card limit increases are bank-initiated.

Figure 4 shows the share of bank-initiated credit limit increases, weighted by the size of the increase (the blue line) and by the number of credit cards undergoing a limit increase (the red line). In both cases, about 75-80% of limit increases have been bank-initiated in

after origination and about \$2700 by 8 years after origination.

[8]This figure includes only credit cards originated since January 2014.

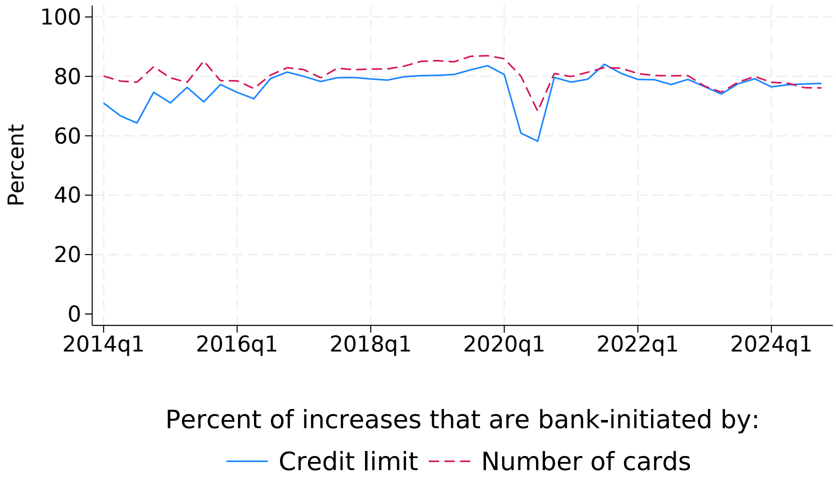

Figure 4: Percent bank-initiated limit increases

recent years. The high level of bank-initiated increases likely relates to the importance of information and relationships in banking. When a new customer applies for a credit card, the lender must decide whether to approve and what size of credit limit to grant the borrower using only information on the borrower's credit score, credit report, and self-reported income. However, after the borrower begins using the card, the lender has extensive information on usage, spending, and payment behavior, which it can then use to adjust the credit limit.

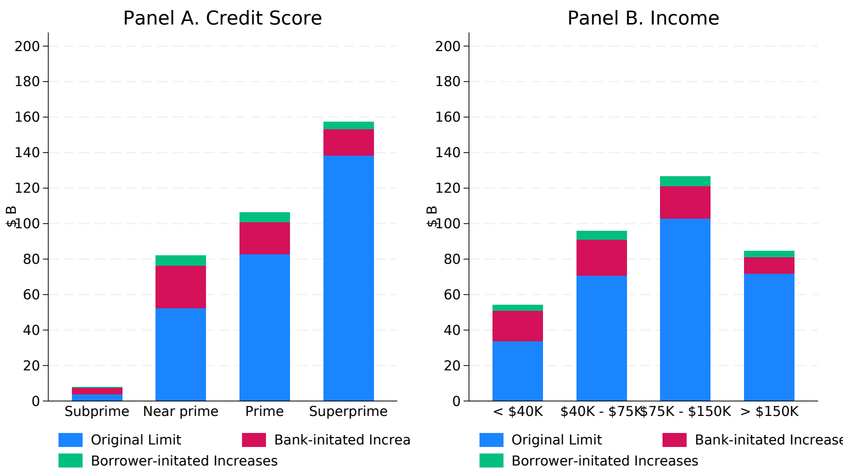

Figure 5: Credit Limit

Figure 5 disaggregates the total credit limit outstanding into three components: the credit limit at origination; bank-initiated limit increases; and borrower-initiated limit increases.[9] Panel A shows that the credit limit at origination comprises the majority of

[9]The figure includes cards originated since January 2014. For simplicity, we ignore limit decreases, which we find are relatively rare (see Figure 9).

total credit limit in each credit score group, and the share of the total limit that is due to increases after origination is higher for subprime and near prime borrowers. Panel B shows a qualitatively similar pattern by income.

Fact 4. The share of revolving balances made possible by limit increases is higher for lower-credit-score borrowers.

A credit card has two main functions: a convenience function which allows the account holder to use it as a method of payment and a credit function which allows the account holder to use the credit card as a source of credit. Figure 6 examines the extent to which limit increases make possible the revolving function of credit cards. To estimate this, we split revolving balances—the part of total credit card balances that is carried from one month to the next, not fully paid off and thus accrues interest—into the component made possible by the card's original credit limit and the component made possible by subsequent limit increases. We use a conservative approach and only assign the part of revolving balances that exceed the original credit limit as supported by limit increases.

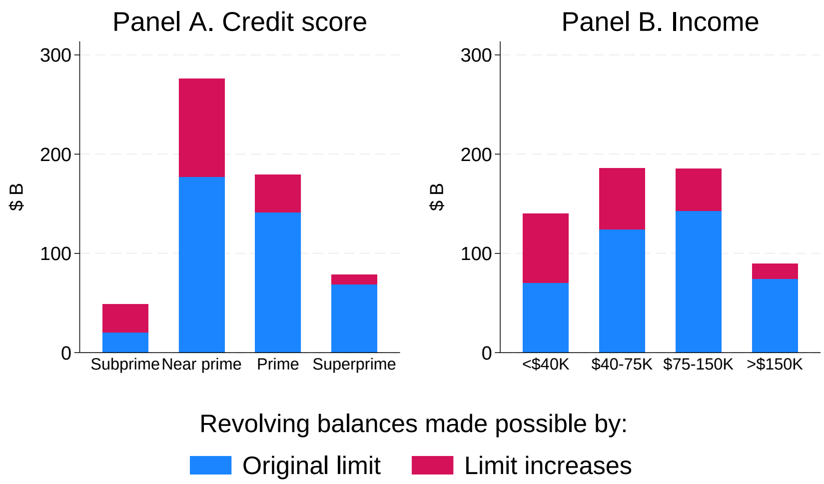

Figure 6: Outstanding revolving balances and limit increases

In this figure, we assign revolving balances as either supported by the original credit limit or by subsequent credit limit increases. We follow a conservative approach and assign only the portion of revolving balances that exceed the original credit limit, we assign the portion to limit increases. Credit scores are as of origination; revolving balances are as of the end of 2024.

Panel A of Figure 6 shows revolving balances as of December 2024, split by whether they were made possible by the original credit limit or subsequent increases, and by the credit score at origination. Overall, about 30% of revolving balances are made possible by limit increases, but this figure varies with credit score. The share of revolving balances supported by credit limit increases is 60% for cards of subprime borrowers, 35% for near prime borrowers, and only 20% and 12% for prime and superprime borrowers, respectively. Panel B splits revolving balances by income at origination. As with Panel A, higher-income groups have a lower share of revolving balances that are made possible by subsequent limit increases.

Fact 5. Lenders that mention AI more have a higher share of revolving balances made possible by limit increases.

Figure 7 splits lenders by the median of the distribution of the number of times they refer to 'machine learning' (ML) or 'artificial intelligence' (AI) in their 2024 10K filings. Panel A shows that for all credit score groups, lenders with above median mentions of AI have a higher share of revolving balances made possible by credit limit increases. AI and machine learning technology lend themselves easily to models of limit increases. Lenders are likely able to use data on customers' credit usage to determine to whom they should give limit increases (Botella, 2022). Further, because revolvers comprise the majority of the profits of a bank's portfolio (Adams et al., 2022), giving limit increases to support revolving balances may help to boost lender profits, subject to the risk of default. Panel B similarly a similar pattern by income. In Figure 23 of the Appendix, we show that these results are not driven by differences in credit limits issued at origination (which can mechanically constrain the revolving debt supported prior to any limit increases).

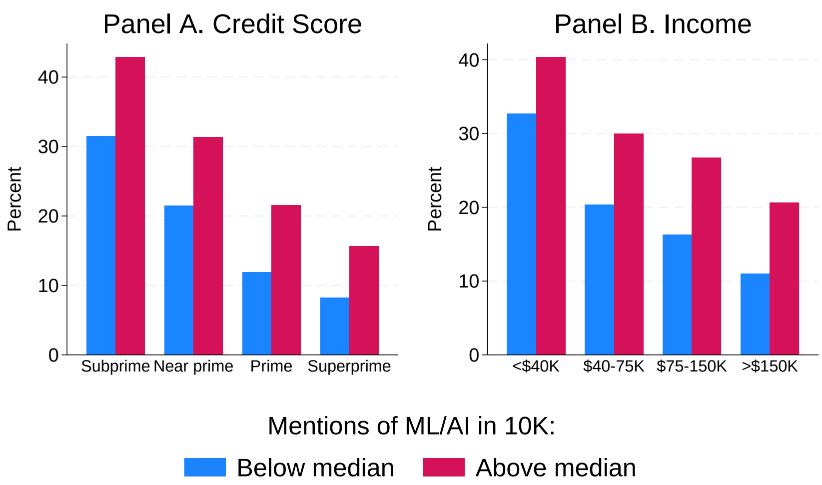

Figure 7: Percent of revolving balances supported by limit increases and ML/AI

This figure splits lenders by the number of mentions of artificial intelligence and machine learning in their 10K reports. The share of revolving balances supported by limit increases is calculated as in Figure 6. Credit scores are as of origination; revolving balances are as the end of 2024.

Fact 6. Credit card limit increases are more prevalent among revolvers.

Figure 8 shows the incidence of limit increases among "revolving accounts" that had revolved at least once in the previous three months and accounts with no recent revolving. Panel A shows that bank-initiated increases are about 1.5-2 times more prevalent among revolving accounts, relative to non-revolving accounts. Panel B shows that for borrower-initiated limit increases, there is essentially no difference in the prevalence of increases among revolvers and non-revolvers. Taken together, the two panels of Figure 8 suggest that an account's revolving status is an important factor in whether the account is granted a bank-initiated limit increase by the lender.[10]

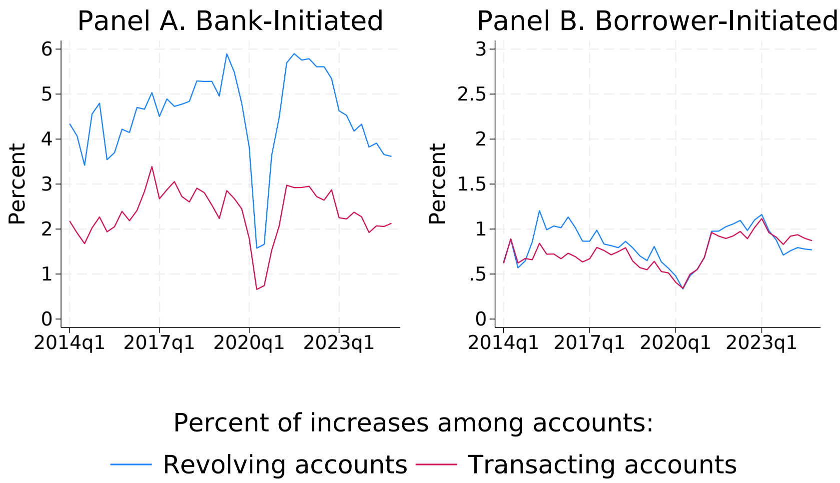

Figure 8: Limit increases and revolving status

Fact 7. Limit decreases are much less common than increases.

Figure 9 examines the relative prevalence of limit increases and decreases.[11] Panel A shows the change in total available credit from limit increases and decreases each quarter, as a share of the total credit limit on open cards. The aggregate amount of available credit increases by about 1-1.5% each quarter from limit increases, and decreases only about 0.2% from limit decreases. The only period when there were more limit decreases than increases was the 2nd quarter of 2020, when the number of limit increases plummeted

[10] In Appendix A.1, we examine how this finding relates to Fact 2, which shows that a high share of limit increases occur within the first 6 months. We find that a large share of accounts begin revolving soon after origination, suggesting that lenders can likely predict future revolving behavior from revolving behavior in the first few months, and may target limit increases accordingly.

[11] Figure 9 examines limit increases and decreases among active cards. About one third of bank-initiated limit decreases are among inactive cards–that is cards, that have not been used in at least one year. Since limit changes on inactive cards are unlikely to have a large effect on borrowers, we focus on active cards throughout this paper.

because of concerns about consumer financial health due to the COVID-19 pandemic. Notably, while limit decreases doubled in that one quarter, they remained at only 0.5% of all cards. Similarly, Panel B shows that about 4% of credit cards undergo a limit increase each quarter, while only 0.5% undergo a limit decrease.[12] Because limit decreases are rare—both in absolute terms and compared to limit increases–we focus on increases in this paper and leave an examination of limit decreases for future research.

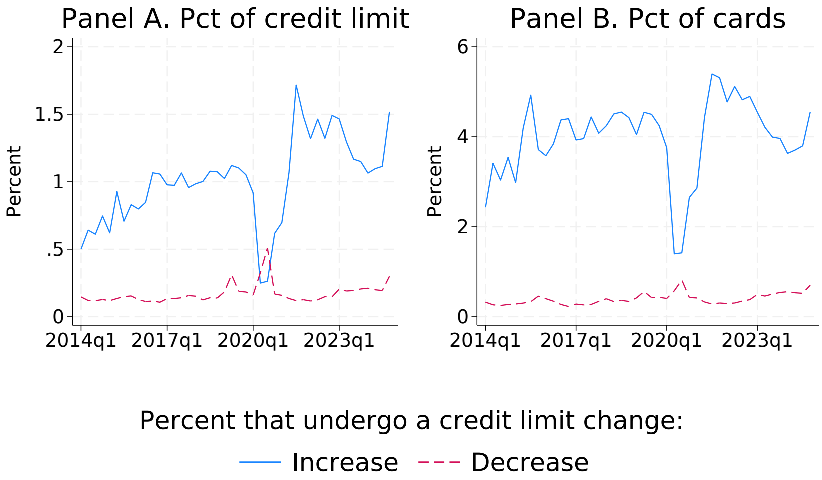

Figure 9: Credit limit increases and decreases

2.3 Regulations on Proactive Credit Limit Increases

Having established the importance of credit card limit increases for consumer credit in the United States, we next turn to a discussion of the regulations surrounding limit increases, both in the United States and internationally.

In the United States, there is very little regulation of bank-initiated credit limit increases. Two main types of regulations apply to any granting of credit, including limit increases and originations of new credit cards. First, the Equal Credit Opportunity Act stipulates that when decreasing a credit limit or rejecting an application to increase a credit limit, lenders need to provide a reason; they similarly need to provide a reason when rejecting an application for a new credit card. Notably, this regulation applies to consumer-initiated credit increases, but not bank-initiated increases. Second, when increasing limits——and approving applications for new cards——lenders need to abide by "ability-to-pay" rules and only make loans that borrowers can reasonably pay given their income and other obligations.[13]

By contrast, several countries have passed legislation that regulates proactive credit

[12]Because there is some autocorrelation in limit increases, aggregating to the yearly level yields that more than 12% of cards undergo a limit increase each year, adding about 5% to total credit limit.
[13]Fulford and Stavins (2025) find that these rules are generally non-binding.

limit increases. The legislation comes in two broad varieties: bans of unsolicited limit increases to some or all borrowers (United Kingdom, Australia) and requirements that lenders obtain customer consent before any increase is implemented (Canada, Singapore, New Zealand, and European Union).

Bans: In the United Kingdom, a lender is prohibited from increasing limits for revolving borrowers who have been in persistent revolving debt for 12 months or longer, where persistent is defined as when the amount paid towards interest and fees exceeds the amount of principal repaid.

In Australia, since 2018, the National Consumer Credit Protection Act has prohibited bank-initiated credit limit increases in all cases. Whereas lenders cannot initiate or suggest a credit limit increase, consumers can request one.

Approvals: In Canada, a lender must obtain a customer's consent at the time of each credit limit increase. If the customer gives oral consent to a credit increase, the lender must confirm it in writing by the next account statement. Similar legislation is in effect in Singapore and New Zealand.[14]

The EU has passed similar legislation requiring member states to implement a ban on "unilateral" credit increases by November 20, 2026. Lenders will not be allowed to increase credit limits without the customer's approval or request.

3 Empirical Analysis

Having established several stylized facts about credit limit increases, we next turn to our empirical analysis. Expanding on Fact 6, we show that revolving behavior appears to play a role in the likelihood of a bank-initiated limit increase. In addition, consistent with existing literature, we find that borrowers' revolving utilization grows after limit increases.

3.1 Who receives credit card limit increases?

In this section, we investigate the account and borrower characteristics that correlate with receiving a credit limit increase. Lenders use machine learning (ML) algorithms and artificial intelligence (AI) approaches to choose which borrowers to give credit limit increases. Since we are unable to observe these proprietary algorithms, we investigate in the data which account features are the most important in receiving a limit increase. This investigation documents a revealed-preference for giving limit increases to revolving borrowers—that is borrowers who do not fully pay off their balances.

[14] The Monetary Authority of Singapore also has rules for the maximum credit limit on each card which depends on borrower income.

Analysts at credit card lenders use a two stage process of first identifying who should receive a limit increase and then quantifying the size of the increase, according to our conversations with them. We adopt a similar two-stage approach. First, we estimate a linear probability model to determine the factors that affect the probability of receiving a limit increase.[15] Second, we briefly examine the factors that affect the size of the limit increase in Appendix A.5. The intensive margin results are broadly in line with the extensive margin results, therefore we do not discuss them in detail.

To identify the factors that affect the likelihood of receiving a limit increase, we estimate an OLS regression of the form:

$$Y_{it} = \beta Utilization_{it} + \gamma Controls_{it} + \lambda_t + \eta_g + \epsilon_{it} \tag{1}$$

where, in most of our regressions, Y_{it} is an indicator for whether account i received a limit increase at time t+1. $Controls_{it}$ are a set of controls and include log of credit card limit, the interest rate margin, an indicator for whether the borrower has other credit cards with the same lender, and indicators for a bank-initiated limit increase, borrower-initiated increase, and change in credit score in the previous 3 months. We include month fixed effects, λ_t. Other fixed effects, η_g include state, lender, credit score deciles, income deciles, and fixed effects for the types of other relationships the consumer has with the same lender.[16] In our main specification, we also include card portfolio group fixed effects. We define card portfolio groups as a combination of the lender, type of card, and annual fee.[17] Including card group fixed effects allows us to compare credit cards that have similar characteristics and thus are targeted to similar types of borrowers. Throughout, standard errors are clustered at the card portfolio group level.

We focus our analysis on the role that $Utilization_{it}$ plays in bank-initiated limit increases. We measure $Utilization_{it}$ as the statement balances divided by credit limit, averaged over the previous 3 months. Because trends in utilization may also affect the likelihood of a limit increase, we also include the change in utilization from the prior three months in our main regressions (see Appendix A.3).

In addition, we further split utilization into two components. Revolving utilization—defined as revolving balances as a share of the credit limit—is the portion of overall utilization due to balances that the borrower revolves and pays interest on. Transacting

[15]We use a linear probability model to allow us to use the multitude of fixed effects discussed below. However, the results are qualitatively similar using other estimation approaches.

[16]Relationships include deposit, investment account, mortgage, home equity, auto, student loan, installment loan, and multiples of these relationships.

[17]The type of card includes a combination of product type—whether the card is a co-branded credit card, an oil and gas credit card, an affinity card, a student card or other—the reward type–whether the card earns cash, miles, points, another reward, or nothing—and whether the credit card is secured.

utilization, by contrast, is the credit card utilization that arises from purchases. This distinction allows us to differentiate between accounts that have high utilization because they have large purchases but pay them off and accounts that have high utilization from using the credit feature of credit cards. It is also important to note that the two types of utilization are intricately linked—a high transacting utilization one month can become revolving utilization the next month if the borrower does not pay it down.

For brevity, we examine two questions below: how does the probability of bank-initiated limit increases vary with revolving and transacting utilization; and how do these results differ for consumer-initiated limit increases. In Appendix A.3, we present how other account characteristics including credit score and income affect the likelihood of bank- and consumer- initiated increases, and examine heterogeneity by credit score and age.

Utilization and bank-initiated limit increases

We begin by examining how the correlation between utilization and credit limit increases varies across the utilization distribution.

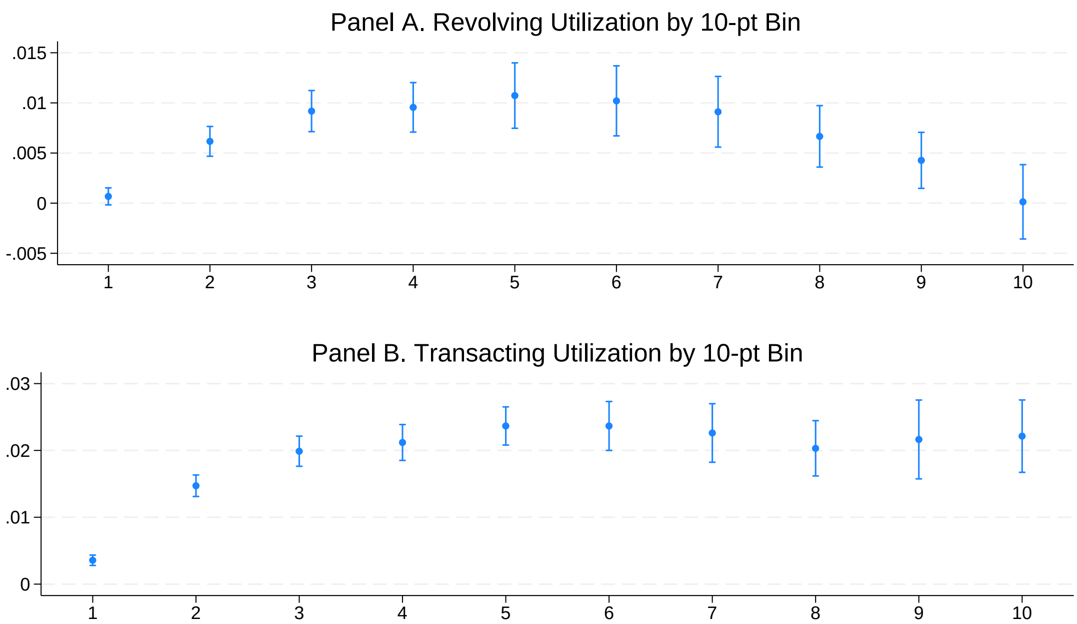

Figure 10: Utilization and the likelihood of a credit limit increase

Figure 10 plots the regression coefficients from 10-point bins of revolving utilization and transacting utilization. The bins include the upper bound but not the lower bound so the first bin includes accounts with revolving utilization $\in (0,0.1]$. We omit the coefficient on cards with zero utilization—which comprise about half of the observations for revolving utilization and a quarter of the observation for transacting utilization—so the coefficients should be interpreted as relative to an account with zero utilization. Panel A shows that the coefficients on the bins of revolving utilization follow an upside down U-shaped curve. Accounts with very high revolving utilization above 0.9 are just as likely likely to receive a limit increase as an account with zero utilization, but all other bins are more likely to, with the probability of an increase rising with utilization until the 5th bin. Panel B repeats the analysis for transacting utilization and shows a pattern similar to that of a logistic growth curve. The probability of receiving a limit increase is lowest for accounts with transacting utilization below 0.1, but rises monotonically until the 4th bin. For transacting utilization above 0.4, the likelihood of receiving a limit increase does not vary with utilization.

Although the coefficients we estimate are small—due to the sparseness of limit increases in the data—they are economically meaningful. In Figure 11, we compare the coefficient on the (0.2,0.3] bin of revolving utilization (the dashed line) to the effect of a recent change in credit score on the likelihood of a limit increase (the solid line).[18] The figure suggests that being in the (0.2,0.3] bin of revolving utilization has approximately the same positive effect on the likelihood of a limit increase as a 66-point increase in credit score. Since Figure 10 suggests that the effect from being in bins corresponding to revolving utilization [0.3-0.7) is slightly larger, these bins would be equivalent to a slightly larger increase in credit score.

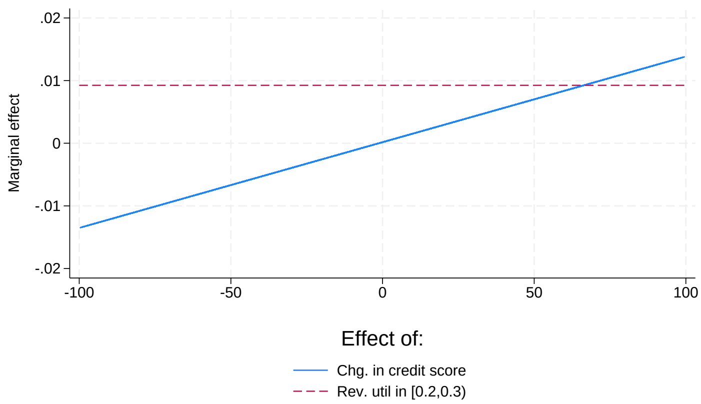

Figure 11: Marginal effect on the likelihood of a limit increase

This figure compares the effect of a change in credit score in the previous 3 months (solid line) to the coefficient on the (0.2,0.3] bin of revolving utilization (the dashed line)

[18] All of our regressions include an indicator for whether the borrower's credit score changed in the previous 3 months and the size of that change. The solid line calculates the marginal effect from a credit score change using these coefficients.

Consumer-initiated increases

Next, we examine whether the drivers of consumer-initiated increases differ from those of bank-initiated increases. Figure 12 shows that for transacting utilization, the patterns across the bins of utilization are similar for bank-and consumer-initiated increases. Revolving utilization, however, exhibits a much flatter upside down U-shaped pattern than in Figure 10, with most of the coefficients not statistically different from zero. In addition, the coefficients on deciles 7-10 are negative. The magnitudes of the coefficients on revolving utilization are also much smaller for consumer-initiated increases than bank-initiated increases, suggesting that for consumer-initiated increases, lenders do not weight revolving utilization as highly in the approval process. We discuss our findings on consumer-initiated increases further in Section A.4 of the Appendix.

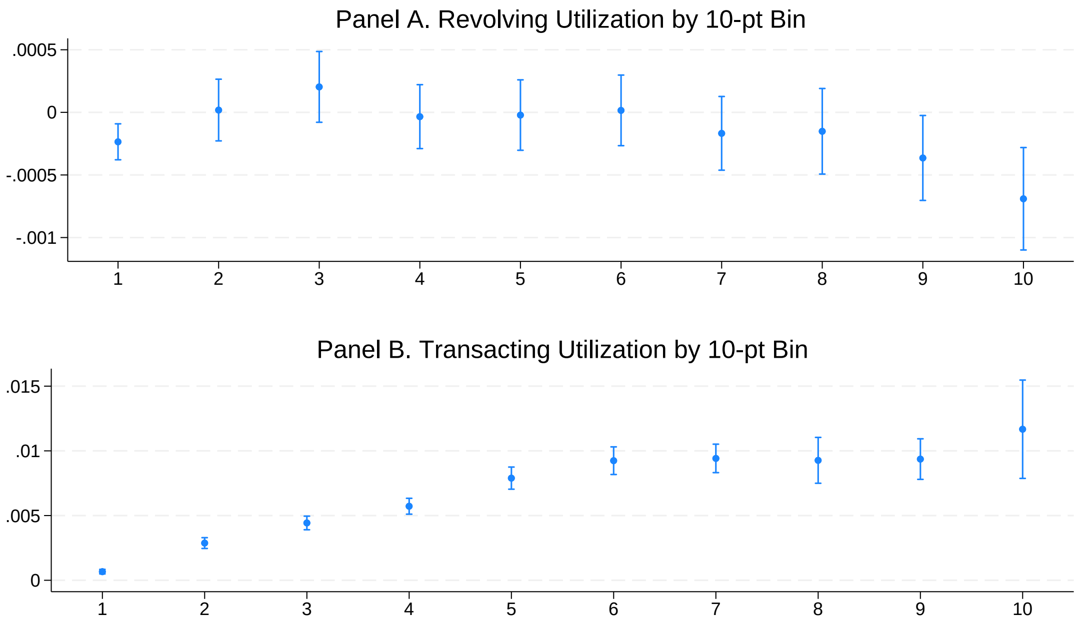

Figure 12: Utilization and consumer-initiated limit increases

3.2 What happens after limit increases?

Having examined the relationship between account utilization and limit increases, we next perform a simple event study to examine how revolving utilization and debt change following a limit increase. Our findings confirm the results established by prior literature that consumers generally borrow more after a limit increase (Aydin, 2022; Gross and Souleles, 2002). In addition, the following discussion should not be interpreted causally; it is meant simply to illustrate borrowers' behavioral responses to limit increases and provide a potential reason for why lenders target revolvers for increases.

Figure 13 presents the results of the event study, using the same sample as in section

3.1. The dependent variable is revolving utilization and we include all accounts, so the coefficients should be interpreted as relative to a credit card account that is not within 12 months of a limit increase.[19] The panel on the left estimates the effect of a bank-initiated increase, and we remove from the sample any accounts within 12 months of a consumer-initiated increase. The panel on the right estimates the effect of a consumer-initiated increase.

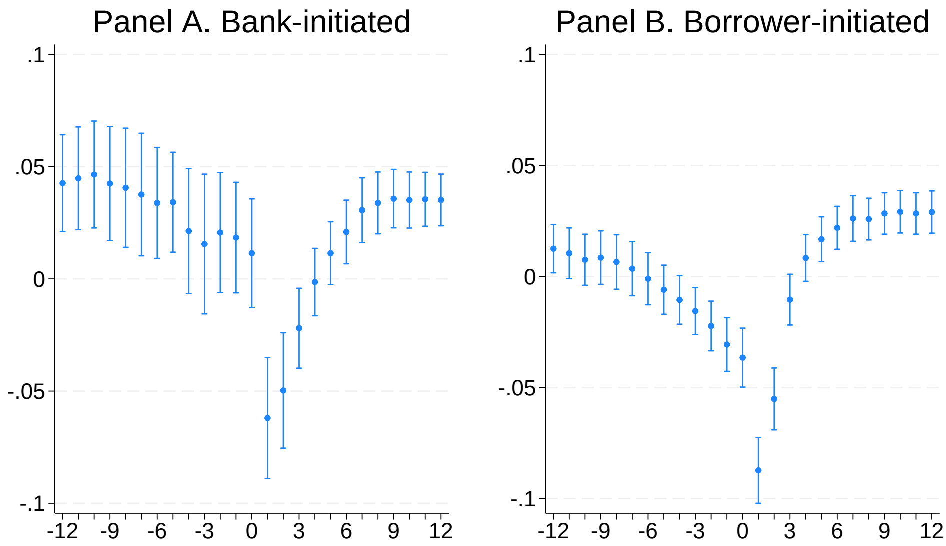

Figure 13: Event study: revolving utilization around limit increases

Three points stand out from figure 13. First, consistent with the results of section 3.1, revolving utilization is higher at accounts before they experience a limit increase than at accounts that are not within 12 months of a limit increase. Moreover, revolving utilization is particularly higher prior to a bank-initiated increase, consistent with banks targeting revolving borrowers for increases. Second, revolving utilization drops sharply when the limit increase takes effect—this is mechanical since utilization is calculated as a share of the credit limit. In addition, utilization experiences a downward trend even before the limit increase, especially for consumer-initiated increases. Third, revolving utilization rebounds in the months after the limit increase. For bank-initiated increases, revolving utilization returns to its previous level by 8 months after the increase. For consumer-initiated increases, revolving utilization rises above its pre-increase level.

Overall, figure 13 suggests that after a limit increase, revolving balances at the average credit card account increase about 40%, with 30% of the limit increase going towards revolving balances.[20] These results are somewhat larger than Gross and Souleles (2002),

[19]Using only "treated accounts" that receive a limit increase at some point gives qualitatively similar results. We exclude observations that are within twelve months of a limit decrease. We include month, state, and card-group fixed effects; including all the controls from section 3.1 does not change the results.

[20]These statistics are calculated using summary statistics from section A.2: average revolving utiliza-

who find that total debt increases by about 13% of the limit increase, and Aydin (2022), who finds that total debt grows by about 18%. Our larger coefficient likely occurs for several reasons. First, our findings are not causal —in so far as lenders target borrowers who are likely to increase their debt, this would lead to higher coefficients. Second, utilization is generally higher in our sample than in the Gross and Souleles (2002) sample, which would lead to larger debt response even if consumers have the same utilization targeting behavior. Third, we focus on revolving utilization, rather than total utilization. In Appendix A.6, we show that the trends in total utilization are similar to revolving utilization but total utilization does not rebound as much, suggesting that total debt increases less, in percentage terms, than revolving utilization.

Importantly, we find that the rebound in revolving utilization in figure 13 is not driven by constrained consumers. In Appendix A.6, we show that limiting to consumers who have a revolving utilization below the mean of 0.28 immediately prior to the increase yields very similar results.

Relying on these empirical findings and the stylized facts in the previous section, we next build a quantitative model of household behavior in the credit card market which reflects the following facts established in this section. First, as facts 1 and 2 showed that limit increases are an important source of credit, in our model, we allow borrowers' credit limits to increase over time. While in the stylized facts, we examined differences by both credit score and income, for tractability, in the model we only focus on borrower income since we find that heterogeneity by income is similar to that by credit score. Second, since facts 3 and 7 established that most limit increases are bank-initiated, we only consider bank-initiated increases in the model and abstract away from consumer-initiated increases and from decreases. Third, fact 6 and the empirics in section 3.1 established that banks target revolvers for increases, likely because these increases lead to higher revolving balances as shown by facts 4 and 5 and the event study in section 3.2. We incorporate these findings into the model by allowing lenders to target borrowers based on their revolving utilization, among other factors. We next turn to a discussion of this model.

4 A life-cycle model with credit card debt

As shown in our empirical analysis, credit card companies frequently and proactively raise credit limits based on observable household behavior. While we do not know banks' exact algorithms for raising credit limits, we observe in our empirical analysis that banks often do so based on the consumer's revolving borrowing. In this section, we develop a novel quantitative model of household behavior that allows us to analyze the positive

tion just prior to limit increase is 28% and the average limit rises 40% after a bank-initiated increase.

and normative implications of banks' approach to proactively raising credit limits, and to assess the costs and benefits of counterfactual polices that may limit banks' ability to raise limits using their preferred algorithm.

We begin with a life-cycle consumption-saving model with credit card borrowing, potential defaults, and preference heterogeneity, in the spirit of Nakajima (2017) and Fulford and Schuh (2024). We include two types of households with heterogeneous preferences, where some share of households have self-control issues and the remaining share do not, similar to Nakajima (2017). The theoretical literature has long predicted that agents with self-control problems should exhibit higher credit card borrowing (Laibson, 1997, Angeletos et al., 2001, Laibson et al., 2017), and empirical evidence strongly supports this prediction. Meier and Sprenger (2010) find that individuals with self-control issues are significantly more likely to have credit card debt. Gathergood (2012) shows that self-control issues are positively correlated with over-indebtedness in the credit card market, and that self-control issues play a stronger role than financial illiteracy. Kuchler and Pagel (2021) document that revolving borrowers exhibit self-control problems by systematically failing to follow through on their own debt paydown plans. Given this empirical evidence, we believe that self-control issues are of first-order importance when thinking about credit card regulation.[21]

This empirically-motivated modeling choice introduces an important economic tradeoff between giving households too much versus too little credit. In a model without self-control issues, greater access to credit is almost always beneficial, as it makes markets more complete and helps households smooth consumption over adverse shocks and over the life-cycle. This standard economic intuition fails to explain the growing prevalence of policies restricting credit limit increases that we observe worldwide. In contrast, in a model with self-control issues, too much credit may be detrimental if it leads some households to over-borrow and over-consume. Our model nests both possibilities, allowing us to quantitatively assess the welfare implications of credit restriction policies that would appear puzzling in standard frameworks without behavioral frictions.

In addition to including heterogeneous preferences, we also extend our model to capture the fact that banks often give low credit limits at origination, which they then proactively increase based on borrower behavior.[22] In our model, we adopt a reduced-form approach to giving bank-initiated credit limit increases to consumers, guided by our empirical analysis. While this abstracts from the richer details of banks' confidential data and learning algorithms, it gives us a rough approximation of how banks give

[21] While other behavioral biases may also be relevant, they are outside the scope of the current paper.

[22] In contrast, most models of credit card lending assume that the credit limit is fixed and/or exogenous. One notable exception is Fulford and Schuh (2024), who model a credit limit that grows exogenously based on age, income, and default. We believe our model is the first to allow the credit limit to increase in response to observed credit utilization.

limit increases, and captures banks' revealed preference of giving more limit increases to revolving borrowers.[23] In our counterfactual analysis, we then restrict the banks' ability to: a) give credit limit increases to revolving borrowers, and b) give credit limit increases without borrowers' consent.

As mentioned above, our model captures a critical trade-off between extending additional credit to individuals who benefit from relaxed credit constraints versus protecting individuals with self-control issue who may be tempted to increase borrowing in response to greater access to credit. From the perspective of firms, however, giving credit to individuals with self-control issues may be highly profitable if these consumers are willing to carry persistent, high-cost revolving balances. To the best of our knowledge, our model is the first to be used to evaluate the pros and cons of allowing banks to increase credit based on observed borrower behavior and to assess potential alternative regulations.

4.1 Modeling Framework

Demographics. The economy is populated by overlapping generations of households. Each generation consists of an equal number of households that are born at age 1 and live up to age T, working during the initial W years of their lives before mandatory retirement. The total measure of households remains constant over time.

Preferences. The model features two types of households with distinct preferences that coexist: a proportion ϕ_{standard} of households with standard preferences, and ϕ_{tempted} with temptation preferences ($\phi_{\text{standard}} + \phi_{\text{tempted}} = 1$). Households maximize the sum of their expected discounted lifetime utility, subject to a budget constraint we define later:

$$\max \mathbb{E}_t \sum_{t=0}^{T} \beta_j^t U_t \tag{2}$$

In this equation, β_j represents the discount factor for households of type j, which is either standard or tempted (i.e., $j \in \{\text{standard}, \text{tempted}\}$). Here, $\mathbb{E}_t$ denotes the expectation operator at time t, and U_t is the utility function in period t. For standard households, the utility function represents the felicity from consuming nondurables:

$$U_t = u(c_t) \tag{3}$$

where c_t is the chosen level of nondurable consumption and the felicity function $u(\cdot)$ is concave and increasing in c_t. For tempted households, we use a modified utility function

[23] In reality, banks' approach to proactive credit limit increases is proprietary, and we do not observe the exact algorithms they use to determine their optimal policy for proactively raising the credit limits of some consumers and not others. For this reason, we approximate bank's policy function for credit limit increases using their observed decisions to give limit increases recorded in the Y-14M data.

that represents temptation preferences developed by Gul and Pesendorfer (2001, 2004):

$$U_t = u(c_t) - \lambda\Big[u(\tilde{c}_t) - u(c_t)\Big] \tag{4}$$

Here $\tilde{c}_t$ represents the most tempting consumption alternative in period t, defined as:

$$\tilde{c}_t = \arg\max_{c_t \in \mathcal{B}_t} u(c_t), \tag{5}$$

where $\mathcal{B}_t$ represents the current period budget set.

The key feature of temptation preferences is that utility $U()$ depends not only on actual consumption, but also the most tempting alternative available in the choice set. The term in square brackets represents the utility cost of temptation – the difference between the felicity one could get focusing only on the present, $u(\tilde{c}_t)$, and the felicity actually enjoyed given chosen consumption, $u(c_t)$. λ captures the degree of temptation. When $\lambda = 0$ the model reduces to standard preferences, eliminating the effects of temptation.[24]

Temptation preferences are not the only framework that gives rise to self-control problems and present biases. Another prominent strand of the literature focuses on models with hyperbolic discounting, as developed by Strotz (1956), Phelps and Pollak (1968), Laibson (1997), Harris and Laibson (2001), and Angeletos et al. (2001). We believe that temptation preferences offer two main advantages in our setting. First, unlike hyperbolic discounting, temptation preferences are time-consistent, facilitating model-consistent welfare analysis. Second, they model self-control problems differently: individuals are influenced by tempting alternatives regardless of their final choices, whereas hyperbolic discounting only considers chosen consumption. This distinction is empirically relevant, as Toussaert (2018) finds that many individuals are willing to pay to remove the tempting alternative even when they do not ultimately succumb to temptation, consistent with temptation but not hyperbolic discounting.

Assets. Households can save or borrow through a fully liquid financial asset, a_t. When saving (positive a_t), households earn a fixed return r^a per period. When borrowing (negative a_t), households can access credit card debt at a constant interest rate r^c, which exceeds the return on savings ($r^c > r^a$). The model incorporates heterogeneous credit limits across k distinct levels, allowing us to evaluate policy interventions related to household-specific credit limit increases. Initially, all households are assigned to the lowest credit limit, with the potential to receive credit limit increases through a stochastic process, which we describe next.

Credit Limit Increases. Motivated by our empirical findings that most limit increases

[24] In our baseline model, we assume that households are fully aware of temptation and thus perfectly sophisticated. Section 5.4 examines robustness when some share of tempted households are naïve.

are bank-initiated, we focus exclusively on bank-initiated credit limit increases in our model, abstracting from consumer-initiated requests for computational tractability. In the baseline model, we adopt a reduced-form approach to modeling the bank's decision to extend a credit limit increase (CLI) to a given household. This approach enables us to flexibly capture the bank's credit decisions using the Y-14M data, while remaining agnostic about the specific algorithms used by the bank in making these decisions. Later, in Section 5.5, we allow banks to solve their own profit maximization problem and confirm the robustness of our results. For now, we adopt the following latent-variable formulation:

$$CLI^* = b_0 + \boldsymbol{x'b} + \varepsilon, \quad \varepsilon \mid \boldsymbol{x} \sim U(-a, a). \tag{6}$$

Here, $\boldsymbol{x}$ represents observable characteristics that influence the probability of receiving a credit limit increase (e.g., income, credit utilization, etc.) The binary credit limit increase (CLI) indicator is then defined as:

$$CLI = \begin{cases} 1 & \text{if } CLI^* > 0 \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

This formulation is convenient as it aligns with the linear probability model used in Section 3.1; the straightforward mapping of regression coefficients allows us to calibrate the coefficients of this function using actual data.[25] This approach delivers a close mapping between revolving utilization and credit limit increases, ensuring that our calibrated model closely matches the data along this important dimension (see figure 14).

Bankruptcy. The model adopts a bankruptcy procedure similar to Chatterjee et al. (2007) and Nakajima (2017), resembling a Chapter 7 bankruptcy filing. When a household defaults on credit card debt, a specific sequence of financial consequences unfolds. The household first incurs a fixed default filing cost of ξ. At this point, the existing debt is entirely eliminated, with no future repayment obligations. During the default period, the household is prohibited from saving, as any potential savings would be immediately garnished by the lender. A proportion η of current labor income is garnished by the lender. Simultaneously, the household's credit history transitions from good ($h = 0$) to bad ($h = 1$). While in bad credit standing ($h = 1$), the household is completely excluded from the credit market, with their borrowing limit reduced to zero. Ultimately, with probability π_h^0, the bad credit history is expunged, returning h to 0 and removing any further negative consequences of the past default.

Income. Households receive an exogenous income stream with explicitly modeled employment and unemployment risk, similar to O'Dea (2018). We assume that employment

[25]The empirical coefficients (β_0, $\boldsymbol{\beta}$) can be mapped to the latent variable coefficients (b_0,$\boldsymbol{b}$) as follows: $b_0 = 2a\beta_0 - a$ and $\boldsymbol{b} = 2a\boldsymbol{\beta}$, when using the same explanatory variables.

($e_{i,t}$) evolves through a first-order Markov process where the probability of employment π_e depends on the previous period's employment status, detailed in Appendix B.1.4. For employed households, log income comprises a deterministic and a stochastic component:

$$\ln y_{i,t} = g_t + z_{i,t}, \qquad g_t = d_0 + d_1 t + d_2 t^2 + d_3 t^3 \tag{8}$$

where g_t represents a deterministic age profile approximated by a third-order age-polynomial and $z_{i,t}$ captures an idiosyncratic shock to log income from an AR(1) Markov process:

$$z_{i,t} = \rho z_{i,t-1} + \varepsilon_{i,t}, \qquad \varepsilon_{i,t} \sim N(0, \sigma_\varepsilon^2), \qquad \varepsilon_{i,0} \sim N(0, \sigma_0^2)$$

The parameter ρ captures income shock persistence, with $\varepsilon_{i,t}$ represents independent and identically distributed shocks across households. The initial variance of income innovations, $\varepsilon_{i,0}$, might differ from subsequent periods to account for initial income heterogeneity. Unemployment results in lower income: when unemployed, households receive unemployment benefits, $\ln y_{i,t} = \ln b$. During the working life, household income is subject to a progressive tax schedule approximating the U.S. tax code. This results in an after-tax household income given by

$$\widetilde{y}_{i,t} = y_{i,t} - \tau(y_{i,t})$$

After retirement, which we assume to happen deterministically at age W, households receive two sources of income: progressive social security income and annuitized disbursements from a mandatory retirement account. Details on taxation and retirement are in Appendices B.1.1-B.1.3.

Households' problem can be formulated recursively and solved using numerical techniques. The problem's state variables are (a_t, z_t, u_t, h_t), where a_t represents asset position, z_t is the persistent labor income shock, u_t is unemployment status and h_t is credit history. In Appendix B.2 we provide a detailed characterization of the problem.

Information Structure. We assume that the credit card company perfectly observes household income, age, and borrowing decisions, but does not observe household type. We assume that the household has a full understanding of the economic environment in which they operate, apart from knowing the firm's policy function for credit limit increases.[26] This latter assumption is helpful for computational simplicity, as it allows us to treat credit limit increases as unexpected from the perspective of the customer.

Firm Profits. We consider a representative credit card company that extends loans to consumers at an interest rate r^c while facing a capital cost of r^a. The firm's profit is

[26]Indeed, it has taken us considerable time and research to begin to understand the firm's policy function for credit limit increases.

defined as the sum of interest expenditure on revolving balances minus provisions in the case of defaults, for each household i and time period t.

$$\pi = \sum_{i,t} \Big((r^c - r^a)(a^d_{i,t})(1 - \mathcal{I}^{\text{def}}) - ((1 + r^a)(a^d_{i,t}) - \eta \widetilde{y}_{i,t})\mathcal{I}^{\text{def}} \Big) \tag{9}$$

Where $\mathcal{I}^{\text{def}}$ is an indicator function denoting default status (1 for default, 0 otherwise), and $a^d_{i,t} = \max\{-a_t, 0.0\}$ represents the household's revolving balance.

When household i does not default in period t, the firm generates interest of r^c while incurring a capital cost of r^a on the outstanding balance, resulting in a net revenue of $(r^c - r^a)a^d_{i,t}$. Conversely, when household i defaults, the firm sustains a loss of $(1+r^a)a^d_{i,t}$, representing the revolving balance adjusted for capital costs, which is partially mitigated by garnishable income $\eta\widetilde{y}_{i,t}$, similar to Nakajima (2017).

In our model, the credit card company makes one simple decision: whether to adjust credit limits for different households.[27] For computational tractability while estimating the preference parameters in the baseline model, we assume this decision follows a predetermined policy rule based on household-specific characteristics such as income or credit utilization rate. In determining this rule, we adopt the latent variable procedure previously outlined in Section 3.1 and further discussed in Section 4.3. Later, in Section 5.5, we assess the robustness of our results by allowing the firm to optimize this policy rule by choosing the coefficients in equation (6) to maximize their profits in equation (9).

4.2 Model intuition via the Euler equations

To better understand the basic mechanisms driving consumption, saving, and borrowing choices in our model, it is useful to examine the optimality conditions for different household types as represented by their respective Euler equations. These Euler equations hold when households are able to freely adjust their liquid assets and are at the interior of their choice set, i.e. they are not at their credit limit or an interest rate notch.

Households with standard preferences who have access to a liquid asset yielding return r^a smooth consumption according to the following Euler equation:

$$u'(c_t) = (1 + r^a)\beta\mathbb{E}_t\Big[u'(c_{t+1})\Big] \tag{10}$$

This standard Euler equation shows that the marginal cost of giving up one unit of current consumption must equal the marginal benefit of consuming the proceeds of the

[27]We allow the lender to change the credit limit as they learn more about the borrower, but leave the interest rate fixed, consistent with the 2009 CARD Act, which limited credit card lenders' abilities to adjust interest rates based on new information about borrowers (see e.g. Nelson, 2025). In practice, lenders rarely adjust interest rate spreads (Grodzicki, 2020). Similarly, we do not allow for credit limit decreases, consistent with how rarely they occur (Figure 9).

extra liquid saving in the next period.[28]

For households with temptation preferences, we derive the Euler equation following Kovacs et al. (2021), who show that temptation generates testable implications for consumption growth through its effects on the modified Euler equation. In particular, tempted consumers abide by the following modified Euler equation in our model:

$$u'(c_t) = (1 + r^a)\beta \mathbb{E}_t \Big[u'(c_{t+1}) - \frac{\lambda}{1+\lambda} u'(\tilde{c}_{t+1}) \Big] \tag{11}$$

The modified Euler equation shows that for tempted households, the marginal cost of giving up one unit of current consumption must equal the marginal benefit of consuming the proceeds of the extra liquid saving in the next period, minus the marginal cost of resisting additional temptation next period.

This modified condition create a fundamental bias towards present consumption. Temptation incentivizes consumption to be tilted toward the present, making households poorer in the future but sparing them the disutility of resisting the temptation to spend all their resources—temptation they would face if they saved more for the future.

The strength of self-control issues depends on two key factors captured in the Euler equation. First, it depends on the importance of temptation itself, measured by $\frac{\lambda}{1+\lambda}$. Second, it depends on the marginal utility of consuming tempting resources $u'(\tilde{c}_{t+1})$. The incentive to shift resources from the future to the present (increasing current consumption) is stronger when either the temptation parameter is high or when the marginal utility from consuming additional tempting resources in the future is high (i.e., when the household expects to have fewer resources available for temptation in the future).

These contrasting Euler equations have important implications for how different types of households respond to changes in credit card limits. For standard households, equation (10) implies that increased access to credit is welfare-enhancing, as it allows better consumption smoothing across periods and states of the world—the additional borrowing capacity helps households optimize their intertemporal consumption plans without distortion. However, for tempted households, equation (11) reveals a different side of credit access: higher credit limits increase the tempting consumption opportunities in future periods (higher $\tilde{c}_{t+1}$), which through the term $\frac{\lambda}{1+\lambda} u'(\tilde{c}_{t+1})$ creates additional incentives to consume today rather than save for tomorrow. This mechanism can lead tempted households to over-borrow and over-consume when credit limits are raised, potentially making them worse off despite having access to more credit.

[28] When a consumer borrows, the Euler equation is unchanged, except that we replace r^a with r^c due to the interest rate wedge between saving and borrowing. For simplicity we denote the time preference parameter as β, which we later estimate differently for standard and tempted consumers.

4.3 Calibration

We set the institutional parameters to reflect the economic environment faced by U.S. households. Parameters set outside the model are detailed in Table 10 in the Appendix.

Demographics. Decisions in the model take place at an annual frequency. All households enter the labor market at age 25 (t = 1), retire at age 65 (W = 41), and die at age 80 (T = 56).

Preferences. The felicity function follows a constant relative risk aversion (CRRA) specification:

$$u(c_t) = \frac{c_t^{1-\gamma}}{1-\gamma} - \mathcal{I}^{\text{def}}\kappa \tag{12}$$

Where γ represents the coefficient of relative risk aversion. We assume that when households default, they incur both a financial and psychological cost. This psychological effect is captured by κ in equation (12), and $\mathcal{I}^{\text{def}}$ is an indicator function that equals 1 when default occurs. The parameters γ and κ are internally calibrated and shown in Table 1.

The temptation parameter λ, specific to households with temptation preferences, is set to the value $\lambda = 0.28$ estimated by Kovacs et al. (2021) using consumption data from the U.S. Consumer Expenditure Survey. This estimate aligns with the existing literature, similar to the value reported by Kovacs and Moran (2025), and close to the 0.23 value found by Huang et al. (2015).

Assets. The return on savings is set to 2% and the credit card borrowing rate is 14%, following Fulford (2024). Households differ in their credit card borrowing capacity, with credit limits varying between 0 and a maximum credit limit (CL_{max}). We discretize the potential credit card limit range from zero to CL_{max} using eight equally spaced grid points ($k = 8$). It is important to note that households affected by temptation are also influenced by their credit card limits. The credit card limit fundamentally affects the budget set, which in turn determines the most tempting consumption alternatives, as specified in equation (5).

Credit Limit Increases. Initially, each household h is assigned to the lowest credit limit, with the potential to receive credit limit increases over time. The credit card limit of each household evolves stochastically, driven by three factors: current credit card limit (CCLimit), credit card utilization (CCUtil), and household income (Income). As a result, the latent varible formulation represented by equations (6)-(7) in the model becomes:

$$CLI^*_{h,t} = b_0 + b_1 \log(CCLimit_{h,t}) + b_2 CCUtil_{h,t} + b_3 Income_{h,t} + \varepsilon_{h,t}. \tag{13}$$

$$CLI = \begin{cases} 1 & \text{if } CLI^* > 0 \\ 0 & \text{otherwise} \end{cases} \tag{14}$$

This specification is consistent with the empirical equation (1) in Section 3 where the dependent variable is an indicator for whether a credit card receives a limit increase, with controls log of the credit card limit and income. The exact coefficients used in the firm's policy function are presented in Table 11.

The procedure for determining credit limit increases for each household h and time period t in the model follows these steps:

1. For each household h and time period t, randomly draw a shock ($\varepsilon_{h,t}$) from a uniform distribution between (-0.5, +0.5).
2. Use the model's simulated values of of credit card limit, credit card utilization and income ($CCLimit_{h,t}, CCUtil_{h,t}, Income_{h,t}$) to create log credit card limit, utilization and income bins that align with the empirical regression counterpart.
3. Map the empirically estimated coefficients of the explanatory variables (β_0, $\boldsymbol{\beta}$ from Table 11) to the latent variable coefficients (b_0, $\boldsymbol{b}$) as shown in footnote 25.
4. Calculate $CLI^*_{h,t}$ using equation (13).
5. Determine whether credit limit increase occurs for each household h and time period t using equation (14).

To demonstrate that the latent variable algorithm for credit limit increases in the model approximates the empirical credit limit increases well, we plot the probability of a credit limit increase by credit card utilization bin in Figure 14 for both the model and the data. Orange dots represent empirical probabilities, while the dark blue line represents probabilities from the model's latent variable method. As shown, the data and model align closely, both exhibiting an inverted U-shaped relationship with card utilization.[29]

As noted previously, credit limit increases crucially affect households' budget sets. For standard households, credit limit increases have an unambiguously positive effect by enabling them to better smooth consumption. In contrast, tempted households face a trade-off: while expanded access to credit facilitates consumption smoothing, it also increases temptation to overborrow and overconsume, as discussed in Section 4.2.

Bankruptcy. The default process is characterized by three key parameters, following Nakajima (2017). We set the filing cost ξ to \$1117 (at 2015 prices), which reflects the

[29] Note that the figure presented here closely resembles Figure 10 in the Empirical Section. The primary difference is that the empirical regression allows for the inclusion of many more controls related to household characteristics than the simulated model does.

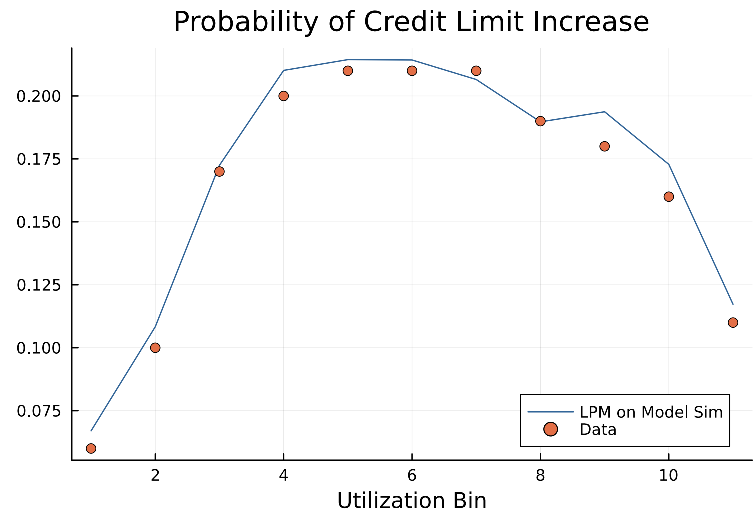

Figure 14: Credit limit increase by utilization

average cost of filing for Chapter 7 bankruptcy. The probability π_h^0 of wiping out a household's bad credit history is calibrated at 0.1, ensuring an average of 10 years before returning to the credit market.[30] The parameter η representing garnished labor income is calibrated within the model and shown below in Table 1.

Income. To establish income-related parameters, we adopt the approach of Kovacs and Moran (2025), who estimate the earnings process using a two-step minimum distance method. For the parameters of the non-linear tax function, we use the estimation results by Keane and Wasi (2016) and convert them to 2015 units. The results of these estimations can be found in Table B.3 in the Appendix.

Internally calibrated parameters. Using the Method of Simulated Moments, we internally calibrate the model parameters that are listed in Table 1 to match aggregate statistics of household behavior, as shown in Table 2. In particular, we target (1) the share of households with revolving credit card debt, (2) the average utilization rate of credit cards, (3) the default probability, (4) and the average debt to income ratio.

Parameter	Symbol	Value
Proportion of standard households	ϕ_{standard}	0.60
Discount factor for standard household	β_{standard}	0.92
Discount factor for tempted household	β_{tempted}	0.96
Risk aversion parameter	γ	0.72
Income garnishment ratio	η	0.30
Utility cost of default	κ	1.76

Table 1: Parameters set within the model

[30] The geometric distribution calculates the average number of periods before exiting the bad credit state (including the successful period) as: $(1-\pi_h^0)/\pi_h^0 + 1$ with this value of 10 implying $\pi_h^0 = 0.1$.

Table 1 reports the estimated values for the calibrated parameters. The results indicate a heterogeneous population: standard households make up roughly 60% ($\phi_{standard}$) of the population, while tempted households account for the remaining 40%. The share of households with self-control issues is consistent with empirical evidence, which finds that a substantial fraction of individuals are present-biased, with estimates ranging from 25-60% across studies (e.g. Ashraf et al., 2006; Meier and Sprenger, 2010; Toussaert, 2018; Balakrishnan et al., 2020; Schneider and Moran, 2024). Further, our sensitivity analysis confirms that $\phi_{standard}$ is pinned down by all four targeted moments, and demonstrates how changes in the targeted moments affect the parameter value (Appendix B.4).

We estimate the time preference parameter β separately for standard and tempted households. The time preference parameter is 0.92 for standard households ($\beta_{standard}$) and 0.98 for tempted households ($\beta_{tempted}$). The higher estimate for tempted households aligns with previous research on self-control problems, as documented by Laibson et al. (2017) and Ganong and Noel (2019).

The rest of the parameters are estimated on the whole population. The coefficient of relative risk aversion γ is 0.65, which is somewhat low compared to studies focusing on standard households. For context, Kovacs et al. (2021) found $\gamma = 0.83$ when estimating a model with tempted households. The income garnishment ratio η is estimated at 0.36, slightly higher than the 0.34 reported by Nakajima (2017). Finally, we estimate the psychological (utility) cost of a default κ to be 2.6.

5 Model results

This section presents the main results from our calibrated model. Section 5.1 outlines the properties of the baseline model, while Sections 5.2 and 5.3 analyze the implications of two counterfactual policies: (i) restricting banks from increasing credit limits for borrowers with revolving debt, and (ii) requiring consumer consent for all limit increases. Sections 5.4 and 5.5 test the robustness of our results by examining scenarios with naïve tempted households and with optimizing firms, respectively.

5.1 Properties of the baseline model

Table 2 presents aggregate statistics from U.S. data (the first column) compared to the baseline model (second column.) In the data, we see that 45 percent of individuals have credit card debt, based on data from the Survey of Consumer Finances and reported by Nakajima (2017). Conditional on having credit card debt, the average utilization rate is around 35 percent, based on data from the CFPB's Consumer Credit Panel (CCP) and reported by Fulford and Schuh (2024).[31] We target all four empirical moments in our calibration routine. Overall, the model gets a relatively good fit of the targeted moments, successfully replicating the large share of consumers with credit card debt and the other targeted moments.

	U.S.	Baseline model		
		All	Standard	Tempted
Proportion in debt	45.0	41.8	18.9	76.0
Utilization rate	35.0	28.9	4.1	73.6
Default probability	0.94	0.94	0.08	2.22
Total debt over income	8.6	6.8	1.1	15.4

Note: All statistics are in percents. The first, second, and fourth targeted moments are from Nakajima (2017). The utilization rate conditional on credit card debt is from Fulford and Schuh (2024). Averages are defined over the working age population.

Table 2: Targeted moments: baseline model vs. data

The third and fourth columns of Table 2 show aggregate statistics for agents without and with temptation. Similar to Nakajima (2017), agents with temptation are much more likely to carry revolving credit card debt and account for the vast majority of borrowing and defaults. In contrast, agents without temptation use credit card borrowing only occasionally, primarily when their income is low and they expect it to rise going forward. As a result, agents without temptation have a much lower utilization rate and debt-to-income ratio than their tempted counterparts. Similarly, agents with temptation have higher credit card utilization rates and debt-to-income ratios, as well as higher probabilities of default.[32]

Figure 15 presents further statistics for the agents without and with temptation in our model. Consistent with the fact that they carry larger revolving balances, we see that agents with temptation pay a substantially higher share of income on credit card interest payments. This is important given that past research shows that revolvers are the primary driver of credit card lenders' income, as interest payments on revolving debt comprise the vast majority of lenders' revenues and profits (Adams et al., 2022). Through the lens of our model, agents with temptation are thus highly attractive customers, given that they are willing to borrow at relatively high interest rates and thus devote a meaningful share of income towards credit card interest payments.

The final panel of Figure 15 shows the probability of a customer receiving a credit

[31] We do not use the utilization rate from the Y-14M since the Y-14M data are at the account-level, rather than the consumer level.

[32] The differences between agents with and without self-control issues are reminiscent of Meier and Sprenger (2010) who find that present-biased individuals are more likely to have credit card debt, Grodzicki and Koulayev (2021) who provide survey evidence that individuals who 'spent more than they wanted' due to credit cards are more than twice as likely to revolve, and Lee and Maxted (2023) who show that present bias induces households to persistently revolve debt.

Figure 15: Baseline model: behavior by type

Note: The variables in this figure are averages of model-simulated data for households aged 25-65.

limit increase based on the firm's policy function specified in equation (13). In our model, tempted agents are much more likely to receive credit limit increases, in large part due to their revolving credit utilization, documented in the previous section. This raises a number of important questions: (i) to what extent do consumers benefit from such limit increases, and (ii) to what extent should policy makers regulate such limit increases? We address both of these questions in due time.

Figure 16 shows the share of individuals with revolving credit card debt by income quintile. As is well-documented, many middle- and high-income consumers maintain revolving credit card debt (see e.g. Exler and Tertilt, 2020). Figure 16 decomposes the contribution of individuals with and without temptation to the aggregate share of revolving borrowers. For agents without temptation, the share of revolving borrowers declines monotonically with income. By contrast, for agents with temptation, the share of revolvers is relatively unaffected by income. This demonstrates the importance of temptation in our model, which plays a key role in explaining why individuals in the top of the income distribution maintain high-interest revolving balances on their credit card.

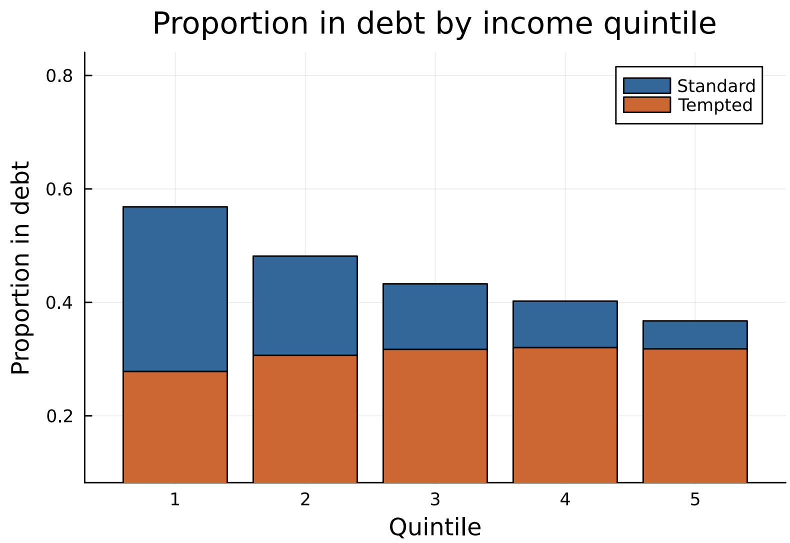

Figure 16: Share of revolvers by income quintile

5.1.1 To what extent are credit limit increases beneficial for consumers?

Figure 17 decomposes the probability of credit limit increases based on revolving utilization. The black dots come from the data and are input into the firm's policy function for credit limit increases. The blue bars show the probability of limit increases in our simulated model. As mentioned previously, the model mechanically generates credit limit increases in line with the empirical evidence. The orange bars in Figure 17 show the probability of limit increases that are beneficial from the perspective of the borrower.[33] The orange bars show that many credit limit increases are not beneficial to the borrower.

What drives the above heterogeneity in whether consumers benefit from higher credit limits? Our model features two opposing channels. On one hand, higher credit limits are beneficial as they relax credit constraints, giving individuals greater flexibility to smooth consumption over the life-cycle, adverse shocks, or unemployment. In our model, agents without temptation almost always benefit from greater credit limits. On the other hand, for agents with temptation, higher credit limits also increase the most tempting consumption alternative, $\tilde{c}_t$ in equation (5), inducing additional temptation to over-consume. In many cases, this may be detrimental for the consumer, if it leads to over-consumption and reduced saving for retirement and precautionary purposes, or if it leads to an above optimal share of income being spent on expensive credit card borrowing.

Figure 17 reveals important heterogeneity across utilization bins. While customers with zero or low utilization almost always benefit from credit limit increases, customers with higher utilization are frequently harmed by credit limit increases in our calibrated

[33] This is computed by comparing the expected value of receiving a credit limit increase relative to the expected value of maintaining your current credit limit.

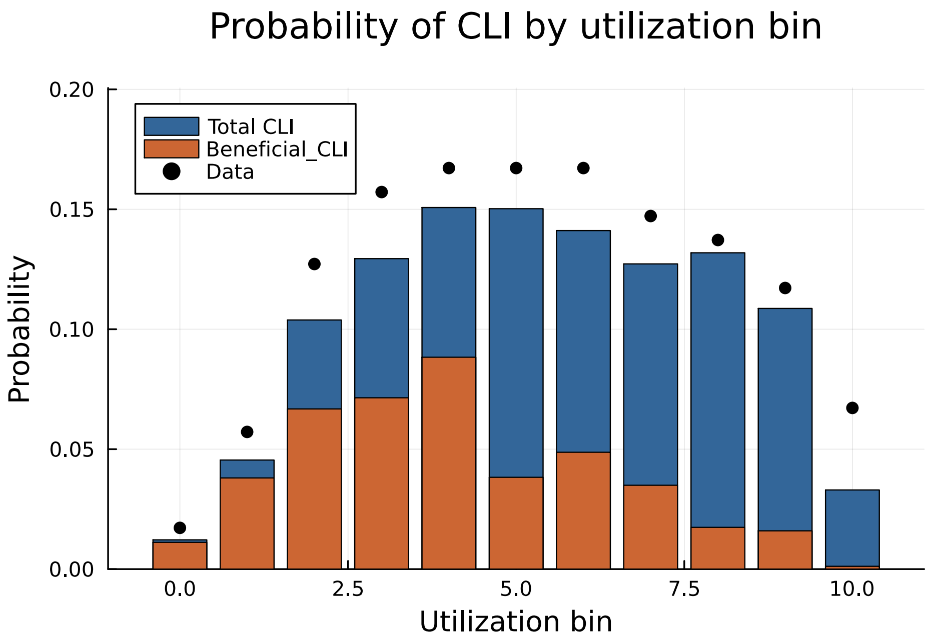

Figure 17: Probability of credit limit increases by card utilization rate

model. For customers with utilization rates above 50 percent, we see that the majority of credit limit increases are detrimental from the perspective of the borrower. This heterogeneity arises because in our model, agents with temptation generally have higher utilization ratios, as shown in Table 2. As a result, credit limit increases going to borrowers with high utilization ratios disproportionately go to agents with temptation, for whom the costs often outweigh the benefits.

5.2 Prohibiting limit increases for revolving borrowers

In this section, we assess a counterfactual policy where credit card companies are restricted from using their preferred algorithm for credit limit increases. More specifically, we assume that firms are prohibited from increasing the credit limits of consumers with revolving debt, despite firms' observed preference to give credit limit increases to these customers. This counterfactual policy analysis is inspired by a recent effort by the UK Financial Conduct Authority (FCA) to regulate bank-initiated credit limit increases.[34] The FCA has recently reached an agreement with UK credit card companies, who have agreed to a new policy where no customer will be given a credit limit increase if they have been in persistent revolving debt for a year or more.

To better understand the effects of this reform on household behavior and well-being,

[34]When discussing credit limit increases, the UK Financial Conduct Authority stated, "An important objective of our regulation of the credit card market is to ensure that consumers, especially those at risk of debt problems and those in financial difficulty, are not given unaffordable credit limit increases and have proper control over their credit limits. This is so they can avoid borrowing more than they intend to and avoid a worsening of their circumstances."

we solve and simulate our calibrated model with a counterfactual policy where banks are prohibited from offering credit limit increases to any consumer with revolving debt. More specifically, we modify equation 7 so that consumers are only given a credit limit increase if $CLI^* > 0$ and $a_t \geq 0$.

While we view this as a first approximation of the effects of the above policy, we note two important limitations. The first is that a period in our model is defined as one year, so even temporary credit card usage results in a temporary restriction against credit limit increases, which is slightly stricter than the UK policy. Second, we do not let the firm adjust the other parameters of the credit limit increase function in response to this policy. Although we examine the impact of partial re-optimization of the firm in Section 5.5, we acknowledge that the welfare results could be further attenuated in a general equilibrium framework. We believe that a fruitful avenue for future research would be to better understand how firms might strategically adapt their broader credit policies in response to these regulatory constraints.

5.2.1 Implications for consumer behavior

Figure 18 presents the effects of the above policy on consumer behavior using our calibrated model. Each bar depicts the contribution of the two household types to the key variables of interest in the baseline model (the first bar) compared to the counterfactual policy (the second bar). The visualization distinguishes between households without temptation (dark blue) and households with temptation (orange). By comparing the Baseline and the UK-style Policy scenarios, we can observe how this regulatory intervention differentially affects these two distinct groups of consumers.

As illustrated in the figure, prohibiting credit card companies from increasing credit limits for revolving borrowers has significant effects on aggregate variables. The most immediate effect is a dramatic reduction in the probability of credit limit increases, which falls from an annual rate of roughly 5.5% to around 1.0%, with substantial declines observed for both standard and tempted households.[35] Simultaneously, default probabilities decrease from about 0.9% to almost zero, predominantly driven by a substantial drop in default probabilities for tempted households. This sharp reduction suggests a crucial mechanism for constraining credit expansion among vulnerable borrowers who are (as default probabilities indicate) mostly the tempted households.

The utilization rate, on the other hand, experiences a slight increase, primarily driven by lower credit limits, which are only partially offset by reduced borrowing. Indeed, the constraint on additional credit for tempted households precipitates a decline in their debt-

[35] Note that the UK policy prohibits limit increases for revolving borrowers regardless of whether the increases are bank- or consumer-initiated. As a result, we believe that our results would be little changed if we were to extend our model to allow consumers to proactively request limit increases.

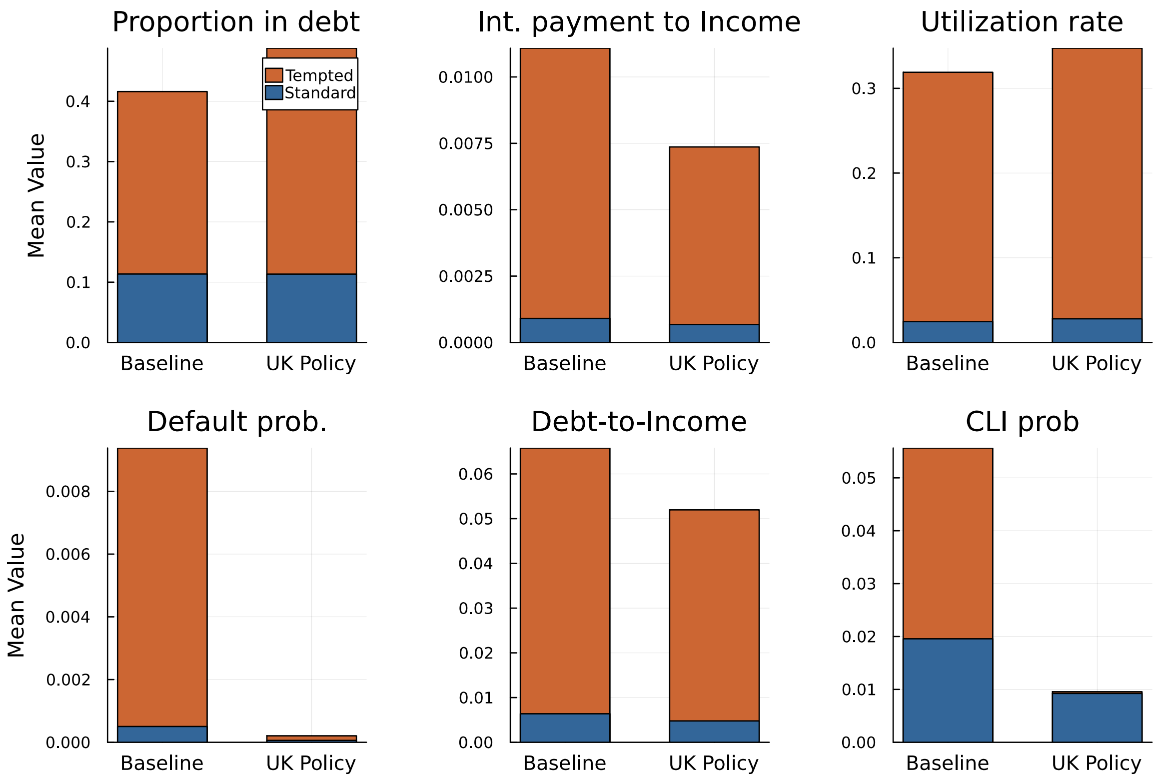

Figure 18: Counterfactual analysis: behavior by type

to-income ratio, which consequently reduces their interest payments relative to income.

5.2.2 Implications for household well-being

In this section, we evaluate the welfare implications of prohibiting credit limit increases to customers with revolving debt. As we pointed out earlier, restricting credit access to households has a trade-off: it might be detrimental for households who would use additional credit for consumption smoothing purposes, while it might be beneficial for households with self-control issues who may over-borrow without credit limit restrictions.[36] To capture these heterogeneous effects, we carry out a welfare analysis that accounts for differences across household types.

To assess the welfare of the counterfactual policy, we compute the consumption equivalent variation (CEV), i.e. the percentage change in consumption that would equate households' utility between the world that prohibits credit limit increases to customers with revolving debt and the world where credit limit increases are allowed unconditionally. We adopt a utilitarian social welfare function, where the social planner gives a weight to the utility of each type proportional to their share in the population.

[36]As John Campbell stated in his 2016 Ely Lecture, "Financial regulators face a difficult tradeoff between the benefits of regulation to households that make mistakes, and the cost of regulation to other financial market participants."

We consider two different measures of welfare. In the baseline, we assume that the social planner has the same preferences as households (thus the planner internalizes the psychological cost of temptation). In the alternative measure, we assume that the social planner ignores the psychological cost of temptation, i.e. the social planner has $\lambda = 0$, even while households still suffer from temptation. The alternative welfare criterion captures the impact of temptation on household choices while ignoring its psychic cost.

Table 3 presents the consumption equivalent variation (CEV) results, which provide insights into the welfare implications of the credit restricting policy. The overall welfare effect is positive, with a 1.12% improvement when psychological costs are considered. This aggregate figure, however, masks significant variations across household types.

	Overall Effect	**Standard**	**Tempted**
With Psychological Costs	1.12	−0.21	3.12
Without Psychological Costs	0.98	−0.05	2.54

Table 3: The effect of prohibiting limit increases for revolving borrowers (CEV, %)

Standard households, who borrow for consumption smoothing purposes only, experience a modest welfare loss of 0.21%. This negative impact stems from the policy's constraint on these households' borrowing capabilities, which limits their ability to smooth consumption. In contrast, tempted households show a remarkably different response. These households experience a substantial welfare gain of approximately 3.12%. This positive outcome suggests that limiting credit access for households prone to impulsive borrowing can actually improve their welfare (i.e., the benefit from reduced temptation outweighs the cost of constrained consumption smoothing for tempted households).

The second row of Table 3 demonstrates the welfare effect when the social planner ignores the psychological cost of temptation in the welfare evaluation. The alternative welfare measure does not fundamentally alter our main conclusions. The benefit from restricting tempted households' access to additional credit still dominates the cost arising from having less available resources for consumption smoothing purposes.[37]

We believe that our main welfare results are robust to alternative assumptions about the source of self-control issues. While we focus on temptation preferences for our quantitative exercise, we believe that hyperbolic discounting would generate qualitatively similar predictions, though with slightly smaller welfare gains. The key difference is that temptation introduces an additional welfare channel: the psychic cost of resisting temptation, which occurs even for households who choose not to borrow. By reducing this psychic

[37]The distribution of welfare changes provides additional insights into the heterogeneous effects of the policy across household types. Figure 36 in Appendix B.5 shows the distribution of welfare effects by households for each type of household. Most standard households, operating at low credit card utilization levels, remain largely unaffected by the policy. Tempted households, however, display more significant and varied responses to the credit limit restrictions.

cost, the policy improves well-being of tempted households that remain debt-free—a mechanism supported by experimental evidence that individuals pay for commitment even without succumbing to temptation (Toussaert, 2018). This is consistent with Table 3, which shows that the psychic costs of temptation amplifies the welfare benefits of prohibiting limit increases for revolving borrowers. Nevertheless, we find that this policy remains beneficial regardless of whether we account for the psychic costs of temptation.

Overall, our findings suggest that carefully designed credit limit policies could potentially enhance aggregate household welfare, particularly in populations with a substantial proportion of tempted borrowers who struggle with self-control problems in financial decision-making.

5.3 Prohibiting limit increases without borrowers' consent

In this section, we analyze a counterfactual policy that restricts credit card companies from increasing credit limits without the borrowers' consent. This analysis is motivated by recent regulatory efforts in several countries, including Canada, Singapore, and New Zealand, aimed at controlling bank-initiated credit limit increases.

To better understand the implications of this reform on household behavior, we solve and simulate our calibrated model under the assumption that banks are prohibited from offering credit limit increases without consumer consent. Specifically, we modify equation 7 to ensure that consumers receive a credit limit increase only if $CLI^* > 0$ and if the consumer's expected value function with the limit increase exceeds that without it.

5.3.1 Implications for consumer behavior

Figure 19 presents the effects of the counterfactual policy on consumer behavior using our calibrated model. Similar to Figure 18, each bar depicts the contribution of the two household types to the key variables of interest under the baseline and counterfactual policy. By comparing the Baseline and the Canada-style policy scenarios, we can observe how this regulatory intervention differentially affects these two distinct consumer groups.

As illustrated in the figure, requiring consumer consent for all limit increases produces significant impacts on aggregate variables. These results, however, are almost identical to those from the previous counterfactual shown in Figure 18, where companies are prohibited from raising credit limits for consumers with revolving debt.

The most significant difference between the two policies is their effects on the probability of credit limit increases. The probability of a credit limit increases goes from 5.5% in the baseline model to roughly 1.9% under the policy requiring consumer consent.[38] For

[38]While it is outside the scope of our model, we do not expect the above policy to result in a large increase in consumer-initiated limit increase requests, since the policy would only prevent limit increases that are not desired by consumers.

Figure 19: Counterfactual analysis: behavior by type

comparison, prohibiting limit increases for revolving borrowers reduces the probability of limit increases even further, to roughly 1.0%. This indicates that requiring consumer consent constrains credit expansion slightly less than the policy targeting revolving debt.

5.3.2 Implications for household well-being

Table 4 presents the implications for household well-being, measured in terms of consumption equivalent variation (CEV), similar to Table 3. The overall welfare effect of the policy is positive, showing a 1.16% improvement when psychological costs are taken into account. Standard households experience a slight welfare reduction of 0.19%, while tempted households benefit significantly, with a welfare gain of approximately 3.19%.

	Overall Effect	**Standard**	**Tempted**
With Psychological Costs	1.16	-0.19	3.19
Without Psychological Costs	1.04	0.00	2.59

Table 4: The effect of prohibiting limit increases without borrowers' consent (CEV, %)

Notably, these welfare results closely resemble those reported in Table 3. Thus, the policy restricting banks from offering credit limit increases without the borrower's consent yields similar welfare implications to the policy that limits credit increases for revolving borrowers, under our baseline assumption that consumers are sophisticated.

5.4 Robustness to the existence of naïves

In our baseline model, tempted households are assumed to be sophisticated, meaning they are fully aware of their self-control issues. However, empirical evidence from several studies suggests that a significant proportion of households exhibit partial or limited awareness of their future self-control problems, with estimates ranging from 50-80% across different experimental settings (see for example DellaVigna and Malmendier, 2006, Augenblick and Rabin, 2018 or Allcott et al., 2021).

For this reason, we evaluate how the presence of naïve tempted households affects our welfare results across the two different policy scenarios. Specifically, we allow a varying fraction of tempted households to be naïve, rather than sophisticated (as in the baseline scenario). We then recalculate the welfare impact of each counterfactual policy. Figure 20 shows the results from this exercise. On the left, we show how the fraction of naïve households affects the welfare impact of prohibiting credit limit increases for revolving borrowers (section 5.2). On the right, we present the same analysis for the policy prohibiting credit limit increases without the borrower's consent (section 5.3).

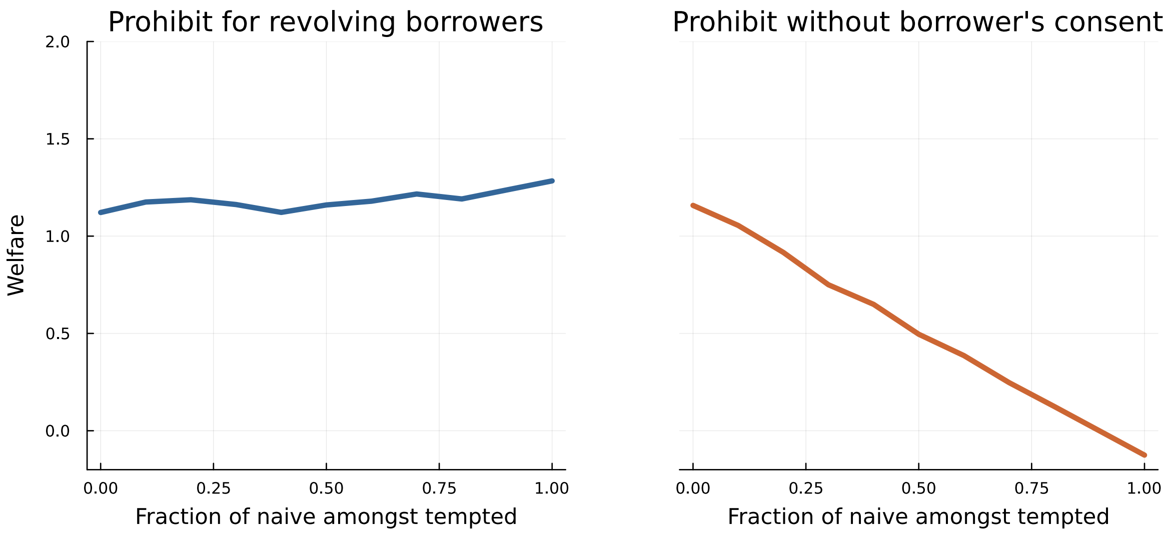

Figure 20: Impact of naïves on welfare across policy scenarios

In each subplot, the x-axis represents the fraction of naïves among tempted households (ranging from zero to one), while the y-axis measures the welfare impact of each policy in terms of consumption equivalent variation (CEV).

There is a striking contrast between how these two policies perform as the fraction of naïves changes. Prohibiting credit limit increases for revolving borrowers remains effective regardless of whether tempted households are sophisticated or naïve, as shown by the relatively flat line in the left subplot. Conversely, prohibiting credit limit increases without the borrower's consent is only welfare-enhancing when most tempted households are sophisticated. The right subplot clearly demonstrates a negative linear relationship between the fraction of naïves among tempted households and the welfare impact of this policy, with welfare gains disappearing entirely as the naive fraction approaches one.

The reason for this difference is straightforward: while the first policy constrains credit limit increases for everyone with revolving credit, the second policy only restricts credit increases for households who actively choose not to have them. Given that the second policy leaves the choice to the household, its effectiveness depends on whether households can accurately judge whether a credit limit increase is beneficial. Sophisticated households, being aware of their own self-control issues, can commit to refusing credit limit increases as a way to restrict their future borrowing. In contrast, naïve households don't recognize that they will have self-control problems in the future, and therefore more frequently opt for credit limit increases that ultimately lead to over-borrowing and over-consumption.

5.5 Robustness to firm re-optimization

In our baseline model, credit card companies follow a policy rule for giving credit limit increases, which we estimate using the Y-14M data to capture the companies' revealed-preference for giving credit limit increases to revolving borrowers, among other factors.[39] While this gives a good approximation of the economic environment currently faced by households in the U.S. and abstracts from variation across time and space, we acknowledge that any change in regulation would likely lead firms to alter their policy function for proactive credit limit increases as they re-optimize their learning algorithms in response to the new regulatory regime.

For this reason, we now assess the sensitivity of our results to allowing the firm to re-optimize its policy function for giving credit limit increases. More specifically, we now allow the firm to optimize the coefficients b_0 and $\boldsymbol{b}$ of the policy function (in equation (6)) for credit limit increase, maximizing its expected profits as defined in equation (9). To do so, we allow the firm to try out alternative policy functions and assess their profitability, similar to the large-scale experiments implemented by banks to optimize credit card profitability (Botella, 2022). In our simple implementation, the firm selects the parameters to maximize profitability using the Nelder–Mead method. And although we cannot capture the full adjustment that would likely be implemented by banks using data and algorithms that we do not observe, we view this as a first approximation.

Figure 21 shows the relationship between credit utilization and the probability of credit limit increases when the firm is allowed to select the CLI parameters optimally. Remarkably, the model delivers a credit limit increase function that has a similar shape

[39] While we do not observe the proprietary learning algorithms used by banks to give credit limit increases, nor the more detailed transaction level data that they are allowed to use, we view our approach as a rough approximation of the companies' preference for giving limit increases to revolvers.

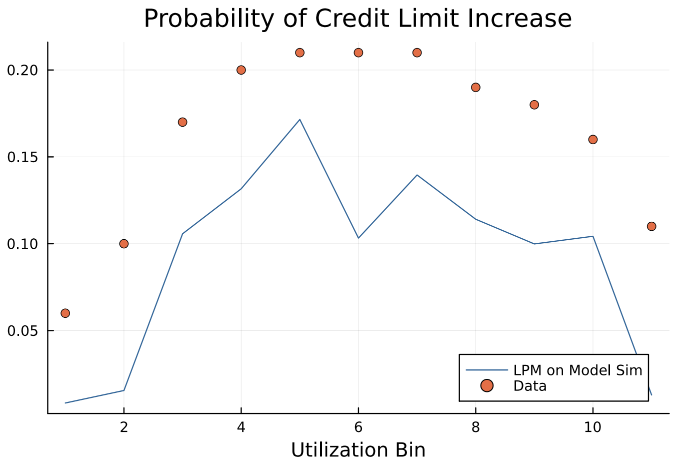

Figure 21: Credit limit increase by utilization with re-optimizing firm

to that observed in the data, with the highest probability of increase being for revolving borrowers who have a moderate degree of credit card utilization. The firm is much less willing to give limit increases to individuals near zero utilization or 100 percent utilization, similar to the empirical results in Figure 10. The main dimension upon which our model fails is that it gives fewer credit limit increases than are observed in the Y-14M data.[40]

Figures 37 and 38 in Appendix B.5 present the effects of both counterfactual policies (prohibiting credit limit increases to revolving borrowers and prohibiting increases without borrower consent) through the lens of our model with an optimizing firm. In this version, we allow the firm to optimize its CLI policy function once in the baseline regime and again in each counterfactual policy regime.

In the baseline with firm optimization, debt-to-income ratios, interest payments, and defaults are all reduced compared to our initial calibration, as the optimizing firm issues fewer CLIs than observed in our baseline data. When we implement the counterfactual policies, we find qualitatively similar implications to our baseline results (presented in Sections 5.2 and 5.3), though with slightly different magnitudes. Under both policy changes, the firm responds by reallocating a higher fraction of limit increases toward consumers without temptation issues compared to the scenario without firm re-optimization. Nevertheless, both counterfactuals still produce reductions in debt-to-income ratios, interest payments, and default probabilities, consistent with fewer credit limit increases being directed toward households with self-control issues.

Turning towards welfare, Table 5 shows the implications of both counterfactual policies when firms are allowed to re-optimize. Similar to our previous results, we again find that the benefits of reduced temptation outweigh the costs of reduced flexibility, leading households to benefit on average from the counterfactual policies. However, the positive

[40] This may be driven by the fact that our model is at an annual frequency, while firms learn monthly.

effects are dampened relative to our baseline calibration. This occurs because our model with the optimizing firm and baseline policy already features fewer limit increases—thus the reduction in limit increases due to the counterfactual policies is smaller than in our previous analysis.

	Overall Effect	Standard	Tempted
UK			
With Psychological Costs	0.21	−0.24	0.89
Without Psychological Costs	0.28	−0.00	0.72
Canada			
With Psychological Costs	0.27	−0.20	0.98
Without Psychological Costs	0.32	−0.01	0.81

Table 5: Welfare results with firm re-optimization, (CEV, %)

6 Conclusion

This paper provides novel evidence on the role and prevalence of bank-initiated credit limit increases. While we do not observe the precise algorithms used by banks, our empirical analysis reveals that banks systematically extend additional credit to revolving borrowers. This pattern suggests a revealed preference for targeting borrowers who are likely to generate interest revenue by carrying revolving balances.

To evaluate the welfare implications of this behavior, we develop a structural model that incorporates borrower heterogeneity and self-control problems. We find that if a significant share of borrowers suffer from self-control issues, restricting banks from unilaterally increasing limits for revolving borrowers could improve welfare.

Our analysis offers a first step toward understanding the costs and benefits of regulating algorithmic credit limit increases. Several directions for future research remain. First, while we focus on self-control problems, other behavioral biases may also influence borrowing behavior. For example, Stango and Zinman (2009) document that many consumers suffer from exponential growth bias—the tendency to underestimate how interest compounds—which can lead consumers to borrow more than they otherwise would. In a related vein, Bäckman et al. (2025) provide evidence that low monthly payments lead to increased borrowing, potentially because some consumers focus on payment size rather than total borrowing costs. Understanding whether banks disproportionately extend credit to consumers exhibiting these biases—and whether alternative policy interventions could offer better protection—remains an important open empirical question.

Second, our model centers on bank-initiated credit limit increases for tractability. A promising direction for future work is to compare bank- versus consumer-initiated increases, particularly in environments where moral hazard may shape banks' responses to consumer requests.

Finally, the regulatory landscape is evolving. Countries such as the UK, Canada, and Australia have introduced policies that restrict or prohibit bank-initiated credit limit increases, and similar regulations are forthcoming in the European Union in 2026. These policy interventions remain largely unstudied. As data become available, empirical work evaluating their effects—both on consumer borrowing behavior and on how banks adjust their credit allocation strategies—would enhance our understanding of the consequences of such regulation.

References

Adams, Robert and Vitaly M. Bord, "The Effects of the COVID-19 Shutdown on the Consumer Credit Card Market: Revolvers versus Transactors," 2020. Feds Notes.

_ , _ , and Bradley Katcher, "Credit Card Profitability," 2022. Feds Notes.

Agarwal, Sumit, Andrea Presbitero, André F. Silva, and Carlo Wix, "Who Pays For Your Rewards? Redistribution in the Credit Card Market," 2023. FEDS Working Paper.

_ , Jessica Pan, and Wenlan Qian, "Age of decision: Pension savings withdrawal and consumption and debt response," *Management Science*, 2019.

_ , Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel, "Do Banks Pass through Credit Expansions to Consumers Who want to Borrow?*," *The Quarterly Journal of Economics*, 07 2017, *133* (1), 129–190.

Alfonso-Sánchez, Sherly, Jesús Solano, Alejandro Correa-Bahnsen, Kristina P Sendova, and Cristián Bravo, "Optimizing credit limit adjustments under adversarial goals using reinforcement learning," *European Journal of Operational Research*, 2024, *315* (2), 802–817.

Allcott, Hunt, Joshua Kim, Dmitry Taubinsky, and Jonathan Zinman, "Are High-Interest Loans Predatory? Theory and Evidence from Payday Lending," *The Review of Economic Studies*, 2021, *89* (3), 1041–1084.

Andrews, Isaiah, Matthew Gentzkow, and Jesse M Shapiro, "Measuring the sensitivity of parameter estimates to estimation moments," *The Quarterly Journal of Economics*, 2017, *132* (4), 1553–1592.

Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg, "The Hyperbolic Buffer Stock Model: Calibration, Simulation, and Empirical Evaluation," *Journal of Economic Perspectives*, 2001, *15* (3), 47–68.

Ashraf, Nava, Dean Karlan, and Wesley Yin, "Tying Odysseus to the mast: Evidence from a commitment savings product in the Philippines," *The Quarterly Journal of Economics*, 2006, *121* (2), 635–672.

Augenblick, Ned and Matthew Rabin, "An Experiment on Time Preference and Misprediction in Unpleasant Tasks," *The Review of Economic Studies*, 2018, *86* (3), 941–975.

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

Bäckman, Claes, Patrick Moran, and Peter van Santen, "Mortgage Design, Repayment Schedules, and Household Borrowing," *Review of Financial Studies (Forthcoming)*, 2025.

Balakrishnan, Uttara, Johannes Haushofer, and Pamela Jakiela, "How soon is now? Evidence of present bias from convex time budget experiments," *Experimental Economics*, 2020, *23* (2), 294–321.

Bertrand, Marianne and Adair Morse, "Information Disclosure, Cognitive Biases, and Payday Borrowing," *The Journal of Finance*, 2011, *66* (6), 1865–1893.

Blattner, Laura, Scott Nelson, and Jann Spiess, "Unpacking the black box: Regulating algorithmic decisions," *arXiv preprint arXiv:2110.03443*, 2021.

Botella, Elena, *Delinquent: Inside America's Debt Machine*, Univ of California Press, 2022.

Chatterjee, Satyajit, Dean Corbae, Makoto Nakajima, and José-Víctor Ríos-Rull, "A Quantitative Theory of Unsecured Consumer Credit with Risk of Default," *Econometrica*, 2007, *75* (6), 1525–1589.

Chava, Sudheer, Rohan Ganduri, Nikhil Paradkar, and Linghang Zeng, "Shocked by Bank Funding Shocks: Evidence from Consumer Credit Cards," *The Review of Financial Studies*, 05 2023, *36* (10), 3906–3952.

DellaVigna, Stefano and Ulrike Malmendier, "Paying Not to Go to the Gym," *American Economic Review*, 2006, *96* (3), 694–719.

Exler, Florian and Michele Tertilt, "Consumer debt and default: A macro perspective," Technical Report, IZA Discussion Papers 2020.

_ , Igor Livshits, James MacGee, and Michèle Tertilt, "Consumer Credit with Over-Optimistic Borrowers," *Journal of the European Economic Association*, 12 2024, p. jvae057.

Fulford, Scott and Joanna Stavins, "Income and the Card Act's Ability-to-Pay Rule in the US Credit Card Market," 2025. Consumer Financial Protection Bureau Office of Research Working Paper No. 24-6.

Fulford, Scott L. and Scott D. Schuh, "Credit cards, credit utilization, and consumption," *Journal of Monetary Economics*, 2024, *148*, 103619.

_ and Scott Schuh, "Revolving versus Convenience Use of Credit Cards: Evidence from U.S. Credit Bureau Data," *Journal of Money, Credit and Banking*, 2023, *55* (7), 1667–1701.

Ganong, Peter and Pascal Noel, "Consumer spending during unemployment: Positive and normative implications," *American Economic Review*, 2019, *109* (7), 2383–2424.

Gathergood, John, "Self-control, financial literacy and consumer over-indebtedness," *Journal of economic psychology*, 2012, *33* (3), 590–602.

Grodzicki, Daniel, "Credit Cards and the Floating Rate Channel of Monetary Policy," 2020. Working Paper.

_ and Sergei Koulayev, "Sustained credit card borrowing," *Journal of Consumer Affairs*, 2021, *55* (2), 622–653.

Gross, David B. and Nicholas S. Souleles, "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data*," *The Quarterly Journal of Economics*, 02 2002, *117* (1), 149–185.

Gul, Faruk and Wolfgang Pesendorfer, "Temptation and Self-Control," *Econometrica*, 2001, *69* (6), 1403–1435.

_ and _ , "Self-control, revealed preference and consumption choice," *Review of Economic Dynamics*, 2004, *7* (2), 243–264.

Harris, Christopher and David Laibson, "Dynamic Choices of Hyperbolic Consumers," *Econometrica*, 2001, *69*, 935–957.

Heidhues, Paul and Botond Kőszegi, “Exploiting Naïvete about Self-Control in the Credit Market,” *American Economic Review*, December 2010, *100* (5), 2279–2303.

_ **and Botond Köszegi**, “On the Welfare Costs of Naiveté in the US Credit-Card Market,” *Review of Industrial Organization*, 2015, *47* (3), 341–354.

Hsu, Joanne W., David A. Matsa, and Brian T. Melzer, “Unemployent Insurance as a Housing Market Stabilizer,” *American Economic Review*, 2018, *108* (1), 49–81.

Huang, Kevin XD, Zheng Liu, and John Qi Zhu, “Temptation and Self-Control: Some Evidence and Applications,” *Journal of Money, Credit and Banking*, 2015, *47* (4), 581–615.

Jolls, Christine and Cass R. Sunstein, “Debiasing through Law,” 2005. NBER Working Paper.

Keane, Michael P. and Nada Wasi, “Labor Supply: The Roles of Human Capital and the Extensive Margin,” *The Economic Journal*, 2016, *126* (592), 578–617.

Kovacs, Agnes and Patrick Moran, “Breaking the Commitment Device: The Effect of Home Equity Withdrawal on Consumption, Saving, and Welfare,” 2021.

_ **and** _ , “Financial Innovation, the Decline in Household Savings, and the Trade-off between Flexibility and Commitment,” *mimeo*, 2025.

_ **, Hamish Low, and Patrick Moran**, “Estimating Temptation and Commitment Over the Life Cycle,” *International Economic Review*, 2021, *62* (1), 101–139.

Kovrijnykh, Natalia, Igor Livshits, and Ariel Zetlin-Jones, “Building Credit Histories,” 2023. Working Paper.

Kuchler, Theresa and Michaela Pagel, “Sticking to your plan: The role of present bias for credit card paydown,” *Journal of Financial Economics*, 2021, *139* (2), 359–388.

Laibson, David, “Golden Eggs and Hyperbolic Discounting,” *Quarterly Journal of Economics*, 1997, pp. 443–477.

_ **, Peter Maxted, Andrea Repetto, and Jeremy Tobacman**, “Estimating Discount Functions with Consumption Choices over the Lifecycle,” *Mimeo*, 2017.

Lee, Sean Chanwook and Peter Maxted, “Credit Card Borrowing in Heterogeneous-Agent Models: Reconciling Theory and Data,” *Working paper*, 2023.

Livshits, Igor, “Regulating Consumer Credit and Protecting (Behavioral) Borrowers,” *Economic Insights*, March 2020, *5* (1), 14–19.

Meier, Stephan and Charles Sprenger, "Present-biased preferences and credit card borrowing," *American Economic Journal: Applied Economics*, 2010, *2* (1), 193–210.

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

Nelson, Scott, "Private information and price regulation in the us credit card market," *Econometrica*, 2025.

O'Dea, Cormac, "Insurance, efficiency and the design of public pensions," *mimeo*, 2018.

Phelps, Edmund and Robert A. Pollak, "On Second-Best National Saving and Game-Equilibrium Growth," *The Review of Economic Studies*, 1968, *35*, 185–199.

Schneider, Patrick and Patrick Moran, "Situational and Behavioral Determinants of Early Withdrawal from Retirement Accounts," 2024.

Schuh, Scott D. Fulford Scott L. , "Credit cards, credit utilization, and consumption," *Journal of Monetary Economics*, 2024, *148.*

Stango, Victor and Jonathan Zinman, "Exponential growth bias and household finance," *the Journal of Finance*, 2009, *64* (6), 2807–2849.

_ **and** _ , "Fuzzy Math, Disclosure Regulation, and Market Outcomes: Evidence from Truth-in-Lending Reform," *The Review of Financial Studies*, 2011, *24* (2), 506–534.

Strotz, Robert H., "Myopia and Inconsistency in Dynamic Utility Maximization," *The Review of Economic Studies*, 1956, *23*, 165–180.

Sutton, R. S. and A. G. Barto, *Reinforcement learning: An introduction*, MIT press, 2018.

Toussaert, Séverine, "Eliciting temptation and self-control through menu choices: A lab experiment," *Econometrica*, 2018, *86* (3), 859–889.

Yin, Xiao, "Learning in the limit: Income inference from credit extension," 2022. Working Paper.

A Appendix

A.1 Stylized Facts

This section presents extra analysis for the stylized facts discussed in section 2.2.

Appendix for Fact 2

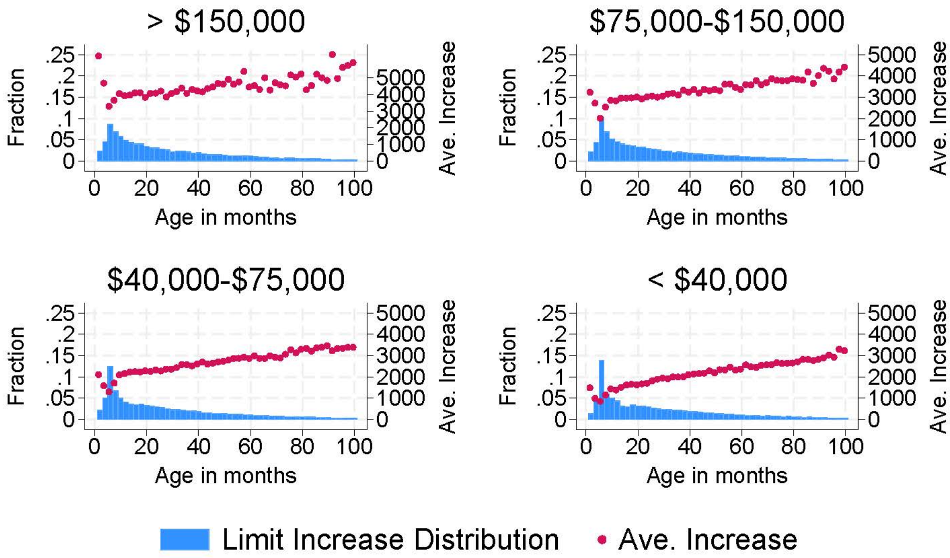

Figure 22: Distribution of limit increases by age and income

This figure includes only credit cards originated since January 2014.

Figure 22 replicates the analysis of Figure 3 by income.

Appendix for Fact 5

Figure 23 compares the average credit limit at origination for credit cards issued by lenders above and below the median of the number of mentions of AI and ML in their 10K statements by credit score (Panel A) and by income (Panel B). The figure shows no systematic differences between the average credit limit at origination for lenders above and below the median in their mentions of ML/AI.

Appendix for Fact 6

In Fact 2, we showed that a high share of limit increases occur within the first 6 months. How does this reconcile with limit increases being more prevalent among revolvers? The blue line in Figure 24 displays the share of accounts that revolve, by age. The figure shows that many accounts begin to revolve balances immediately upon card origination. Indeed, of cards that revolve at some point in their first 10 years of existence, about

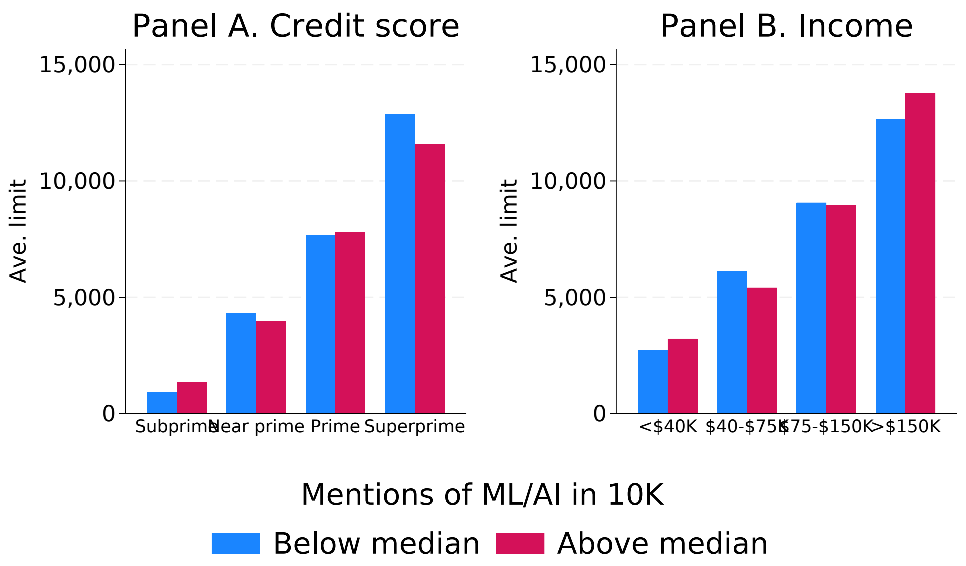

Figure 23: Average Credit Limit

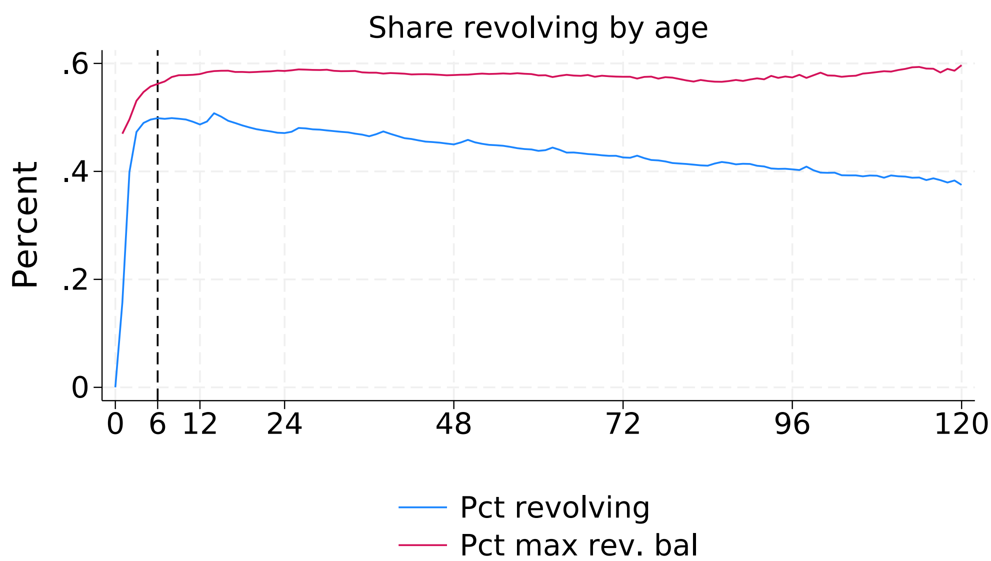

Figure 24: Share of revolving by age

two-thirds revolve in the first 6 months after origination. This finding is consistent with Grodzicki and Koulayev (2021), who find that when accounts start revolving balances they continue to do so for over a year, on average.

The red line in Figure 24 shows the fraction of current revolving balances divided by the maximum revolving balances the account ever carries, averaged across ever-revolving accounts. Conditional on revolving, even young accounts revolve about 60% of their maximum revolving amount, on average. Thus, it is not the case—on the extensive margin or the intensive margin—that young accounts revolve less often or smaller amounts. This suggests that lenders can likely predict future revolving behavior from revolving behavior in the first few months, and may target limit increases accordingly.[41]

A.2 Summary Statistics

This section discusses the summary statistics for the empirical analysis presented in section 3. Table 6 shows summary statistics for our sample using Y-14 data. Columns 1 and 2 present means and standard deviations for the full sample, while columns 3 and 4 focus on the bank-initiated increases and columns 5 and 6 focus on the borrower-initiated increases. For columns 3-6, the data are as of the month prior to the limit increase.

Several points stand out. First, accounts that undergo a bank-initiated credit increase tend to have somewhat lower incomes and credit scores relative to the full sample. 55% of instances of bank-initiated increases have a prime or superprime credit score, compared to 65% for the full sample. Similarly, 29% of bank-initiated increases have reported income above $75,000, compared to 31% for the full sample. Interestingly, borrower-initiated increases tend to be have somewhat lower credit scores than the full sample, but higher incomes: 60% of observations have a prime or superprime credit score and 37% have reported income above $75,000. [42]

Second, accounts that undergo limit increases tend to have higher utilization. While both bank and borrower-initiated accounts tend to have higher overall and transaction utilization, only bank-initiated increases tend to have higher revolving utilization. Indeed, while about 45% of the active cards in the full sample are revolvers–the same share as among accounts that undergo a borrower-initiated increase–60% of accounts with a bank-initiated increase are revolvers.

Third, the average bank-initiated credit limit is about $2,400, while the average borrower-initiated limit increase is a slightly large $2,800. That said, the average bank-initiated increase is larger in percentage terms, as the average credit limit prior to a bank-initiated increase is somewhat smaller.

[41] Of course, lenders learn other information about the account in the first 6 months, not just the account's likelihood to revolve. Importantly, lenders learn about the account's likelihood to default. Of accounts that become 30 or more day delinquent at some point, about a third become delinquent for the first time in the first 6 months after origination.

[42] Credit scores are measured contemporaneously. If a credit score since origination is not available, then credit score at origination is used. We similarly use most recent income when available. However, as discussed in Fulford and Stavins (2025), since 2020, updated income is no longer collected in the data. Limiting to just the 2014-2020 period and using only updated income does not change our results.

Variable	All		Bank-Initiated Increase		Borrower-Initiated Increase	
	Mean	SD	Mean	SD	Mean	SD
CC Limit ($)	9,883.763	8,969.863	7,557.323	7,447.187	8,081.482	8,595.169
Interest rate margin	14.586	5.157	15.299	5.664	15.545	4.873
I{Mult CC with Lender}	0.383	0.486	0.290	0.454	0.359	0.480
Credit score	751.216	80.056	732.718	65.776	739.140	66.640
Superprime	0.442	0.497	0.248	0.432	0.293	0.455
Prime	0.231	0.422	0.326	0.469	0.334	0.472
Near prime	0.253	0.435	0.366	0.482	0.317	0.465
Subprime	0.073	0.261	0.060	0.238	0.056	0.229
Income: > $150,000	0.091	0.287	0.080	0.271	0.125	0.331
Income: $75,000 - $150,000	0.227	0.419	0.224	0.417	0.262	0.440
Income: $40,000 - $75,000	0.279	0.449	0.316	0.465	0.296	0.456
Income < $40,000	0.240	0.427	0.320	0.467	0.249	0.432
Utilization	0.295	0.350	0.411	0.302	0.353	0.312
Revolving Utilization	0.224	0.334	0.278	0.308	0.210	0.296
Transactional utilization	0.073	0.113	0.136	0.143	0.145	0.163
Revolver	0.435	0.496	0.598	0.490	0.449	0.497
Heavy Revolver past 12 Mo	0.253	0.435	0.303	0.460	0.213	0.409
Light Revolver past 12 Mo	0.381	0.486	0.469	0.499	0.473	0.499
Transactor	0.356	0.479	0.219	0.413	0.308	0.462
I{Recent bank-initiated increase}	0.037	0.189	0.016	0.126	0.015	0.122
I{Recent borrower-initiated increase}	0.009	0.092	0.005	0.068	0.036	0.185
I{Recent credit score change}	0.752	0.432	0.790	0.407	0.811	0.391
Recent credit score change	0.077	9.114	0.917	9.561	1.331	9.452
Age <6m	0.046	0.209	0.108	0.310	0.054	0.227
Age 6m-1y	0.062	0.241	0.125	0.331	0.098	0.297
Age 1-2y	0.356	0.479	0.405	0.491	0.489	0.500
Age 3-5y	0.244	0.429	0.268	0.443	0.313	0.464
Age 5-10y	0.230	0.421	0.200	0.400	0.210	0.407
Age 10+	0.306	0.461	0.162	0.368	0.150	0.357
Chg Limit ($)			2,524.417	2,412.428	2,980.702	3,449.598
Pct Chg Limit			0.413	0.337	0.423	0.330
Observations	160059951		2008939		466631	

Table 6: Summary statistics

A.3 Who receives limit increases?

In this section, we present more detailed analysis examining the correlations between credit card account characteristics and the likelihood of limit increases. Specifically, in section A.3.1 we examine the effect of a wide range of account characteristics, as well as the change in account utilization, on the likelihood of receiving a limit increase. In section A.3.2, we show that our results are not driven by the dynamic relationship between transacting and revolving utilization. In section A.3.3 and A.3.4, we examine how our results vary by credit score and age, respectively. In section A.4, we present more analysis on borrower-initiated limit increases, while in section A.5, we briefly discuss the results on the intensive margin.

A.3.1 Account characteristics and limit increases

Table 7 presents the results of equation (1), where the dependent variable is an indicator equal to 1 if the account experienced a bank-initiated increase in the following month. In specification (1), we include controls for age, coarse credit score and income bins, Mult CC with Lender—an indicator for whether the borrower has another credit card with the lender—log of the credit limit, interest rate margin, indicators for recent bank-initiated limit, borrower-initiated limit and credit score increases, and the change in credit score over the previous 3 months. We also include month, state, and bank fixed effects. Specification (2) adds card group fixed effects, while specification (3) adds credit score and income decile fixed effects.

Our main variable of interest in specification (1)-(4) is utilization, which we decompose into revolving and transacting utilization in specifications (5)-(6). In specifications (4) and (6), we also add the change in utilization from the prior 3 months. As expected, utilization is positively correlated with receiving a limit increase, with a 1 percentage point increase in utilization leading to approximately a 1.4 basis point increase in the probability of receiving a limit increase.[43] Decomposing utilization into its revolving and transacting components, we find that both components are positively correlated with limit increases. In addition, growth in both components is negatively correlated with limit increases. The results of specification (6) suggest that a one standard deviation increase in revolving utilization increases the probability of a limit increase by 0.2 basis points, while a one standard deviation increase in transacting utilization increases the probability of a limit increase by 0.5 basis points.

The coefficients on the other variables are generally as expected. Younger cards were

[43] Because of the sparseness of the limit increase indicator at the monthly frequency—as shown in Table 6–both the coefficients and our R-squared are quite low. Fulford and Stavins (2025) find a higher R-Squared of about 0.15, but their sample consists of credit cards younger than 5 years, which are more likely to experience a limit increase.

more likely to receive a credit limit increase, and higher credit score cards were more likely to receive one than subprime cards. Cards with a higher credit limit were less likely to receive a limit increase, as were cards that recently experienced a limit increase. Cards of borrowers with recent growth in their credit score were more likely to receive a limit increase.

	(1)	(2)	(3)	(4)	(5)	(6)
Utilization	0.013***	0.012***	0.014***	0.013***		
	(0.001)	(0.001)	(0.002)	(0.002)		
Chg. in utilization				0.001		
				(0.002)		
Revolving utilization					0.007***	0.008***
					(0.002)	(0.002)
Chg in rev. utilization						-0.015***
						(0.002)
Transactional utilization					0.059***	0.068***
					(0.003)	(0.003)
Chg in tr. utilization						-0.005
						(0.003)
Age <6m	0.021***	0.021***	0.020***	0.023***	0.016***	0.027***
	(0.006)	(0.006)	(0.006)	(0.007)	(0.006)	(0.009)
Age 6m-1y	0.018***	0.017***	0.016***	0.016***	0.016***	0.016***
	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)	(0.002)
Age 2-5y	0.006***	0.006***	0.005***	0.005***	0.005***	0.005***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Age 5-10y	0.004***	0.003***	0.003***	0.003***	0.003***	0.003***
	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Superprime	0.015***	0.017***				
	(0.002)	(0.002)				
Prime	0.021***	0.023***				
	(0.002)	(0.002)				
Near prime	0.016***	0.017***				
	(0.002)	(0.002)				
Income: > $150,000	-0.000	-0.000				
	(0.001)	(0.001)				
Income: $75,000 - $150,000	-0.001	-0.000				
	(0.000)	(0.000)				
Income: $40,000 - $75,000	-0.000	0.000				
	(0.000)	(0.000)				
Mult CC with Lender	-0.003***	-0.003***	-0.003***	-0.003***	-0.002***	-0.002***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Log CC Limit	-0.002***	-0.001***	-0.002***	-0.002***	-0.001	-0.000
	(0.000)	(0.001)	(0.000)	(0.000)	(0.000)	(0.000)
Interest rate margin	-0.000	-0.000	-0.000	-0.000	-0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
I{Recent bank-initiated increase}	-0.014***	-0.016***	-0.017***	-0.017***	-0.019***	-0.020***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
I{Recent borrower-initiated increase}	-0.010***	-0.010***	-0.010***	-0.010***	-0.014***	-0.015***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
I{Recent credit score increase}	0.000***	0.000***	0.000***	0.000*	0.001***	0.000*
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Chg. in Credit Score	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
State, Date, Bank FE	Yes	Yes	Yes	Yes	Yes	Yes
Card group FE		Yes	Yes	Yes	Yes	Yes
CS and Income Decile FE			Yes	Yes	Yes	Yes

Observations	159281782	159281782	159281782	156768808	159281782	154792670
R-Squared	0.009	0.012	0.012	0.012	0.014	0.015

Table 7: Bank-Initiated Limit Increases

Change in Utilization

Our preferred specifications include both the average utilization over the previous 3 months and the change in utilization from the prior 3 months. Figure 25 presents the coefficients of the deciles of the change in revolving utilization in Panel A and the change in transacting utilization in Panel B. The effect of the change in revolving utilization from the prior 3 months is S-shaped. It appears to peak somewhat earlier than the mid-point, with the highest probability of receiving an offer among accounts with a large decline in revolving utilization. Accounts with revolving utilization that declined recently, as well as accounts that saw large growth in utilization, are less likely to receive a limit increase than accounts with zero change in utilization. For the change in transacting utilization, a change of (-0.6, 0.8] appears to have minimal effect on the likelihood of a limit increase, while a change of (-0.8,-0.6] or (0.8,1] reduces the likelihood of a limit increase.

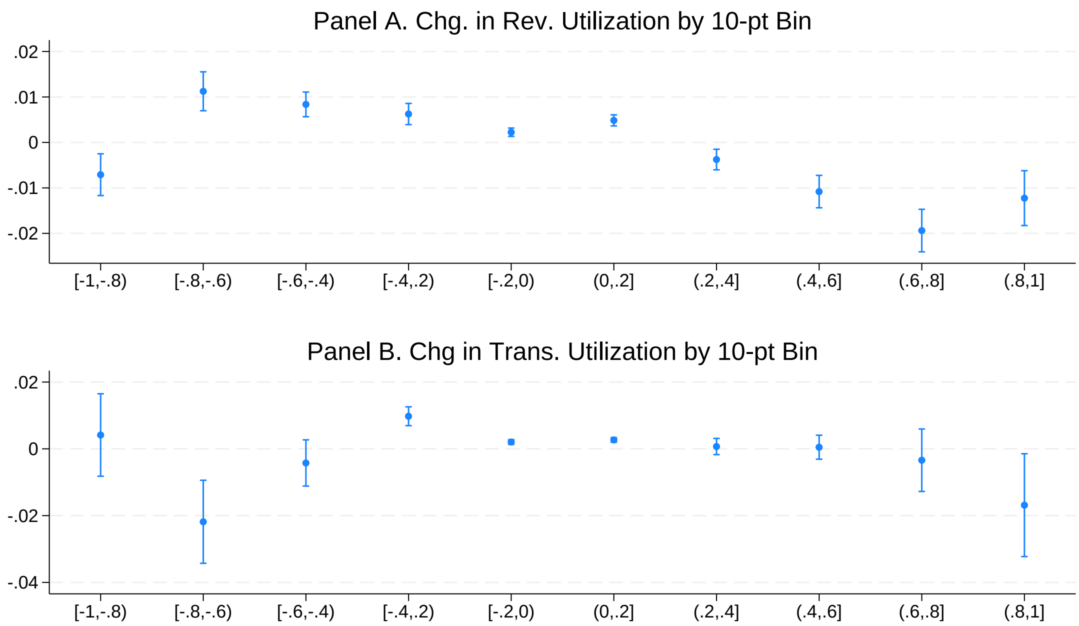

Figure 25: Change in utilization and the likelihood of a credit limit increase

Variable	Revolver		Heavy Revolver		Light Revolver		Transactor	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Revolving Utilization	0.52	0.35	0.61	0.33	0.18	0.30	0.01	0.07
Transacting Utilization	0.07	0.12	0.05	0.09	0.09	0.16	0.06	0.13
Credit Limit ($)	8,328	8,256	8,014	7,773	9,887	9,236	11,343	9,277

Table 8: Revolving and transacting utilization by account type

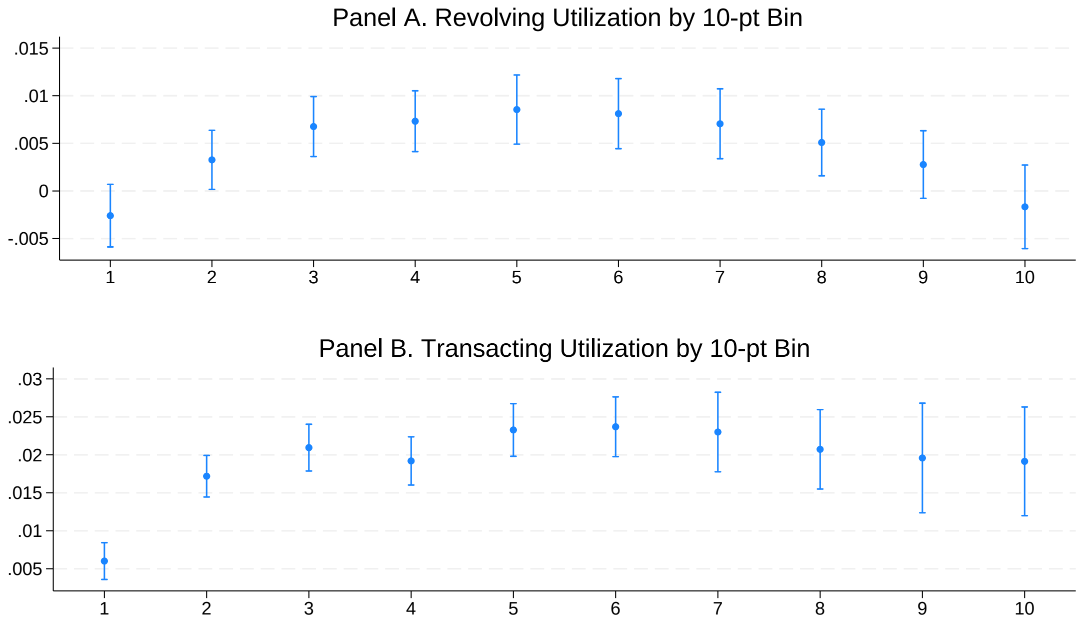

Figure 26: Revolving and transacting utilization among revolvers

A.3.2 Utilization and account type

There are two potential impediments to ascribing the difference between Panels A and B of Figure 10 to a difference in how banks increase limits in response to usage. First, it is possible that the differences in Figure 10 are driven by differences between revolvers and transactors. Accounts that consistently revolve tend to have higher revolving utilization; by contrast, transacting accounts have zero revolving utilization by definition.

To shed more light on the difference between revolvers and transactors, Table 8 presents the mean and standard deviation for revolving utilization, transacting utilization and total credit limit for revolvers (columns 1-2), heavy revolvers (columns 3-4), light revolvers (columns 5-6), and transactors (columns 7-8). For example, Column 1 of Table 8 shows that revolving utilization among revolvers is 0.52, consistent with Lee and Maxted (2023), who point out that empirically, borrowing is lower than one would expect among constrained, "hand-to-mouth" borrowers who would take on debt to their credit limit.

Thus, it is possible that the results of Panel 10 are due to comparing revolvers with an average revolving utilization of 0.52 with transactors with an average revolving utilization of 0. To rule out that this difference in accounts–rather than differences in usage–drive our results, in Figure 26, we rerun the analysis limiting just to accounts that revolved in that month. The figure shows patterns in both revolving and transacting utilization that are very similar to those for the full sample.

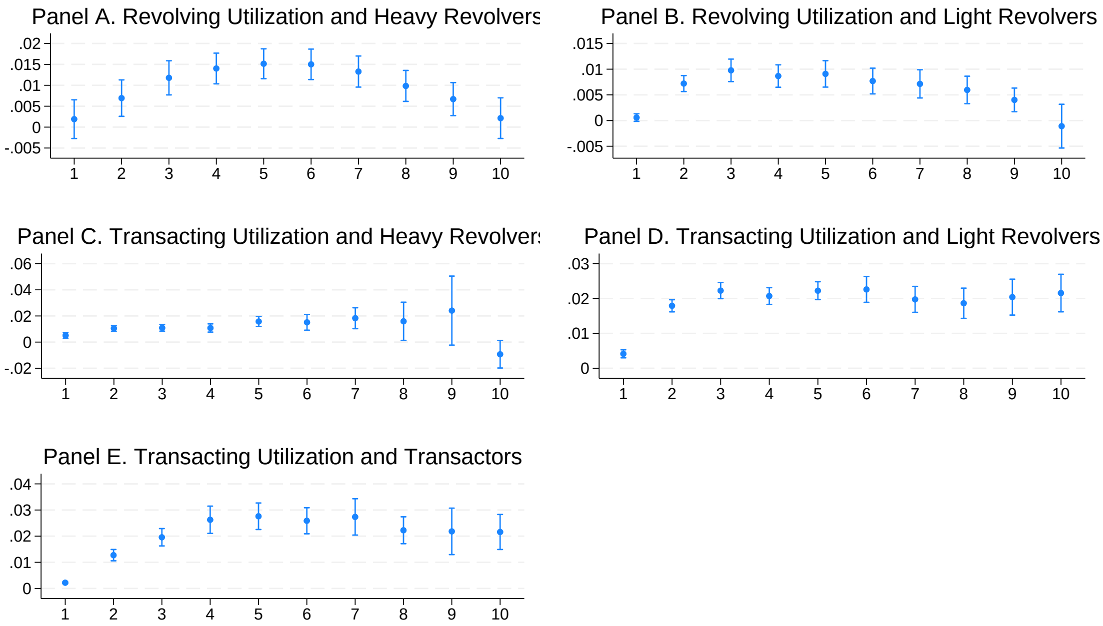

Figure 27: Utilization by revolving status

The second potential issue is that transacting and revolving utilization are linked, since this month's transacting utilization can become next month's revolving utilization. Thus, it is possible that what we ascribe to banks' responses to revolving utilization are lagged responses to transacting utilization. To disentangle this dynamic relationship, we interact revolving and transacting utilization by revolving status. Following Adams and Bord (2020), we examine three types of accounts: heavy revolvers who revolved each of the prior 12 months, light revolvers who revolved 1-11 times in the prior 12 months, and transactors who did not revolve at all in the prior 12 months. Table 8 shows that average revolving utilization for heavy revolvers is 0.61, while their transacting utilization is only 0.05. Thus, the dynamic relationship between transacting and revolving utilization is unlikely to matter much for this group of accounts. Similarly, since transactors have a revolving utilization of zero, their transacting utilization is unrelated to their revolving utilization.

Figure 27 presents the results. Panels A and B show that revolving utilization for both heavy and light revolvers follows similar inverted U-shaped pattern as when coefficients are estimated on the pooled sample. Similarly, Panels C, D, and E show that transacting utilization follows similar patterns for all three types of accounts, although coefficients for transacting utilization at transactors are slightly larger in magnitude than at heavy and light revolvers.[44] This suggests that the dynamic relationship between revolving and transacting utilization do not drive our findings.

A.3.3 Utilization and credit score

Panel A: Revolving utilization

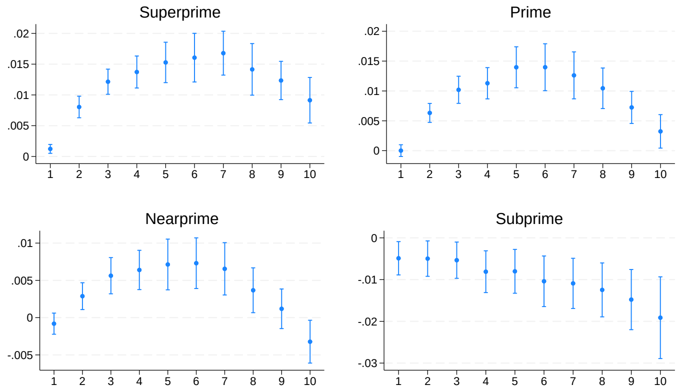

Panel B: Transacting utilization

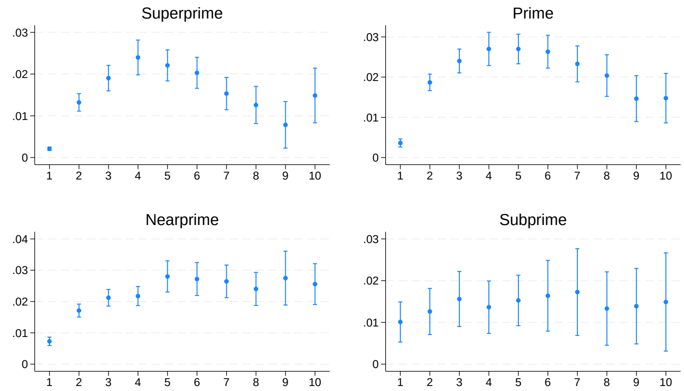

Figure 28: Differences in utilization by credit score

Figure 28 examines how the patterns identified in Figure 10 vary across the contemporaneous credit score distribution. Panel A shows that the U-shaped pattern holds for all credit score groups except for subprime. Subprime borrowers, by contrast, experience the highest probability of receiving a limit increase at low levels of revolving utilization and a monotonic decline with utilization. This likely reflects their higher risk, with lenders unwilling to increase available credit to borrowers whose probability of default is high.

[44]The pattern of the coefficients for transacting utilization at heavy revolvers appears somewhat different from that of revolving utilization at light revolvers. However, this is purely driven by the low coefficient on the 10th bin, and by definition, there are very few heavy revolver accounts with transacting utilization above 90%.

Interestingly, both subprime and nearprime borrowers have a higher probability of a limit increase at very low levels of revolving utilization than at utilization above 0.9, as opposed to prime and superprime borrowers for whom the converse holds. Again, this is likely a reflection of risk. Revolvers are the main source of revenues for credit card operations (Adams et al., 2022). For prime and superprime borrowers, the risk of default even at high levels of revolving utilization is low; therefore the potential profit to granting a limit increase may outweigh the cost. However, for nearprime and subprime borrowers, the risk of default is higher, and so lenders are relatively less willing to increase limits at high levels of revolving utilization.

Panel B similarly shows that the logistic growth pattern of transacting utilization also varies by credit score. The patterns for subprime and nearprime borrowers mirror that of the overall. However, prime and superprime borrowers exhibit a U-shaped pattern with the probability of a limit increase higher for the middle bins of utilization than for the top bins.

A.3.4 Age and credit score

Figure 29 examines how the likelihood of a limit increase varies by credit score and age. The patterns are consistent with the histograms of Figure 3. All credit score bins experience more limit increases during the first year of a credit card, but the effect is particularly pronounced for near prime and subprime.

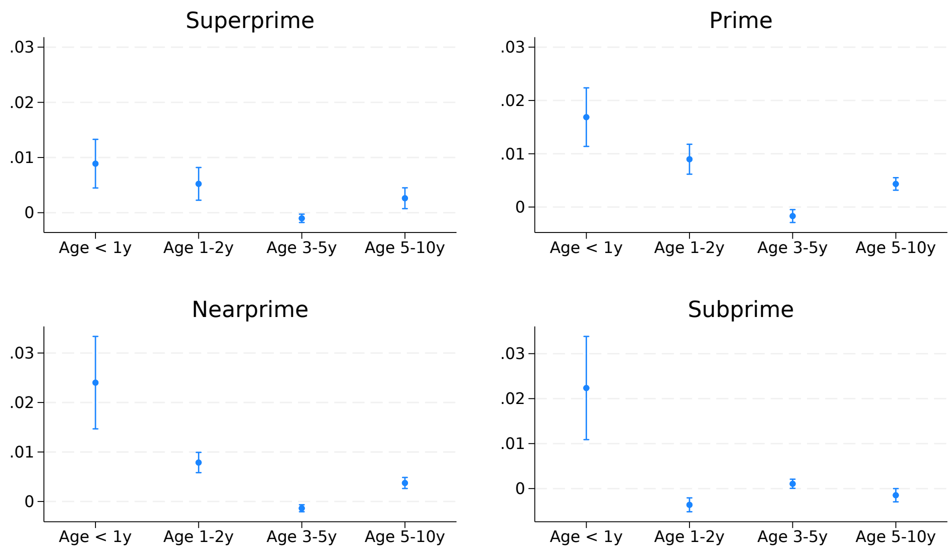

Figure 29: Age and limit increases, by credit score

A.4 Borrower-initiated increases

Table 9 repeats the analysis of Table 7 for borrower-initiated increases. Because borrower-initiated increases are more rare than bank-initiated increases, the magnitudes on all the coefficients are smaller. However, the difference in the magnitudes in Tables 9 and 10 are not just driven by differences in the prevalence of increases. A one-standard deviation in utilization can explain about 4% of the standard deviation in the bank-initiated increase indicator, and only 0.9% of the standard deviation in the borrower-initiated increase indicator. Specification (6) decomposes utilization into its revolving and transacting components and shows that revolving utilization has no correlation with borrower-initiated increases. Moreover, although the coefficient on transacting utilization is statistically significant, it explains less of the variation in borrower-initiated increases than for bank-initiated increases. A one-standard deviation in transacting utilization can explain about 4.5% of the standard deviation in the bank-initiated increase indicator, and 2.7% of the standard deviation in the borrower-initiated increase indicator.

	(1)	(2)	(3)	(4)	(5)	(6)
Utilization	0.001***	0.001***	0.001***	0.001***		
	(0.000)	(0.000)	(0.000)	(0.000)		
Chg. in Utilization				-0.000		
				(0.000)		
Revolving Utilization					-0.001**	-0.000
					(0.000)	(0.000)
Chg in rev. Utilization						-0.005***
						(0.001)
Transactional utilization					0.016***	0.019***
					(0.001)	(0.001)
Chg in tr. utilization						-0.006***
						(0.001)
Age <6m	0.001***	0.001***	0.001***	0.002***	0.000	0.001***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Age 6m-1y	0.003***	0.003***	0.003***	0.003***	0.002***	0.003***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Age 2-5y	0.002***	0.002***	0.002***	0.002***	0.002***	0.002***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Age 5-10y	0.001***	0.001***	0.001***	0.001***	0.001***	0.001***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Superprime	0.002***	0.003***				
	(0.001)	(0.001)				
Prime	0.004***	0.004***				
	(0.001)	(0.001)				
Near prime	0.002***	0.003***				
	(0.000)	(0.000)				
Income: > $150,000	0.002***	0.002***				
	(0.000)	(0.000)				
Income: $75,000 - $150,000	0.001***	0.001***				
	(0.000)	(0.000)				
Income: $40,000 - $75,000	0.000***	0.000***				
	(0.000)	(0.000)				
Mult CC with Lender	-0.000***	-0.000***	-0.000***	-0.000***	-0.000***	-0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Log CC Limit	-0.001***	-0.001***	-0.001***	-0.001***	-0.001***	-0.001***

	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Interest rate margin	-0.000	0.000***	0.000***	0.000***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
I{Recent bank-initiated increase}	-0.003***	-0.003***	-0.003***	-0.003***	-0.003***	-0.004***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
I{Recent borrower-initiated increase}	0.008***	0.007***	0.007***	0.007***	0.005***	0.005***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
I{Recent credit score increase}	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Chg. in Credit Score	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
State, Date, Bank FE	Yes	Yes	Yes	Yes	Yes	Yes
Card group FE		Yes	Yes	Yes	Yes	Yes
CS and Income Decile FE			Yes	Yes	Yes	Yes
Observations	159281782	159281782	159281782	156768808	159281782	154792670
R-Squared	0.002	0.003	0.003	0.003	0.004	0.004

Table 9: Borrower-Initiated Limit Increases

Figure 30 highlights some differences between bank and borrower-initiated increases by credit score. While for bank-initiated increases prime and superprime borrowers exhibited U-shaped patterns for both revolving and transacting utilization, for borrower-initiated increases, the pattern is flatter, suggesting that lenders are relatively more likely to grant these borrowers increases, if they ask for them, at high levels of utilization.

Finally, Figure 31 shows that while bank-initiated increases occur mainly among very young cards in the first year after origination, borrower-initiated increases are also likely to occur in the second year after origination. This likely reflects a demand effect of borrowers asking for higher limits as they continue using the card.

A.5 Intensive margin: the size of the limit increase

For the sake of brevity, we do not present the intensive margin results, which are available upon request and are similar to the extensive margin results. For both bank and borrower-initiated limit increases, revolving utilization generally follows an upside down U-shape, with somewhat lower limit increases for cards with high levels of revolving utilization. The one exception is subprime borrowers, for whom the size of the increase decreases monotonically in revolving utilization. For transacting utilization, lower levels of utilization receive smaller limit increases, with the exception of subprime borrowers.

Panel A: Revolving utilization

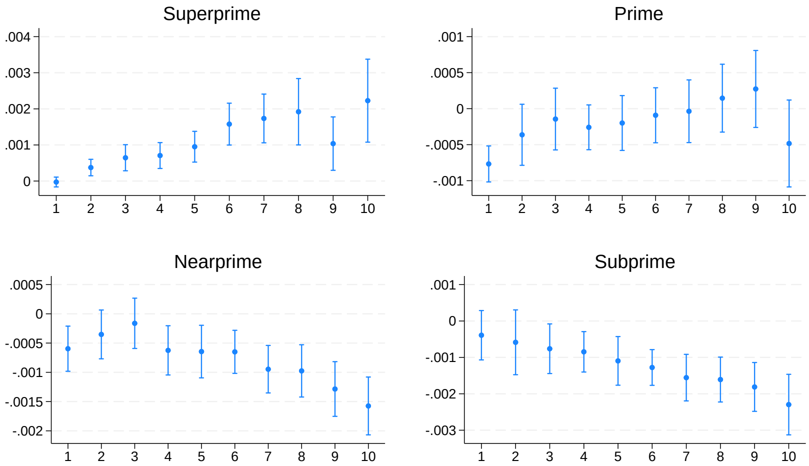

Panel B: Transacting utilization

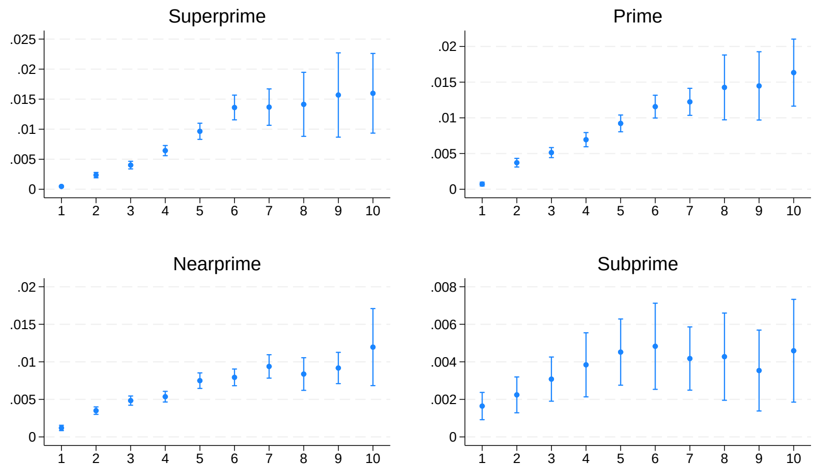

Figure 30: Differences in utilization by credit score, consumer-initiated increases

A.6 Event study: what happens after limit increases

In this section, we present some more details and robustness for the event study in section 3.2. To conduct the event study, we run a regression of the form:

$$Y_{it} = \sum_{\tau=t-12}^{t+12} D_{i,\tau} + \lambda_t + \eta_g + \epsilon_{it} \tag{15}$$

where λ_t are month fixed effects and η_g are state and card group fixed effects. Including credit score and income bin deciles, age, and other controls and fixed effects does not change our results.

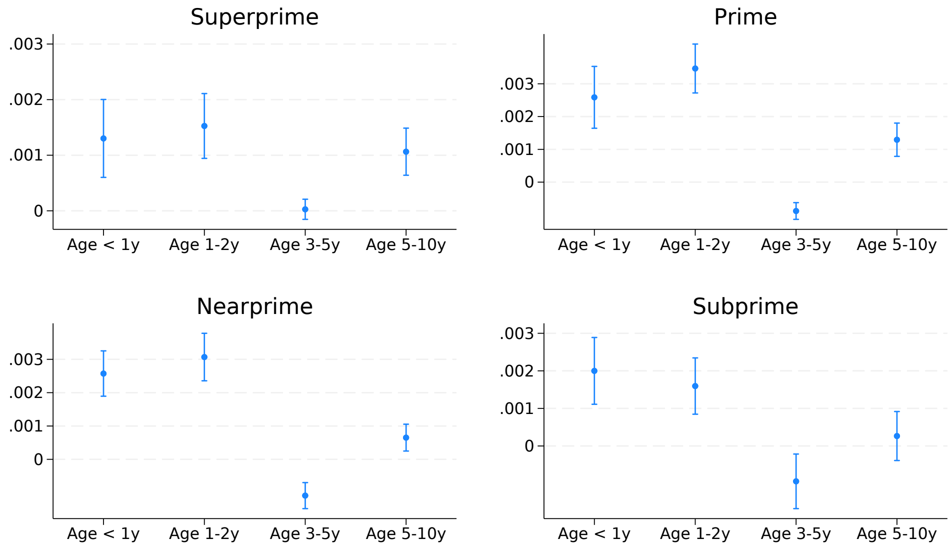

Figure 31: Age and consumer-initiated increases, by credit score

Figure 32 presents how total utilization changes around limit increases, as a contrast to revolving utilization presented in figure 13. For both bank- and consumer-initiated increases, utilization rebounds after the limit increase to a level approximately 2 percentage points lower than before the increase.

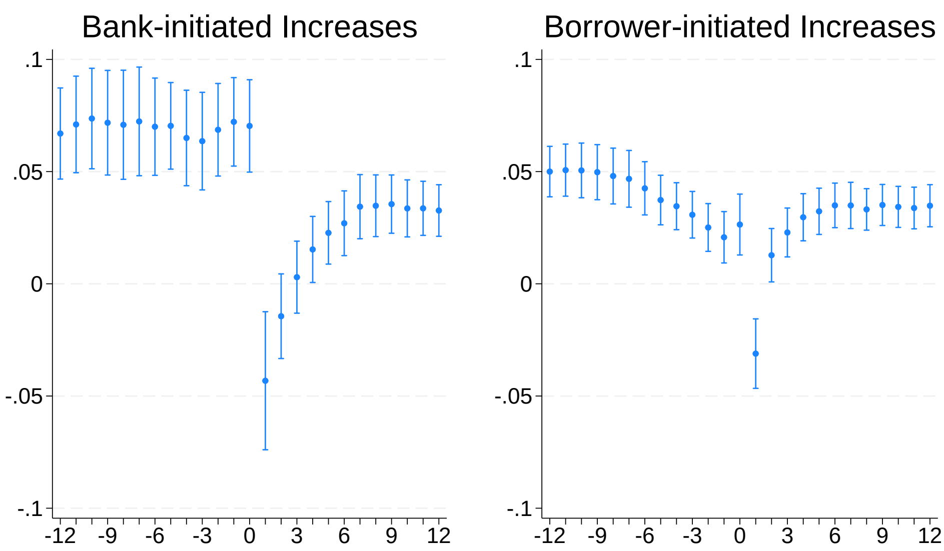

Figure 32: Event study: utilization around limit increases

Furthermore, figure 33 repeats the analysis of figure 13 limiting the accounts that undergo an increase to those that, as of the month before the increase, have revolving utilization below the mean of 27% for bank-initiated increases and 22% for consumer-initiated increases. The trends in figure 33 are very similar to those of figure 13, suggesting that the rise in revolving balances after the increase is not driven purely by constrained accounts for whom the limit is a binding constraint.

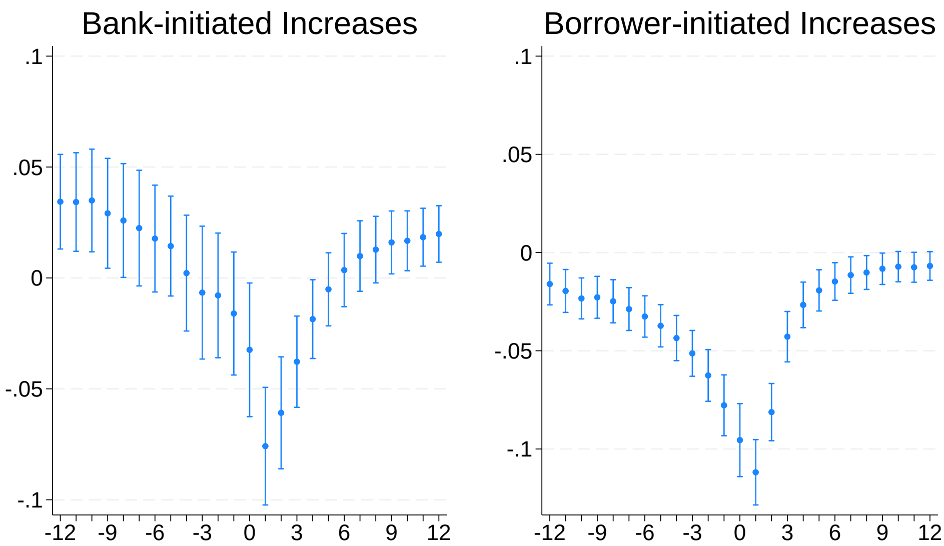

Figure 33: Event study: revolving utilization around limit increases.

This figure limits the sample to accounts with revolving utilization below the mean at time 0.

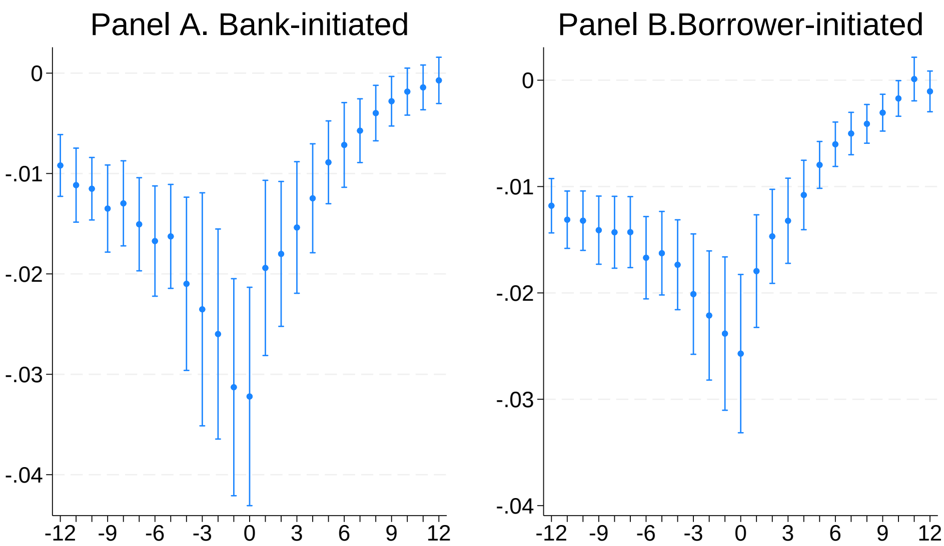

Figure 34: Event study: 30-day delinquency around limit increases

Finally, we repeat the analysis of figure 13 using as the dependent variable whether the account is at least 30 days past due. Figure 34 shows that although accounts that undergo bank- and consumer-initiated limit increases tend to have lower than average delinquency rates–as evidenced by the negative coefficients on the event-time indicators prior to time 0–delinquency rates rise after the limit increase.

B Model Appendix

B.1 Model details

B.1.1 Taxes

Household pre-tax income ($y_{i,t}$) during the working-life is taxed using a progressive tax schedule $\tau(\cdot)$. We follow Keane and Wasi (2016) and assume a nonlinear tax function:

$$\tau(y_{i,t}, a_{i,t}) = e^{\tau_1 + \tau_2 \log(y_{i,t} - \tau_d)} \tag{16}$$

where the parameters τ_1 and τ_2 determine the progressivity of the aggregate tax schedule. These parameters are estimated based on income and tax data from the Current Population Survey, therefore $\tau(y_{i,t})$ represents the sum of federal, state, and municipal taxes, plus mandatory social security contributions.

In addition, τ_d represents the deduction which is subtracted from income before the tax is applied. We define τ_d to be the standard deduction τ_d^{standard}. This results in an after-tax household income given by $\widetilde{y}_{i,t} = y_{i,t} - \tau(y_{i,t})$.

B.1.2 Social security

Following retirement ($\forall\ t > W$), households receive two sources of income: progressive social security income and annuitized disbursements from an individual retirement account:

$$y_{i,t} = y_{i,t}^{SS} + y_{i,t}^{IRA} \quad \forall\, t > W \tag{17}$$

The progressive social security-style pension is determined by the following rule:

$$y_{i,t}^{SS} = \max\Big\{\text{SS Income Floor, Annual PIA}(y_{i,W})\Big\} \tag{18}$$

where Annual PIA($y_{i,W}$) is the annual social security benefit (the primary insurance amount) received upon retirement, based on average indexed monthly earnings (AIME), which we approximate based on the last working period income, $y_{i,W}$.[45] We calibrate the social security income floor and primary insurance amount based on U.S. legislation from 2015.[46]

[45] In reality, to calculate AIME, the worker's wage during the years of employment is first expressed in today's dollars, then the wages of the highest 35 years are summed up. This sum is then divided by 420 (12*35) in order to get the real average monthly earnings.

[46] The PIA is a piecewise linear function with two break points: 90% of AIME up to breakpoint 1, 32% of AIME up to breakpoint 2, and 15% of AIME up to the social security wage base.

B.1.3 Retirement accounts

Households also accumulate wealth for retirement via individual retirement accounts. We assume that retirement contributions are mandatory and must be converted to an annuity at age W. During each year of retirement, households receive annuitized disbursements from the individual retirement accounts that depend on the actuarially fair annuity rate η_1 and the size of their account at retirement $\text{IRA}_{i,W+1}$:

$$y_{i,t}^{IRA} = \eta_1 * \text{IRA}_{i,W+1} \tag{19}$$

The value of these disbursements depend on the age of retirement, life expectancy, and the household's income during its final working period. The actuarially fair annuity rate for a household that purchases an annuity at age j, given the survival probability s_t, is:

$$\eta_1 = \left[\sum_{t=j}^{T} \frac{s_t}{(1+r)^{t-j}} \right]^{-1} \tag{20}$$

This annuity rate is actuarially fair as it enables the purchase of a guaranteed income stream until death, where the price of the annuity is equal to its expected discounted value, conditional on the interest rate and survival probabilities. We require all households to purchase an annuity in the first year of retirement using the entirety of their retirement account.

We assume that the size of the retirement account is a linear function of the household's last working period income.[47] This simplifying assumption allows us to include retirement accounts without the introduction of an additional state variable. The size of the retirement account is given by the following simple formula:

$$\text{IRA}_{i,W+1} = \eta_2 * y_{i,W} \tag{21}$$

The relationship between last period income and the size of retirement accounts (η_2) can then be estimated using PSID data.

B.1.4 Unemployment

Although unemployment benefits in the United States depend on past income, they are capped at a low level, therefore the dependence on past income is relatively low compared to other countries. We calibrate the unemployment benefit to be \$11,270 per year, which is the maximum benefit averaged across states reported in Hsu et al. (2018), adjusted to

[47] In reality, few countries have compulsory retirement accounts. One notable exception is Singapore where working age households are required to put at least 20% of their income into the Central Provident Fund each year. Agarwal, Pan and Qian (2019) provide evidence that this has an important impact on consumption, evidence that is consistent with the view that households suffer from temptation.

2015 dollars. To compute the employment transition probabilities, we use data from the PSID. We restrict our sample to household heads aged 50 or below, in order to avoid the potential effect of early retirement. As our data is biannual, we first compute the two-year probability of transitioning from employment to unemployment or out of the labor force. We then convert this to an annual probability (π_u). We use the same procedure to estimate the probability of re-employment (π_{re}).

B.2 The recursive problem of households

The household's budget sets differ based on credit history. For a household with bad credit history, the budget set is:

$$\mathcal{B}_t^{h=1} = \left\{a_{t+1} \in \mathbb{R}^+ : c_t + a_{t+1} = \tilde{y}_t + (1 + r^a)a_t, c_t \geq 0\right\} \tag{22}$$

For a household with good credit history, the budget set is:

$$\mathcal{B}_t^{h=0} = \{a_{t+1} \in \mathbb{R} : c_t + a_{t+1} = \tilde{y}_t + (1 + r^a)a_t, c_t \geq 0\} \tag{23}$$

These budget sets define the household's choice set and determine the most tempting consumption alternative each period. Households are always tempted to spend all available resources. Notice, households with bad credit history are excluded from the credit market, as reflected in their constrained budget set.

The problem's state variables are (a_t, z_t, u_t, h_t), where a_t represents asset position, z_t is the persistent labor income shock, u_t is unemployment status and h_t is credit history. For households with good credit history ($h = 0$), the value function involves choosing whether to default:

$$V_t(a_t, z_t, u_t, h_t = 0) = \max\left\{V_{\text{non},t}(a_t, z_t, u_t, h_t = 0), V_{\text{def},t}(a_t, z_t, u_t, h_t = 0)\right\} \tag{24}$$

where $V_{\text{non},t}$, and $V_{\text{def},t}$ are the value functions conditional on not defaulting and defaulting, respectively. The Bellman equation for a household with good credit history ($h_t = 0$), conditional on not defaulting, is:

$$V_{\text{non},t}(a_t, z_t, u_t, h_t = 0) = \begin{cases} -\infty & \text{if } \mathcal{B}_t^{h=0} = 0 \\ \max_{a_{t+1} \in \mathcal{B}_t^{h=0}} \left\{U_t + \beta \mathbb{E}_t V_{t+1}(a_{t+1}, z_{t+1}, u_{t+1}, h_{t+1} = 0)\right\} & \text{if } \mathcal{B}_t^{h=0} \neq 0 \end{cases} \tag{25}$$

With the budget constraint:

$$c_t = (1 + r^a)a_t + \tilde{y}_t(z_t) - a_{t+1} \tag{26}$$

, where $\mathbb{E}$ is a expectation operator, taken with respect to (z_{t+1}, u_{t+1}). The first case in equation (25) handles credit constrained scenarios, where an empty budget set leads to involuntarily default.

For households with good credit history choosing to default, the Bellman equation is:

$$V_{\text{def},t}(a_t, z_t, u_t, h_t = 0) = U_t + \beta \mathbb{E}_t V_{t+1}(0, z_{t+1}, u_{t+1}, h_{t+1} = 1) \tag{27}$$

Subject to:

$$c_t = \max\{\tilde{y}_t(z_t)(1 - \eta) - \xi, \bar{c}\} \tag{28}$$

Key differences in the default scenario include wiping out existing credit card debt (a_t), paying a default cost (ξ), setting savings to zero $(a_{t+1} = 0)$, and introducing a consumption floor $(\bar{c})$.

For households with bad credit history $(h_t = 1)$, the value function is:

$$V_t(a_t, z_t, u_t, h_t = 1) = \begin{cases} V_{\text{def},t}(a_t, z_t, u_t, h_t = 1) & \text{if } \mathcal{B}_t^{h=1} = 0 \\ \max_{a_{t+1} \in \mathcal{B}_t^{h=1}} \left\{ U_t + \beta \mathbb{E}_t V_{t+1}(a_{t+1}, z_{t+1}, u_{t+1}, h_{t+1} = \tilde{h}) \right\} & \text{if } \mathcal{B}_t^{h=1} \neq 0 \end{cases} \tag{29}$$

subject to the constraint seen in equation (26):

$$c_t = (1 + r^a)a_t + \tilde{y}_t(z_t) - a_{t+1}$$

The Bellman equation for a household with bad credit history is the same as the Bellman equation for for a household with good credit history, conditional on defaulting, and is is defined as:

$$V_{\text{def},t}(a_t, z_t, u_t, h_t = 1) = U_t + \beta \mathbb{E}_t V_{t+1}(0, z_{t+1}, u_{t+1}, h_{t+1} = 1) \tag{30}$$

$\mathbb{E}$ is the expectation operator, taken now with respect to $(z_{t+1}, u_{t+1}, \tilde{h})$. Notice that households with bad credit history can only default involuntarily when their budget set is empty. Bad credit history is stochasically wiped out with probability π_0^h and remains with probability π_1^h.

B.3 Parameters set outside the model

Table 10 shows the parameters that we calibrate externally. The estimates for the income process come from Kovacs and Moran (2021), who estimate the income process parameters using data from the Panel Study of Income Dynamics between 1999 and 2015.

Parameter	Symbol	Value	Source
Income Persistence	ρ	0.97	PSID 1999-2015
Std Dev Income Shocks	σ_ε	0.180	PSID 1999-2015
Initial Income	σ_0	0.410	PSID 1999-2015
Income Constant	d_0	8.2007	PSID 1999-2015
Income *Age* Effect	d_1	0.1378	PSID 1999-2015
Income Age^2 Effect	d_2	-0.0019	PSID 1999-2015
Income Age^3 Effect	d_3	0.000007	PSID 1999-2015
Unemployment probability	π_u	0.053	PSID 1999-2015
Re-employment probability	π_{re}	0.397	PSID 1999-2015
Unemployment benefit	b	\$11,270	Hsu et al. (2018)
Liquid asset return	r	0.02	Annual risk free interest rate
Credit card rate	r^C	0.14135	Fulford (2024)
Credit recovery probability	π_h^0	0.10	Nakajima (2017)
Cost of a bankruptcy filing	ξ	\$1,117	Nakajima (2017)
Income garnishment ratio	η	0.335	Nakajima (2017)
Share with zero initial assets	a_0^{zero}	0.433	PSID 1999-2015
Cond. mean initial assets	μ_{a_0}	7.117	PSID 1999-2015
Cond. std dev initial assets	σ_{a_0}	1.972	PSID 1999-2015

Table 10: Parameters set outside the model

Table 11 shows the empirical estimates of the linear probability model, which we then feed into equation (13) of the quantitative model. The coefficients are estimated from a yearly version of equation (1) with an indicator for a limit increase during the calendar year regressed on deciles of average utilization, the log of the credit limit at the start of the year, and bins for income at the start of the year. The regression is estimated without the fixed effects used in Table A.3.1, although our results are little changed when we use alternative controls.

B.4 Sensitivity of the standard share to targeted moments

Figure 35 presents the sensitivity of our estimated parameter ϕ_{standard} to changes in our targeted moments using the sensitivity matrix derived by Andrews et al. (2017). Following their approach, the sensitivity matrix captures how the estimated parameter value varies due to small changes in the targeted moments. The four targeted moments in our estimation are: the share of households with revolving credit card debt (revolver_share), the average utilization rate of credit cards (util_rate), the default probability (default_prob),

Coefficient	Estimate
Lowest Utilization Decile	0.036
Utilization Decile 2	0.110
Utilization Decile 3	0.141
Utilization Decile 4	0.151
Utilization Decile 5	0.153
Utilization Decile 6	0.148
Utilization Decile 7	0.138
Utilization Decile 8	0.124
Utilization Decile 9	0.103
Top Utilization Decile	0.049
Log Limit	-0.006
Income: > \$150,000	0.002
Income: \$75,000 - \$150,00	0.006
Income: \$40,000 - \$75,000	0.011
Constant	0.06

Table 11: Coefficients for the firm's policy function

and the average debt to income ratio (mean_debt_y). We normalize the sensitivity measure so that each point represents the level change in parameter ϕ_{standard} to a 1 percent increase in the respective targeted moment.

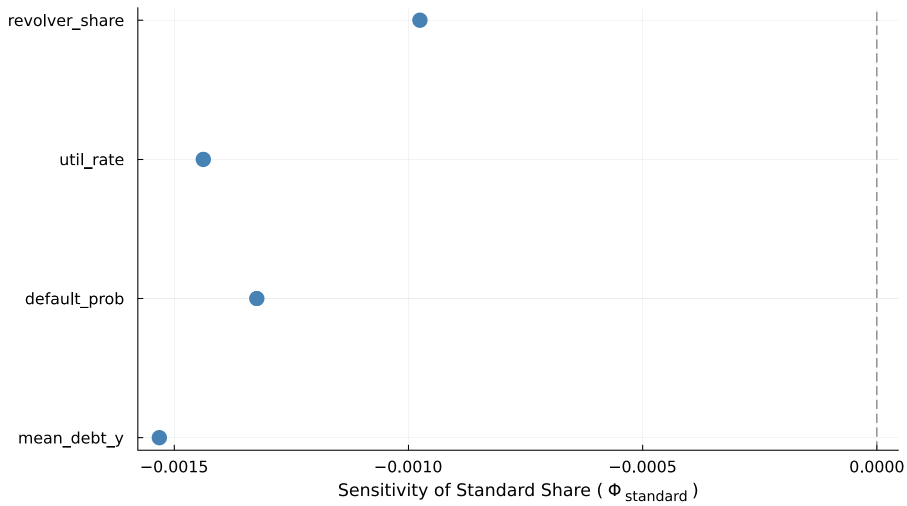

Figure 35: Sensitivity of the standard share to changes in targeted moments

The results show that all targeted moments play a role in pinning down the share of standard consumers, with the same sign and slightly varying magnitudes. Specifically, increases in any of the four targeted moments lead to decreases in the estimated share of standard households, which aligns with our theoretical expectations. For instance, when the share of households with revolving credit card debt increases by 1 percent (from 0.45 to 0.4545), the estimated share of standard households falls by approximately 0.001 (from 0.6 to 0.599). This negative relationship reflects the intuitive result that higher levels of revolving credit card usage in the data require a larger share of tempted households in the model to match these patterns.

B.5 Additional model figures

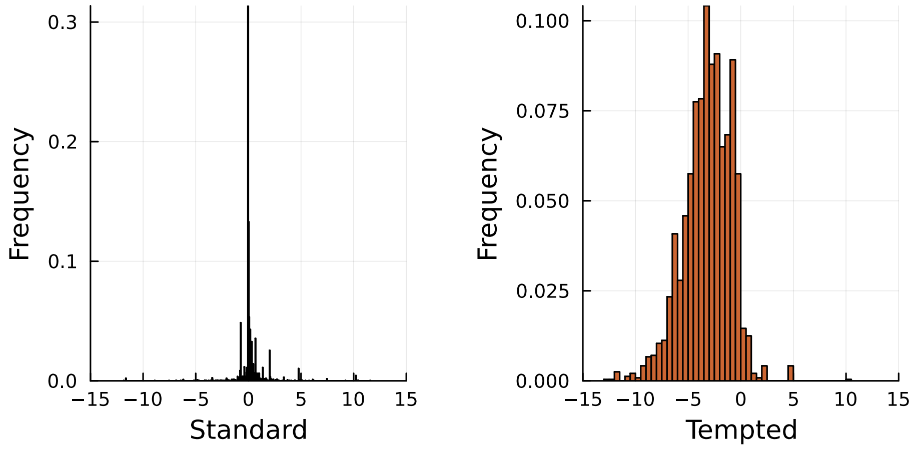

Figure 36: Consumption Equivalent Variation (%)

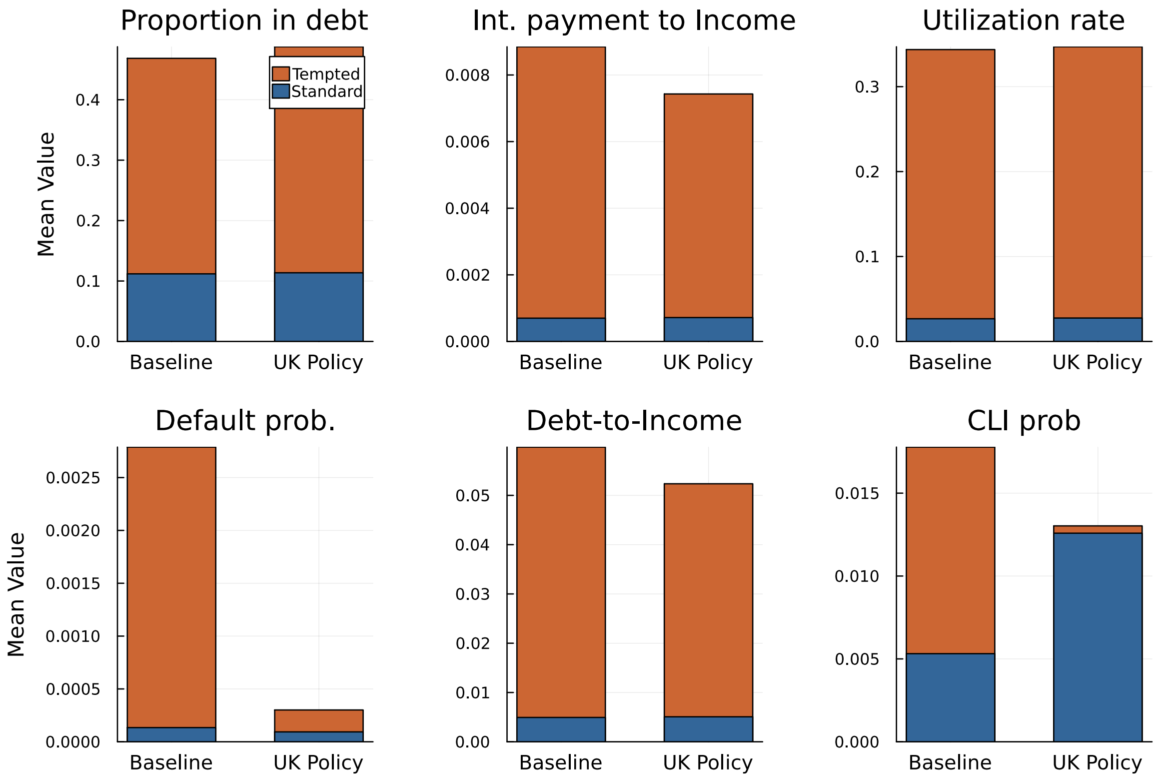

Figure 37: Behavior by type with firm re-optimization, no CLI to revolving borrower

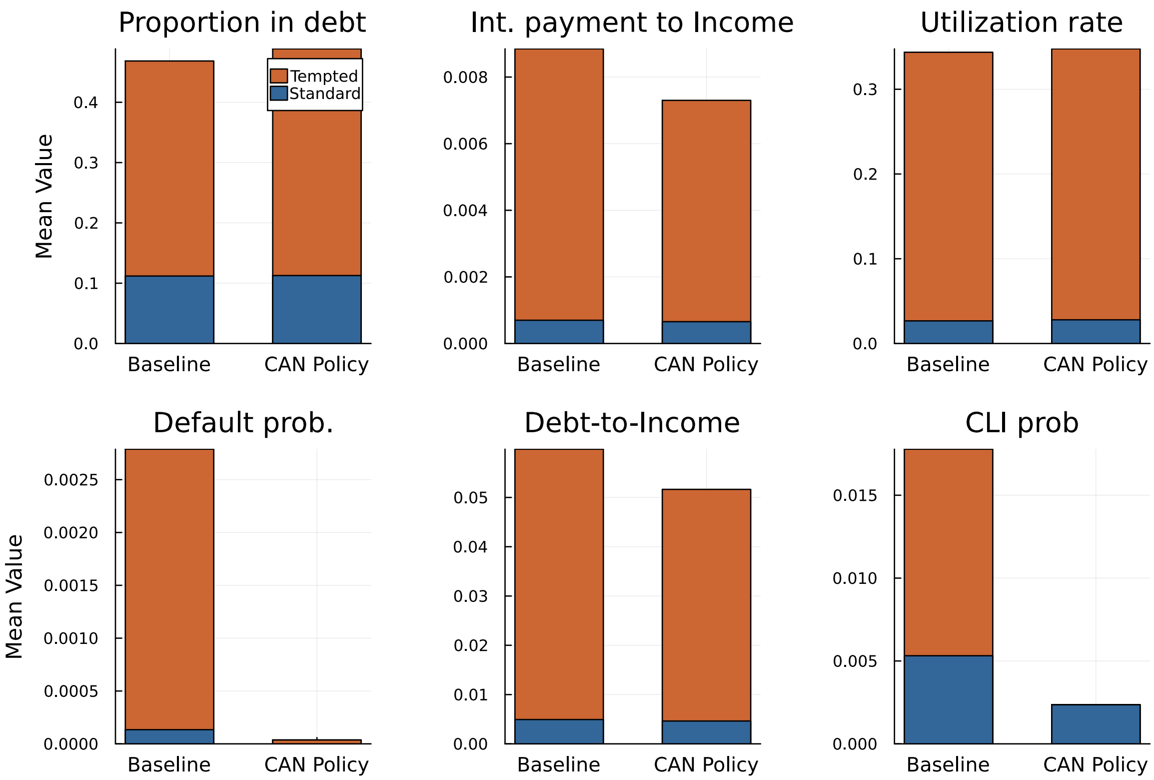

Figure 38: Behavior by type with firm re-optimization, no CLI without borrower's consent