The Impact of Social Insurance on Household Debt *

Gideon Bornstein† Sasha Indarte‡

March 24, 2026

Abstract

This paper investigates how the expansion of social insurance affects households' accumulation of debt. Insurance can reduce reliance on debt by lessening the financial impact of adverse events such as illness and job loss. But it can also weaken the motive to self-insure through savings, and households' improved financial resilience can increase access to credit. Using data on 10 million people and a quasi-experimental research design, we estimate the causal effect of expanded insurance on household debt, exploiting ZIP-code level heterogeneity in exposure to the staggered expansions of one of the largest US social insurance programs: Medicaid. We find that a 1 percentage point increase in a ZIP code's Medicaid-eligible population increases credit card borrowing by 0.46%. Decomposing this effect in a model of household borrowing, we show that increased credit supply in response to households' improved financial resilience drives the rise in borrowing and contributed 32% of the net welfare gains of expanding Medicaid.

*First version: December 15, 2019. This version: March 24, 2026. For helpful comments and suggestions, we thank Nathan Blascak (discussant), Corina Boar (discussant), Claire Célérier (discussant), Taha Choukhmane (discussant), Dean Corbae, Marty Eichenbaum, Simon Freyaldenhoven, Paul Goldsmith-Pinkham (discussant), Dan Grodzicki (discussant), Kyle Herkenhoff, Ben Keys, Olivia Mitchell, Kurt Mitman, Arna Olafsson (discussant), Gordon Phillips, and Giorgia Piacentino (discussant), as well as conference and seminar participants at LSE, Imperial, Penn-Wharton brown bag, SED, EM3C, Rochester, INSPER, Minnesota, SFS Cavalcades, Columbia, UVA, EAGLS, Bank of Canada, UIUC Gies, Chicago Booth, Columbia GSB, RUNI, Wisconsin, Berkeley, Duke, NYU, NBER Micro Data and Macro Models, MIT Sloan, Goethe University Frankfurt, Bonn, Northwestern, and UChicago. We are grateful to Michael Boutros, Joyce Chen, Tanvi Jindal, Zoe Perez, and Paige Stevenson, who provided excellent research assistance. We gratefully acknowledge financial support for this project from the Rodney L. White Center for Financial Research.

†Wharton. Email: **gideonbo@wharton.upenn.edu**.

‡Wharton. Email: **aindarte@wharton.upenn.edu**.

1 Introduction

Unsecured debt is an important tool that households rely on to smooth their consumption—not only across time but also across states of the world. For example, of the 4 in 10 US adults that report anticipating difficulty in meeting an unexpected $400 expense, credit cards are the most-cited tool that they expect to rely on.[1] And when experiencing an income shortfall, 43% of US households report turning to borrowing, including credit cards.[2] When unsecured debt functions as a de facto source of insurance, changes to social insurance could significantly affect households' desire and ability to access credit.

Whether social insurance crowds in or crowds out household debt is theoretically ambiguous. The total effect depends on the strength of competing channels. On the one hand, improved access to social insurance can reduce households' reliance on debt to cope with adverse shocks such as job loss and illness, crowding out debt. On the other hand, credit demand and supply channels may also work in the opposite direction. Social insurance can increase the demand for debt by weakening the precautionary savings motive. Additionally, when social insurance makes households less likely to default, this reduction in risk can make creditors more willing to lend. Expanded insurance can improve households' financial resilience, which can in turn increase access to credit.

This paper sheds new light on how credit markets shape the impact of social insurance on household debt and welfare. We first estimate the causal effect on credit card borrowing of the expansion of Medicaid under the Affordable Care Act (ACA), one of the largest changes to the US social safety net in recent years. We find that a 1 percentage point increase in the share of a ZIP code's population eligible for Medicaid increases revolving credit card balances by 0.46 percentage points. Next, we develop a heterogeneous-agent model in which households face both income and medical expenditure shocks while having the ability to borrow and default on unsecured debt. Simulating the expansion of Medicaid in the model, we find that increased credit supply is the main force behind the overall increase in household debt. Moreover, the credit supply response to expanding Medicaid accounts for 32% of the net welfare gain.[3] Our

[1] Source: 2019 SHED.

[2] Source: 2016 SCF.

[3] Our measure of net welfare reflects improvements stemming from the financial benefits of the Medicaid ex-

findings demonstrate that expanding social insurance can *crowd in* the supply of unsecured debt and that this improvement in credit access can lead to first-order welfare gains. We conclude that policymakers may significantly underestimate the benefits of social insurance programs if they overlook potential increases in credit access spurred by improved financial resilience.

Our focus on health insurance is motivated by the importance of Medical expenses as a source of financial distress for US households. Unexpected medical bills recently ranked first among expenses that US residents report worrying most about affording, ahead of rent or mortgages, transportation, and food.[4] Among bankruptcy filers, 29% named medical bills as the reason for their bankruptcy (Himmelstein, Thorne, Warren and Woolhandler, 2009). Health expenditures also constitute a large and growing share of US GDP, reaching 18.3% in 2021.[5] Despite the importance of health expenses, they have received less attention in the macroeconomics literature compared to other sources of risk such as job loss and housing.[6]

We begin by estimating the causal effect of Medicaid eligibility on a variety of credit outcomes using credit bureau data on 10 million US borrowers over 2010-2021 in a continuous difference-in-difference (DID) analysis. We exploit granular variation in the *size* of eligibility changes induced by states' staggered expansions of Medicaid under the ACA. This variation arises from two sources. The first comes from states' pre-expansion Medicaid policies. Expanding Medicaid under the ACA required states to implement a common set of eligibility requirements. As a result, ZIP codes with similar distributions of income saw larger increases in eligibility if their state previously had stricter requirements. The second source of variation comes from within-ZIP differences in the distribution of income among low-income households. Eligibility increased more in ZIP codes that had more low-income households clustered between the pre- and post-expansion income eligibility thresholds.

Our DID estimator compares credit outcomes before and after expanding across ZIP codes with bigger versus smaller changes in eligibility. The within-state variation in the im-

pansion (net of the tax costs of expansion). We do not model either mortality or utility as a function of health. Therefore, our welfare estimates do not reflect gains from improved health or reduced mortality.

[4]Source: Kaiser Family Foundation Health Tracking Poll (February 2020). Link: https://www.kff.org/health-costs/poll-finding/data-note-public-worries-about-and-experience-with-surprise-medical-bills/.

[5]Source: Centers for Medicare and Medicaid Services NHE Fact Sheet.

[6]Notable exceptions include Livshits, MacGee and Tertilt (2007); Chatterjee, Corbae, Nakajima and Ríos-Rull (2007); Attanasio, Kitao and Violante (2010); De Nardi, French and Jones (2010).

pact of expansions makes it possible to use county-time fixed effects in our analysis, which help account for other unobserved state- or county-level policy changes that may have coincided with expansions and also affected credit outcomes. Identification requires that the change in eligibility is not correlated with other shocks coinciding with the expansions.

We find that expanded Medicaid eligibility increases credit card borrowing. A 1 percentage point increase in eligibility leads to a 0.11 percentage point increase in the share of households with at least one credit card and raises revolving (i.e., unpaid) credit card balances by 0.46 percentage points. Consistent with a positive credit supply response, we document that expanded eligibility also increases credit limits and overall reduces credit card utilization. On the demand side, we find that the number of credit card inquiries rises. In line with our financial resilience hypothesis, we show that expanding Medicaid also reduces severe forms of default (debt in collections), minor forms of default (delinquency) either decrease or remain stable, and credit scores increase.

We then construct a heterogeneous-agent model in which households have access to credit card debt. Households face idiosyncratic income shocks as well as idiosyncratic expenditure shocks. They can save using a risk-free asset or borrow via credit card debt, which they can decide not to repay. Households incur debt both by choosing credit card borrowing and as a result of experiencing expenditure shocks that they are unwilling or unable to cover on impact. The interest rate paid on credit card debt is endogenous and depends on the probability households do not repay their debt.

The nature of credit card debt in our model is a hybrid of one-period and long-term debt used in the literature on consumer bankruptcy and sovereign debt.[7] When households are not delinquent, they must roll over their debt that period. That is, credit card debt must be repaid in full to avoid delinquency. However, households have the option not to repay their debt and enter a delinquent state. In that state, they cannot take on more debt before paying their current debt in full.[8] Households in a delinquent state get a stochastic haircut on their debt, which can help them exit the delinquency state. Since financial intermediaries

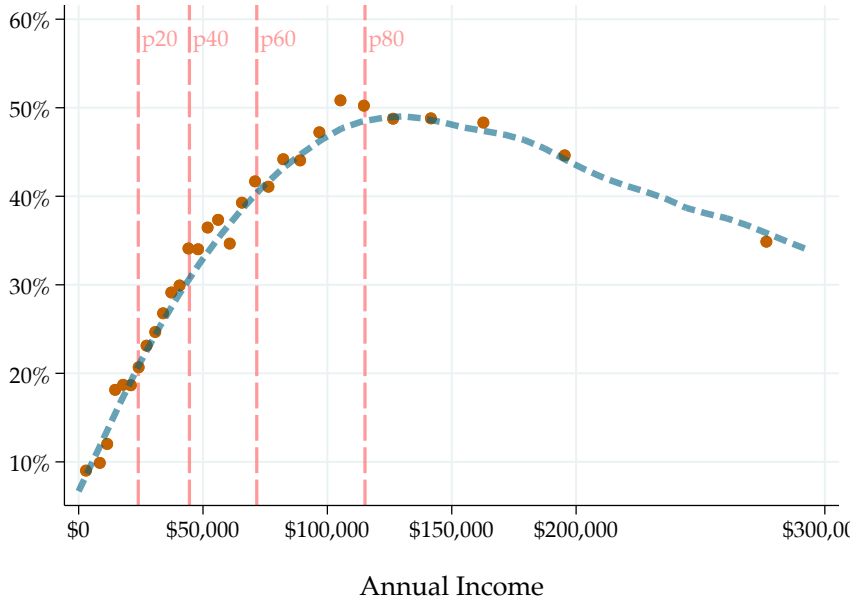
[7]Notable examples of such models include Chatterjee et al. (2007); Livshits et al. (2007); Arellano (2008) for short-term debt and Hatchondo and Martinez (2009); Chatterjee and Eyigungor (2012) for long-term debt.

[8]While households cannot choose to take on more debt in the delinquent state, uninsured expenditure shocks add to households' total debt.

who hold claims on delinquent debt are not repaid immediately, the pricing of debt has a long-term component to it.

The delinquency option on debt allows us to capture a key aspect of the data. The relationship between having unpaid/revolving credit card balances and income follows an inverse U-shape, as displayed in Figure 1. Less than 25% of households with annual income below $25,000 have any credit card debt. This share rises to 50% for households with an annual income of $100,000 and declines as household income rises. Delinquency, together with endogenous credit supply, restricts low-income households' access to credit.

Figure 1: Share of Households with Credit Card Debt



Notes: This figure plots a binscatter of the share of households with a non-zero amount of unpaid (i.e., revolving) credit card balances. Data come from the PSID (2011–2019). Annual income is in real 2019 USD.

Capturing the inverse U-shape relationship between income and credit card debt is important for studying policies that target low-income households, such as Medicaid. This is because means-tested social insurance policies generally target households whose default risk limits their credit access. By reducing this risk, social insurance can lead to better credit access. We model Medicaid as a policy that covers a fraction of households' expenditure shocks, where eligibility for Medicaid depends on household income.

Health insurance policies affect the aggregate level of credit card debt through three chan-

nels. First, the direct effect of more generous health insurance is increasing households' disposable income. Households can achieve the same consumption levels while borrowing less. However, higher disposable income also reduces default, which can raise overall debt as less is being discharged. A priori, the direct effect of health insurance on credit card debt is ambiguous.

The second channel is through credit demand. Health insurance affects demand for credit even if credit terms remain unchanged. Lower medical expenditures reduce households' precautionary savings motive and, as a result, increase their borrowing. However, households are more likely to repay their debt in the future. This results in a higher level of *expected* repayment, which discourages borrowing. These competing forces mean that the effect of the credit demand channel is also theoretically ambiguous.

The third channel is the credit supply channel. The reduction in delinquency rates leads to lower interest rate spreads in equilibrium. Lower interest rates induce households to take on more credit, leading to an increase in the aggregate level of credit card debt.

While we ground our model in the context of health insurance, the economic channels we study can arise in many other contexts. Our model and approach can be readily adapted to study the impact of other types of social insurance (unemployment insurance, minimum wages, disability insurance, etc.). Our model could also be adapted to study how changes to insurance/risk exposure affect sovereign and corporate credit outcomes. Quantitatively, one could find that expanded insurance either increases or decreases borrowing, as the channels can work in opposing directions, and which dominates in equilibrium may vary with context.

Our model abstracts away from several features of health— notably, "physical" health that enters the utility function, price-sensitive demand for health services ("moral hazard"), and correlations between income and health expenditures. As we discuss in detail in Section 3, this likely makes our counterfactual analysis conservative regarding the welfare and borrowing responses.

To estimate our model, we use data on credit card usage, health expenditures and insurance, and our reduced-form estimate of the causal effect of the expansion on revolving credit card debt. The model predicts that expanding Medicaid led to a 4.1% overall increase in credit card debt. We find that the direct channel has a small impact on the aggregate level of credit

card debt and that the credit demand channel reduces the aggregate level of credit card debt substantially. Together, these two channels decrease credit card debt by 10.2%. The overall increase in credit card debt is primarily driven by the credit supply channel, which increases aggregate credit card debt by 14.3%.

We use our model to study the welfare benefits associated with the expansion of Medicaid. We find that the policy is equivalent to a permanent increase in consumption of 41 basis points. Of the welfare gains, 32% is a result of the reduction in interest rate spreads households face on their debt. This result suggests policymakers should take into account the effect social insurance has on credit supply. Disregarding the credit supply channel could substantially underestimate the welfare gains of social insurance programs.

Related Literature: This paper contributes to several strands of literature by bringing a new macroeconomic perspective to the effects of social insurance. First, our model builds on the macroeconomic literature on consumer bankruptcy and default (e.g., Chatterjee, Corbae, Nakajima and Ríos-Rull, 2007; Livshits, MacGee and Tertilt, 2007; Mitman, 2016). This literature focuses on the drivers of default—in particular, the role of bankruptcy policy. We study a policy that does not directly target bankruptcy or default: public health insurance. This policy can reduce default and increase both credit access and borrowing.

Second, we add to a large macroeconomic literature on heterogeneous agent models with uninsurable risk (Bewley, 1986; Aiyagari, 1994; Huggett, 1993). This literature highlights how precautionary savings play a key role in shaping the macroeconomy. Our work studies how insurance provision can reduce the precautionary savings motive and traces the macroeconomic implications. We take a partial equilibrium approach—that is, the risk-free rate in the economy is taken as given. In this sense, our approach is similar to the work of Imrohoroğlu (1989), Zeldes (1989), Deaton (1991), and Hubbard, Skinner and Zeldes (1995), who also study how social insurance affects the precautionary savings motive.

We build on these works by taking into account the impact social insurance has on the endogenous debt pricing schedule. We also build on Krueger and Perri (2011). Their model features contingent contracts, resulting in no default occurring in equilibrium and credit supply *decreasing* when social insurance expands (financial autarky becomes less costly and thus

default more tempting). Our model instead features non-contingent defaultable debt and leads to a positive credit supply response, which is consistent with our empirical estimates. Additionally, our analysis of the impact of social insurance on credit markets complements concurrent work studying the "opposite" question: the impact of credit access on the design of optimal UI (Braxton, Herkenhoff and Phillips, 2022).

Third, we build on an empirical microeconomic literature studying the consumer finance consequences of expanding insurance. We present new evidence on revolving credit card debt and, using our model, quantify the welfare effects and decompose the underlying channels. Prior work on Medicaid finds that expansions reduced bankruptcy, medical debt, missed debt and bill payments, reliance on payday loans, and debt in collections (Gross and Notowidigdo, 2011; Allen, Swanson, Wang and Gross, 2017; Mazumder and Miller, 2016; Miller, Hu, Kaestner, Mazumder and Wong, 2018; Gallagher, Gopalan and Grinstein-Weiss, 2019b; Goldsmith-Pinkham, Pinkovskiy and Wallace, 2020). Consistent with credit supply playing an important role, Brevoort, Grodzicki and Hackmann (2017) finds that expanded Medicaid eligibility reduced interest rates on credit card offers and Célérier and Matray (2017) finds that mortgage application approval rates increased. Our focus on borrowing complements Gallagher, Gopalan and Grinstein-Weiss (2019a), which finds that health insurance reduces savings among households experiencing financial hardship. Less saving does not immediately imply higher gross or net borrowing.[9] Our work also builds on Hu, Kaestner, Mazumder, Miller and Wong (2018), which estimates null effects of the Medicaid eligibility on credit card balances in the first following expansion; we study *unpaid* credit card debt, use ZIP-level variation in eligibility, and find that there are positive effects on unpaid balances emerging two years post-expansion. Our empirical analysis also builds on Finkelstein et al. (2012), which finds lower rates of either borrowing or skipping bills to pay for medical expenses among participants enrolled in Medicaid via the Oregon health experiment. Our analysis differs by studying all credit card borrowing (rather than only borrowing to pay for medical expenses) and by measuring borrowing with administrative credit bureau data rather than surveys.

[9] In contrast to standard models of consumption and saving, household borrowing presents a "credit card puzzle" in that households tend to hold both high-interest credit card debt and low-interest savings simultaneously Gross and Souleles (2002). This behavior is consistent with a motive to maintain a liquidity buffer in the presence of incomplete markets (Telyukova, 2013; Druedahl and Jørgensen, 2018).

Research on unemployment insurance also finds evidence of improved financial resilience (lower mortgage default) and expanded credit access in the form of lower interest rates and higher application approval rates (Hsu, Matsa and Melzer, 2018; Matsa, Melzer and Zator, 2022). Our paper also relates to Aaronson, Agarwal and French (2012), which documents strong and positive spending and collateralized debt increases among minimum wage workers following rises in the minimum wage. We build on their findings by documenting that revolving credit card debt also rises after the expansions of another part of the social safety net (Medicaid). We show that this rise is accompanied by a decline in default, improvements in credit scores, and increases in measures of credit supply (e.g., credit card approval rates).

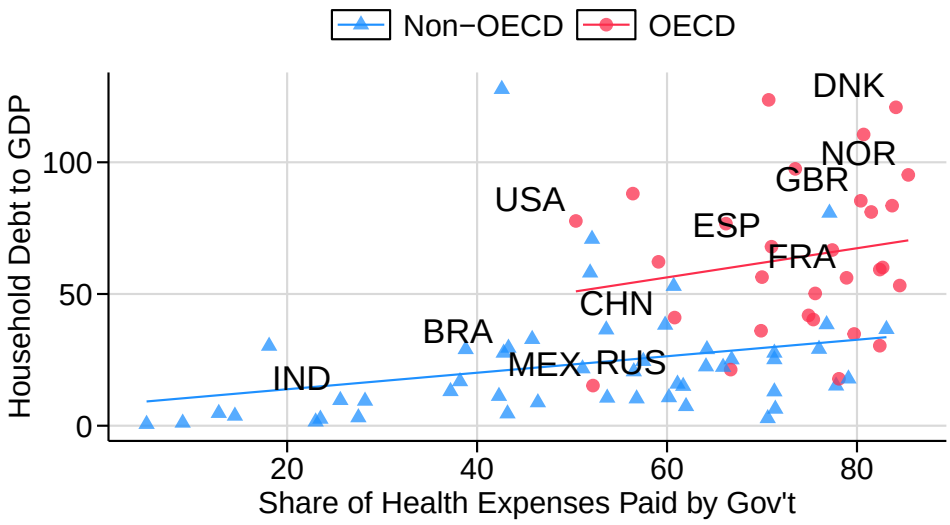
This paper is organized as follows. Section 2 presents background on the Medicaid expansions and the empirical analysis of the impact of expansions on credit card debt. Next, Section 3 presents the model. Section 4 analyzes policy counterfactuals on debt (its distribution and aggregate level) as well as welfare. We decompose the effect of expanding health insurance on borrowing into its direct impact on debt for households experiencing adverse shocks and its general equilibrium impacts through credit demand and supply. Section 5 concludes.

2 The Effect of Social Insurance on Household Debt: Evidence from Medicaid Expansions

At the macro level, cross-country comparisons show that the generosity of social insurance is *positively* related to household debt. Figure 2 plots the ratio of household debt to GDP versus the share of total health expenses paid by the government, a proxy for social insurance provision, across countries. Countries with more generous insurance have significantly higher levels of household debt. At the top right corner are the Scandinavian countries, which have both a high provision of social insurance and high levels of household debt. Viewed through the lens of our hypothesis, one reason Scandinavian households take on higher levels of debt is because they are more financially resilient—they are better insured against adverse events.

This section estimates the effect of expanded social insurance on household debt. We study one of the largest changes to the US social safety net in recent decades: the expansion

Figure 2: Household Debt and Social Insurance across Countries



Notes: Data come from the IMF Global Debt and WHO Global Health Observatory data repository.

of Medicaid under the Affordable Care Act (ACA). Our empirical strategy exploits both the staggered nature of the Medicaid expansions across states and within-state granular heterogeneity in the impact of the expansions. Examining a variety of credit outcomes, we find a positive relationship between credit card debt and Medicaid eligibility, as well as supporting evidence consistent with credit supply playing an important role in shaping the equilibrium response.

2.1 Institutional Background: Medicaid

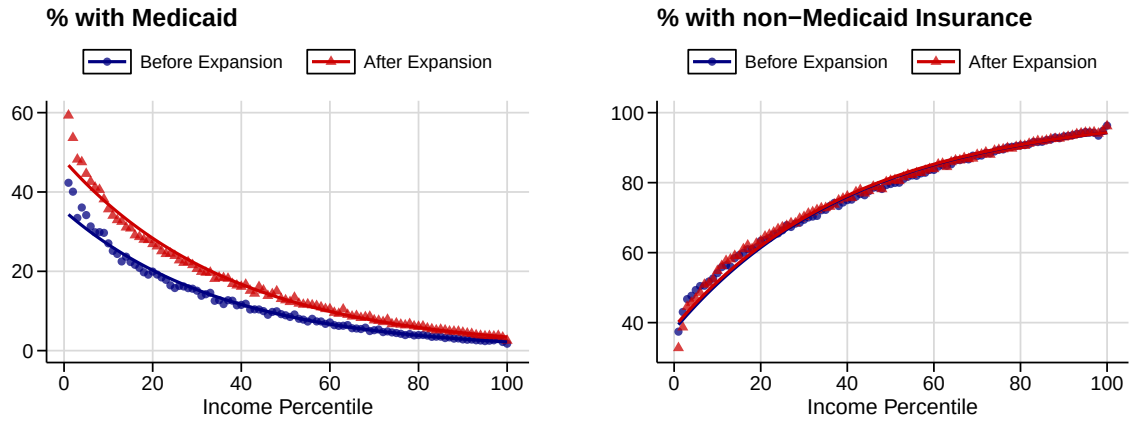
Medicaid is a program, administered jointly by state and federal governments, that offers low-income households free or low-cost health insurance. By the end of 2019, 64.3 million individuals—nearly 20% of the US population—were receiving health insurance through Medicaid. In the wake of COVID-19, the number of enrollees steadily climbed to 85.3 million as of December 2022.[10] Medicaid spending is a large fraction of total US health care expenses, totaling $597.4 billion in 2018 (or comprising 16% of aggregate health expenditures). To qualify for Medicaid, a household must have income below a specified threshold.[11]

[10] We obtain enrollment count data from the Kaiser Family Foundation: https://www.kff.org/other/state-indicator/medicaid-and-chip-monthly-enrollment.

[11] Several states also use an asset-based means test in addition to the income threshold to determine eligibility. To expand Medicaid under the ACA, states were required to remove any asset-based means tests.

The ACA expanded Medicaid in participating states by requiring these states to set the income eligibility threshold to at least 138% of the federal poverty level (FPL) for *all* adults. Participating states receive federal funds to support the costs of the Medicaid expansion. Prior to the ACA, only a handful of states offered Medicaid to adults aged 64 or under without dependents. After expanding, the uninsured population decreased 30% on average among participating states. The share of the population enrolled in Medicaid rose 4.99 percentage points in expanding states.[12] This average change in enrollment is within range of the estimated *causal* effect of the expansions on enrollments, which is 3-5 percentage points.[13] The crowding out of private and other sources of government insurance appears limited; the fraction receiving non-Medicaid sources of insurance is similar before and after the expansions. Figure 3 plots insurance type by income both before and after expanding.

Figure 3: Health Insurance Type by Income (Pre- and Post-Expansion)



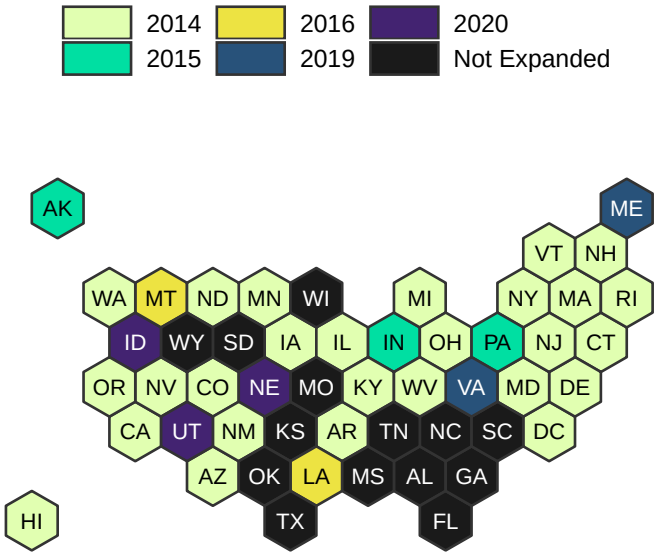
Notes: The graphs plot the fraction of people with either Medicaid (left panel) or other non-Medicaid sources of insurance (right panel), such as private insurance or Medicare. We report averages both before (blue circles) and after (red triangles) states expanded Medicaid. We include non-expanding states in the "before" group. Insurance type is measured at the end of the year, while income is measured over the entire year. Note that some "high-income" households may qualify for Medicaid if, for example, they lost their job during the year and were eligible by the end of the year. This, and aggregation across household sizes, time, and states, is likely why the share with Medicaid is smooth with respect to income (as opposed to being discontinuous at eligibility thresholds). Source: American Community Survey (ACS), 2010 to 2020.

Participating in the expansion is optional, and the timing of adoption varied significantly across time (see Figure 4). The first, and most common, year of adoption was 2014. As of

[12] These figures are calculated using American Community Survey (ACS) data from 2010 to 2020.

[13] See, for example, Sommers et al. (2014); Frean et al. (2017); Kaestner et al. (2017); Courtemanche et al. (2019).

Figure 4: Medicaid Expansion Dates



Notes: Expansion dates come from the Kaiser Family Foundation.

2021, 36 states (and D.C.) have expanded Medicaid eligibility under the ACA.

The ACA was primarily financed by cuts to federal government spending on health care, taxes on insurers, and taxes on individuals. Congressional Budget Office (CBO) estimates projected the ACA to *reduce* budget deficits overall during the period 2013–2022 by $109 billion (Congrssional Budget Office, 2012). Over this period, the ACA's projected costs to the federal government of expanding Medicaid, the Children's Health Insurance Program (CHIP), health insurance exchanges, and other non-coverage provisions totaled $1,455 billion. Health care-related government spending cuts were projected to save $741 billion over this same time period, accounting for 50.9% of the total cost. The primary source of these cuts were reductions in payment rates for hospital services rendered to Medicare and Medicare Advantage patients and reductions in Disproportionate Share Hospital (DSH) payments. DSH payments help compensate hospitals for the cost of uncompensated care (i.e., missing or partial payments owed by patients).

2.2 Data

This section describes the two main datasets used in our reduced-form analysis. The first is credit bureau data. The second dataset records Medicaid eligibility at the ZIP-code level.

Throughout, we deflate nominal variables to real 2020 dollars.

Experian Data. Our credit bureau data comes from Experian and covers a random sample of 10 million US residents. This Experian sample is an annual panel spanning the period 2010 to 2021 and contains over 90 million borrower-level observations. A subset of outcomes are measured at a quarterly frequency over 2010–2020. Our sample is geographically representative in the sense that individuals are randomly sampled from ZIP codes in proportion to their ZIP code's population share. In our empirical analysis, we aggregate the borrower-level data to the ZIP-code level. We drop ZIP codes with fewer than 150 borrowers reporting data and ZIP codes that cannot be matched to our Medicaid eligibility data (these are mainly small ZIP codes). After these restrictions, we have 12,511 unique ZIP codes over 12 years.

Experian Summary Statistics. Table 1 presents summary statistics for our main credit outcomes of interest in the ZIP-code level sample. On average, 85% of borrowers in a ZIP have at least one credit card, and 22% obtained a new credit card in the past year. Average borrower-level credit card balances are $4,299. We measure revolving (i.e., unpaid) balances as the difference between average monthly balances and average monthly payments. These monthly averages are available at a quarterly frequency in the Experian data. In our sample, borrowers carry $3,680 on average in revolving balances each month. Within this sample, 60% of the borrowers have unpaid balances at least once per quarter; the rate is 71% among credit cardholders.

Our measure of unpaid balances may be subject to measurement error that depends on when balances are measured relative to the consumer's scheduled payment date.[14] Average use of revolving credit card debt in our sample is similar to estimates from regulatory data, suggesting that measurement error is close to mean zero.[15] Importantly, this measurement error will *unlikely* bias the regression analysis. Revolving balances will be a response variable

[14] Ideally, we would subtract payments from the *statement* balance at the end of the billing cycle. Average current balances will tend to understate statement balances when measured after the end of the billing cycle and overstate it when measured shortly before.

[15] Using regulatory data, Adams, Bord and Katcher (2022) report that around 68% of active credit card accounts are used to revolve in a given year and that approximately 75% of credit card balances are revolving at a given point in time. Alternatively, the share of revolvers is 52% in the 2019 Survey of Household Economics and Decisionmaking (SHED) and 45% in the 2019 Survey of Consumer Finances (SCF). However, Zinman (2009) argues that the SCF significantly under-reports credit card debt.

in our regressions, and the timing of a consumer's billing cycle is likely unrelated to their state's Medicaid expansion.

We consider several additional variables that reflect credit access. In our sample, average credit card utilization is 31%, and average limits are $18,013. Credit limits vary significantly across ZIP codes; the standard deviation is $8,740. The average ratio of new accounts per credit card inquiry is 0.52, indicating that typically two inquiries are necessary to obtain one credit card. Credit card inquiries per year average 0.44 per borrower.

Table 1: Summary Statistics for ZIP-Code Level Credit Outcomes

	Mean	SD	25th %	50th %	75th %	N
1[Has CC] (%)	84.73	8.73	79.55	86.28	91.46	117,629
1[New CC] (%)	21.75	4.63	18.66	21.83	24.83	117,629
CC Bal. (All)	4,299.26	1,744.38	3,055.66	3,986.69	5,215.71	372,226
CC Rev. Bal.	3,680	1,448.77	2,656.52	3,436.17	4,445.41	372,226
1[Revolver] (%)	60.46	11.48	52.53	60.77	69.08	372,226
CC Util. (%)	31.27	12.11	21.87	30.03	39.38	117,629
CC Lim.	18,013.08	8,740.22	11,637.48	15,975.89	22,635.94	117,629
New CC to Inq.	0.52	0.12	0.44	0.52	0.61	117,629
CC Inq. (#)	0.44	0.14	0.34	0.43	0.53	117,629
90+ Delinq. (%)	7.15	3.22	4.81	6.82	9.06	117,629
90+ Utilities Delinq. (%)	4.46	3.59	2.10	3.43	5.71	117,629
1[Any Col.]	23.61	11.81	14.5	21.89	30.92	117,629
Col. Bal.	664.6	436.73	351.17	566.75	872.22	117,629
Col. Bal./Income (%)	1.19	1.11	0.43	0.88	1.59	117,629
Vantage Score	683.92	36.95	659.03	684.65	710.26	117,629

Notes: This table reports ZIP-code level summary statistics for the main credit outcomes of interest from the Experian data. All variables are reported annually (June of each year) except for credit card balances (total and revolving) and the percentage of people that have revolving balances ("revolvers"). The latter are reported quarterly. We calculate revolving (i.e., unpaid) balances by subtracting credit card payments from total credit card balances. Nominal variables are CPI adjusted to be in terms of 2020 dollars. Throughout, "CC" denotes "credit card." The "Has CC" variable indicates the fraction of borrowers with at least one credit card, while the "New CC" variable captures people that obtained at least one new credit card in the past year. Balances, limits, and utilization are calculated across all credit cards. Utilization and limit calculations include households with zero credit cards. The "inquiries" variables reflect total activity in the past year. The delinquency variables reflect whether borrowers had any debt 30+ or 90+ days past due (excluding collections). The collections variables indicate the fraction of households with some debt in collections (non-medical and medical debt are tabulated separately). Our credit score measure is the Vantage Score.

For default, we examine both measures of delinquency and whether consumers have debt in collections. On average, 7% of consumers have some debt that is 90 or more days past due. Delinquency on utilities bills is fairly common, with on average 4.5% of consumers being 90 more days behind on at least one utility bill. These measures of delinquency exclude debt in

collections, which is even more common. On average, 24% of consumers have at least some debt in collections.

Default is highly persistent. Among people that either become newly delinquent on credit card debt or have non-medical debt enter into collections, 61.6% are in default the following year. 19.5% of delinquency spells last at least 5 years. Appendix Figure A.5 and Table A.10 report the persistence of delinquency for households in the Experian sample for a 10-year horizon.

Medicaid Eligibility Data. We combine several data sources to estimate Medicaid eligibility at the ZIP-code level. Adult household members are eligible for Medicaid if their income falls below a threshold, where the income threshold depends on household size and the number of dependents. Directly calculating eligibility requires individual-level data containing ZIP codes, income, and household composition. Because such data are difficult to access, we instead estimate the eligible population share. Our estimation approach first estimates the income distribution for each ZIP-year using data from the IRS Statistics of Income (SOI). Next, relying on the law of iterated expectations and Bayes' rule, we use the joint distribution of income and household size from the 2010 American Community Survey (ACS) to estimate the share of households with income below their corresponding Medicaid threshold. For details on our procedure, see Appendix B. We use IRS SOI data from 2009 to 2020 in our calculations. We obtain annual ZIP-code level data on Medicaid enrollment from the ACS.

Eligibility Summary Statistics. Table 2 reports summary statistics for our eligibility and income data. In the median ZIP code, throughout our sample period, approximately 14% of households are eligible for Medicaid. Conditional on expanding, the newly eligible share of the population on average increased by 17 percentage points within ZIP codes.

2.3 Empirical Strategy: Continuous Difference-in-Difference

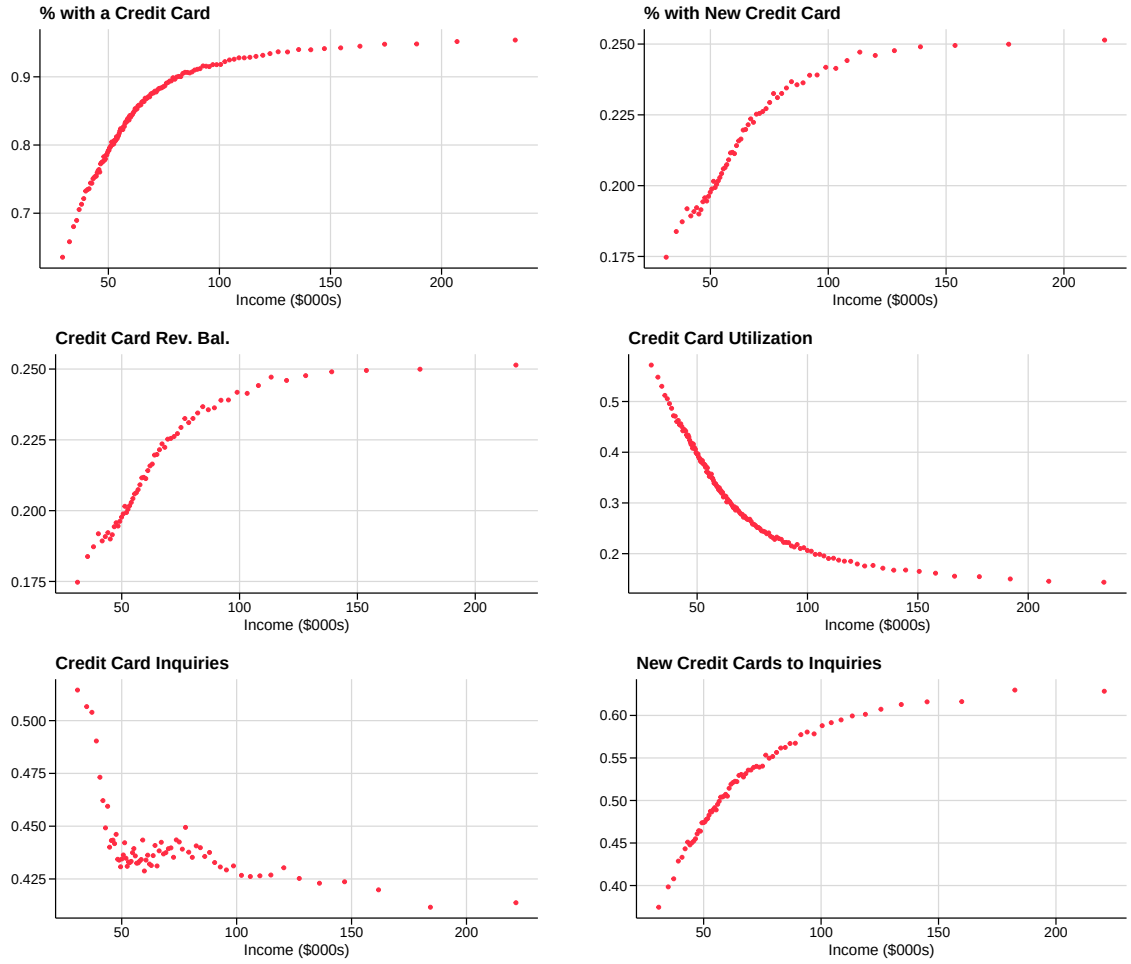
Identifying the causal effect of Medicaid eligibility is challenging. Eligibility depends on household income, which is likely correlated with omitted variables. Figure 5 shows that income is generally positively related to credit access and use. Comparing credit outcomes of eligible and non-eligible people would likely *understate* the causal effect of Medicaid eligibil-

Table 2: Summary Statistics for ZIP-Code Level Eligibility and Income Data

	Mean	SD	25th %	50th %	75th %	N
Elig. (%)	17.66	11.93	6.43	13.7	27.84	117,629
ΔElig.	16.55	9.86	8.54	18.28	23.81	7,828
Avg. HH Inc.	79.86	63.18	51.22	63.1	85.5	117,629

Notes: This table reports ZIP-code level summary statistics for the eligibility measures and income. Eligibility is calculated as described in the text, with additional details provided in Appendix B. The first row reports average eligibility across all ZIP codes and years. The second row reports the average change in eligibility in the first year of expansion (this excludes ZIP codes in states that did not expand). Average household income is the average adjusted gross income (AGI) reported in the ZIP code. Nominal variables are CPI-adjusted to be in terms of 2020 dollars.

Figure 5: ZIP-Code Level Credit Access and Income Facts



Notes: This figure plots binscatters of credit outcomes against average income, with all variables measured as the average within a ZIP code for a given year. All nominal variables deflated to 2020 dollars.

ity, as such comparisons would conflate the causal effect of eligibility with negative effects of low-income.

To overcome these identification challenges, we exploit quasi-experimental variation in the *size* of Medicaid expansions in a difference-in-differences (DiD) analysis. We draw on two sources of variation. First, we exploit cross-state variation arising from states' *pre-existing* Medicaid eligibility rules. Expanding under the ACA required states to raise the income eligibility limit to a common threshold of 138% of the federal poverty level for all adults aged 64 or less. States with previously stricter eligibility rules, such as lower income thresholds or rules precluding adults without dependents from eligibility, tended to experience larger changes in eligibility when expanding. Second, we exploit within-state variation due to differences in the *distribution* of income among low-income households. Specifically, ZIP codes where more households had income lie between the pre- and post-expansion income eligibility thresholds experienced a larger increase in eligibility.

Because we have a continuous measure of treatment intensity, we estimate a continuous DiD. The DiD compares ZIP codes with larger versus smaller shares of people that became newly eligible as a result of their state expanding Medicaid. This approach is similar in spirit to that of Goodman-Bacon (2018, 2021b), which uses state-level variation in exposure to Medicaid expansions to estimate the impact on mortality. An advantage of using ZIP-code level variation is that we can include state- or county-time fixed effects to account for the impact of state- or county-level events coinciding with expansions. Motivated by similar concerns, Dranove et al. (2016) and Garthwaite et al. (2019) also use a continuous ZIP-code level DiD to study the impact of Medicaid on health care utilization. We further refine their approach by constructing an eligibility measure that takes into account how Medicaid eligibility cutoffs vary with family size.

We estimate

$$Y_{zcst} = \beta \left(\text{Post}_{st} \times \text{NewElig}_{zs} \right) + \phi_z + \phi_{ct} + \gamma X_{zcst} + \varepsilon_{zcst}. \tag{1}$$

Above, Y_{zcst} is an average credit outcome in ZIP code z (located in county c of state s) in year

t.[16] Post_{st} equals one if state s has expanded Medicaid as of year t. We capture the intensity of treatment with a continuous variable, NewElig_{zs}, which corresponds to the fraction of ZIP z's population that became newly eligible for Medicaid over the first year of state s's expansion (relative to the year before its expansion). ZIP code fixed effects, ϕ_z, account for persistent differences across ZIP codes. We also include county-time fixed effects, ϕ_{ct}, to account for the effects of other unobserved time-varying county and state-level shocks. We include time-varying controls (X_{zcst}); our main specification controls for logged ZIP-code level income per household and deciles of the ZIP code's income distribution.[17] The coefficient of interest is β on the interaction term, which reflects the causal effect of an expansion that induced a 1 percentage point increase in eligibility. Throughout, we cluster our standard errors at the state level because the decision to expand Medicaid occurs at the state level.

When does OLS estimation of the above equation identify the causal effect β? The key identifying assumption is that household credit outcomes would have evolved in parallel—across locations with high versus low changes in eligibility—if Medicaid had not expanded. Phrased differently, we assume that the size of the newly eligible population is uncorrelated with other factors *changing* at the time of the expansion. This assumption would fail if other policies were implemented along with the expansions that also targeted the newly eligible population. Our design also mitigates this concern through the use of granular variation. The ZIP-level variation in the intensity of Medicaid expansions makes it possible to include county-time fixed effects. These fixed effects help account for the impact of other state or county-level policy changes on credit outcomes.

2.4 Empirical Results: The Effects of Medicaid Eligibility on Household Debt

Credit Card Debt. We first examine the response of credit card debt; Table 3 reports DiD estimation results. We start by documenting positive effects of increased eligibility on the extensive margin of credit card use. A 1 percentage point increase in a ZIP code's Medicaid-eligible population leads to a 0.11 percentage point increase in the fraction of borrowers that

[16] Most of our outcome variables are measured annually, but a subset are observed at a quarterly frequency. The quarterly variables are credit card balances (both total and revolving).

[17] Our results are robust to including fewer or more moments from the local income distribution, such as only controlling for logged average income or for 19 percentiles of the distribution. Appendix Table A.1 reports these results for our main variable of interest: revolving credit card balances.

have a credit card. The fraction of households that received a new credit card in the last year grows by 0.17 percentage points in response to the same shock. Per newly eligible household, these correspond to an 11 and 17 percentage point increase (respectively).

In terms of total credit card balances, average ZIP-code level balances increase 0.81 percentage points in response to a 1 percentage point increase in eligibility. This corresponds to approximately a $34.95 increase in average balances per capita.

Next, examining *revolving* balances, we find that a 1 percentage point increase in eligibility increases credit card borrowing by 0.46 percentage points. The borrowing response constitutes a $16.74 increase in average balances. Per newly eligible household, this corresponds to a $1,674 increase in credit card borrowing, which is approximately 23% of average household credit card balances. [18] Additional heterogeneity analyses reveal that both the extensive and intensive margin responses are strongest in low-income ZIP codes (see Appendix Table A.6). In particular, revolving balances barely change in high-income ZIP codes and exhibit the strongest response in low-income ZIP codes. The statistical significance of this result is also robust to a Fisher randomization test, which yields p-values below 0.001.[19]

Table 3: DiD Results for Credit Card Debt

	1[Has CC] (1)	1[New CC] (2)	log(CC Bal.) (3)	log(CC Rev. Bal.) (4)
$\text{NewElig}_{zs} \times \text{Post}_{st}$	0.105***	0.171***	0.813***	0.455***
	(0.017)	(0.016)	(0.083)	(0.076)
Obs	117,628	117,628	372,222	372,218
R2	0.979	0.865	0.970	0.957
Mean	85%	22%	$4,299	$3,680

Notes: This table reports results from estimating the DiD in Equation (1). Each specification uses county-time and ZIP code fixed effects and controls for logged ZIP-code level average adjusted gross income (AGI) and logged deciles of the ZIP code's income distribution (coefficients omitted). Note that data are annual in columns 1-2 and quarterly in columns 3-4 (due to data availability). Standard errors are clustered by state. Nominal variables are CPI-adjusted to be in terms of 2020 dollars. The dependent variable is labeled above the column number, and its mean is reported in the bottom row. Statistical significance: 10%*, 5%**, and 1%***.

[18] The 23% figure comes from dividing $1,674 by $7,210, where the $7,210 is average household-level credit card balances as of 2019. We calculate this figure using the aggregate credit card balance data from the Federal Reserve Bank of New York's Consumer Credit panel and household population count data from the American Community Survey.

[19] Because treatment varies both in timing and intensity, we randomize both of these aspects for our test. We permute the newly eligible share within counties across ZIP codes and the timing of adoption across states among those that ever adopted.

While large, the estimated borrowing response is plausible. First, this large response is consistent with households going from little to no credit card usage to newly having access to a credit card, and beginning to catch up to more typical levels of credit card balances.

Second, the DiD estimates the credit card response for up to 12 years before and after expanding. For each household made newly eligible *at the time* of expansion (our explanatory variable), more households likely end up eligible in the *following years* as a result of adverse events such as job loss. Putting our estimates in terms of *newly* eligible households likely overstates the impact in terms of the number of households that access Medicaid at some point *over* the next 12 years.

Third, credit demand and supply may also change for borrowers that never become eligible during the period of study. For example, after Medicaid expands, households at risk of adverse events may expect ex ante that they are more likely to qualify for Medicaid if an adverse event occurs. Even if these particular households do not experience an adverse event ex post, credit supply may respond positively for non-eligible households if underwriting models suggest that they present a lower risk of default after the expansion. In our structural analysis, our model allows these effects to be present.

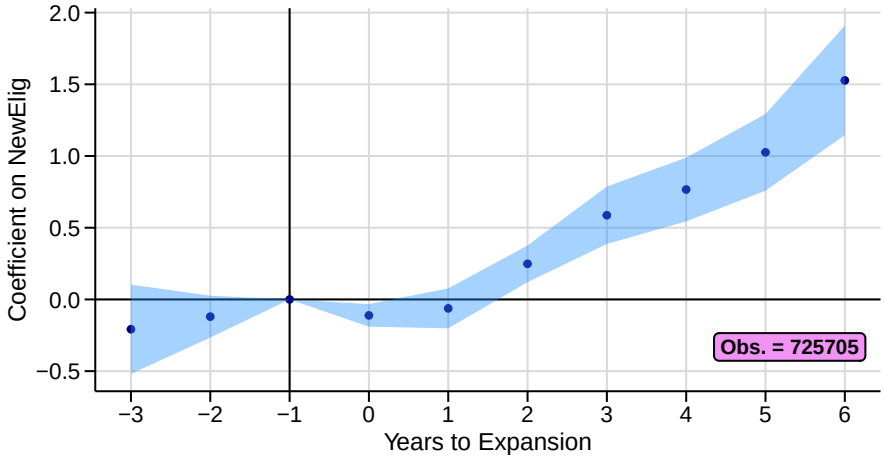
Fourth, our large estimates are consistent with other empirical evidence on spending responses to social insurance. Notably, Aaronson et al. (2012) estimate strong spending responses to minimum wage hikes. They find that a $1 rise in the minimum wage increases quarterly income and spending by $250 and $800 (respectively) among households receiving over 20% of their income from minimum wage labor.

Pre-trend Testing and Dynamic Estimates. We estimate a dynamic version of Equation (1), where we replace the post-treatment indicator with a set of categorical variables reflecting years relative to expansion. Figure 6 plots these estimates. Prior to expanding, we do not see economically or statistically significant pre-trends.[20] That is, the subsequent change in the share of the population that becomes newly eligible does not predict differential trends in credit card debt prior to expansion.

[20]While the point estimates in the pre-period are small in magnitude and not statistically different from zero, we nonetheless further assess robustness to violations of parallel trends. We report confidence intervals under the "honest DiD" approach of Rambachan and Roth (2023) for all outcome variables in Appendix Tables A.3 and A.4. Overall, we find our results robust to even large violations of parallel trends.

After expanding, revolving balances gradually increase over time. The gradual response is likely due to several factors. First, financial resilience can take time to improve, delaying the responses of credit supply and demand. Indeed, we find that on impact utilization decreased and repayment increased—two important determinants of credit scores. Credit scores begin improving one-year post expansion. The responses of these variables evolve gradually over time as well (see Appendix Figure A.2). Second, the Card Act of 2009 made it difficult for credit card lenders to update existing credit cards' interest rates. So most of the new and improved terms would come from new credit cards—delaying the reaction of credit card borrowing. Third, as more households experience medical expenditure shocks over the time, the share of people that have directly benefited from Medicaid will continue to grow.

Figure 6: Dynamic Estimates for Revolving Balances.



Notes: This figure plots dynamic DiD estimates using the stacked estimator of Cengiz et al. (2019) multiple horizons before and after Medicaid expansions. The shaded area is a 95% confidence interval for standard errors clustered at the state-level. The standard two-way fixed effect estimator yields similar results (see Appendix Figure A.1).

The gradual response we find also aligns with results from Hu et al. (2018), which finds a null effect of the 2014 Medicaid expansion on total credit card balances in 2014–2015. Our analysis reveals that studying the longer-run impact of Medicaid expansions is important, as some impacts on financial outcomes can emerge gradually.

Credit Demand and Supply Proxies. Next, to shed further light on the role of supply and demand, we examine proxies for these two forces in Table 4. Our proxy for credit card demand is the number of credit card inquiries per borrower in the past year.

Table 4: DiD Results for Credit Card Demand and Supply Proxies

	Demand	Supply		
	CC Inq. (#)	log(CC Lim.)	CC Util. (%)	New CC to Inq.
	(1)	(2)	(3)	(4)
$\text{NewElig}_{zs} \times \text{Post}_{st}$	0.420***	0.452***	-0.302***	0.046
	(0.036)	(0.099)	(0.025)	(0.033)
Obs	117,628	117,628	117,628	117,628
R2	0.941	0.987	0.980	0.848
Mean	0.44	$18,013	31%	0.52

Notes: This table reports results from estimating the DiD in Equation (1). Each specification uses county-year and ZIP code fixed effects and controls for logged ZIP-code level average adjusted gross income (AGI) and logged deciles of the ZIP code's income distribution (coefficients omitted). Standard errors are clustered by state. Nominal variables are CPI-adjusted to be in terms of 2020 dollars. The dependent variable is labeled above the column number, and its mean is reported in the bottom row. Statistical significance: 10%*, 5%**, and 1%***.

We estimate that credit card inquiries increase by 0.42 per newly eligible person. This suggests that demand may be contributing to the equilibrium rise in credit card debt.

For our supply proxies, we consider credit limits, the ratio of credit card balances to credit limits (utilization), and the ratio of new credit cards to credit card inquiries (over the past year). Overall, we find evidence consistent with a positive credit supply response to expansion. A 1 percentage point increase in a ZIP code's eligibility leads to a 0.45 percentage point increase in credit limits, corresponding to a $81.06 increase. In dollar terms, this is larger than the implied change in either total or revolving balances. Indeed, utilization decreases 0.30 percentage points per 1 percentage point increase in the eligible population. Thus, despite the increase in credit card borrowing, credit constraints are overall relaxed in the sense of credit limits becoming less binding.

Lastly, we estimate a small and statistically insignificant effect of eligibility on the ratio of new credit cards to inquiries. However, when separately examining ZIP codes with below-median income, we find a statistically and economically significant increase of 0.1 percentage points in the likelihood of a credit card inquiry resulting in a new credit card per 1 percentage point increase in eligibility (Appendix Table A.7). These findings are also consistent with

an expansion of credit supply. First, note that the flat, average response of credit card application success rates happens *despite* coinciding with a rise in both the number of inquiries and indebtedness. Additionally, if marginal applicants are less likely to be successful than the average applicant, the flat response of the average success rates would indicate increased success rates for marginal applicants (relative to what they would experienced absent expanding Medicaid). Moreover, we do find that applicants from low-income ZIP codes do experience higher success rates in response to expanded Medicaid.

These proxies have several limitations. First, we note that inquiries are an equilibrium outcome that may be influenced by marketing practices of credit card lenders and households' perceptions of lending standards, making them an imperfect proxy for demand. Second, we would ideally also observe the response of credit card interest rates to gauge the plausibility that the rise in credit card debt is supply driven. While interest rates are not available in our sample, other work using credit card offer data finds evidence of reduced credit card APRs. Brevoort et al. (2017) estimate that Medicaid expansions led to improvements in credit card offer terms worth $520 million per year. Ultimately, we rely on our model to quantify the roles of supply and demand.

Financial Resilience. Finally, we turn to measures of default and find evidence indicating that increased eligibility improved households' financial resilience in Table 5. We first examine the likelihood a household is 90 or more days delinquent on any debt, finding a small and statistically insignificant effects of Medicaid eligibility. This indicates that household "debt tolerance" has increased in the sense that, despite households increasing credit card debt overall, they do not become more likely to enter delinquency.

In order to better gauge the direct effect of insurance on delinquency, separate from indirect effects through changes in debt, we examine 90-day delinquency on utility bills (gas, electric, etc.). In response to expanded Medicaid, households are less likely to adjust consumption of utilities, compared to other goods. This makes utility delinquency a better measure of the direct effect of insurance on financial resilience. In line with improved financial resilience, we find that a 1 percentage point increase in eligibility lowers the rate of 90-day delinquency on utilities by 0.14 percentage points. This decrease is about three times larger

in below-median income ZIP codes: 0.16 versus 0.05 percentage points (Appendix Table A.8).

Next, we study more severe forms of default and find that Medicaid eligibility reduces debt in collections. A 1 percentage point increase in eligibility reduces the likelihood of having debt in collections by 0.13 percentage points. Similarly, the amount of debt in collections relative to average income (AGI) also shrinks by 0.04 percentage points in response to a 1 percentage point rise in eligibility. In relative terms, this constitutes a 3.08% reduction in this ratio. Note also that some the borrowers, who would have otherwise have had debt enter collections if not for Medicaid, may not entirely avoid default and instead experience milder forms of default such as delinquency. This substitution towards milder forms of default could explain the muted response of delinquency. To sum up, households appear to become more financially resilient as a result of increased Medicaid eligibility.

Table 5: DiD Results for Default and Credit Scores

	1[Delinquency]		Debt in Collections		Credit Score
	1[90+ Days]	1[90+ (Util.)]	1[Any Col.]	Col./Inc.	Vantage Score
	(1)	(2)	(3)	(4)	(5)
$\text{NewElig}_{zs} \times \text{Post}_{st}$	-0.006	-0.137***	-0.131***	-0.037***	0.230**
	(0.011)	(0.050)	(0.026)	(0.010)	(0.103)
Obs	117,628	117,628	117,628	117,628	117,628
R2	0.902	0.923	0.982	0.940	0.988
Mean	7.2%	4.5%	24%	1.2%	684

Notes: This table reports results from estimating the DiD in Equation (1). Column 1's delinquency measure reflects delinquency on any kind of debt, not only credit cards. Column 2's measure specifically reflects delinquency on utility bills. The collections measures also reflect outstanding collections for any kind of debt. Each specification uses county-year and ZIP code fixed effects and controls for logged ZIP-code level average adjusted gross income (AGI) and logged deciles of the ZIP code's income distribution (coefficients omitted). Standard errors are clustered by state. Nominal variables are CPI-adjusted to be in terms of 2020 dollars. The dependent variable is labeled above the column number, and its mean is reported in the bottom row. Statistical significance: 10%*, 5%**, and 1%***.

The improved financial resilience of borrowers may also translate into higher credit scores. We estimate an economically significant increase in credit scores in response to expanded Medicaid. Our point estimate implies that a 1 percentage point increase in eligibility leads to a 0.23 increase in consumers' average Vantage Score. This corresponds to a 23 point score increase per newly eligible person.

To sum up, we find that expanding Medicaid eligibility increased both the prevalence

of credit cards and revolving credit card balances. We find evidence suggesting both supply and demand contributed to the increase in equilibrium credit card usage. Lastly, we show that households appear more financially resilient in terms of reduced default, which could drive an increase in credit supply in response to expanded Medicaid eligibility.

2.5 Robustness and Additional Results

Treatment Effect Heterogeneity Robustness. Recent econometric work highlights that when the timing of treatment is staggered and treatment effects are heterogeneous, standard DiD estimators can identify a weighted average treatment effect (e.g., Goodman-Bacon, 2021a). These weights can sometimes be *negative*, biasing estimates in the opposite direction of the true (unweighted) average treatment effect. Settings like those in this paper, in which there are many *early* treated units, are especially vulnerable to having negative weights (Roth, Sant'Anna, Bilinski and Poe, 2022). This suggests that the estimates above may tend to understate the true treatment effects.

To address this concern, we verify that our results are robust to using a "stacked" DiD estimator (as in Cengiz, Dube, Lindner and Zipperer, 2019). The idea of this approach is to restructure the data and regression in order to estimate the treatment effect using only "good comparisons." The problematic negative weights that may arise for the standard two-way fixed effect (TWFE) estimator are due to its use of "bad comparisons." The "bad comparisons" occur when newly-treated units are compared to previously-treated units.[21] The stacked approach instead relies only on "good comparisons" (or "clean controls") of never-treated units to treated units.[22]

We implement the stacked DiD as follows. We first split our sample into several groups based on their expansion year (2014, 2015, 2016, 2019, or 2020). For each group of expanders, we append rows for the non-expanders. We create a group-identifier for each group and then "stack" observations for each group into a single data table. We then estimate a modified DiD

[21] The problem with such comparisons is that already-treated units may not provide a valid counterfactual for newly-treated units. For example, if treatment effects grow over time, the change in outcomes for already-treated units reflected both their potential outcomes (absent treatment) and any subsequent changes in treatment effects. Therefore, in this scenario, subtracting the change in outcomes for already-treated units from the change for newly-treated units would tend to understate the true treatment effect.

[22] For a helpful, more detailed discussion of the "stacked" approach, see Baker, Larcker and Wang (2022).

regression, specified below:

$$Y_{zcst\mathbf{g}} = \beta \left(\text{Post}_{st\mathbf{g}} \times \text{NewElig}_{zs} \right) + \phi_{z\mathbf{g}} + \phi_{ct\mathbf{g}} + \gamma X_{zcst} + \varepsilon_{zcst}.$$

The regression above modifies our baseline, Equation (1), by allowing the ZIP and county-time fixed effects ($\phi_{z\mathbf{g}}$ and $\phi_{ct\mathbf{g}}$, respectively) to vary by group $\mathbf{g}$. Estimating the modified equation results in effectively estimating the treatment effect for each group separately and then taking an average of these estimated treatment effects.

Overall, we find that the stacked specifications yield quantitatively similar estimates. Appendix Tables A.3, A.4, A.5 report estimation results. Generally, the stacked regressions yield estimates within two to three decimal points of those obtained via TWFE estimation on the same sub-sample and with similar statistical significance. We create "balanced panels" by restricting the stacked sub-samples to include the same number of pre- and post-expansion years. Since only one and two years of post expansion data are available for the 2019 and 2020 expansions (respectively), in this robustness analysis we use only data for 2014, 2015, and 2016 expanders (constituting 94% of our original sample). Because including later expansions comes at the cost of including fewer years of data, and treatment effects appear to grow over time (Figure 6), we present three sets of results that vary in whether they include the 2015 and 2016 expanders. To gauge the bias induced by "bad comparisons," the tables also display TWFE estimates for the same sub-samples.

Notably, our stacked DiD estimate for revolving balances is 0.471 for the set of 2014 expanders, which is close to its TWFE counterpart obtained on the same sub-sample (0.425), and close to our original estimate of 0.455. This suggests that bias related to treatment effect heterogeneity is small in our setting. Our estimates become smaller (but remain positive and highly statistically significant) when adding additional expanders and shortening the horizon over which we estimate. This is because the shorter horizon excludes the later years, where the estimated treatment effects are largest. We also find that our dynamic estimates are extremely similar under either TWFE or the stacked estimator (Appendix Figure A.1).

Income Placebo Test. An important advantage of our granular ZIP-level measure of exposure to Medicaid expansions is that it allows us to control for county-time fixed effects, which absorb regional differences in time trends of credit card borrowing. Additionally, our exposure measure exploits cross-ZIP variation in a narrow range of the income distribution, while controlling for many moments of the ZIP's income distribution.[23] Together, these features allow us to plausibly control for other policies or local economic developments that tended to coincide with states' decisions to expand Medicaid, which may have also affected credit card borrowing. To better examine our success in avoiding bias from such confounders, we conduct a placebo test. We interact a post expansion indicator with either one of two measures of ZIP-level income: a binary indicator for whether a ZIP had above median average AGI or simply the ZIP's logged average AGI. The idea of the test is to see whether it is specifically expanded Medicaid eligibility that explains differences in credit card borrowing across ZIP codes, or is it due to another factor that is correlated with ZIP-level income.

Appendix Table A.2 reports results from this test for revolving balances. Our findings here yield two useful insights. First, our estimate of the causal effect of increased eligibility remains similar (and highly statistically significant) even when including these additional interactions (0.440 and 0.518, versus our original 0.455). This suggests, in the sense of a mediation analysis, that our original findings are not simply an artifact from failing to control for other factors related to ZIP codes' differences in income. Second, we also find that the interactions of either income measure with a post expansion indicator has a precisely estimated null coefficient. The point estimates imply that being in an above-median ZIP code predicts a (statistically insignificant) -0.2 pp lower level of revolving balances; a 1 pp increase in the ZIP's average AGI predicts a (statistically insignificant) 0.008 pp lower level. The small estimates here indicate that higher- versus lower-income ZIP codes did not generally have revolving balances start trending differently around states' Medicaid expansions (conditional on our controls). Taken together, these results are consistent with our measure of Medicaid exposure successfully using variation in a narrow income range that is unrelated to broader income-related confounders, and that the risk of such confounders is low given the absence

[23] For example, the pre-ACA limit in California was 106% of the FPL and 138% post-ACA. For a family of three in 2020, this corresponds to incomes of $23,023 to $29,974.

of differential trends between high- and low-income ZIP codes around expansion timing.

Additional Results: Other Forms of Borrowing. Our analysis centers on card debt because of its widespread use and because it is a more natural form of credit to smooth consumption across states of the world (versus across time). But given the positive impact of expanded Medicaid eligibility on credit scores, it is worth examining whether borrowing changes for other forms of debt, even if they are primary used to finance the purchase of durables. To this end, we examine the impact of expanding Medicaid on auto and mortgage loans in Appendix Table A.9. We find statistically significant and economically meaningful effects on auto loans, while effects on mortgages are small and statistically indistinguishable from zero. A 1 pp increase in the eligible population share leads to a 0.176 pp increase in the fraction of people with an auto loan and a 1.11% increase in auto loan balances. Two features can account for the difference in responses between auto and mortgage loans. First, risk tends to be more strongly reflected in auto loan interest rates than in mortgage rates.[24] This may be due to the greater prevalence of negative equity for vehicles versus homes, and the tendency of borrowers to prioritize mortgage payments over auto loan payments (Experian, 2026). Thus, improvements in creditor scores can translate into a larger reduction in auto and credit card rates, while mortgage rates react less. Second, the people most affected by the Medicaid expansions will tend to have lower income. These households may be closer to the margin of getting an auto loan (or a credit card), but not as close to the margin of getting a home loan.

3 Model

In this section, we develop an incomplete-markets heterogeneous-agents model with health expenditure shocks, health insurance, and delinquency on unsecured debt. Households face idiosyncratic income shocks as well as idiosyncratic health expenditure shocks. They can save and borrow using a one-period non-state-contingent asset and can choose not to repay their debt obligations. After reneging on their debt obligations, their debt enters a delinquent state, and the household is excluded from financial markets and suffers a utility loss. In every

[24] Auto loan rates vary much more with borrower credit scores than mortgage rates. For example, typical APRs on used auto loans recently ranged from 7.43% to 21.60% across the credit score distribution (Experian, 2025) In contrast, the typical mortgage rate varied from 7.20% to 7.85% across the credit score distribution (Bankrate, 2025).

period, delinquent debt receives a stochastic haircut (i.e., it is reduced by some percentage). Households can exit the delinquency state by repaying their debt.

Credit to households is supplied by credit card companies. These companies have access to funds at the risk-free interest rate, which we assume is exogenous. The assumption that the risk-free rate is constant, rather than the aggregate net supply of debt, allows us to study how different policies affect the aggregate stock of household debt in the economy. We assume that credit card companies are risk neutral and behave competitively. As a result, the interest rate spread a household pays on its loans is such that lenders' expected profits equal zero.

We use the model to study the effects of different insurance policies on the aggregate stock of debt, wealth inequality, and welfare.

3.1 Household problem

There is a continuum of measure one of households in the economy, denoted by $i \in (0,1)$. Household's i income at time t is denoted by y_{it} and evolves according to a compound Poisson process:

$$\ln y_{it} = \begin{cases} \rho \ln y_{it-1} + \epsilon_{it}^{y} & \text{w.p.} \quad \lambda_y, \\ \ln y_{it-1} & \text{w.p.} \quad 1-\lambda_y, \end{cases} \tag{2}$$

where λ_y is the probability that an income shock arrives in a period. If such shock does not arrive, the household's income does not change. Given an income shock, the household income follows an AR(1) process, where ρ is the degree of persistence and ϵ_{it}^{y} is an idiosyncratic income shock with mean zero and variance $\sigma_{\epsilon,y}^2$.[25]

In addition to income risk, households are subject to stochastic medical expenditure shocks, denoted by m_{it}. The stochastic expenditure shocks follow a log-normal distribution with mean μ_e and variance σ_e^2.

Medical bills are partially covered by health insurance. We allow the type of insurance to vary exogenously with household income. Health insurance covers a share of the medical bill due. We denote the share of medical bills a household has to pay out-of-pocket as $o(y_{it})$. Note that the out-of-pocket share depends only on the current household income, as we assume

[25]The assumption that income shocks occur with probability λ_y allows us to match the high kurtosis of income changes observed in the data. The high kurtosis is quantitatively important for matching delinquency rates in the data.

insurance varies only with household income.

Each household has access to non-state-contingent one-period debt, denoted by b_{it}. A negative value of b_{it} represents household savings. In the beginning of each period, the household can choose to repay or renege on its debt obligations. The household's decision to repay debt depends on its total debt, its level of income, and the size of its current medical bill. Households face an interest rate schedule for the amount of borrowing they choose, which we denote by $r(b', y)$.[26] The price of debt is denoted by $q(b', y) = \frac{1}{1+r(b',y)}$. We denote the total amount of repayments owed by the household by $\tilde{b}_{it} = b_{it} + o(y_{it})m_{it}$. Note that $\tilde{b}$ is the relevant state variable from the perspective of the household.

The household's present discounted value is denoted by $V(\tilde{b}, y)$, where $\{\tilde{b}, y\}$ are its individual state variables. This value function is given by the max between the present discounted value of repaying debt obligations, denoted by $V^r(\tilde{b}, y)$, and the value of reneging and declaring delinquency, denoted by $V^d(\tilde{b}, y)$:

$$V(\tilde{b}, y) = \max \left\{ V^r(\tilde{b}, y), V^d(\tilde{b}, y) \right\}. \tag{3}$$

Conditional on the decision to repay its debt obligations, the recursive problem of the household with total debt obligations $\tilde{b}$ and income y is given by

$$\begin{aligned} V^r(\tilde{b}, y) &= \max_{c, b'} u(c) + \beta \mathbb{E} V\left(b' + o(y')m', y'\right), \\ \text{s.t.} \qquad & c + \tilde{b} \leq y + q(b', y)b', \end{aligned} \tag{4}$$

where $u(\cdot)$ is the utility of the household from consumption, which is assumed to be strictly increasing, concave, and continuously differentiable. The household's discount factor is β.

When a household reneges on its debt obligations, the debt moves into a delinquency state. A household with debt in a delinquency state cannot save or borrow and suffers a utility cost ξ. At the end of the period, the household receives a stochastic haircut on its debt obligations. The haircut, denoted by δ, is assumed to follow a Power distribution with shape parameter α. That is, the probability that the share of debt forgiven is less than δ is given

[26] The size of medical bill does not appear in the interest rate schedule as it is independent across periods.

by δ^{α}. This haircut process may be interpreted as a reduced-form approximation of either bankruptcy or debt forgiveness obtained through negotiations with creditors.

The value of the household in the delinquent state is given by

$$V^d(\tilde{b}, y) = u(y) - \xi + \beta \mathbb{E} V\left((1-\delta)\tilde{b} + o'(y')m', y'\right). \tag{5}$$

The timeline in every period is as follows. First, the household learns its current income and medical bill. Then, it decides whether or not to repay its outstanding debt obligations. If it decides to repay its debt obligations, it chooses the level of consumption and borrowing. If it decides to renege on its debt obligations, it consumes its income, suffers a utility loss, and at the end of the period draws a stochastic haircut rate.

We denote the delinquency policy function of a household with total debt obligations $\tilde{b}$ by

$$d(\tilde{b}, y) = \mathbb{1}\left[V^r(\tilde{b}, y) < V^d(\tilde{b}, y)\right], \tag{6}$$

where $\mathbb{1}$ is the indicator function. The function $d(\tilde{b}, y)$ equals 1 when the household defaults. When indifferent, we assume the household repays its debt obligations.

3.2 Credit supply

Credit to households is supplied by risk-neutral credit card companies. Perfect competition among these companies ensures an expected zero-profit condition holds in equilibrium. We assume credit card companies have unlimited access to funds at the risk-free interest rate, which we denote by r^f.

Debt in the economy is a hybrid of short-term and long-term debt. Debt is of short maturity in nature, as households need to repay their debt obligations in the following period.[27] However, when debt becomes delinquent, credit card companies do not receive payments within that period. Instead, they need to wait until the household decides to repay its debt—either because it received a sufficiently large haircut or because its income level changes so that it decides to repay.

[27] Note that households can potentially roll over their debt by borrowing to cover their current debt obligations.

The zero-profit condition that pins down the price of debt, $q(b', y)$, is

$$q(b', y)(1 + r_f) = \left[1 - \mathbb{E}\left(d(b' + o(y')m', y')\right)\right] + \mathbb{E}\left[d(b' + o(y')m', y')(1 - \delta')q\left((1 - \delta')(b' + o(y')m'), y'\right)\right]. \quad (7)$$

That is, the credit companies equate the cost of their loan (LHS) with its return (RHS). The return from the loan is the sum of two parts. If the debt is repaid, the credit company receives its full face value. If the debt is delinquent, which happens with probability $\mathbb{E}\left(d(b' + o(y')m', y')\right)$, the creditor is left with a claim on a delinquent debt. The worth of such claim is given by $(1 - \delta')q\left((1 - \delta')(b' + o(y')m'), y'\right)$. The fixed point function $q(\cdot)$ that solves equation (7) is the debt pricing schedule. As the following proposition shows, this function exists and is unique.[28]

Proposition 1. *Given a default policy function, $d(\tilde{b}, y)$, there exists a unique pricing schedule $q(b', y)$ that satisfies equation (7).*

3.3 Stationary Equilibrium

We denote the joint distribution of households across total debt obligations and income levels at the beginning of the period, after the expenditure shock realization, by $\Lambda\left(\tilde{b}, y\right)$. Four components characterize the law of motion for this joint distribution: (i) the borrowing decision of households, $b'(\tilde{b}, y)$, (ii) the debt pricing schedule, $q(b', y)$, (iii) the household's default decision, $d(\tilde{b}, y)$, and (iv) the exogenous processes of income, expenditure shocks, and haircuts. The law of motion for the joint distribution is defined as follows. For all Borel sets $\mathcal{B} \times \mathcal{Y} \subset \mathbb{R} \times \mathbb{R}^+$,

$$\Lambda'\left(\mathcal{B} \times \mathcal{Y}\right) = \int_{m'} \int_{y' \in \mathcal{Y}} \int_{\mathscr{B}\left(\tilde{b}, y, m'\right)} d\Lambda(\tilde{b}, y) dF(y'|y) dG(m') + \int_{\delta} \int_{m'} \int_{y' \in \mathcal{Y}} \int_{y, ((1-\delta)\tilde{b} + o(y')m') \in \mathcal{B}} \mathbb{1}\left[d\left(\tilde{b}, y\right)\right] d\Lambda(\tilde{b}, y) dF(y'|y) dG(m') dH(\delta), \quad (8)$$

and

$$\mathscr{B}\left(\tilde{b}, y, m'\right) = \left\{(\tilde{b}, y) \quad \text{s.t.} \quad d\left(\tilde{b}, y\right) = 0 \quad \text{and} \quad b'\left(\tilde{b}, y\right) + o(y')m' \in \mathcal{B}\right\},$$

[28] The proof is in Appendix C.

where $F(\cdot)$, $G(\cdot)$, and $H(\cdot)$, are the CDF of the different exogenous variables. In the stationary equilibrium, the distribution $\Lambda(\tilde{b}, y)$ is constant over time. The definition of the stationary Markov-perfect equilibrium is as follows.

Definition 1 (Equilibrium)**.** *A stationary Markov-perfect equilibrium is given by a default policy function $d(\tilde{b}, y)$, a borrowing policy function $b'(\tilde{b}, y)$, a debt pricing schedule $q(b', y)$, and a joint distribution of households across total debt and income levels, $\Lambda(\tilde{b}, y)$, such that*

1. *The default and borrowing policy functions solve the household's problem given the debt pricing schedule.*
2. *The debt pricing schedule satisfies the zero-profit condition, (7).*
3. *The joint distribution of households across debt and income levels is stationary.*

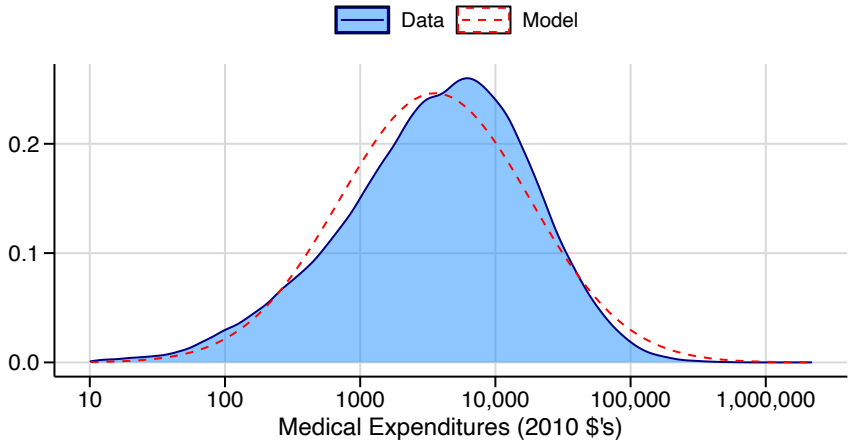
3.4 Calibration

We calibrate the risk-free interest rate in the economy to 2% and assume that households have log utility in consumption, $u(c) = \ln(c)$. We can divide our calibration into five sets of structural parameters. First, there are two preference parameters—the discount factor β and the disutility of delinquency ξ. Second, there is the haircut process parameter—the shape parameter of the Power distribution, α. Third, there are three parameters governing the income process—the arrival rate of income shocks λ_y, the persistence of income ρ, and the variance of income shocks $\sigma^2_{\epsilon,y}$. Fourth, there are two parameters governing the health expenditure process—its log mean μ_e and variance σ^2_e. Finally, we need to calibrate the out-of-pocket share function $o(y)$.

We proceed as follows. We use income data from the Panel Study of Income Dynamics (PSID) to calibrate the income process. We use the Medical Expenditure Panel Survey (MEPS) to calibrate the parameters governing health expenditure shocks and the out-of-pocket function. We then calibrate the preference parameters and haircut shape parameter to match several features of credit card debt in the data.

Income Process Parameters. We calibrate the income process parameters to match key moments of household income dynamics from the PSID. Our sample includes households whose

Figure 7: Distribution of annual medical expenditure



Notes: Data source: Medical Expenditure Panel Survey (MEPS).

head is between the ages of 25–55. The sample period is 1999–2019. The three moments we target are the two-year autocorrelation of log household income (0.837), the standard deviation of two-year log-income changes (0.536), and the kurtosis of two-year log-income changes (5.586).[29] The calibration yields an arrival rate of income shocks $\lambda_y = 0.348$, so that income shocks arrive on average every 2.8 years. The calibrated persistence of income conditional on an income shock is $\rho = 0.756$, and the variance of income shocks is $\sigma^2_{\epsilon,y} = 0.378$.[30]

Health Expenditure and Insurance Parameters. We calibrate the mean log expenditure and its variance to match the distribution of annual medical expenditures of households in the MEPS data over 2000 to 2017. We find a mean expenditure of 8% of median income and a variance of 2.62. Figure 7 compares the calibrated distribution to its empirical counterpart. The calibrated parameters imply that a medical expenditure one standard deviation above the average equals 40% of median income. Recall that the expenditure shocks, in both the data and the model, are not the amounts that households have to pay out-of-pocket.

The MEPS dataset also contains information on household income, insurance type, and out-of-pocket (OOP) expenditures. To construct the out-of-pocket share as a function of in-

[29] The data moments are in line with the income moments in Guvenen et al. (2015), who use Social Security Administration data on the annual income of males in the US.

[30] In Appendix C, we show how the moments map to these three parameters. Note that while the persistence of income conditional on an income shock is 0.756, the unconditional income persistence is 0.915.

come, we proceed in two steps. First, we split insurance types into three categories: Medicaid, non-Medicaid insurance, and uninsured. Non-Medicaid insurance includes households with private insurance (including employer sponsored) and Medicare. We compute the average OOP share for each insurance type. The average OOP share in Medicaid is 6.8%, 27.5% under non-Medicaid insurance, and 63% for the uninsured.[31]

Second, using a log-linear function, we approximate the share of households with each insurance type. We compute these approximations both before and after the expansion of Medicaid. Figure 3 displays the empirical shares alongside our approximation. The stationary distribution of our benchmark model uses the pre-expansion approximations, while the post-expansion counterfactual economy uses the post-expansion approximations.

We combine the relationship between insurance type and income along with the OOP share by insurance type to impute the average OOP share by income. Specifically, the OOP share in our model is

$$o(y) = \Pr(\text{Medicaid}|y) \times 6.8\% + \Pr(\text{Other}|y) \times 27.5\% + \Pr(\text{Uninsured}|y) \times 62.7\%. \quad (9)$$

A key advantage of this two-step approach, relative to directly estimating the OOP share along the income distribution, is that we can conduct counterfactuals in which we change the share of households under each insurance type.

Preference and Haircut Parameters. There are three remaining parameters: (i) the discount factor β, (ii) the disutility of delinquency ξ, and (iii) the haircut distribution shape parameter α. We use an indirect inference procedure to calibrate these parameters.

We target four credit card moments in the pre-expansion period. These moments include the aggregate level of revolving credit card debt in the data, the share of revolving credit card debt owed by the bottom 40% of households, the share of delinquent households, and the persistence of delinquency.[32]

[31] Uninsured patients can be required to pay less than 100% of their expenses when either they negotiate lower payments with hospital or when other sources cover their payments. Other sources of payment include non-health sources of insurance (e.g., homeowners or auto insurance in relation to injuries), the Veteran's Administration, and Worker's Comp.

[32] The revolving credit card data at the household level are obtained from the PSID, while the delinquency rate as well as the hazard rate of delinquency are obtained from the Experian dataset. The data counterpart

Note that not all debt in our model corresponds to credit card debt in the data. Consider, for example, a low-income household who can only borrow at very high interest rate spreads, and as a result decides to have no debt. Suppose this household suffers a large medical shock, which it chooses not to repay. As a result, this household has delinquent debt. We do not consider this form of debt to be credit card debt, as no lender extended any loans to this household. In particular, credit card debt in the model corresponds to any non-delinquent debt households owe, as well as delinquent debt that was given to the household in the previous period by a credit card lender.

In addition to the four credit card moments, we also target the main regression coefficient from our empirical analysis: a 1 percentage point increase in a ZIP code's Medicaid-eligible population increases credit card borrowing by 0.46%. To construct a model-implied regression coefficient, we proceed as follows. Starting from the pre-expansion ergodic distribution, we divide households into ZIP codes according to their income.[33] We then study the transition dynamics following the expansion of Medicaid.[34] We construct the estimand in the model in three steps. First, for each ZIP code, we calculate the share of newly Medicaid-eligible households based on the income distribution in that ZIP code as well as the difference between the estimated pre- and post-expansion Medicaid insurance across household income.[35] Second, for each ZIP code, we track the impact the expansion has on credit card debt for the first five years following the expansion. We do so by comparing the log amount of credit card debt in the ZIP code following the expansion of Medicaid, relative to what would be the log amount of credit card debt absent the implementation of the expansion.[36] Finally, we compute the average treatment effect across ZIP codes.

The estimation targets as well as the model fit are reported in Table 6. All model param-

of delinquency is an individual with any delinquent debt that is over 90-days due or any non-medical debt in collection.

[33]To do so, we use IRS SOI data to obtain the income distribution in every ZIP code in the US. We then draw households from the ergodic distribution so that we exactly match the observed income distribution within each ZIP code and across ZIP codes.

[34]Details on the counterfactual experiment of expanding Medicaid are relegated to Section 4.

[35]Since our model only uses data on the insurance type individuals have, we need to take a stand on the take-up rate in order to obtain Medicaid eligibility. For the analysis, we use a take-up rate of 50%.

[36]Note that while the national ergodic distribution of income does not vary over time in the steady state, the income distribution within each ZIP code does slowly converge toward the national stationary distribution. For that reason, we do not compare the log amount of debt post-expansion to the same variable before the expansion, but rather to the log amount of debt along the transition dynamics in each ZIP code absent the expansion of Medicaid.

eters are reported in Table 7. The model does a reasonable job at matching the data targets, considering we are using 3 parameters to match 5 targets. The estimated annual discount factor β is estimated to be 0.85. This relatively low discount factor is needed for households to borrow despite the non-negligible interest rate spreads.[37]

Table 6: Estimation Targets and Model Fit

	Data	Model
Aggregate credit card debt (% of median income)	5.3%	4.8%
Credit card debt share held by bottom two quintiles	19.3%	16.3%
Delinquency share	25.7%	29.9%
Delinquency hazard rate	61.6%	62.6%
Eligibility effect on credit card debt	0.46	0.39

Notes: This table presents the targeted data moments as well as the model-implied moments of the estimated model. The last row reports the main regression coefficient from Section 2 and the estimand in the model.

The estimated value for the disutility of delinquency is $\xi = 0.137$. This value implies that households are indifferent between the periodic non-pecuniary cost of delinquency and a 13.7% drop in consumption.[38] This sizeable cost is powerful delinquency deterrent, which enables households to maintain positive levels of debt in equilibrium without going delinquent. Finally, we estimate the shape parameter of the haircut distribution to be 0.865. This value implies that, on average, about 45% of delinquent debt will be forgiven within one year. This haircut rate helps the model match the persistence of delinquency, which is about 60% from one year to the next (see Appendix Figure A.5). Moreover, it is consistent with lenders selling their delinquent debt to debt collectors for a low price relative to the face value of debt.

3.5 Equilibrium properties

We solve the model globally and compute the stationary distribution across income and debt levels. Panel A of Figure 8 presents the regions of the state space where households choose

[37]The estimated value of the discount factor is within the range of impatient and patient households found in Chatterjee et al. (2020) and Aguiar et al. (2020). Chatterjee et al. (2020) construct a quantitative model of credit scores and estimate the discount factor of impatient households to be 0.81 and that of patient households to be 0.93. Aguiar et al. (2020) use the PSID to study the behavior of hand-to-mouth households and estimate preference heterogeneity across households. They estimate that the discount factor for impatient households is 0.72, while that for patient households is between 0.94–0.97.

[38]This is approximately 25-50% of the consumption-equivalent cost of personal bankruptcy for marginal filers estimated in Indarte (2023). This smaller cost is plausible, as delinquency is likely less stigmatized than bankruptcy.

Table 7: Model Parameters

Parameter description	Parameter	Value	Calibration Source
Risk-free rate	r_f	2%	Standard value
Arrival rate of income shocks	λ_y	0.348	PSID income moments
Persistence of income shocks	ρ	0.756	PSID income moments
Volatility of income shocks	$\sigma_{\epsilon,y}$	0.615	PSID income moments
Expenditure shock mean	μ_e	0.08	MEPS expenditure moments
Expenditure shock volatility	σ_e	1.61	MEPS expenditure moments
Discount factor	β	0.854	Indirect inference (Table 6)
Disutility of delinquency	ξ	0.137	Indirect inference (Table 6)
Haircut power law shape	χ	0.865	Indirect inference (Table 6)

Notes: This table presents the targeted data moments as well as the model-implied moments of the estimated model. The last row reports the main regression coefficient from Section 2 and the estimand in the model.

to repay or renege on their debt obligations.[39] In general, lower income implies a lower debt threshold, above which households do not repay their debt. Because the utility function is concave, low-income households are more tempted not to repay their debt and increase their contemporaneous consumption. As a result, low-income households cannot maintain a high level of debt.

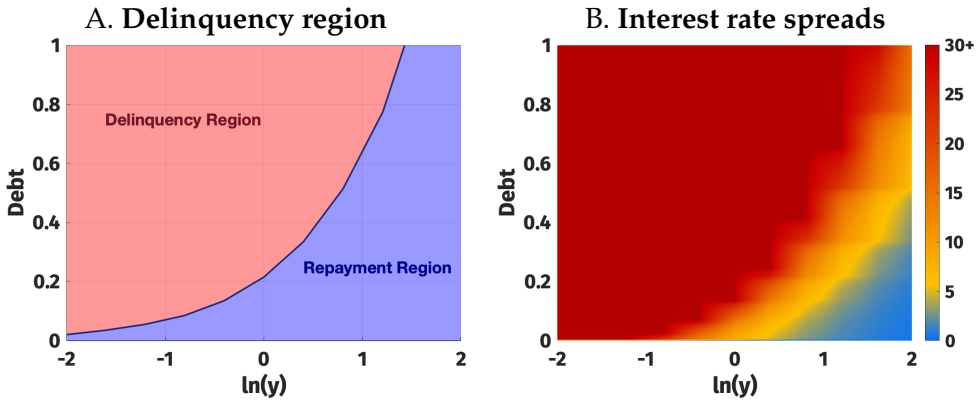
The probability of delinquency affects the interest rate spreads households face. Panel B of Figure 8 presents the interest rate spreads households face in equilibrium. The horizontal axis is the current income level of the households, and the vertical axis represents the debt obligations promised to be repaid in the following period. Because high-income households are less likely to renege on their debt obligations, they face lower interest rate spreads. In a similar fashion, low-income households face very high interest rate spreads.

The interest rate schedule faced by low-income households effectively limits their access to credit. Policies that reduce the delinquency probability of these households expand their credit access by lowering these interest rate spreads. In the next section we study the effects of one such policy: the expansion of Medicaid.

Model Abstractions. Our model abstracts away from several aspects of health, which will generally make our policy counterfactual analysis conservative. First, we do not model house-

[39] While the upper limit of debt in Figure 8 is 1, we solve the model with the upper limit of debt being 6. The share of households with debt level above 1 is very small.

Figure 8: Equilibrium properties



Notes: The left panel presents the regions where households choose to repay their debt or go delinquent. The right panel presents the interest rate spreads households pay as a function of promised repayment. In the left panel, the debt level is capped at 1 for illustration purposes. The delinquency threshold for the highest income household is 5. In the right panel, the interest rate spread is capped for illustration purposes at 30%, so the dark share of red represents interest rate spreads above that cap.

holds' "health." Therefore, our welfare calculations reflect improvements in *financial* health rather than *physical* health. In reality, the welfare benefits from improved physical health are also likely large, as Medicaid expansions significantly reduced both child and adult mortality (Goodman-Bacon, 2018, 2021b; Miller et al., 2021).

Second, our model does not feature moral hazard in terms of risky health behaviors or use of medical services. Prior analyses of the ACA's Medicaid expansion do not find effects on risky behaviors (e.g., smoking and lack of exercise, Simon, Soni and Cawley, 2017). However, research generally estimates moderate moral hazard responses in terms of overall health expenditures (Manning et al., 1987; Finkelstein et al., 2012; Garthwaite et al., 2019), although a recent analysis of the Medicaid expansion under the ACA finds no impact on total (pre-insurance) medical expenditures (Shupe, 2023). Because households do not increase total medical expenditures in our model, the model likely understates the borrowing response. While the additional costs of the expansion that would arise from increased expenditure are not included in our welfare calculation, these costs should be weighed against the welfare gains from improved physical health.

Third, income and medical expenditure shocks are uncorrelated in our model. In reality, the health events resulting in large medical expenditures may reduce income and employ-

ment (Dobkin et al., 2018; Stepner, 2019). If expanding Medicaid improves access to care and reduces the likelihood of major illnesses, then our model would understate the financial benefits from higher income associated with these health improvements. It likely also understates the borrowing response, as avoiding low-income states of the world improves credit access.

We also abstract away from explicitly modeling other collateralized forms of debt such as auto loans and mortgages. As discussed earlier in Section 2.5, we find that auto debt does react to the Medicaid expansion but mortgage debt does not. This suggests excluding mortgage debt is likely innocuous for the two main applications of our model: studying its impact on household borrowing and welfare. Regarding auto loans, this abstraction would tend to make our model analysis conservative, at worst understanding the increase in total household debt and welfare.

4 The Effect of Medicaid Expansion on Debt and Welfare

In this section, we study how the expansion of Medicaid shapes the distribution of debt across households in the economy and study its welfare implications. We begin by examining the channels through which health insurance policy can affect households' accumulation of debt. We then study the effect of expanding Medicaid in our model.

4.1 Theoretical analysis

We model the expansion of health insurance as a change in the out-of-pocket share of medical expenditure that households of different income levels face, $o(y)$. A more generous health insurance policy corresponds to a reduction in the out-of-pocket share that different households pay for medical shocks.

Expanding health insurance affects households' accumulation of credit card debt in several ways, both directly and indirectly. There are two direct effects of the policy. First, a more generous health insurance increases households' disposable income. As a result, household debt levels decline. Second, the insurance against medical expenditure shocks reduces delinquency rates. When a household in delinquency receives a medical expenditure shock, its additional debt is not counted as credit card debt but rather as medical debt. When a household is not in delinquency, receiving a medical shock may result in more borrowing to cover

that expenditure. The decline in delinquency rates leads to a substitution from medical debt to credit card debt. Overall, the direct effects of the policy reduce overall debt levels, but may result in larger levels of credit card debt.

To study the indirect channels of the policy, consider the household's optimality condition with respect to debt accumulation, which is given by

$$u'(c)\frac{\partial\left(q(b',y)b'\right)}{\partial b} = \beta\mathbb{E}\mathbb{1}_{V^r\geq V^d}u'\left(c(b'+o(y')m',y')\right) + \beta\mathbb{E}\mathbb{1}_{V^r<V^d}V_1^d\left(b'+o(y')m',y'\right) \quad (10)$$

The household equates the benefits from borrowing (LHS) to the utility costs of borrowing (RHS). By increasing debt obligations, b', the household increases its current funds by $\frac{\partial(q(b',y)b')}{\partial b}$. Note that the household internalizes how its borrowing decision affects the interest rate it pays on debt. There are two potential costs of borrowing. If the household repays its debt in the following period ($V^r \geq V^d$), the marginal cost of debt obligations is simply the marginal utility of consumption. Alternatively, if the household goes delinquent ($V^d > V^r$), the household's cost of debt obligations is $\frac{\partial V}{\partial b}$.

The first indirect channel through which social insurance affects household debt is a reduction in the *precautionary savings motive*. A reduction in $o(y')$ reduces the volatility of out-of-pocket medical expenditure, $o(y')m'$. This results in a lower volatility of future consumption. If the utility function features prudence ($u'''(\cdot) > 0$), as it does in our calibration, then this reduction in volatility results in a smaller expected utility cost of borrowing. The reduction in the marginal cost of borrowing induces households to take on more debt. That is, through the precautionary savings channel, social insurance raises household debt levels.

The second indirect channel we label the *debt aversion* channel. Debt leads to higher utility loss in the states where households repay their debt obligations. Debt is less costly in the delinquency state as households expect to pay it only in the future, and after a haircut. A more generous insurance policy raises the probability of repayment, $\mathbb{E}\mathbb{1}_{V^r\geq V^d}$, as medical expenditures that would have pushed households into delinquency are now partially insured. The increase in the repayment probability deters households from taking on debt, as they are more likely to repay it. Therefore, through the debt aversion channel, social insurance reduces

household debt levels.

Both the precautionary savings motive and the debt aversion channels do not depend on lenders changing their behavior. That is, they do not depend on the supply side of loans. We refer to the combined effect of these two channels as the *credit demand* channel.

The final indirect channel is the *credit supply* channel. The reduction in delinquency probability induces lenders to lower interest rate spreads, $q(b', y)$. This raises the benefits from taking on more debt obligations b'. For each unit of consumption promised to be repaid in the following period, households receive more units of consumption in the current period. This induces households to increase their debt obligations. So, through the credit supply channel, social insurance increases household debt levels.

Overall, the effect of social insurance on the aggregate level of household debt is ambiguous. We now turn to study the expansion of Medicaid in our model and quantitatively assess the strength of the different channels.

4.2 Medicaid expansion

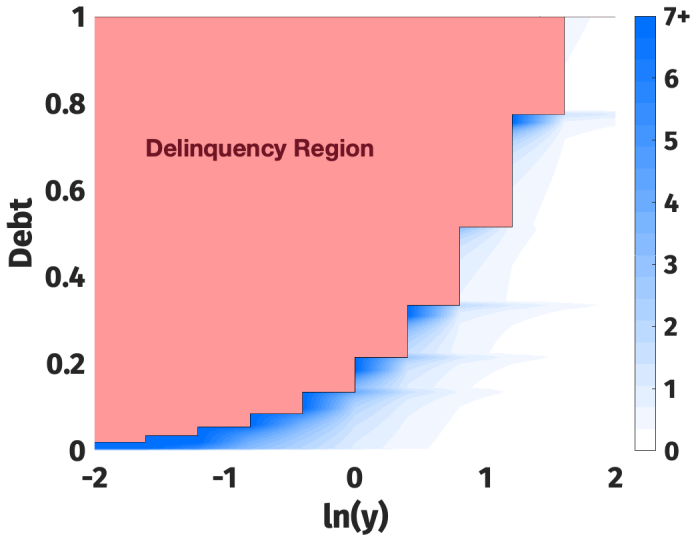
Our benchmark specification assumes that the share of households covered by Medicaid insurance is log-linear in household income percentile. Low-income households are more likely to be covered by Medicaid. In this section, we consider a policy that mimics the expansion of Medicaid as part of the ACA in the data. Following the results in the left panel of Figure 3, we change the coefficients of the Medicaid coverage function so that the share of households insured by Medicaid follows the post-expansion pattern across states. We maintain the same share of households insured by non-Medicaid insurance.

The resulting Medicaid insurance approximation implies that an additional 4.99 percent of households are covered by Medicaid. We solve the model and compute the stationary distribution, as well as the transition dynamics towards the new stationary distribution. We assume that the out-of-pocket share of the Medicaid policy remains unchanged at a rate of 6.8%.

The expansion of Medicaid reduces the delinquency probability of households, as health expenditure shocks that would push households into the delinquency region are not partially covered by their health insurance. This results in lower interest rate spreads in equilibrium.

The reduction in equilibrium spreads as a result of the policy is plotted in Figure 9. Households who are close to the delinquency region are now facing lower interest rate spreads relative to the pre-expansion interest rate spreads.

Figure 9: Reduction in interest rate spreads due to policy



Notes: This figure presents the reduction in equilibrium interest rate spreads in percentage points.

The reduction in interest rate spreads affects low-income households, who tend to be close to the default frontier. High-income households, who often do not hold any debt and are not covered by Medicaid, are less affected by the expansion of Medicaid.

The expansion of Medicaid in our model leads to a long-run increase of 4.1% in credit card debt. This increase in debt levels is not uniform across households. Low-income households, who are more directly exposed to the policy, respond by raising their credit card debt levels more relative to other households. The bottom income quintile of households raise their debt level by 16.7%. The response of the second and third income quintiles is about 8%. The top two income quintiles do not change their debt levels, on average. The response of each income quintile is presented in Appendix Figure A.6.

Our model allows us to decompose the aggregate increase in credit card debt into the three theoretical channels we laid out in the previous subsection: the direct channel, the credit demand channel, and the credit supply channel. The decomposition results are summarized in the top panel of Table 8.

Table 8: Decomposing the effect of Medicaid expansion

Not budget neutral

	Total impact	Direct impact	Credit demand	Credit supply
Credit card debt	**+4.08%**	**+0.91%**	**-11.12%**	**+14.29%**
Total welfare	**+0.41%**	**+0.30%**	**+0.004%**	**+0.11%**

Budget neutral

	Total impact	Direct impact	Credit demand	Credit supply
Credit card debt	**+4.36%**	**+0.91%**	**-11.04%**	**+14.50%**
Total welfare	**+0.35%**	**+0.23%**	**+0.004%**	**+0.11%**

To get the direct impact of Medicaid expansion, we keep the debt pricing schedule as well as the policy functions of households as they are prior to the expansion. The only change we consider here is the change to the out-of-pocket share of medical expenditure, $o(y)$. The direct impact of the policy on debt is an increase of 0.9% in the aggregate level of credit card debt.

The credit demand channel is computed by keeping the debt pricing schedule at its levels prior to the expansion of Medicaid but allowing households to re-optimize given the original pricing schedule. The credit demand channel reduces the aggregate-level credit by 11.1%. This implies that the debt aversion channel is much stronger than the precautionary savings channel.

Finally, the credit supply channel is computed by also updating the debt pricing schedule so that financial intermediaries make zero profits. The reduction in interest rates, which can be seen in Figure 9, leads to a large increase in credit card debt. The credit supply channel raises the aggregate level of credit card debt by 14.3%, leading to an overall increase of aggregate credit card debt across households.

Our model also allows us to study the welfare impact of the policy. We first consider the policy when budget is not neutral, so that no tax is levied on households. The model is useful in studying how important the different channels are in driving welfare, as well as comparing the welfare benefits across different households. Following Chatterjee et al. (2007), we calculate welfare by computing the percentage drop in consumption in all periods following the expansion of Medicaid that would make households indifferent between implementing and not implementing the policy. On average, the policy leads to a welfare benefit of 41 basis points in consumption-equivalent terms. That is, the average household in the economy

is willing to incur a 0.41% drop in consumption in all periods so that Medicaid expansion remains in place.

We decompose the effects of the Medicaid expansion to the three channels. The results are presented in the second row of Table 8. The reduction in out-of-pocket medical expenditure accounts for the majority of welfare benefits. This effect is expected: we assume households do not pay any cost to implement the policy. The direct channel accounts for 73% of the welfare gains.

The credit demand channel has only a negligible welfare effect. This is simply the envelope theorem. Households were optimizing their borrowing decision prior to the policy. So adjusting their borrowing decision following the policy can only lead to second-order welfare gains.

Unlike the credit demand channel, the credit supply channel leads to sizable welfare benefits: 0.11% out of a total of 0.41%. That is, 27% out of the total welfare gains of the expansion of Medicaid can be attributed to the reduction in interest rate spreads households pay. This result suggests policymakers should take into account the impact of social insurance policies on the supply of credit. Disregarding the effect of social insurance on the supply of credit understates its welfare benefits.

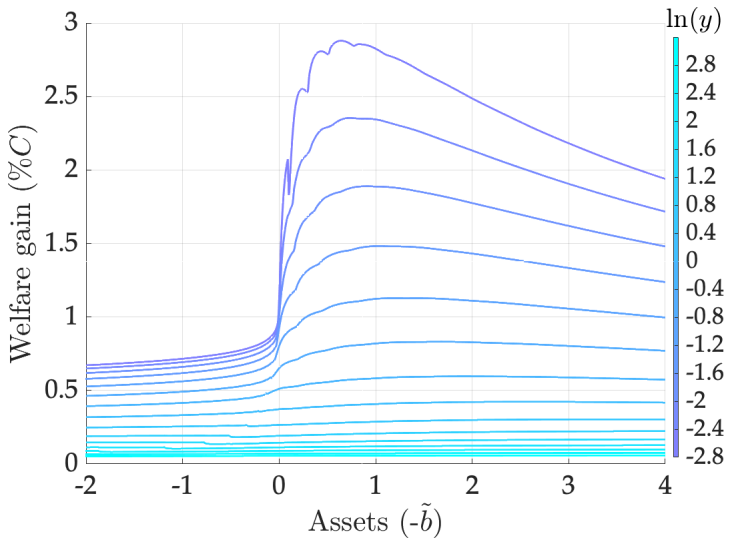
Finally, we consider an alternative in which the policy is financed through a uniform tax levied on all households. Medicaid expansion reduces uncompensated medical care. We choose the uniform tax rate so that the tax revenues are equal to the cost of enacting the policy net of the reduction in uncompensated medical care.[40] The resulting uniform tax is 0.043%. The bottom panel of Table 8 presents the effect of the budget-neutral Medicaid expansion on debt and welfare.

The aggregate impact on credit card debt is slightly larger under the budget-neutral policy. Households' disposable income is lower as a result of the tax, so their debt level goes up. The direct welfare benefits from the program fall from 0.41% to 0.35%. Note, however, that the welfare benefits attributed to the credit supply channel remain approximately the same. The credit supply channel amplifies the direct welfare benefits of Medicaid expansion

[40] Uncompensated medical care is defined as the sum of unpaid out-of-pocket medical bills net of the market value of medical debt due to unpaid medical bills.

by almost 50%.

Figure 10: Heterogeneous welfare effect of budget-neutral Medicaid expansion



Notes: This figure presents the welfare gains (in consumption equivalence terms) from Medicaid expansion across the state space. The lighter shades of blue correspond to higher income households. Negative levels of assets correspond to positive debt levels.

Figure 10 presents the net welfare gains from the budget-neutral policy. We see that the majority of the welfare gains come from low-income non-delinquent households. The less income a household has, the more health insurance they receive, and the higher their welfare gains from the policy. It is worth noting that even high-income households do not suffer welfare losses despite the increase in their income tax. The top 5% of the income distribution obtain welfare gains of between 0.05%–0.1%. These high-income households value the insurance they gain from the policy more than the income tax they pay to finance it.

5 Conclusion

This paper investigates how social insurance affects household debt. We exploit heterogeneity in the impact of the staggered expansions of Medicaid as a source of quasi-experimental variation in households' access to health insurance. Using a difference-in-difference design, we estimate that a 1 percentage point increase in health insurance coverage leads to a 0.46% increase in credit card debt. Consistent with credit supply increasing in response to expanded Medicaid, we find that credit limits rise, overall utilization decreases, and in below-median

income ZIP codes credit card inquiries become more successful. In support of improved financial resilience driving this rise in credit supply, we show that debt in collections falls, and delinquencies either fall or remain stable despite the increase in debt. Credit scores increase as well.

Our paper builds on prior empirical work by focusing on general equilibrium channels and both macroeconomic and distributional outcomes. We develop a heterogeneous-agent model in which households face permanent and transitory differences in their income, health expenditure shocks, and incomplete markets. Households incur debt both by choosing how much to borrow on a credit card and as a result of health expenditure shocks. Using the model, we explore the impact of expanding health insurance.

While insurance can help households avoid taking on debt when experiencing adverse events such as job loss and illness, it can also increase borrowing by enhancing their financial resilience. Insurance softens the financial impact of adverse events, making it easier to avoid default or states of the world in which consumption is extremely low. In doing so, insurance can dampen households' precautionary savings motive and raise lenders' expected returns, increasing both credit demand and supply. Our findings suggest that credit supply drives the equilibrium rise in borrowing induced by expanded health insurance. Public insurance can have an important impact on the quantity and distribution of household debt, and the reaction of credit markets can be an large part of the overall welfare benefits of these policies.

References

Aaronson, Daniel, Sumit Agarwal, and Eric French, "The Spending and Debt Response to Minimum Wage Hikes," *American Economic Review*, 2012, *102* (7), 3111–3139.

Adams, Robert, Vialty Bord, and Bradley Katcher, "Credit Card Profitability," *FEDS Notes*, 2022.

Aguiar, Mark A, Mark Bils, and Corina Boar, "Who are the Hand-to-Mouth?," Technical Report, National Bureau of Economic Research 2020.

Aiyagari, Rao S., "Uninsured Idiosyncratic Risk and Aggregate Saving," *The Quarterly Journal of Economics*, 1994, *109* (3), 659–684.

Allen, Heidi, Ashley Swanson, Jialan Wang, and Tal Gross, "Early Medicaid Expansion Associated with Reduced Payday Borrowing in California," *Health Affairs*, 2017, *36* (10), 1689–1776.

Arellano, Cristina, "Default risk and income fluctuations in emerging economies," *American economic review*, 2008, *98* (3), 690–712.

Attanasio, Orazio, Sagiri Kitao, and Giovanni L Violante, "Financing Medicare: A general equilibrium analysis," in "Demography and the Economy," University of Chicago Press, 2010, pp. 333–366.

Baker, Andrew C, David F Larcker, and Charles CY Wang, "How much should we trust staggered difference-in-differences estimates?," *Journal of Financial Economics*, 2022, *144* (2), 370–395.

Bankrate, "Mortgage Rates by Credit Score," https://www.bankrate.com/real-estate/what-credit-score-do-you-need-to-buy-a-house/ 2025.

Bewley, Truman, "Stationary Monetary Equilibrium with a Continuum of Independently Fluctuating Consumers," *Contributions to Mathematical Economics in Honor of Gérard Debreu*, 1986, *79*.

Braxton, J. Carter, Kyle Herkenhoff, and Gordon M. Phillips, "Can the Unemployed Borrow? Implications for Public Insurance," 2022.

Brevoort, Kenneth, Daniel Grodzicki, and Martin B. Hackmann, "Medicaid and Financial Health," Working Paper, National Bureau of Economic Research 2017.

Célérier, Claire and Adrien Matray, "Public Health Insurance, Access to Credit, and Wealth Accumulation," 2017.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer, "The effect of minimum wages on low-wage jobs," *The Quarterly Journal of Economics*, 2019, *134* (3), 1405–1454.

Chatterjee, Satyajit and Burcu Eyigungor, "Maturity, indebtedness, and default risk," *American Economic Review*, 2012, *102* (6), 2674–2699.

_ , Dean Corbae, Kyle P Dempsey, and José-Víctor Ríos-Rull, "A quantitative theory of the credit score," Technical Report, National Bureau of Economic Research 2020.

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

Congrssional Budget Office, "Letter to John Boehner regarding H.R. 6079, the Repeal of Obamacare Act," 2012, https://perma.cc/Q9VQ-C63D.

Courtemanche, Charles J, James Marton, and Aaron Yelowitz, "Medicaid coverage across the income distribution under the Affordable Care Act," Technical Report, National Bureau of Economic Research 2019.

Deaton, Angus, "Saving and Liquidity Constraints," *Econometrica*, 1991, *59* (5), 1221–48.

Dobkin, Carlos, Amy Finkelstein, Raymond Kluender, and Matthew J Notowidigdo, "The economic consequences of hospital admissions," *American Economic Review*, 2018, *108* (2), 308–52.

Dranove, David, Craig Garthwaite, and Christopher Ody, "Uncompensated care decreased at hospitals in Medicaid expansion states but not at hospitals in nonexpansion states," *Health Affairs*, 2016, *35* (8), 1471–1479.

Druedahl, Jeppe and Casper Nordal Jørgensen, "Precautionary Borrowing and the Credit Card Debt Puzzle," *Quantitative Economics*, 2018, *9* (2), 785–823.

Experian, "State of the Automotive Finance Market: Q3 2025," https://www.experian.com/content/dam/noindex/na/us/automotive/finance-trends/2025/Experian-SAFM-Q3-2025.pdf September 2025.

_ , "2026 State of the U.S. Housing Market Report," https://www.experian.com/content/dam/marketing/na/thought-leadership/business/documents/2026-state-of-housing-report.pdf January 2026.

Finkelstein, Amy, Sarah Taubman, Bill Wright, Mira Bernstein, Jonathan Gruber, Joseph P Newhouse, Heidi Allen, Katherine Baicker, and Oregon Health Study Group, "The Oregon health insurance experiment: evidence from the first year," *The Quarterly journal of economics*, 2012, *127* (3), 1057–1106.

Frean, Molly, Jonathan Gruber, and Benjamin D Sommers, "Premium subsidies, the mandate, and Medicaid expansion: Coverage effects of the Affordable Care Act," *Journal of health economics*, 2017, *53*, 72–86.

Gallagher, Emily A., Jorge Sabat Radhakrishnan Gopalan, and Michal Grinstein-Weiss, "Medicaid and Household Savings Behavior: New Evidence from Tax Refunds," *Journal of Financial Economics*, 2019.

_ **, Radhakrishnan Gopalan, and Michael Grinstein-Weiss**, "The Effect of Health Insurance on Home Payment Delinquency: Evidence from ACA Marketplace Subsidies," *Journal of Public Economics*, 2019, *172*, 67–83.

Garthwaite, Craig, John A Graves, Tal Gross, Zeynal Karaca, Victoria R Marone, and Matthew J Notowidigdo, "All Medicaid expansions are not created equal: the geography and targeting of the Affordable Care Act," Technical Report, National Bureau of Economic Research 2019.

Goldsmith-Pinkham, Paul, Maxim Pinkovskiy, and Jacob Wallace, "The Geography of Financial health," 2020.

Goodman-Bacon, Andrew, "Public insurance and mortality: evidence from Medicaid implementation," *Journal of Political Economy*, 2018, *126* (1), 216–262.

_ , "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

_ , "The long-run effects of childhood insurance coverage: Medicaid implementation, adult health, and labor market outcomes," *American Economic Review*, 2021, *111* (8), 2550–93.

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

Gross, Tal and Matthew J. Notowidigdo, "Health Insurance and the Consumer Bankruptcy Decision: Evidence from Expansion Medicaid," *Journal of Public Economics*, 2011, *95* (7), 767–778.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song, "What do data on millions of US workers reveal about life-cycle earnings risk?," Technical Report, National Bureau of Economic Research 2015.

Hatchondo, Juan Carlos and Leonardo Martinez, "Long-duration bonds and sovereign defaults," *Journal of international Economics*, 2009, *79* (1), 117–125.

Himmelstein, David U, Deborah Thorne, Elizabeth Warren, and Steffie Woolhandler, "Medical bankruptcy in the United States, 2007: results of a national study," *The American journal of medicine*, 2009, *122* (8), 741–746.

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

Hu, Luojia, Robert Kaestner, Bhashkar Mazumder, Sarah Miller, and Ashley Wong, "The Effect of the Affordable Care Act Medicaid Expansions on Financial Wellbeing," *Journal of Public Economics*, 2018, *163*, 99–112.

Hubbard, R Glenn, Jonathan Skinner, and Stephen P Zeldes, "Precautionary saving and social insurance," *Journal of political Economy*, 1995, *103* (2), 360–399.

Huggett, Mark, "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies," *Journal of Economic Dynamics and Control*, 1993, *17* (5–6), 953–969.

Imrohoroğlu, Ayşe, "Cost of business cycles with indivisibilities and liquidity constraints," *Journal of Political economy*, 1989, *97* (6), 1364–1383.

Indarte, Sasha, "Moral hazard versus liquidity in household bankruptcy," *Journal of Finance (forthcoming)*, 2023.

Kaestner, Robert, Bowen Garrett, Jiajia Chen, Anuj Gangopadhyaya, and Caitlyn Fleming, "Effects of ACA Medicaid expansions on health insurance coverage and labor supply," *Journal of Policy Analysis and Management*, 2017, *36* (3), 608–642.

Krueger, Dirk and Fabrizio Perri, "Public versus private risk sharing," *Journal of Economic Theory*, 2011, *146* (3), 920–956.

Livshits, Igor, James MacGee, and Michele Tertilt, "Consumer Bankruptcy: A Fresh Start," *American Economic Review*, 2007, *97* (1), 402–418.

Manning, Willard G, Joseph P Newhouse, Naihua Duan, Emmett B Keeler, and Arleen Leibowitz, "Health insurance and the demand for medical care: evidence from a randomized experiment," *The American economic review*, 1987, pp. 251–277.

Matsa, David A., Brian T. Melzer, and Michael Zator, "Dual Credit Markets: Income Risk, Household Debt, and Consumption," *Working Paper*, 2022.

Mazumder, Bhashkar and Sarah Miller, "The effects of the Massachusetts health reform on household financial distress," *American Economic Journal: Economic Policy*, 2016, *8* (3), 284–313.

Miller, Sarah, Luojia Hu, Robert Kaestner, Bhashkar Mazumder, and Ashley Wong, "The ACA Medicaid Expansion in Michigan and Financial Health," Working Paper 2018.

_ , Norman Johnson, and Laura R Wherry, "Medicaid and mortality: new evidence from linked survey and administrative data," *The Quarterly Journal of Economics*, 2021, *136* (3), 1783–1829.

Mitman, Kurt, "Macroeconomic Effects of Bankruptcy and Foreclosure Policies," *American Economic Review*, 2016, *106* (8), 2219–2255.

Nardi, Mariacristina De, Eric French, and John B Jones, "Why do the elderly save? The role of medical expenses," *Journal of political economy*, 2010, *118* (1), 39–75.

Rambachan, Ashesh and Jonathan Roth, "A more credible approach to parallel trends," *Review of Economic Studies*, 2023, *90* (5), 2555–2591.

Roth, Jonathan, Pedro HC Sant'Anna, Alyssa Bilinski, and John Poe, "What's trending in difference-in-differences? A synthesis of the recent econometrics literature," *arXiv preprint arXiv:2201.01194*, 2022.

Shupe, Cortnie, "Public Health Insurance and Medical Spending: The Incidence of the ACA Medicaid Expansion," *Journal of Policy Analysis and Management*, 2023, *42* (1), 137–165.

Simon, Kosali, Aparna Soni, and John Cawley, "The impact of health insurance on preventive care and health behaviors: evidence from the first two years of the ACA Medicaid expansions," *Journal of Policy Analysis and Management*, 2017, *36* (2), 390–417.

Sommers, Benjamin D, Genevieve M Kenney, and Arnold M Epstein, "New evidence on the Affordable Care Act: coverage impacts of early Medicaid expansions," *Health affairs*, 2014, *33* (1), 78–87.

Stepner, Michael, "The Insurance Value of Redistributive Taxes and Transfers," 2019.

Telyukova, Irina A., "Household Need for Liquidity and the Credit Card Debt Puzzle," *Review of Economic Studies*, 2013, *80* (3), 1148–1177.

Zeldes, Stephen P, "Optimal consumption with stochastic income: Deviations from certainty equivalence," *The Quarterly Journal of Economics*, 1989, *104* (2), 275–298.

Zinman, Jonathan, "Where is the missing credit card debt? Clues and implications," *Review of income and wealth*, 2009, *55* (2), 249–265.

Online Appendix

A Additional Figures and Tables

Table A.1: Adding Flexible Income Controls (Revolving Credit Card Balances)

	(1)	(2)	(3)	(4)
$\text{NewElig}_{zs} \times \text{Post}_{st}$	0.557***	0.592***	0.455***	0.434***
	(0.116)	(0.102)	(0.076)	(0.073)
Obs	372,218	372,218	372,218	372,218
R2	0.957	0.957	0.957	0.958
Percentile Controls	None	Quartiles	Deciles	5th, 10th,...95th

Notes: This table reports results from regressions that control for additional moments from the ZIP code's income distribution for our main outcome of interest: log revolving credit card balances. All columns control for log average AGI (coefficients omitted). Column 2 adds controls for logged quartiles (i.e., the 25th, 50th, and 75th percentile of the ZIP code's income distribution. Column 3 instead uses deciles (10th, 20th, ..., 90th) percentiles. Column 4 uses 19 percentiles ranging from the 5th to 95th percentile. Standard errors are clustered by state. Nominal variables are CPI-adjusted to be in terms of 2020 dollars. Statistical significance: 10%*, 5%**, and 1%***.

Table A.2: Placebo Test: Income Interactions

	(1)	(2)	(3)
$\text{NewElig}_{zs} \times \text{Post}_{st}$	0.455***	0.440***	0.518***
	(0.073)	(0.063)	(0.111)
$\text{HighInc}_{zst} \times \text{Post}_{st}$		-0.002	
		(0.003)	
$\log(\text{AGI}^{\text{avg}}_{zst}) \times \text{Post}_{st}$			0.008
			(0.010)
Obs	369,728	369,728	369,728
R2	0.957	0.957	0.957

Notes: The dependent variable in all columns is log revolving credit card balances. Column 1 here reproduces our baseline estimate from column 4 of Table 3 to facilitate comparison. Each specification uses county-time and ZIP code fixed effects and controls for logged ZIP-code level average AGI and logged deciles of the ZIP code's income distribution (coefficients omitted). Column 2 adds an interaction of the post-expansion indicator with a binary indicator HighInc_{zst} which equals one if ZIP code z had average AGI above the median of all ZIP codes' average AGI in year t. An un-interacted HighInc_{zst} is included too but its coefficient is omitted for brevity. Column 3 adds an interaction of the post-expansion indicator with ZIP z's logged average AGI in year t (recall that un-interacted log AGI is already included as a control). Standard errors are clustered by state. Nominal variables are CPI-adjusted to 2020 dollars. Statistical significance: 10%*, 5%**, 1%***.

A.1 Treatment Effect Heterogeneity Robustness

Table A.3: Comparison of Non-Stacked and Stacked DID Estimates (1/3)

	Expand = 2014		Expand $\leq$ 2015		Expand $\leq$ 2016	
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: Outcome = 1[Has CC]					
NewElig$_{zs}\times$ Post$_{st}$	0.131***	0.132***	0.108***	0.108***	0.073***	0.073***
	(0.026)	(0.027)	(0.023)	(0.024)	(0.023)	(0.023)
Obs	60,871	100,505	62,525	133,885	57,223	152,386
R2	0.973	0.979	0.974	0.981	0.977	0.982
	Panel B: Outcome = 1[New CC]					
NewElig$_{zs}\times$ Post$_{st}$	0.171***	0.171***	0.160***	0.153***	0.148***	0.138***
	(0.017)	(0.019)	(0.018)	(0.020)	(0.017)	(0.020)
Obs	60,871	100,505	62,525	133,885	57,223	152,386
R2	0.843	0.87	0.849	0.879	0.863	0.888
	Panel C: Outcome = log(CC Bal.)					
NewElig$_{zs}\times$ Post$_{st}$	0.817***	0.864***	0.735***	0.755***	0.656***	0.665***
	(0.079)	(0.088)	(0.064)	(0.078)	(0.055)	(0.069)
Obs	192,430	318,319	207,331	450,667	188,578	525,943
R2	0.970	0.971	0.970	0.971	0.972	0.972
	Panel D: Outcome = log(CC Bal. Rev.)					
NewElig$_{zs}\times$ Post$_{st}$	0.425***	0.471***	0.363***	0.378***	0.297***	0.293***
	(0.075)	(0.085)	(0.058)	(0.073)	(0.053)	(0.067)
Obs	192,428	318,317	207,329	450,665	188,576	525,941
R2	0.956	0.959	0.956	0.959	0.958	0.960
Window	-3 to +6 years		-3 to +5 years		-3 to +4 years	
Stacked	No	Yes	No	Yes	No	Yes

Notes: This table reports results from estimating the "stacked" version of the DID in Equation (1). Columns 1 and 2 use the sub-sample of observations whose states expanded in 2014 using data from 3 years before to 6 years after expansion (the window). Columns 3 and 4 expand the sample to observations whose states expanded in 2015 or earlier but use a narrower window of 3 years before to 5 years after. Columns 5 and 6 expand the sample to include observations whose states expanded in 2016 or earlier, but restrict the window to 3 years before to 4 years after expanding. Each panel these six regressions for different outcome variables. To facilitate comparisons, the odd-numbered columns display non-stacked regression results for the sub-samples used in the stacked regressions. Each specification uses county-time and ZIP code fixed effects and controls for logged ZIP-code level average adjusted gross income (AGI) and logged deciles of the ZIP code's income distribution (coefficients omitted). Standard errors are clustered by state. Nominal variables are CPI-adjusted to be in terms of 2020 dollars. Statistical significance: 10%*, 5%**, and 1%***.

Table A.4: Comparison of Non-Stacked and Stacked DID Estimates (2/3)

	Expand = 2014		Expand $\leq$ 2015		Expand $\leq$ 2016	
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: Outcome = CC Inquiries (#)					
$\text{NewElig}_{zs} \times \text{Post}_{st}$	0.411***	0.419***	0.383***	0.384***	0.350***	0.341***
	(0.038)	(0.037)	(0.038)	(0.039)	(0.036)	(0.040)
Obs	60,871	100,505	62,525	133,885	57,223	152,386
R2	0.931	0.939	0.934	0.943	0.938	0.946
	Panel B: Outcome = log(CC Limits)					
$\text{NewElig}_{zs} \times \text{Post}_{st}$	0.451***	0.478***	0.383***	0.360***	0.325***	0.275***
	(0.094)	(0.103)	(0.073)	(0.090)	(0.066)	(0.087)
Obs	60,871	100,505	62,525	133,885	57,223	152,386
R2	0.986	0.987	0.987	0.988	0.988	0.989
	Panel C: Outcome = CC Utilization (%)					
$\text{NewElig}_{zs} \times \text{Post}_{st}$	-0.305***	-0.313***	-0.281***	-0.277***	-0.254***	-0.244***
	(0.028)	(0.030)	(0.023)	(0.027)	(0.022)	(0.028)
Obs	60,871	100,505	62,525	133,885	57,223	152,386
R2	0.976	0.981	0.978	0.982	0.981	0.984
	Panel D: Outcome = New CC to Inquiries Ratio					
$\text{NewElig}_{zs} \times \text{Post}_{st}$	0.059	0.037	0.073**	0.04	0.082**	0.047
	(0.035)	(0.038)	(0.035)	(0.038)	(0.034)	(0.036)
Obs	60,871	100,505	62,525	133,885	57,223	152,386
R2	0.840	0.847	0.839	0.846	0.847	0.851
Window	-3 to +6 years		-3 to +5 years		-3 to +4 years	
Stacked	No	Yes	No	Yes	No	Yes

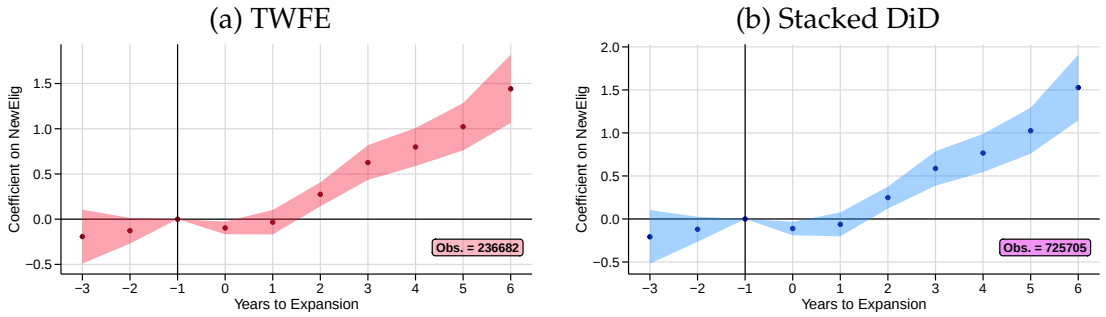
Notes: This table reports results from estimating the "stacked" version of the DID in Equation (1). Columns 1 and 2 use the sub-sample of observations whose states expanded in 2014 using data from 3 years before to 6 years after expansion (the window). Columns 3 and 4 expand the sample to observations whose states expanded in 2015 or earlier but use a narrower window of 3 years before to 5 years after. Columns 5 and 6 expand the sample to include observations whose states expanded in 2016 or earlier, but restrict the window to 3 years before to 4 years after expanding. Each panel these six regressions for different outcome variables. To facilitate comparisons, the odd-numbered columns display non-stacked regression results for the sub-samples used in the stacked regressions. Each specification uses county-time and ZIP code fixed effects and controls for logged ZIP-code level average adjusted gross income (AGI) and logged deciles of the ZIP code's income distribution (coefficients omitted). Standard errors are clustered by state. Nominal variables are CPI-adjusted to be in terms of 2020 dollars. Statistical significance: 10%*, 5%**, and 1%***.

Table A.5: Comparison of Non-Stacked and Stacked DID Estimates (3/3)

	Expand = 2014		Expand ≤ 2015		Expand ≤ 2016	
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: Outcome = 1[90+ DPD]					
$\text{NewElig}_{zs} \times \text{Post}_{st}$	-3e-3	-0.006	0.006	4e-5	0.008	0.002
	(0.013)	(0.013)	(0.011)	(0.011)	(0.010)	(0.010)
Obs	60,871	100,505	62,525	133,885	57,223	152,386
R2	0.893	0.902	0.894	0.903	0.899	0.907
	Panel B: Outcome = 1[Utilities 90+ DPD]					
$\text{NewElig}_{zs} \times \text{Post}_{st}$	-0.130**	-0.133**	-0.127**	-0.127**	-0.116**	-0.114**
	(0.059)	(0.060)	(0.049)	(0.050)	(0.045)	(0.046)
Obs	60,871	100,505	62,525	133,885	57,223	152,386
R2	0.917	0.923	0.918	0.929	0.927	0.938
	Panel C: Outcome = 1[Any Col.]					
$\text{NewElig}_{zs} \times \text{Post}_{st}$	-0.141***	-0.153***	-0.133***	-0.143***	-0.136***	-0.148***
	(0.027)	(0.028)	(0.023)	(0.026)	(0.022)	(0.026)
Obs	60,871	100,505	62,525	133,885	57,223	152,386
R2	0.978	0.983	0.978	0.984	0.980	0.984
	Panel D: Outcome = Collections/Income (%)					
$\text{NewElig}_{zs} \times \text{Post}_{st}$	-0.042***	-0.044***	-0.039***	-0.041***	-0.040***	-0.041***
	(0.010)	(0.010)	(0.009)	(0.010)	(0.008)	(0.010)
Obs	60,871	100,505	62,525	133,885	57,223	152,386
R2	0.929	0.940	0.929	0.941	0.932	0.942
	Panel E: Outcome = Vantage Score					
$\text{NewElig}_{zs} \times \text{Post}_{st}$	0.231**	0.264**	0.186**	0.204*	0.163**	0.178*
	(0.108)	(0.117)	(0.086)	(0.103)	(0.078)	(0.099)
Obs	60871	100,505	62,525	133,885	57,223	152,386
R2	0.986	0.988	0.986	0.988	0.987	0.989
Window	-3 to +6 years		-3 to +5 years		-3 to +4 years	
Stacked	No	Yes	No	Yes	No	Yes

Notes: This table reports results from estimating the "stacked" version of the DID in Equation (1). Columns 1 and 2 use the sub-sample of observations whose states expanded in 2014 using data from 3 years before to 6 years after expansion (the window). Columns 3 and 4 expand the sample to observations whose states expanded in 2015 or earlier but use a narrower window of 3 years before to 5 years after. Columns 5 and 6 expand the sample to include observations whose states expanded in 2016 or earlier, but restrict the window to 3 years before to 4 years after expanding. Each panel these six regressions for different outcome variables. To facilitate comparisons, the odd-numbered columns display non-stacked regression results for the sub-samples used in the stacked regressions. Each specification uses county-time and ZIP code fixed effects and controls for logged ZIP-code level average adjusted gross income (AGI) and logged deciles of the ZIP code's income distribution (coefficients omitted). Standard errors are clustered by state. Nominal variables are CPI-adjusted to be in terms of 2020 dollars. Statistical significance: 10%*, 5%**, and 1%***.

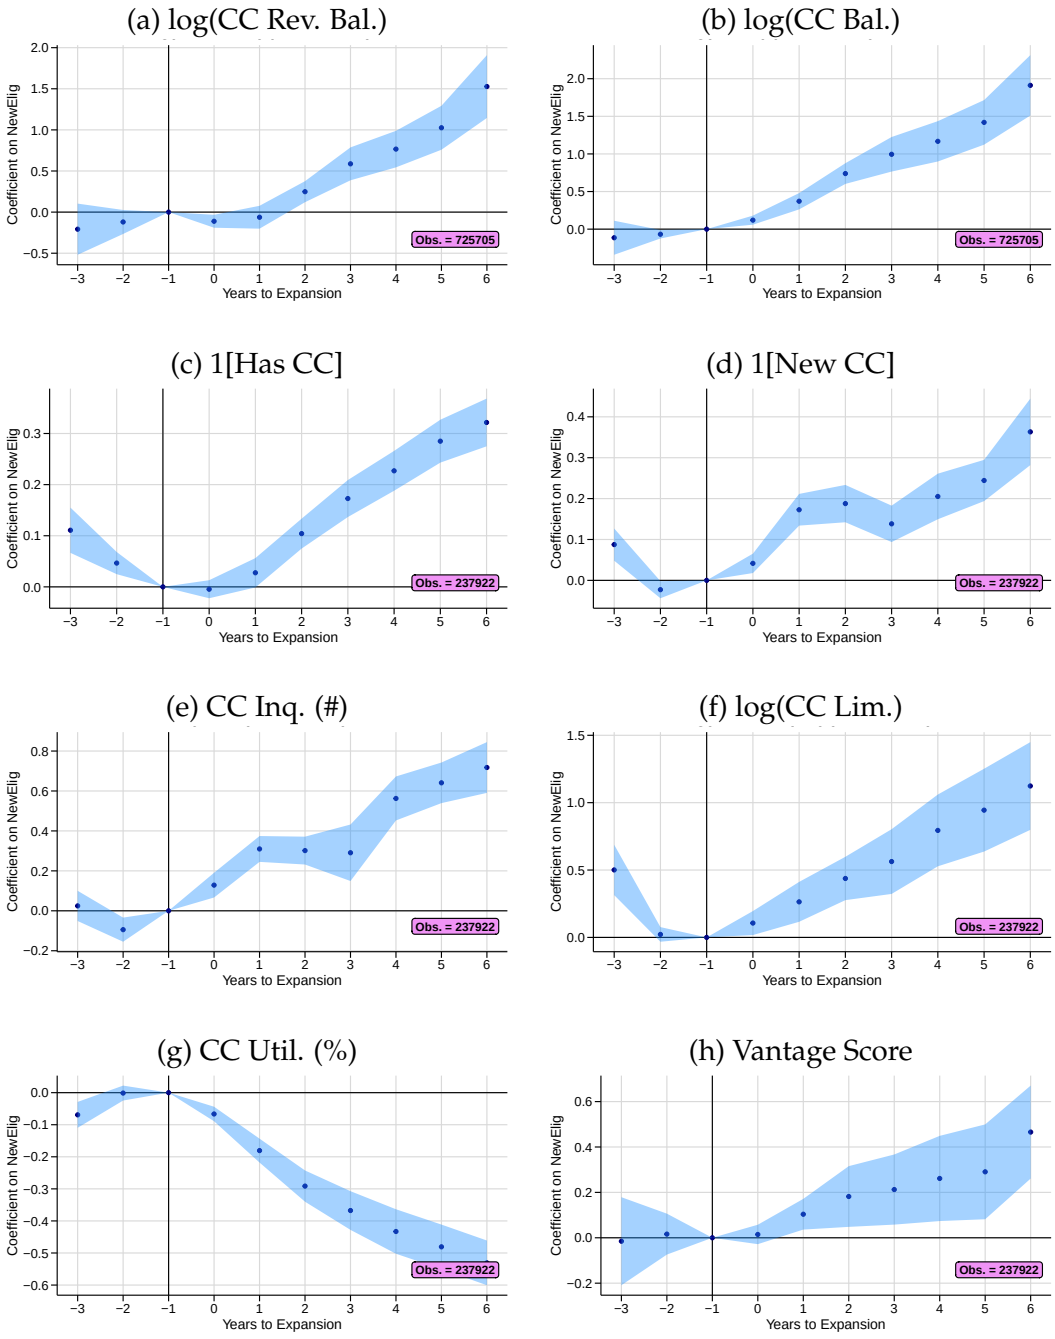
Figure A.1: Comparison of TWFE and Stacked DiD Estimates: log(CC Rev. Bal.)



Notes: Panel (a) reports dynamic difference-in-differences estimates from a standard two-way fixed effects specification. Panel (b) reports stacked difference-in-differences estimates following Cengiz et al. (2019), which addresses potential bias from heterogeneous treatment effects across cohorts. Shaded areas represent 95% confidence intervals for standard errors clustered at the state-level. The sample size discrepancy arises because ZIPs from "never treated" states appear multiple times when stacking data by expansion year cohort.

A.2 Dynamic DiD (additional outcomes)

Figure A.2: Credit Card Outcomes and Credit Score



Notes: Each panel reports stacked difference-in-differences estimates following Cengiz et al. (2019). The shaded area represents the 95% confidence interval for standard errors clustered at the state-level. The sample size discrepancy across panels is due to quarterly availability for credit card balances (total and revolving), whereas the remaining outcomes are measured annually. Note that the rise in total balances in early periods when revolving balances have not yet risen indicates that borrowers are making larger payments (such that their revolving balances do not increase despite increased card usage).

A.3 Heterogeneity by ZIP Code Income

Table A.6: Borrowing Outcomes (Low versus High Income)

	1[Has CC] (1)	1[New CC] (2)	log(CC Bal.) (3)	log(CC Rev. Bal.) (4)
Below-Median Income				
$\text{NewElig}_{zs} \times \text{Post}_{st}$	0.049*	0.104***	0.548***	0.507***
	(0.025)	(0.024)	(0.077)	(0.075)
Obs	61,258	61,258	192,918	192,915
R2	0.975	0.898	0.951	0.940
Mean	79%	20%	$3,193	$2,818
Above-Median Income				
$\text{NewElig}_{zs} \times \text{Post}_{st}$	0.128***	0.079***	0.498***	0.066
	(0.013)	(0.023)	(0.070)	(0.096)
Obs	56,370	56,370	179,304	179,303
R2	0.956	0.823	0.956	0.938
Mean	91%	24%	$5,489	$4,607

Notes: This table reports results from estimating the DID in Equation (1) separately within sub-samples (below- and above-median income ZIP codes). Each specification uses county-time and ZIP codes fixed effects and controls for logged ZIP-code level average adjusted gross income (AGI) and logged deciles of the ZIP code's income distribution (coefficients omitted). Nominal variables are CPI-adjusted to be in terms of 2020 dollars. The dependent variable is labeled above the column number, and its mean is reported below the R2. Statistical significance: 10%*, 5%**, and 1%***.

Table A.7: Demand and Supply Proxies (Low versus High Income)

	Demand		Supply	
	CC Inq. (#)	log(CC Lim.)	CC Util. (%)	New CC to Inq.
	(1)	(2)	(3)	(4)
Below-Median Income				
$\text{NewElig}_{zs} \times \text{Post}_{st}$	0.305***	0.459***	-0.143***	0.100***
	(0.055)	(0.113)	(0.036)	(0.035)
Obs	61,258	61,258	61,258	61,258
R2	0.960	0.974	0.977	0.862
Mean	0.448	$12,156	39%	0.475
Above-Median Income				
$\text{NewElig}_{zs} \times \text{Post}_{st}$	0.160***	0.136	-0.281***	-0.046
	(0.050)	(0.160)	(0.030)	(0.039)
Obs	56,370	56,370	56,370	56,370
R2	0.934	0.978	0.960	0.821
Mean	0.429	$24,378	23%	0.571

Notes: This table reports results from estimating the DID in Equation (1) separately within sub-samples (below- and above-median income ZIP codes). Each specification uses county-time and ZIP codes fixed effects and controls for logged ZIP-code level average adjusted gross income (AGI) and logged deciles of the ZIP code's income distribution (coefficients omitted). Nominal variables are CPI-adjusted to be in terms of 2020 dollars. The dependent variable is labeled above the column number, and its mean is reported below the R2. Statistical significance: 10%*, 5%**, and 1%***.

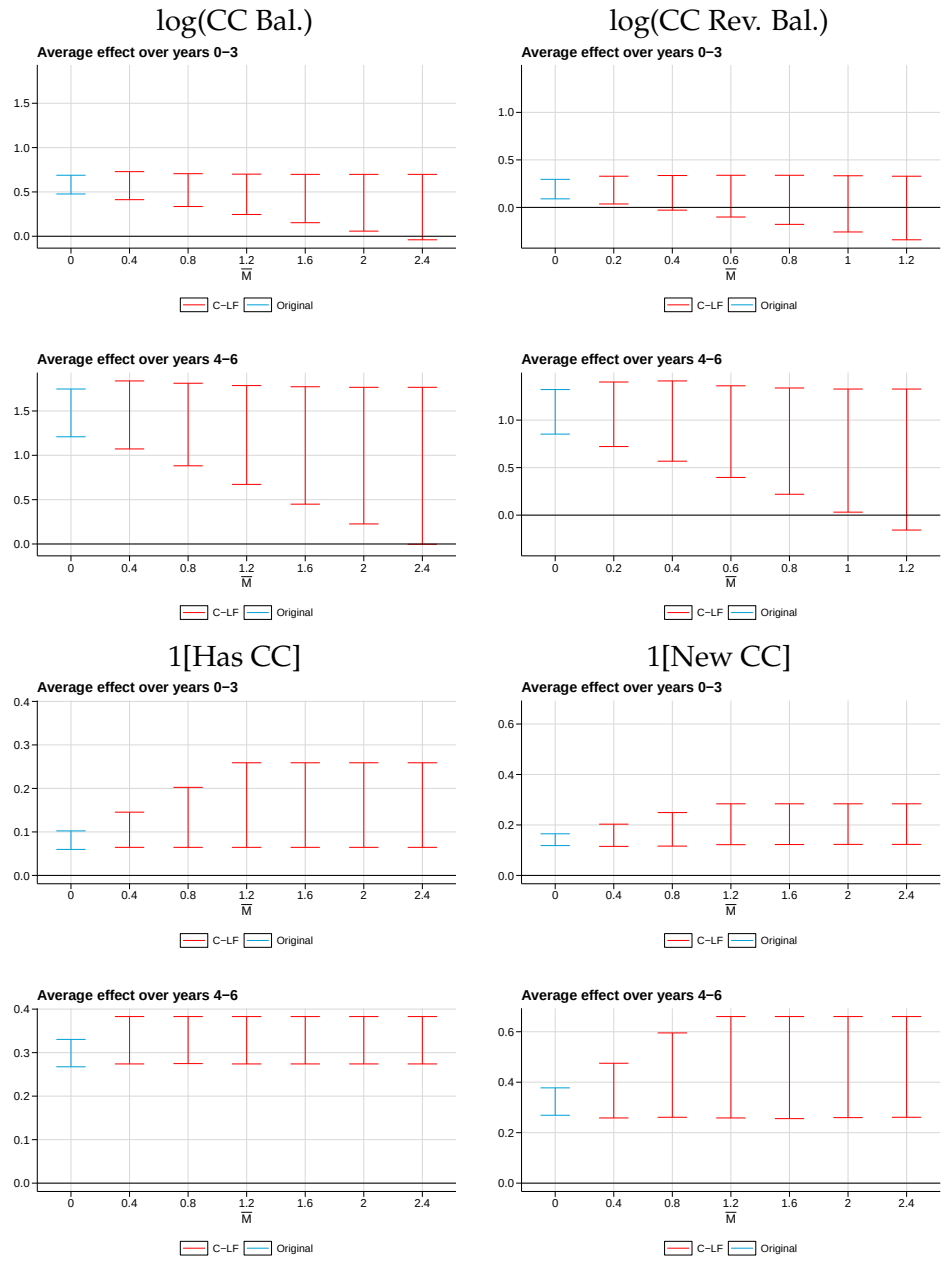
Table A.8: Default and Credit Scores (Low versus High Income)

	1[Delinquency]		Debt in Collections		Credit Score
	1[90+ Days]	1[90+ (Util.)]	1[Any Col.]	Col./Inc.	Vantage Score
	(1)	(2)	(3)	(4)	(5)
Below-Med. Inc.					
$\text{NewElig}_{zs} \times \text{Post}_{st}$	0.019	-0.157**	-0.097***	-0.024**	0.101
	(0.020)	(0.069)	(0.032)	(0.009)	(0.069)
Obs	61,258	61,258	61,258	61,258	61,258
R2	0.884	0.942	0.942	0.979	0.983
Mean	8.9%	5.4%	31%	1.8%	659.65
Above-Med. Inc.					
$\text{NewElig}_{zs} \times \text{Post}_{st}$	-0.033***	-0.047**	-0.106***	-0.023***	0.433***
	(0.009)	(0.022)	(0.240)	(0.006)	(0.111)
Obs	56,370	56,370	56,370	56,370	56,370
R2	0.853	0.918	0.969	0.895	0.979
Mean	5.2%	3.5%	16%	0.5%	710.31

Notes: This table reports results from estimating the DID in Equation (1) separately within sub-samples (below- and above-median income ZIP codes). Each specification uses county-time and ZIP codes fixed effects and controls for logged ZIP-code level average adjusted gross income (AGI) and logged deciles of the ZIP code's income distribution (coefficients omitted). Nominal variables are CPI-adjusted to be in terms of 2020 dollars. The dependent variable is labeled above the column number, and its mean is reported below the R2. Statistical significance: 10%*, 5%**, and 1%***.

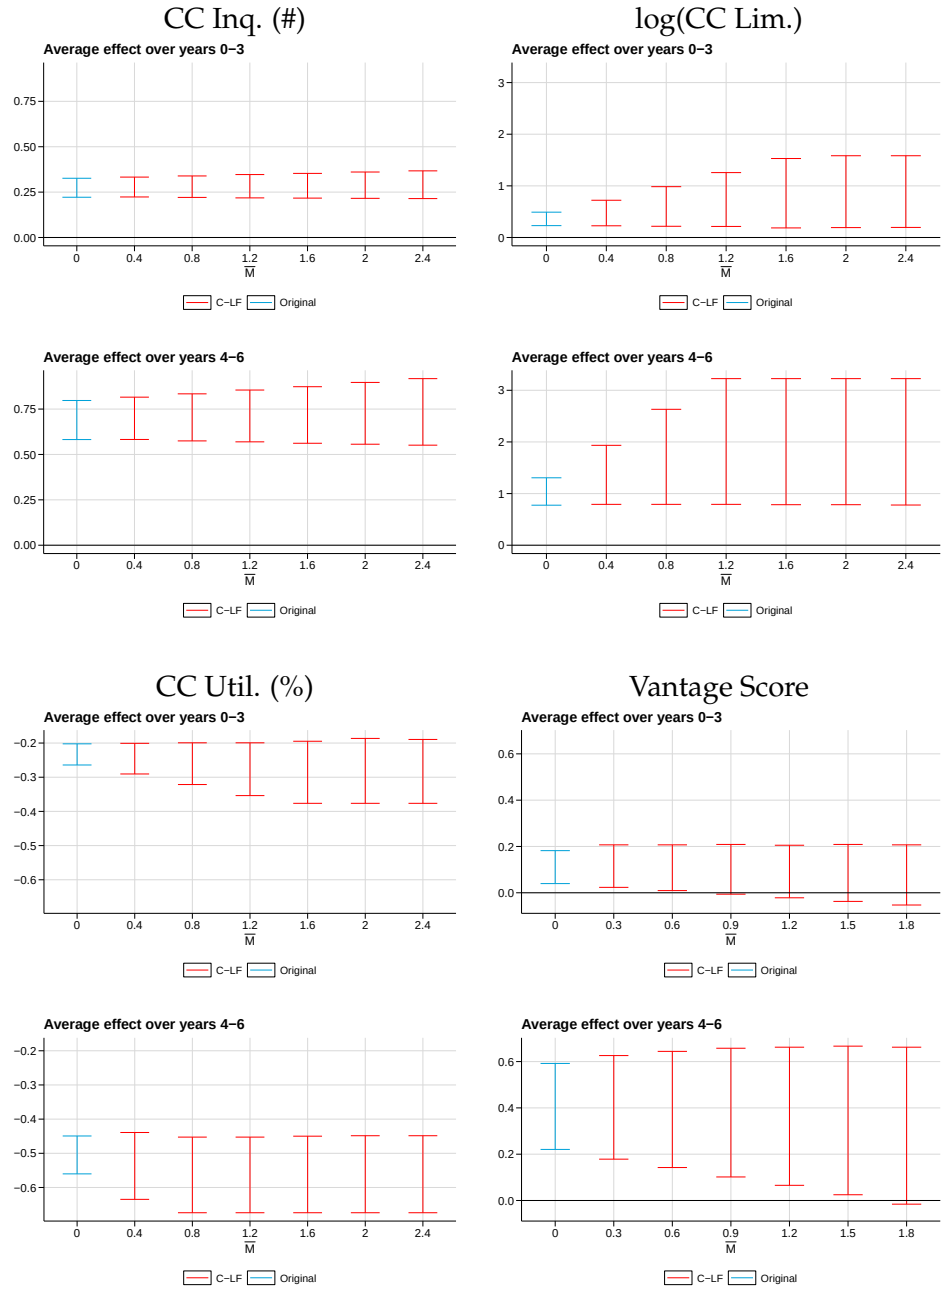
A.4 Honest DiD Robustness

Figure A.3: Honest DiD Estimation Results (1/2)



Notes: For each variable, we report 90% confidence intervals for the estimated average treatment effect over two horizons: early (years 0-3) and later (years 4-6). Red "least favorable" confidence intervals ("C-LF") are obtained via the "relative magnitudes" approach of Rambachan and Roth (2023), which allows various violations of the parallel trends assumption. Above, each $\bar{M}$ corresponds to a scenario allowing parallel trends to be $\bar{M}$-times the size of the average pre-trend. Estimates and confidence intervals resulting from our baseline regression are in blue ("original").

Figure A.4: Honest DiD Estimation Results (2/2)



Notes: For each variable, we report 90% confidence intervals for the estimated average treatment effect over two horizons: early (years 0-3) and later (years 4-6). Red "least favorable" confidence intervals ("C-LF") are obtained via the "relative magnitudes" approach of Rambachan and Roth (2023), which allows various violations of the parallel trends assumption. Above, each $\bar{M}$ corresponds to a scenario allowing parallel trends to be $\bar{M}$-times the size of the average pre-trend. Estimates and confidence intervals resulting from our baseline regression are in blue ("original").

A.5 Reactions of Other Forms of Debt

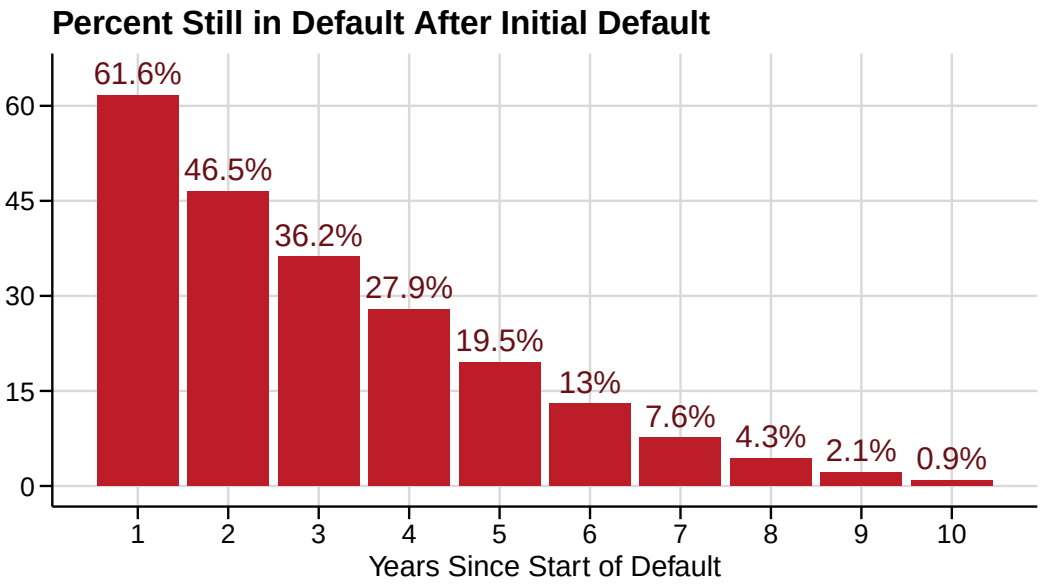
Table A.9: DiD Results for Auto and Mortgage Debt

	1[Has Auto] (1)	1[Has Mortgage] (2)	log(Auto Loan) (3)	log(Mortgage Bal.) (4)
$\text{NewElig}_{zs} \times \text{Post}_{st}$	0.176***	-0.045	1.109***	-0.068
	(0.016)	(0.041)	(0.076)	(0.145)
Obs	140,941	140,941	140,941	140,941
R2	0.954	0.970	0.964	0.983
Mean	35%	30%	$6,812	$69,611

Notes: This table reports results from estimating the DiD in Equation (1) for auto and mortgage loans. Each specification includes county-time and ZIP code fixed effects. Standard errors are clustered by state. The dependent variable is labeled above the column number, and its mean is reported in the bottom row. Statistical significance: 10%*, 5%**, and 1%***.

A.6 Delinquency Duration

Figure A.5: Duration of Delinquency



Notes: This figure plots the fraction of people that are delinquent one to ten years later (after an initial delinquency). Here, delinquency is measured as having either credit card debt 90-180 days past due or non-medical debt in collections.

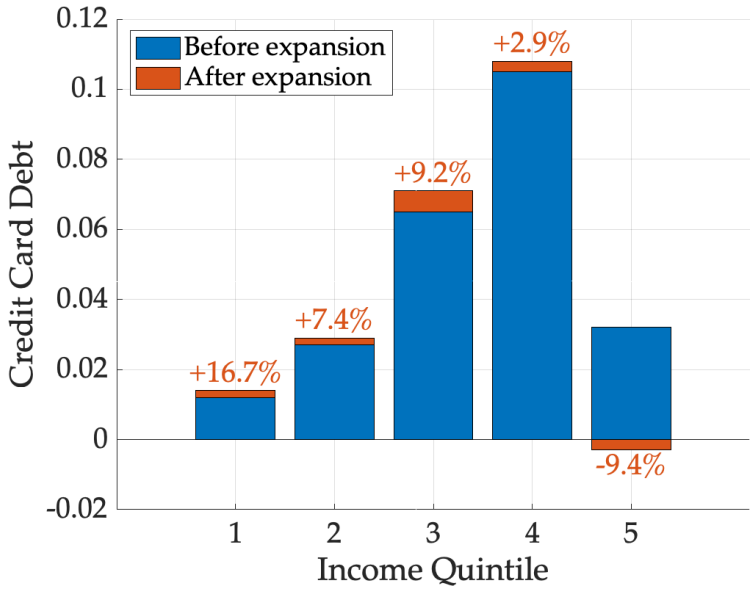
Table A.10: Persistence of Delinquency

Years	% of People Still Delinquent
1	61.62%
2	46.54%
3	36.19%
4	27.86%
5	19.53%
6	12.98%
7	7.62%
8	4.32%
9	2.10%
10	0.88%

Notes: This table reports the fraction of people that are delinquent zero to 10 years later (after an initial delinquency). Here, delinquency is measured as having either credit card debt 90-180 days past due or non-medical debt in collections.

A.7 Additional Model Figures

Figure A.6: Credit Card Debt Along the Income Distribution



Notes: This figure presents the average credit card debt before and after the expansion of Medicaid across different income quintiles. The figure shows that the rise in credit card debt is larger for lower income households, who are more exposed to the policy.

B Eligibility Estimation

This appendix describes how we generate a measure of Medicaid eligibility at the ZIP $\times$ year level.

Data: We use the following datasets from 2009 to 2016:

1. IRS SOI: ZIP-code level individual income tax statistics
2. KFF: Medicaid eligibility limits
3. ACS: joint distribution of income and household size
4. KFF: Medicaid expansion date data

B.1 Estimand

We want estimates of the fraction of adults eligible for Medicaid in year t and ZIP z. Let E denote the event a household is eligible for Medicaid. Denote the probability that a household is eligible in (t, z) by $P_{tz}(E)$. $P_{tz}(E) = P(E|t, z)$. The probability that a household of size n is eligible depends on the probability that its income is below the relevant cutoff c_{tz}^n. We can write out the probability above into the sum of probabilities for these different cases using the law of iterated expectations:

$$P_{tz}(E) = P_{tz}(y \leq c) = \sum_{n=1}^{N} P_{tz}(y \leq c_{tz}^n | c = c_{tz}^n) P(c = c_{tz}^n)$$

where y is income and c_{tz}^n is the eligibility cutoff for a household of size n in $\{z, t\}$. We can rewrite each term using Bayes' rule to get

$$P_{tz}(y \leq c_{tz}^n | c = c_{tz}^n) P(c = c_{tz}^n) = P_{tz}(c = c_{tz}^n | y \leq c_{tz}^n) P(y \leq c_{tz}^n).$$

The first term is the probability that a household is of a given size:

$$P_{tz}(c = c_{tz}^n | y \leq c_{tz}^n) = P_{tz}(\tilde{N} = n | y \leq c_{tz}^n) = g_{tz}(n | y \leq c_{tz}^n)$$

where g_{tz} is the probability mass function of a household size conditional on income being below the relevant threshold. The second term $P(y \leq c_{tz}^n)$ is the probability that the household's income is below the threshold. Let $F_{tz}(\cdot)$ denote the cumulative distribution function of income in $\{t, z\}$. Then

$$P(y \leq c_{tz}^n) = F_{tz}(c_{tz}^n).$$

Household Size Notation: Households of size two may comprise either a head and a dependent or a head an a non-dependent (e.g., a parent and child versus a childless couple). The eligibility thresholds can differ across these situations. To capture both, we use 2 to indicate a household with a head and a dependent and 2^* to indicate a two-person household with no dependents. Let n denote household size. We write $n \in n\{1, 2, 2^*, 3, 4, ..., N\}$, where N is the largest observed household size.

Note on Cutoffs: Income cutoffs vary over time because of changes in the threshold and federal poverty line. They also differ across states; specifically, Alaska and Hawaii generally have higher thresholds in each year.

B.2 Estimation

We estimate $P_{tz}(E)$ using the following steps:

1. We procure the number of tax returns and annual gross income by income bins for each ZIP code and year using IRS SOI.

2. Then, we interpolate the CDF $\hat{F}_{tz}(\cdot)$ of income between the known bins. With that, we estimate $\hat{F}_{tz}(c_{tz}^n)$ where c_{tz}^n is obtained from Medicaid eligibility limits. Note that this equals the second term of the probability, $P(y \leq c_{tz}^n)$. We generate this for $n \in N$ household sizes where N is defined above.

3. We keep one observation per ZIP and year.

4. We use the ACS data to calculate $P(\tilde{N} = n | y \leq c_{tz}^n)$ for each n . We do this using income and household size information available by ZIP code and year. Note that this equals the first term of the probability, $g_{tz}(n | y \leq c_{tz}^n)$.

5. We can calculate the probability that a household of size n is eligible for Medicaid by multiplying the first and second terms for each n.

6. Then, $P_{tz}(E)$ is equal to the sum of the probabilities calculated in the step above.

C Mathematical Appendix

C.1 Proof of Proposition 1 - Uniqueness of Debt Pricing Schedule

The debt pricing schedule is given by

$$\begin{aligned} q(b',y) =& \frac{1}{1+r_f}\left[1-\mathbb{E}\left(d(b'+o(y')m',y')\right)\right] \\ &+\frac{1}{1+r_f}\mathbb{E}\left[d(b'+o(y')m',y')(1-\delta')q\left((1-\delta')(b'+o(y')m'),y'\right)\right]. \end{aligned}$$

Let $B(\mathbb{R}^2)$ be the set of bounded functions $q:\mathbb{R}^2\to\mathbb{R}$. Define the operator $T:B(\mathbb{R}^2)\to B(\mathbb{R}^2)$ as follows:

$$\begin{aligned} T(q)(b',y) =& \frac{1}{1+r_f}\left[1-\mathbb{E}\left(d(b'+o(y')m',y')\right)\right] \\ &+\frac{1}{1+r_f}\mathbb{E}\left[d(b'+o(y')m',y')(1-\delta')q\left((1-\delta')(b'+o(y')m'),y'\right)\right]. \end{aligned} \tag{11}$$

We will use Blackwell's sufficient conditions to show that T is a contraction mapping with modulus $\frac{1}{1+r_f}<1$.

Monotonicity. Suppose $q_1(b',y)\geq q_2(b',y)$ for all b' and y. Then,

$$\begin{aligned} T(q_1)(b',y) &= f(b',y)+\frac{1}{1+r_f}\mathbb{E}\left[d(b'+o(y')m',y')(1-\delta')q_1\left((1-\delta')(b'+o(y')m'),y'\right)\right] \\ &\geq f(b',y)+\frac{1}{1+r_f}\mathbb{E}\left[d(b'+o(y')m',y')(1-\delta')q_2\left((1-\delta')(b'+o(y')m'),y'\right)\right] \\ &= T(q_2)(b',y) \end{aligned}$$

where $f(b',y)=\frac{1}{1+r_f}\left[1-\mathbb{E}\left(d(b'+o(y')m',y')\right)\right]$. The first inequality follows from the fact that $d(b'+o(y')m',y')(1-\delta')\geq 0$ for all (b',y').

Discounting. We have that

$$
\begin{aligned}
&T(q+a)(b',y) \\
&\quad = f(b',y) + \frac{1}{1+r_f}\mathbb{E}\left[d(b'+o(y')m',y')(1-\delta')\left[q\left((1-\delta')(b'+o(y')m'),y'\right)+a\right]\right] \\
&\quad = T(q)(b',y) + \frac{1}{1+r_f}\mathbb{E}\left[d(b'+o(y')m',y')(1-\delta')\right]a \\
&\quad \leq T(q)(b',y) + \frac{1}{1+r_f}a,
\end{aligned}
$$

where the last inequality follows from $d(b'+o(y')m',y')(1-\delta') \in [0,1]$ for all (b',y'). Therefore, T is a contraction mapping with modulus $\frac{1}{1+r_f}$. Hence, the debt pricing schedule $q(b',y)$ exists and is unique, given a delinquency policy function $d(\tilde{b}',y)$.

C.2 Moments of Income Process

Suppose that the income process takes the following form:

$$
\ln y_t = \begin{cases} \rho \ln y_{t-1} + \epsilon_t & \text{w.p.} \quad \lambda, \\ \ln y_{t-1} & \text{w.p.} \quad 1-\lambda, \end{cases} \tag{12}
$$

where ϵ_t follows a Normal distribution with standard deviation σ. We calibrate the structural parameters $\{\rho, \sigma, \lambda\}$ to match the persistence of income, as well as the standard deviation and kurtosis of income changes. Since the PSID reports household-level income every two years, we match the 2-year auto-correlation of income, as well as the standard deviation and kurtosis of 2-year income changes.

Proposition 2. *The income process admits the following moments:*

$$
\begin{aligned}
\text{acorr}_2(\ln y_t) &= \frac{\mathbb{E}\left[\ln y_t \ln y_{t+2}\right]}{\text{var}(\ln y_t)} = (1-\lambda+\lambda\rho)^2, \\
\text{var}\left(\Delta \ln y_t\right) &= \mathbb{E}\left[(\ln y_{t+1} - \ln y_t)^2\right] = 2\left(1-(1-\lambda+\lambda\rho)\right)\text{var}(\ln y_t), \\
\text{var}\left(\Delta_2 \ln y_t\right) &= \mathbb{E}\left[(\ln y_{t+2} - \ln y_t)^2\right] = 2\left(1-(1-\lambda+\lambda\rho)^2\right)\text{var}(\ln y_t), \\
\text{kurt}\left(\Delta \ln y_t\right) &= \frac{3}{4}\frac{1}{\lambda}\left((1-\rho)^4 + 2(1-\rho)^2(1-\rho^2) + (1-\rho^2)^2\right)\frac{(1+\rho)^2}{(1-\rho^2)^2}, \\
\text{kurt}\left(\Delta_2 \ln y_t\right) &= \frac{2(1-\lambda)\,\text{K}_\Delta\,\text{V}_\Delta^2 + \lambda^2(3(\rho^2-1)^4\,\text{V}_y^2 + 3(1+\rho^4)\sigma^4 + 6\rho^2\sigma^4 + 6(\rho^2-1)^2(1+\rho^2\,\text{V}_y\,\sigma^2)}{\text{V}_{\Delta_2}^2},
\end{aligned}
$$

where $\mathrm{V}_y = \mathrm{var}(\ln y_t)$, $\mathrm{V}_\Delta = \mathrm{var}(\Delta \ln y_t)$, $\mathrm{V}_{\Delta_2} = \mathrm{var}(\Delta_2 \ln y_t)$ *and* $\mathrm{K}_\Delta = \mathrm{kurt}(\Delta \ln y_t)$.

Proof. 1. **Proof of 2-year auto-correlation**.

$$\mathbb{E}\left[\ln y_t \ln y_{t+2}\right] = (1-\lambda)^2 \mathrm{var}(\ln y_t) + 2\lambda(1-\lambda)\rho \,\mathrm{var}(\ln y_t) + \lambda^2\rho^2 \,\mathrm{var}(\ln y_t).$$

Rearranging we have

$$\frac{\mathbb{E}\left[\ln y_t \ln y_{t+2}\right]}{\mathrm{var}(\ln y_t)} = (1-\lambda)^2 + 2(1-\lambda)\rho\lambda + \lambda^2\rho^2 = (1-\lambda+\lambda\rho)^2.$$

2. **Proof of 1-year income change variance**.

$$\mathbb{E}\left[(\ln y_{t+1} - \ln y_t)^2\right] = 2\,\mathrm{var}(\ln y_t) - 2\mathbb{E}\left(\ln y_{t+1} \ln y_t\right).$$

Similarly to how we proved the 2-year auto-correlation, we have that

$$\mathbb{E}\left(\ln y_{t+1} \ln y_t\right) = (1-\lambda+\lambda\rho)\,\mathrm{var}(\ln y_t).$$

So we obtain

$$\mathbb{E}\left[(\ln y_{t+1} - \ln y_t)^2\right] = 2\left(1-(1-\lambda+\lambda\rho)\right)\mathrm{var}(\ln y_t).$$

3. **Proof of 2-year income change variance**.

$$\mathbb{E}\left[(\ln y_{t+2} - \ln y_t)^2\right] = 2\,\mathrm{var}(\ln y_t) - 2\mathbb{E}\left(\ln y_{t+2} \ln y_t\right).$$

We have that

$$\begin{aligned}\mathbb{E}\left(\ln y_{t+2} \ln y_t\right) &= (1-\lambda)\mathbb{E}\left(\ln y_{t+1} \ln y_t\right) + \lambda\rho\mathbb{E}\left(\ln y_{t+1} \ln y_t\right) \\ &= (1-\lambda+\lambda\rho)\mathbb{E}\left(\ln y_{t+1} \ln y_t\right).\end{aligned}$$

Combining everything we have

$$\mathbb{E}\left[(\ln y_{t+2} - \ln y_t)^2\right] = 2\left(1 - (1 - \lambda + \lambda\rho)^2\right) \text{var}(\ln y_t).$$

4. **Proof of 1-year income change kurtosis**. Let's start from the unconditional variance, $var(\ln y_t)$. From equation (12), we have

$$\text{var}(\ln y_t) = \lambda\left(\rho^2 \text{var}(\ln y_t) + \sigma^2\right) + (1 - \lambda) var(\ln y_t).$$

Rearranging we obtain

$$\text{var}(\ln y_t) = \frac{1}{1 - \rho^2}\sigma^2.$$

Now we derive the kurtosis:

$$\begin{aligned}
\mathbb{E}\left(\Delta \ln y_t\right)^4 &= \lambda \mathbb{E}\left((\rho - 1)\ln y_t + \sigma\epsilon\right)^4 \\
&= \lambda\left[(\rho - 1)^4 \mathbb{E}(\ln y_t)^4 + 6(1 - \rho)^2 \mathbb{E}(\ln y_t)^2 \sigma^2 + 3\sigma^4\right] \\
&= \lambda\left[(1 - \rho)^4 + 2(1 - \rho)^2(1 - \rho^2) + (1 - \rho^2)^2\right] \frac{1}{(1 - \rho^2)^2} 3\sigma^4
\end{aligned}$$

where we've used

$$\mathbb{E}(\ln y_t)^2 = \text{var}(\ln y_t) = \frac{1}{1 - \rho^2}\sigma^2,$$

and we can obtain $\mathbb{E}(\ln y_t)^4$ as follows

$$\mathbb{E}(\ln y_t)^4 = (1 - \lambda)\mathbb{E}(\ln y_t)^4 + \lambda\left[\rho^4 \mathbb{E}(\ln y_t)^4 + 6\rho^2 \frac{1}{1 - \rho^2}\sigma^4 + 3\sigma^4\right]$$

so that

$$\mathbb{E}(\ln y_t)^4 = \frac{1}{1 - \rho^4}\left(\frac{2\rho^2}{1 - \rho^2} + 1\right) 3\sigma^4 = \frac{1}{(1 - \rho^2)^2} 3\sigma^4$$

Finally we get that

$$Kurt\left(\Delta \ln y_t\right) = \frac{\mathbb{E}\left(\Delta \ln y_t\right)^4}{\text{var}^2(\Delta \ln y_t)} = \frac{3}{4\lambda}\left[(1 - \rho)^4 + 2(1 - \rho)^2(1 - \rho^2) + (1 - \rho^2)^2\right] \frac{(1 + \rho)^2}{(1 - \rho^2)^2}$$

where we've used $Var(\Delta \ln y) = 2\lambda \frac{(1-\rho)}{1-\rho^2}\sigma^2 = 2\lambda \frac{1}{1+\rho}\sigma^2$.

5. **Proof of 2-year income change kurtosis**.[41]

$$\begin{aligned}\mathbb{E}\left(\ln y_{t+2} - \ln y_t\right)^4 &= 2(1-\lambda)\lambda \mathbb{E}\left((\rho-1)\ln y_t + \sigma\epsilon\right)^4 + \lambda^2 \mathbb{E}\left((\rho^2-1)\ln y_t + \rho\sigma\epsilon_1 + \sigma\epsilon_2\right)^4 \\ &= 2(1-\lambda)\mathbb{E}(\Delta \ln y)^4 + \lambda^2 \mathbb{E}\left((\rho^2-1)\ln y_t + \rho\sigma\epsilon_1 + \sigma\epsilon_2\right)^4\end{aligned}$$

Let's expand the second term:

$$\mathbb{E}\left((\rho^2-1)\ln y_t + \rho\sigma\epsilon_1 + \sigma\epsilon_2\right)^4 = (\rho^2-1)^4 3V_y^2 + (1+\rho^4)3\sigma^4 + 6\rho^2\sigma^4 + 6(\rho^2-1)^2(1+\rho^2)V_y\sigma^2,$$

where we've used $\text{kurt}(\ln y_t) = 3$. Combining this expression with the first term we obtain

$$\begin{aligned}&\mathbb{E}\left(\ln y_{t+2} - \ln y_t\right)^4 \\ &\quad = 2(1-\lambda)\,\mathrm{K}_\Delta\,\mathrm{V}_\Delta^2 + \lambda^2(3(\rho^2-1)^4\,\mathrm{V}_y^2 + 3(1+\rho^4)\sigma^4 + 6\rho^2\sigma^4 + 6(\rho^2-1)^2(1+\rho^2\,\mathrm{V}_y\,\sigma^2).\end{aligned}$$

Together with the definition of kurtosis, $\text{kurt}(\Delta_2 \ln y_t) = \frac{\mathbb{E}(\ln y_{t+2} - \ln y_t)^4}{\left(\mathbb{E}(\ln y_{t+2} - \ln y_t)^2\right)^2}$, we proved the desired expression.

□

[41] Since this moment involves extensive algebra, we confirmed via simulations that our formula is correct.