

The Credit Consequences of Unpaid Medical Bills*

Kenneth Brevoort [†] Daniel Grodzicki [‡] Martin B. Hackmann [§]
Sergei Koulayev [¶]

July 31, 2018

Abstract

This paper quantifies the costs of leaving medical bills unpaid and what these costs imply for the value of health insurance to beneficiaries. We argue that most unpaid medical bills are sent to third-party collections and reported to credit bureaus, with detrimental effects on patients' credit outcomes. Combining a large panel of credit records with data on credit offers, we find that the ACA Medicaid expansion reduced newly reported medical collections by \$5.89 billion and led to better terms of credit valued at \$670 million annually. We calculate the financial benefits of Medicaid double when including this indirect credit channel.

JEL Codes: D14, H51, I13

*This is an updated version of our paper formerly titled "Medicaid and Financial Health". We thank seminar participants at the CFPB, the FDIC, the NBER Summer Institute, the NTA, The Ohio State University, UCLA, the University of Mannheim, and the University of Chicago as well as Mark Duggan, Liran Einav, Bradley Heim, Jon Kolstad, Amanda Kowalski, Neale Mahoney, David Molitor, Tomas Philipson, and Pietro Tebaldi for thoughtful comments and suggestions. The views expressed are those of the authors and do not necessarily reflect those of the Consumer Financial Protection Bureau or the United States. Authors are solely responsible for all remaining errors.

[†]Consumer Financial Protection Bureau, kenneth.brevoort@cfpb.gov.

[‡]The Pennsylvania State University and Consumer Financial Protection Bureau, djg39@psu.edu.

[§]The University of California at Los Angeles, CESifo, and NBER, mbhackmann@gmail.com.

[¶]Consumer Financial Protection Bureau, sergei.koulayev@cfpb.gov.

1 Introduction

A fundamental goal of health insurance is to provide financial risk protection against large and unforeseen medical expenses. As such, the existing literature on the financial benefits of health insurance highlights consumer welfare gains arising from reductions in the mean and variance of out-of-pocket medical expenses (Zeckhauser, 1970). However, the uninsured typically pay only 20 percent of their overall health care utilization out-of-pocket (Coughlin, 2014). This suggests that the financial benefits of health insurance to beneficiaries may be relatively small (Finkelstein, Hendren and Shepard, 2017). At the same time, a majority of the uninsured report making substantial sacrifices to pay for medical care, including significant changes to their financial situation, lifestyle, and or employment (Hamel et al., 2016), suggesting otherwise.

Reconciling these conflicting points, this paper examines the costs of leaving medical bills unpaid and what these costs imply for the value of health insurance to beneficiaries. We argue that a large fraction of unpaid medical bills is sent to third-party collection agencies, with detrimental consequences for patients' future terms of credit. By guarding against unpaid medical bills, health insurance thus provides additional indirect financial benefits through its impact on beneficiaries' credit market experiences. Complementing previous landmark studies on the benefits of insurance (Finkelstein and McKnight, 2008; Finkelstein, Hendren and Luttmer, 2015), we highlight these indirect financial benefits from protection against unpaid medical bills.

We begin by extending the textbook model of insurance to examine the role of unpaid medical bills in consumer welfare. In our conceptual framework, uninsured individuals derive utility from consumption and face a disutility from leaving medical bills unpaid. They then choose what portion of their medical expenses to leave unpaid, trading off greater consumption with the disutility of not paying their bills. With this model, we decompose the financial benefits of health insurance into two parts: (1) the *direct* gains from insurance against out-of-pocket spending and (2) the reduction in disutility from fewer unpaid bills, which operates through the *indirect* credit channel.

We quantify these direct and indirect financial benefits of health insurance in the context of the Patient Protection and Affordable Care Act (ACA), which was signed into law in 2010. One of the ACA's marquee provisions sought to expand Medicaid eligibility to all individuals earning less than 138% of the federal poverty level (FPL). While this expansion was intended to apply nationwide, in 2012 the Supreme Court ruled that states must be allowed to decide individually whether they would adopt the expanded Medicaid eligibility rules. As of the end of 2016, 31 states and the District of Columbia have adopted the Medicaid expansion and 19

states have chosen not to sign on to the expansion. This provides us with quasi-experimental variation in the Medicaid expansion, which we exploit in a difference-in-differences research design.

Our analysis combines state adoption decisions and Census tract-level variation in eligibility from the Medicaid expansion with administrative data from the Consumer Financial Protection Bureau’s Consumer Credit Panel (CCP), a nationally representative panel of over 5 million de-identified credit records. An important advantage of this CCP, when compared to other panels, is that it contains information on individual credit obligations (trade-lines). In particular, this includes the source of each debt (bank, third-party collection agency, etc.) and the date the obligation was credited. As a result, we are able to separately identify unpaid medical bills that are in collection and the dates in which they were credited.

We find that the Medicaid expansion directly reduced newly-accrued medical debt by 35 percent, with a disproportionately greater effect for larger medical debts. On average, the reform led to an annual decline in accrued medical debt of \$54 per person, or \$1,227 per treated person per year. This translates into an aggregate reduction of medical debt of \$5.89 billion between the beginning of 2014 and the end of 2016. When compared to overall health care utilization and out-of-pocket spending, our estimates indicate that about 51 percent of overall utilization and 64 percent of unpaid medical bills (uncompensated care) of the uninsured go into collection. We also find that collection agencies are able to recover about 9 percent of the face value of these debts in the first two years, providing a financial incentive for health care providers to sell uncompensated medical claims to third-party collection agencies.

The CCP also makes it possible to identify movements into repayment delinquency for various debts as well as changes in individuals’ overall credit risk, measured by their credit score. Using an event study approach, we first provide direct evidence on the relationship between medical debt and these measures of creditworthiness. Specifically, we document a sudden, sharp, and persistent drop in an individual’s credit score immediately following her first medical collection. Building on this evidence, we return to the Medicaid expansion and document a reduction of new delinquencies and an increase in credit scores in the post-reform years following the reduction in medical debt. Moreover, and consistent with earlier work on Medicaid expansions, we find that credit scores were little moved at first but increased substantially in the second and third post-reform year, most noticeably for the middle two quartiles of the credit score distribution.

To put these improvements in creditworthiness into perspective, we quantify the implied interest rate savings on outstanding debt. For this purpose, we use novel data on direct-mail credit offers from Mintel Comperemedia (Mintel) in conjunction with aggregated lender

rate sheets collected by the Fair Isaac Corporation (FICO) to measure effects of the policy on the pricing of credit offered to consumers. We first document that the reform led to a decline in the price of offered credit. We then calculate a dollar value of this price decline by simulating a refinancing of debt by these individuals under the new (lower) interest rates. Our simulation suggests large annual savings to consumers, which come predominantly from credit card and unsecured personal loan debt. On average, we calculate annual savings of \$14.60 per person, or \$332 per treated person. This translates into about \$670 million in aggregate annual savings.

Finally, we return to our conceptual framework and quantify the relative importance of the indirect credit channel of insurance on consumer welfare. We do so using two alternative approaches. In the first approach, we calibrate individuals' consumption utility, and recover their disutility over unpaid medical bills by leveraging information on observed out-of-pocket spending. This approach builds on the idea that out-of-pocket payments are informative about the implicit disutility from higher medical debt. Intuitively, out-of-pocket spending would reduce to zero in the absence of utility costs from higher medical debt. We refer to this as the revealed preference approach. In addition to obtaining closed-form expressions of the compensating variation for a mean reduction in medical bills, we derive the risk premium and assess the value of risk protection from a reduction in the variance of medical expenditures. Our estimated compensating variation and risk premium exceed those from a benchmark model that only considers the benefits from reductions in out-of-pocket spending by a factor of 2.5 and 2.8, respectively. While the revealed preference approach relies on strong functional form assumptions, we note that the main insights are robust to a variety of alternative assumptions concerning the patient's risk aversion and out-of-pocket spending.

In the second approach, which we call the direct approach, we add our calculated interest savings to the direct benefits of reduced out-of-pocket spending. Using this method, we find that the financial benefits of a mean reduction in medical bills increases by 69 percent when considering the indirect benefits in addition to reduced out-of-pocket spending. We view this as a conservative estimate since it ignores numerous other benefits of insurance such as less hassling by debt collectors, diminished risk of legal action by creditors, a reduced bankruptcy risk, as well as the consumer gains from changes in borrowing. Overall, the findings from both approaches suggest that the financial benefits of a mean and variance reduction in medical bills double when considering the indirect financial benefits of insurance.

Our paper contributes to three main literatures. First, our analysis complements recent studies on the value of Medicaid (Finkelstein, Hendren and Luttmer, 2015) and the value of public insurance more generally (Kowalski, 2015; Cabral and Cullen, 2016; Finkelstein and McKnight, 2008). These studies investigate the overall consumer benefit of public insur-

ance, taking financial and health related benefits into account. In the context of Medicaid, [Finkelstein, Hendren and Luttmer \(2015\)](#) find that beneficiaries value the program by only \$0.2 to \$0.4 per dollar of government spending, mostly stemming from reduced out-of-pocket spending. Our approach abstracts away from changes in health care utilization as uninsured individuals gain Medicaid insurance and shifts the focus to the financial benefits of Medicaid insurance. To this end, we complement the analysis of financial benefits in [Finkelstein, Hendren and Luttmer \(2015\)](#) by adding and quantifying the indirect benefits from a reduction in unpaid medical bills through improved terms of credit.

Second, our findings add to a growing body of work studying the link between Medicaid, or insurance expansions more generally, and measures of financial health ([Finkelstein et al., 2012](#); [Mazumder and Miller, 2016](#); [Gross and Notowidigdo, 2011](#); [Hu et al., 2016](#); [Sojourner and Golberstein, 2017](#); [Caswell and Waidmann, 2017](#); [Gallagher, Gopalan and Grinstein-Weiss, 2018](#); [Argys et al., 2017](#)). To the best of our knowledge, we are the first to document the relationship between unpaid medical bills and the terms of credit, which has important implications for the value of health insurance. To this end, we provide novel direct evidence on the negative relationship between new medical debt in collection and credit scores using an event study approach. We corroborate this relationship by providing new evidence on a reduction in medical debt and improvements in credit scores following the Medicaid expansion. Our analysis of medical debt improves upon limitations in other credit panels that are not able to distinguish between medical debt and other debt in collection. Our evidence on credit scores contributes to mixed existing evidence in the existing literature and is supported by extensive robustness checks.¹ Most importantly, we develop two novel empirical approaches that allow us to quantify the consumer welfare implications of reductions in unpaid medical bills. Leveraging novel data on interest rates for credit card debt, personal loans, automobile loans, and mortgages, we are able to quantify the financial benefits of a reduction in unpaid medical bills in dollars.²

Third, our results shed new light on the incidence of uncompensated care. Several recent studies document the important role of uncompensated care for health care delivery (e.g., ([Coughlin, 2014](#)) and ([Dranove, Garthwaite and Ody, 2016](#))). Notably, [Garthwaite, Gross and Notowidigdo \(2015\)](#) document that hospitals act as "insurers of last resort," as the uninsured pay only a small fraction of their medical bills out-of-pocket providing a potential explanation for the low willingness-to-pay among low-income individuals, [Finkelstein,](#)

¹For instance, ([Mazumder and Miller, 2016](#)) find that the Massachusetts health reform led to an increase credit scores, whereas ([Finkelstein et al., 2012](#)) find no significant effect of Medicaid insurance on credit scores in the first year.

²Consistent with our evidence on improved access to credit markets, ([Allen et al., 2017](#)) find that California's early Medicaid expansion led to a reduction in pay-day loan borrowing.

Hendren and Shepard (2017). We contribute to these studies by shedding new light on the incidence of uncompensated care. We use tradeline-level variation in credits and subsequent repayment of medical debt in collection to study the incidence of uncompensated care. Specifically, we examine the likelihood with which providers seek repayment through third-party collections, the rate at which new medical collections are repaid, and how these debts affect low-income uninsured patients through their subsequent interaction with broader credit markets. Our findings suggest that the incidence of uncompensated care at least partially falls on the low-income uninsured patients themselves, through this indirect credit channel.

The remainder of this paper is organized as follows. Section 2 presents our conceptual framework, which formalizes the credit channel of health insurance. Section 3 discusses institutional details surrounding the Medicaid expansion and unpaid medical bills. We describe the data and lay out our difference-in-difference approach in Section 4. In Sections 5 and 6 we present our main empirical findings on medical debt and financial distress, respectively, as well as examine the impact of improved financial health on the prices of offered credit and the implied dollar value of this benefit. Returning to the effects on consumer welfare, we present our overall financial benefit estimates in Section 7. Section 8 concludes with a discussion of our main findings.

2 The Credit Channel of Health Insurance

To assess the role of the indirect credit channel, we begin by extending the textbook model on the value of risk protection provided by insurance to include the impact to consumers' welfare from a reduction in paid *and* unpaid medical bills. Our conceptual framework focuses on the financial benefits of health insurance. To this end, we treat health care utilization as exogenous and abstract away from changes in utilization following the Medicaid expansion. As a result, we do not model utility over health care consumption.

We consider an environment in which an uninsured individual derives positive utility from consumption, $g(c)$, and faces a utility loss from unpaid bills, $-h(D)$. Utility losses from unpaid bills capture a number of factors such as costs of worsening credit options, hassles from dealing with debt collectors, and legal complications related to unpaid bills and bankruptcy. Let an individual's utility be of the form

$$U = g(c) - h(D), \tag{1}$$

with $g'(\cdot) > 0, g''(\cdot) < 0$ and $h'(\cdot) > 0, h''(\cdot) \geq 0$. The marginal utility of consumption is decreasing while the marginal disutility of leaving medical bills unpaid is weakly increasing.

An individual earns income Y and is exposed to random medical bills $\epsilon_{MB} \sim G$, where G denotes the underlying distribution function. Given an existing stock of unpaid medical bills $\bar{D} \geq 0$, an individual decides on the optimal amount of new medical bills $0 \leq b \leq \epsilon_{MB}$ that goes unpaid. This decision triggers an inherent trade off in utility from greater consumption and disutility from leaving bills unpaid. Upon incurring a medical expenditure, ϵ_{MB} , consumers' problem is given by

$$\max_{0 \leq b \leq \epsilon_{MB}} g(Y - (\epsilon_{MB} - b)) - h(\bar{D} + b), \quad (2)$$

where in optimality

$$g'(Y - (\epsilon_{MB} - b^*)) - h'(\bar{D} + b^*) = 0. \quad (3)$$

Introducing this trade-off between consumption and disutility from unpaid bills changes the consumer welfare implications of reductions in both the mean and the variance in incurred medical costs, which we discuss separately below.

2.1 Mean Reduction and the Compensating Variation

We first analyze the effect of a mean reduction in medical bills on consumer welfare and ignore uncertainty in medical expenditures. We evaluate the financial harm of a fixed medical bill (ϵ_{MB}), the key implications of which are illustrated in Figure 1. In the Figure, consumption is plotted on the horizontal axis and marginal (dis)utility on the vertical axis. The downward sloping line represents the marginal utility of consumption (MU_C), and the upward sloping line is the marginal disutility of unpaid bills (MU_D).

Absent any medical expenses, an individual consumes her income Y . When facing a medical bill of size ϵ_{MB} , she decides on the amount that she is willing to pay out-of-pocket, $\epsilon_{MB} - b^*$. In an optimum, the marginal utility of an additional dollar of consumption must equal the marginal disutility of an additional dollar in unpaid medical bills. This is depicted in point B^* . We can then decompose the welfare loss resulting from a medical bill as the sum of two effects: the direct effect on out-of-pocket spending and the indirect effect, or the *credit channel*.

In the figure, the red area D bounded by the marginal utility of consumption, the individual's baseline income Y , and her final consumption, $Y - (\epsilon_{MB} - b^*)$, captures the direct effect, or the utility loss from reduced consumption due to increased out-of-pocket payments. The indirect, or *credit channel*, effect is then the blue area, bounded by the marginal disutility of unpaid medical bills, final consumption, $Y - (\epsilon_{MB} - b^*)$, and final consumption minus the borrowed amount $Y - \epsilon_{MB}$. As described above, this *credit channel* highlights

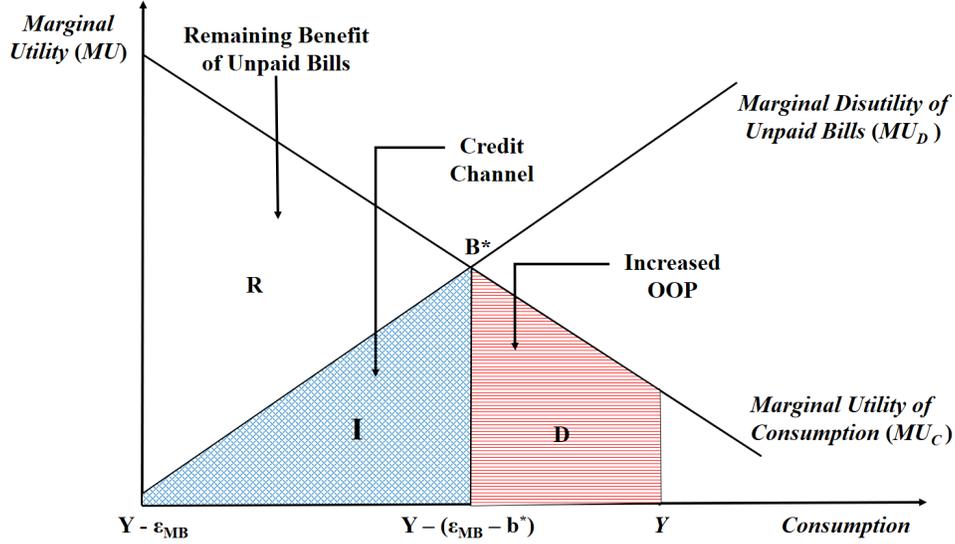


Figure 1: Welfare Benefits of Mean Reduction: Example

the potentially adverse consequences of unpaid bills on access to and the price of credit as well as other costs associated with not paying bills. The sum of the two areas captures the overall utility loss from the medical bill shock ϵ_{MB} . Finally, the white area (R) captures any remaining net benefit from unpaid medical bills. To see this, note that were the individual to pay the entire amount out-of-pocket, the utility loss would be the entire area underneath the marginal utility of consumption between: ($= R + I + D$).

To translate the effects on consumer utility into dollars, we analyze the compensating variation (CV). In this context, the CV describes the amount of income a person is willing to forgo if the medical bill of the amount ϵ_{MB} is removed:

$$CV = e(p_0, u_0) - e(p_1, u_0) = e(\epsilon_{MB}, u_0) - e(0, u_0) . \quad (4)$$

Here, $e(\cdot)$ denotes the underlying expenditure function. Naturally, we have $CV = \epsilon_{MB}$ if the person pays the entire bill out-of-pocket. Conversely, if only a portion of the medical bill is paid out-of-pocket, then we have $\Delta OOP \leq CV \leq \epsilon_{MB}$, where ΔOOP denotes the counterfactual savings in out-of-pocket payments. Looking back at Figure 1, $Y - CV$ corresponds to the point on the horizontal axis where the area underneath the marginal utility of

consumption curve bounded by $Y - CV$ from the left and Y from the right equals the sum of the blue and the red area (D+I).

2.2 Variance Reduction and the Risk Premium

As we allow individuals to be risk averse, we further analyze the effects of a reduction in the variance of medical expenditures on consumer welfare. For this, we reintroduce uncertainty to the model and consider the risk premium RP , which isolates the financial benefits of a reduction the variance of incurred medical bills. Given that only a fraction of the cost of medical services is paid out-of-pocket, the risk premium combines the benefits from a variance reduction in out-of-pocket spending and unpaid medical bills.

We capture these benefits by decomposing the overall risk premium into an out-of-pocket RP, RP^{oop} , and a unpaid medical bill RP, RP^D , respectively. These are implicitly defined by the following equations:

$$E[g] = g(Y - (\bar{\epsilon}_{MB} - \bar{b}^*) - RP^{OOP}) \quad (5)$$

$$E[h] = h(\bar{D} + \bar{b}^* + RP^D) \quad (6)$$

where $E[g]$, $E[h]$ denote expected consumption utility and expected disutility over unpaid medical bills, respectively. $\bar{\epsilon}_{MB}$ denotes the average medical bill, and \bar{b}^* is the average increase in unpaid bills. Finally total expected utility is given by $E[U] = E[g] - E[h]$.

Note that RP^D is denominated in dollars of unpaid medical bills. Using a first order approximation to the utility functions, we translate RP^D into consumption dollars and express the overall risk premium as follows:

$$RP = RP^{OOP} - \frac{h'(\bar{D} + \bar{b}^*)}{g'(Y - (\bar{\epsilon}_{MB} - \bar{b}^*))} \cdot RP^D \geq RP^{OOP} . \quad (7)$$

Like for the mean transfer effect above, it becomes evident from this representation that a naive consideration of benefits from a variance reduction in out-of-pocket spending, captured by RP^{OOP} understates the overall value of risk protection denoted by RP . The difference largely depends on the curvature of $h(\cdot)$, the disutility over unpaid bills. Moreover, this difference is increasing in the share of medical bills left unpaid, which could be large given that a majority of health care costs are not repaid by the uninsured. In all, broadening the textbook analysis to include the disutility from leaving medical bills unpaid illustrates that restricting attention to changes in out-of-pocket spending can substantially understate the full financial benefit of insurance. We return to this claim and provide a quantitative analysis of it in Section 7.

3 Institutional Details and Context

3.1 The Medicaid Expansion

Signed into law in 2010, the Patient Protection and Affordable Care Act (ACA) was one of the most sweeping health care reforms in U.S. history. Among its most important and controversial provisions was its expansion of the Medicaid program to cover all individuals earning less than 138% of the federal poverty level. Before the ACA, Medicaid’s principal beneficiaries were low-income children, their parents, and people with disabilities. Childless adults between the ages of 18 and 65 were for the most part ineligible to receive insurance in nearly all states. The ACA forced each state government to choose between expanding its Medicaid program and losing its federal Medicaid funding altogether. Twenty-six states filed suit to challenge this provision of the ACA and, in 2012, the Supreme Court found it to be unconstitutional. The Court required that states not expanding their Medicaid programs be allowed to retain the federal funding for their existing Medicaid programs.³

By January 1, 2014, on the eve of the expansion’s intended rollout, only 24 states plus the District of Columbia had adopted the measure. Of these, 19 states chose to expand their Medicaid programs on January 1, 2014. The other 5 states and the District of Columbia had already expanded their programs.⁴ Another 7 states expanded coverage at various points after January 1, 2014.⁵ This left 19 non-adopting states as of the end of 2016, the final year of our analysis. Figure 2 illustrates states’ adoption decisions since passage of the ACA. In our empirical analysis, we exclude consumers in the early- and late-adopting states and focus on trends in the 19 states that expanded Medicaid on January 1, 2014 (which we refer to throughout as the adopting or treatment states) and the 19 non-adopting states (control states).

Following the reform, health care coverage increased substantially in adopting states. According to the *Medicaid and Children Health Insurance Program (CHIP) Enrollment Reports*, there were 4.9 million more people enrolled in Medicaid in the 19 on-time adopting states in December 2014 than the average enrollment in these same states from July-September 2013, an increase of 26%. In the 19 non-adopting states, enrollment was up by only 1.6 million people or 9%.⁶ Hence, we attribute a differential Medicaid enrollment increase of 3.3 mil-

³The court case is known as *National Federation of Independent Business v. Sebelius*, 567 U.S. 519 (2012). Also see [Kaiser Family Foundation \(2012a\)](#) for more detail.

⁴Early adopting states include CA, CT, MN, NJ, WA and the District of Columbia.

⁵Alaska expanded on September 1st, 2015. Indiana expanded on February 1st, 2015. Louisiana expanded on July 1st, 2016. Michigan expanded on April 1st, 2014. Montana expanded on January 1, 2016. New Hampshire expanded on August 15th, 2014. Pennsylvania expanded on January 1st, 2015.

⁶See <https://www.medicaid.gov/medicaid/program-information/medicaid-and-chip-enrollment-data/monthly-reports/index.html>, last accessed on June 26, 2017. Enrollment figures for the control

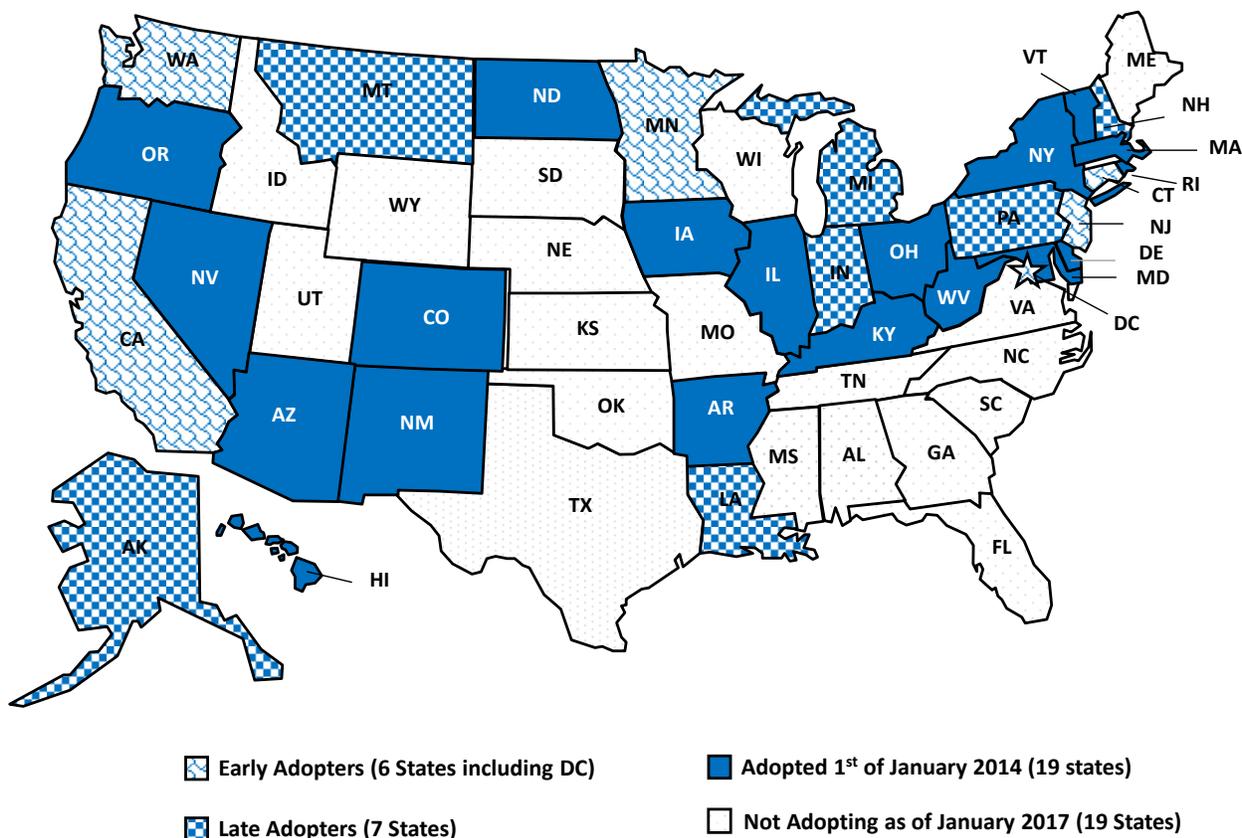


Figure 2: Medicaid Adoption Across States

Source: Kaiser Family Foundation, see

<http://kff.org/health-reform/slide/current-status-of-the-medicaid-expansion-decision/> for further details.

lion, about 4.1% of the non-elderly population, to the Medicaid expansion, which is roughly consistent with estimates from the literature.⁷ Repeating the analysis for December 2016, we find a cumulative increase of about 4.6 percentage points over the three post-expansion years. This suggests that the coverage gains were disproportionately larger in the first post-expansion year. In what follows, we assume that the Medicaid expansion led to an average increase in insurance coverage of 4.4 percentage points among non-elderly adults over the post-expansion years 2014-2016. In what follows, we assume that the Medicaid expansion led to an average increase in insurance coverage of 4.4 percentage points among non-elderly adults over the post-expansion years 2014-2016.

states exclude Maine, for which data are unavailable. The increase in enrollment was concentrated among adults. We find only small changes in CHIP enrollment over this period.

⁷ Most closely related to our context, [Courtemanche et al. \(2016\)](#) find a coverage increase of 5.9 percentage points among the non-elderly adults in Medicaid expansion states by the end of 2014. In contrast, coverage increased by only 3 percentage points in non-expansion states suggesting an additional 2.9 percentage point increase due to the Medicaid expansion. [Frean, Gruber and Sommers \(2016\)](#) find that the ACA Medicaid expansion increased insurance coverage by 9 percentage points among individuals who were newly eligible for Medicaid with no evidence that the expansion crowded out private insurance.

3.2 Unpaid Medical Bills in Uninsured’s Balance Sheet

Recent survey evidence from the Kaiser Family Foundation (KFF) ([Hamel et al., 2016](#)) notes that about a quarter of non-elderly adults in the U.S. have difficulties paying their medical bills, with that figure rising to more than half among the uninsured. Not surprisingly, previous studies have found that the uninsured pay up to 20% of medical bills out-of-pocket ([Finkelstein, 2007](#)), or \$480 out of about \$2,400 in overall annual health care spending according to recent estimates based on data from the Medical Expenditure Panel Survey (MEPS). The remaining cost is left as uncompensated care ([Coughlin, 2014](#)).

This uncompensated care can be decomposed into charity care and uninsured care, or bad debt. According to the American Hospital Association (AHA), charity care comprises services for which the hospital never received but also never expected payment, possibly because of the patient’s inability to pay. Bad debt consists of services for which the hospital anticipated but did not receive payment. While charity care is not billed to consumers, ‘bad debt’ is billed to consumers through third-party collection agencies. Collection accounts reported to credit bureaus can severely impact the debtors’ creditworthiness, reducing the quality of credit options available to them. Conceptually, we view charity care as “free” care from the point of view of the patient as it is not billed to her. Bad debt, on the other hand, is not free care as it is sent to collection agencies with potentially detrimental consequences for future terms and access to credit.

In practice, the distinction between charity care and bad debt is blurry and hospitals often struggle to draw the distinction. Not surprisingly, there is little empirical evidence on the relative magnitudes of charity care and bad debt. Instead, studies have focused on quantifying the prevalence of uncompensated care in general and how it is affected by the Medicaid expansion. For example, ([Bachrach, Boozang and Lipson, 2015](#)) find that the Medicaid expansion led to a net reduction in uncompensated care in hospitals of about \$2.6 billion per year in expansion states. This translates into a reduction in total uncompensated care of about \$4.3 billion considering that hospitals provide about 60% of uncompensated care to the uninsured ([Coughlin, 2014](#)).

4 Data and Empirical Strategy

4.1 Consumer Credit Panel

Data on debts, delinquencies, and credit scores used in this study come from the Consumer Financial Protection Bureau’s Consumer Credit Panel (CCP), a nationally representative, 1-in-48 random sample of de-identified credit records drawn quarterly from a nationwide

credit reporting company (NCRC). Each credit record in the sample includes information about each of the individual debt obligations, or tradelines, reported on that record. This information includes each tradeline’s origination date, the source of the debt, its current balance, and its past payment history. Although de-identified, credit records in the CCP can be linked over time, allowing us to study the evolution of debts for consumers in our sample.

We define medical debts as those reported directly by a medical provider or by third-party debt collectors as unpaid medical bills.⁸ In what follows, we refer to medical debt and medical collections interchangeably. We note that our definition of medical debt is somewhat narrow by necessity. For example, credit card balances that are generated by paying for medical services could be considered a type of medical debt. However, while credit records contain information about outstanding credit card balances, the information is insufficient to determine the portion of those balances derived from medical services. Consequently, we exclude these from our definition of medical debt.

Our measure of medical debt is the flow of new debts incurred in each quarter. We focus on the flow of debt because we believe it more precisely captures the impact of the Medicaid expansion than the stock of outstanding medical debt more commonly used in previous studies. The stock of debt often includes older lingering obligations, many of which remain on a record for several years, and portions of which have been repaid. Because we observe the date an obligation is incurred, its origination date, we can be more precise about the timing of the debt. Moreover, this measure allows us to separate the original obligation from any subsequently repaid amount. In a novel and complementary exercise, we calculate collections repayment rates and incorporate these into our analysis of overall savings to consumers.

We measure financial distress as a deterioration in repayment status. For each credit account, the CCP includes up to 84 months of payment history. Using this information, we determine whether a tradeline transitioned into a higher state of delinquency during each quarter. This includes a transition from current to 30 days or more past due or from 60 into 90 or more days past due.⁹ We define transitions of a given category on any loan as a new delinquency.¹⁰ Like for medical debt, isolating flows into missed repayments, rather than

⁸Our definition does not distinguish debts reported directly by medical providers from those reported by third-party debt collectors. Nevertheless, nearly all medical debts (>> 99%) accrue in the form of unpaid bills sent to collections. For a broader discussion of medical collections see [Brevoort and Kambara \(2015\)](#).

⁹We consider any account that starts a quarter as 90 days past due or worse to be in default and do not include further transitions, such as charge-offs or repossessions, which often reflect lender-initiated actions, as instances of financial distress.

¹⁰Later in Section 6.5 we separate out serious (90 days or more) delinquency by loan type when considering net indirect effects of the expansion. Details are in Appendix Section D.

contemporaneous payment status of all outstanding debts, allows us to better capture the timing of delinquency and to more cleanly identify changes in distress due to the reform.

Creditworthiness is often summarized via a credit score, such as the FICO or Van-tageScore. Each quarterly archive of CCP data includes a widely-used, commercially available credit score that was produced for that record. Although payment history is generally the most important determinant of credit scores, other factors that are associated with future default, such as utilization rates on revolving credit accounts, are incorporated into the scoring models. This provides a comprehensive assessment of the consumer’s financial health and likelihood of future financial distress that is used extensively by lenders to assess creditworthiness when underwriting and pricing credit. We use these scores as a measure of the direct link between the Medicaid expansion and the financial health of the newly insured.

Our analysis is based on a balanced sample of adults aged 18-64 residing either in states that expanded Medicaid on January 1, 2014, our treatment group, or in states that had not expanded by the end of our sample period (end of 2016), our control group. We exclude individuals residing in early-adopting states because the information necessary to determine which third-party collection accounts from medical debt is not available in the CCP for archives before September 2011. Given that early expansions began around the passage of the ACA in 2010, this means we cannot measure pre-treatment outcomes for residents of these states. Moreover, early expansions were generally implemented more gradually across the state, with a complete roll out not occurring until 2014 ([Kaiser Family Foundation, 2012b](#)). Were our data to extend far enough in time, it would still be difficult to ascribe a treatment period for these states.¹¹

We also exclude late adopting states for two reasons. First, we lack sufficient post-expansion data for some of these states as our sample period goes only until the end of 2016. Second, and more importantly, we adopt the synthetic control method proposed by [Abadie and Gardeazabal \(2003\)](#) to address potential concerns about pre-existing trends in the treatment states. Naturally, this method can be applied to a simple difference-in-differences framework but it is unclear how the framework should be extended to multiple treatment periods (events) when considering the late adopters. Hence, to simplify the analysis we exclude the late adopters in the main analysis and consider their inclusion in additional robustness exercises.¹² Finally, we balance the sample to clean out records flagged as fraudulent or belonging to individuals just entering the formal credit sector, mostly very young

¹¹We further note that the timing of early expansions coincides with the peak of the Great Recession. Including these states may therefore confound effects of the recession on household balance sheets.

¹²We included the late adopters in an earlier version of the paper, which did not adopt the synthetic control method. The old findings are very similar to the results presented below and are available upon request.

borrowers. In all, our baseline sample consists of about 5 million quarterly records tracked from the third quarter of 2011 to the end of 2016.¹³

Unfortunately, our data do not allow us to identify Medicaid eligibility at the individual level. To help assuage concerns related to this, we supplement the CCP with data on within-Census-tract income distributions and pre-reform statewide income eligibility criteria. Income data are from the American Community Survey’s (ACS) 2009-2013 5-year averages, and statewide eligibility criteria are compiled by the Kaiser Foundation.¹⁴ We define and calculate the proportion of newly eligible adults in each tract as the maximum of zero and difference in the fraction of residents eligible under the new benchmark of 138% of the federal poverty level and under the prior statewide rules. This provides us with rich variation in the share of newly eligibles between Census tracts, ranging from 0% to 100% of the entire Census tract population. We leverage this variation in supplemental analyses to corroborate our baseline findings.

4.2 Loan Offers and Pricing (Intel and MyFico)

Data on credit offers and pricing are from Intel Comperemedia (Intel), for credit cards and personal loans, and aggregated rate-sheet data from Fair Isaac Corporation’s MyFico website, for auto loans and mortgages. These data sources provide information on loan offers and pricing, allowing us to measure the effect of improved financial health on consumer’s credit options.

The Intel data provide information on credit card and personal loan offers. These data are generated via a nationally representative monthly survey of approximately 2,000 households, or 4,000 individuals. Participants are asked to provide Intel with all mail solicitations they received during the month. These include offers of new credit from all lenders in the marketplace.¹⁵ Despite the rise of internet search sites, direct mail remains one of the most popular and effective channels by which lenders advertise both credit cards and personal loans to potential customers. More than 1 billion offers are sent to consumers each year. Intel combines offer information from the mailings with the demographic profiles of respondents, including the county in which they reside. Because the Intel data provide extensive information on the supply decisions of nearly all lenders in the marketplace, as well

¹³Frequently there is a significant lag between when debts are incurred and when they are reported to the NCRCs. To account for this lag, we identify new tradelines opened in a calendar quarter by looking at the CCP archive from the following quarter. Because the reporting delay does not affect the tradeline’s reported opening date, we can assign later-reported medical debts to the quarters in which they were incurred.

¹⁴Specifically, for parsimony, and to highlight that the main beneficiaries of the expansion are childless adults, we use eligibility criteria by state for childless adults as of January 1, 2013.

¹⁵These include nearly all marketing solicitations and are not restricted to direct credit offers.

as demographic information on recipients, they are uniquely suited for exploring changes in the supply of credit to consumers following the Medicaid expansion. In our analysis, we focus on credit card and personal loan offers that have been pre-screened.¹⁶ Pre-screened offers are made to potential customers whose credit record, and score, has been previously checked and as a result are targeted toward specific risk types.¹⁷ We use data on repeated cross sections of respondents from the third quarter of 2011 to the end of 2016.

Mortgages and auto loans are less commonly offered through direct mail. However, in pricing mortgage and auto loans, lenders often set rates uniformly within credit score ranges. *Rate sheets*, which are often set by each lender statewide or nationally, translate credit score ranges into the interest rates available from a lender. Fair Isaac Corporation uses aggregated rate-sheet information to estimate the prevailing interest rates on mortgage and auto loans for credit score ranges that are widely used by lenders for both products. They publish this information on their educational *MyFico* website. In our analysis, we use these rate sheets to assign to each consumer the auto loan and mortgage interest rate they would have qualified for in that quarter based on their credit score. We then estimate the impacts of the reform on available auto loan and mortgage rates using these imputed values.¹⁸

4.3 Empirical Strategy

To identify the effect of the Medicaid expansion on medical debt accruals, creditworthiness, and lenders' pricing of credit products to consumers, we exploit the states' decisions whether and when to adopt in a simple difference-in-differences framework. Our primary specification is as follows:

$$y_{ct}^k = \alpha_c^k + \eta_t^k + \sum_{j=1}^3 \beta_j^k \times Exp_c \times Post_j + \epsilon_{ct}^k, \quad (8)$$

where, y_{ct}^k denotes outcome measure k in Census tract c in year-quarter t . Specifically, k is either a measure of medical debt, delinquency, credit score, or the interest rate on different types of loans. α_c^k and η_t^k denote Census-tract and quarter-year fixed effects, respectively. Exp_c is an indicator variable that turns on if Census tract c is located in an expansion state. $Post_j$ is an indicator variable that turns on in the j^{th} year following a states' adoption of the Medicaid expansion. The coefficients of interest are, β_j , $j = 1, 2, 3$, which map out the

¹⁶Pre-screened offers are identified via a flag for the presence of a pre-screen opt-out disclosure. Opt-out disclosures are required by law for pre-screened mail out offers.

¹⁷Often mail-out offers are made without screening consumers. These occur with the roll out of new products in an effort to learn their profitability.

¹⁸Consumers with credit scores below the bottom price tiers are excluded from calculations concerning automobile loans and mortgages, as they likely cannot qualify for a loan.

full dynamic effects of the policy during its first three years.¹⁹ We interpret the coefficients β_j , $j = 1, 2, 3$ as the intent-to-treat (ITT) effect of the Medicaid expansion. To construct an estimate of the average treatment effect on the treated (ATT), we first estimate an average intent-to-treat effect over the three post-reform years. Specifically, we estimate a simplified version of Equation 8 in which we replace the three post-reform indicators by a single indicator. We then divide the parameter estimate by the average increase in insurance coverage of 4.4 percentage points over the three post-reform years, see Section 3.

One empirical challenge is that medical debt and measures of distress differ in treatment states from national pre-reform averages. While we can control for differences in levels via Census tract fixed effects, one may be concerned that our reform effects may be confounded by differences in trends between treatment and control states, which are unrelated to the reform. We address these concerns in two ways. First, we apply the synthetic control method proposed by [Abadie and Gardeazabal \(2003\)](#). Specifically, we construct weights for the non-expansion states such that they match pre-reform trends and levels in medical collections, measures of financial distress, and offered interest rates in the expansion states. This data-driven procedure equates the trends in the key endogenous variables between adopting and non-adopting states allowing us to isolate the impact of the reform. We provide a graphical inspection of the "parallel trend" assumption for each outcome variable over the pre-reform period.²⁰

Second, we also exploit within-state variation across Census tracts that differ in the fraction of newly eligible adults. To this end, we split the sample into quantiles of Census tracts based on their fraction of newly Medicaid eligible adults. We then estimate Equation 8 within each subsample allowing for a non-parametric comparison of the reform effects between Census tracts.

5 Direct Effects of Medicaid on Medical Debt

5.1 Mean Effects

Table 1 shows the effects of the expansion on medical debt accrual. The dependent variable is the average value of newly accrued medical debt in each tract-quarter-year (ct). The table

¹⁹We acknowledge the potential for serial correlation in the unobservables. Accordingly, we cluster standard errors at the Census tract level.

²⁰Our unit of analysis is the Census-tract-quarter. To construct the synthetic control weights, we first aggregate observations at the state-quarter level and calculate state-specific weights following [Abadie and Gardeazabal \(2003\)](#). We present the corresponding weights in Appendix Section B. We then scale the weights to Census tract level by multiplying the state weight with the fraction of the state population aged 18-64 that lives in the respective Census tract.

shows the full dynamic effects of the reform in each of its first three years ($Exp \times Post_j$) as well as the average impact over the period. Column 1 of the table shows results for the

Table 1: Decline in Value of New Medical Collections

	By Proportion of Newly Medicaid Eligible Adults				
	All	1 st Quartile	2 nd Quartile	3 rd Quartile	4 th Quartile
	(1)	(2)	(3)	(4)	(5)
1 st Year ($Exp \times Post_1$)	-10.07 (1.31)	-2.99 (1.35)	-8.25 (2.11)	-10.58 (2.55)	-18.94 (4)
2 nd Year ($Exp \times Post_2$)	-17.17 (1.52)	-4.51 (1.39)	-10.87 (1.96)	-21.12 (3.78)	-33.06 (3.97)
3 rd Year ($Exp \times Post_3$)	-13.37 (1.14)	-6.42 (1.15)	-10 .8 (1.83)	-13.54 (2.35)	-23.59 (3.3)
<i>Average</i>	-13.54 (1.01)	-4.64 (0.9)	-9.97 (1.58)	-15.08 (2.18)	-25.19 (2.83)
<i>Pre-Expansion Ave. Value</i>	38.54	18.91	35.41	56.97	80.16
<i>Percent Newly Eligible Adults</i>	0-100	0-10	10-19	19-32	32-100
<i>Mean Tract CCP Population</i>	2,030	2,231	2,118	1,981	1,529
Observations	1,425,565	436,874	365,817	314,352	308,522

Notes: The table shows effects of the Medicaid expansion on the value of newly accrued medical debt from a regression using Equation 8. Data are from the CFPB’s Consumer Credit Panel described in section 4. The unit of observation is a Census-tract in a quarter-year (ct). As discussed in Section 4.3, all regressions include Tract and quarter-year fixed effects and are weighted by the synthetic control weight. Standard errors (in parentheses) are clustered by tract.

entire sample, while Columns 2-4 show these effects by quartile of newly eligible individuals.

As shown in the table, the Medicaid expansion led to a substantial decline in the value of newly accrued medical debt on the order of \$13.50 per quarter, or 35 percent, over its first three years. This effect grew in magnitude over time as the reform gained steam in its second and third year. The decline was also more pronounced in areas with a higher proportion of newly eligible adults, mostly poorer communities. The intent-to-treat effect was nearly 5 times larger in high (Column 5) relative to low (Column 1) eligibility Tracts. For the most vulnerable consumers, the average quarterly decline in newly accrued medical debt amounted to more than \$25. These average effects imply substantial aggregate reductions in medical debt across treatment states. Scaling our per-capita estimates with population weights from the CCP, Column 1 in the top panel of Table 2 shows that the expansion prevented $\$1.46 + \$2.39 + \$2.04 = \5.89 billion in medical collections from being debited to households balance sheets over the first three post-reform years.

Table 2: Reduction and Repayment of Medical Debt

	All (1)	1 st Quartile (2)	2 nd Quartile (3)	3 rd Quartile (4)	4 th Quartile (5)
<i>Annual Aggregate Decrease in Accrued Medical Collections (\$ Millions)</i>					
1 st Year - 2014	1,457	226	417	360	455
2 nd Year - 2015	2,390	341	548	717	785
3 rd Year - 2016	2,037	486	543	458	550
<i>Proportion of New Medical Collections Repaid (p.p)</i>					
<i>After One Year</i>					
Repaid	7.53	9.67	8.79	8.16	5.46
Repaid or Removed	38	35.85	37.34	38.55	38.69
<i>After Two Years</i>					
Repaid	8.57	10.59	10	9.36	6.3
Repaid or Removed	53.02	50.22	52.53	53.87	53.58
<i>Lower Bound of Medical Collection Repayments (\$ Millions)</i>					
1 st Year - 2014	110	22	37	29	25
2 nd Year - 2015	180	33	48	58	43
3 rd Year - 2016	153	47	48	37	30

Notes: This table presents our calculations of aggregate annual reduction in medical debt, repayment rates, and decline in out-of-pocket expenditures due to the Medicaid expansion. Accrued savings are calculated by multiplying estimates from Table 1 by the CCP population in respective Census tracts during each respective quarter-year. Repayment rates are calculated directly from the data. Percent repaid is the proportion of new medical collections accrued in year-quarter t that were repaid one and two years later, respectively. Percent removed is the proportion of new medical collections removed from the credit report one- and two-years later, respectively. The lower bound on the decline in out-of-pocket expenditures is calculated as the product of decline in accrued debts and the proportion of debts paid off within one year. The CCP Population is calculated by multiplying the number of records in each respective Census tract-quarter-year by 48, the inverse of the population sampling rate (Section 4)

A direct financial benefit from fewer medical debts in collection is a reduction in repayments to collection agencies. Collection repayments may or may not be included in self-reported out-of-pocket expenditures and have received very little attention in the existing literature. To identify collection repayments, we track these debt obligations in the CCP for two years after they originated. In doing so, we observe that some obligations continue to exist on the record one or two years after being debited and that, in some instances, the balances on those debts has declined. We also observe that some debt obligations are removed altogether from the credit record. These removals may represent instances in which the outstanding amount was repaid in full or in which the debt obligation was removed because the debt was returned by the debt collector to the original creditor.

We interpret partial reductions in the face value of collections as consumer repayments. As shown in the middle panel of Table 2, this suggests that out of every dollar sent to collections about 8 (9) cents are repaid within one (two) years. This proportion declines slightly in Census tracts with a high proportion of newly eligible individuals, tempering any effects on repayment savings to these consumers. We note that our estimated repayment rate falls into the ballpark of estimated purchasing prices for medical debt, paid by collection agencies. In 2009, according to [Federal Trade Commission \(2013\)](#), debt buyers paid about 5 cents per dollar of medical debt. Taken at face value, this provides collection agencies with a margin of 4 cents per dollar of medical debt to cover remaining operating expenses. Combining partial repayments and removed debt obligations, Table 2 also shows that out of every dollar sent to collections about 38 (53) cents are no longer owed after one (two) year(s) of the collection being debited.

In the bottom panel of the table, we quantify direct savings in repayments by multiplying the aggregate decline in newly accrued debt by the proportion of collections partially repaid after one year of being debited. We view this estimate as a conservative lower bound as we attribute no removals of debt obligations to consumer repayments. Even with this conservative measure, our calculations suggest substantial aggregate decreases in medical debt repayments on the order of $\$110 + \$180 + \$153 = \443 million between 2014 and 2016.²¹

To put our estimated reduction in medical debt into perspective, we benchmark our results to estimates from the landmark Oregon Health Insurance Experiment. ([Finkelstein et al., 2012](#)) find that Medicaid insurance reduced medical debt by \$390 (standard error 177) per treated person per year in its first year. Our findings suggest a per-capita decline in newly accrued medical debt of \$10.07 per quarter in the first post-reform year, see Table 1. Divided by the estimated coverage increase in 2014, we find an annual reduction of $\frac{\$10.07 \times 4}{0.041} = \982 per newly insured person. When accounting for differences in the measurement of medical collections resulting from attrition (e.g. $\sim 38\%$ of debts are no longer owed after one year) we find a debt reduction per treated person per year of approximately \$609. Although the Oregon experiment focused on a small and geographically concentrated sample of consumers, we find its estimated savings to be remarkably close to our national averages. Unlike in the Oregon experiment, we do not observe individual treatment for Medicaid. Therefore, we

²¹Of course, our repayment calculations to collection agencies are subject to various caveats. Most importantly, we cannot ascertain that reductions in the face value actually reflect repayments. Alternatively, the collection agency may forgive a fraction of the face value and or the debt was made in error. In our view, correcting errors and or providing discounts might lead to larger reductions in medical debt but we acknowledge that we cannot address these concerns completely. To provide a conservative assessment on the costs of leaving medical bills unpaid, we do not include these repayments in our consumer welfare calculations presented in Section 7.

interpret this congruence as further validation of our intent-to-treat approach for identifying the exogenous effects of the reform.

Our estimates also provide evidence on the relative significance of uninsured care or bad debt in uncompensated care, an estimate that, to the best of our knowledge, is not readily available from the literature. As outlined in Section 3.2, we take as given that the uninsured pay about 20% of overall health care utilization, worth \$2,400 per year, out-of-pocket (Hamel et al., 2016). This suggests that uncompensated care equals about \$1,920 per uninsured person and year.²² Considering the average effects of the post-reform period and dividing by the average coverage increase of 4.4 percentage points, we find a reduction in medical debt in collections of about $\frac{\$13.5 \times 4}{0.044} = \$1,227$ per treated person in the first year, which is about 51% of overall health care utilization or about 64% of uncompensated care. We conclude that almost two thirds of uncompensated care is sent to collections.

Robustness: As discussed in section 4.3, we consider various robustness checks to support our main findings. First, we complement the regression results with graphical evidence to support our identifying assumptions. To this end, we plot the time-series of the value of new medical debt in collection for initial adoption and the (synthetic) control states to inspect the "parallel trends" assumption. The trends are very similar in the pre-reform period and clearly diverge following the Medicaid expansions, which lends support to our empirical strategy, see Figure A.1.

Further, we investigate the robustness of our findings with respect to the concurrent opening of the private individual insurance exchanges, which may directly affect the accumulation of medical debt. To address this concern, we repeat our analysis using only those states that opted for a federal platform. We find that our findings are robust to this restriction, with the caveat that the smaller sample size introduces more noise into the results, see the left graph of Figure A.2.

Lastly, we consider that our findings may be driven by systematic changes in collections activities coinciding with the expansion of Medicaid. As aforementioned, our data are unique in that they allow us to distinguish between third-party collection originating from Medical providers and all other non-medical third-party collections. We exploit this feature of our data to analyze potential effects of the reform on non-medical collections. We find no evidence of changes in non-medical collections across treatment and control groups and conclude that systematic changes in collections activities are likely not important determinants of our

²²This is roughly consistent with the evidence from Garthwaite, Gross and Notowidigdo (2015), who find that each additional uninsured person costs a local hospital about \$900 annually in uncompensated care, given that hospitals only provide about 60% of the overall uncompensated care to the uninsured, see (Coughlin, 2014).

measured effects, see the right graph of Figure A.2. For details on these robustness exercises see the Appendix Section A.

5.2 Variance Effects

As outlined in Section 2, financial benefits arise from both a mean *and* variance reduction in unpaid medical bills. To assess the potential benefits from a variance reduction in paid and unpaid medical bills, we now turn to the distributional effects of the Medicaid expansion. To this end, we analyze the policy’s impact separately on the incidence of accruing large medical bills ($> \$1,000$) as compared to small bills ($\leq \$1,000$) in collection.²³ We then measure its overall effect on the distribution of medical debt accrual. For ease of exposition, we abstract away from the dynamic treatment effects and focus on the average impact over the adoption period.

The top panel of Figure 3 shows the effects of the reform separately by value of collections and across eligibility quartiles. The dependent variable is the proportion of individuals receiving a small (large) collection in a Tract-quarter-year. As compared to small collections, we note a substantially greater reduction in the incidence of large bills. While the propensity to accrue large unpaid medical collections is less than a third of that for small medical collections, the decline in accrual due to the reform is substantially greater. Moreover, we find that this difference is monotonically increasing with the proportion of newly eligible adults in the community. For example, in a community at the bottom quartile of eligibility, the incidence of a large medical collection declines by approximately 34 percent, as compared to 4 percent for a small collection. In communities with high rates of newly eligible adults, the incidence of small debts declines by about 11 percent. For these same communities, large debts drop by more than half, suggesting that the policy went a long way toward eliminating costly bills for the uninsured in these communities.

We refine the analysis of small and large collections in the bottom graphs of Figure 3, which illustrate the effects of the policy on the distribution of newly accrued medical debt. The bottom left panel summarizes the results of a separate regression at each percentile, whereby the dependent variable is simply the natural logarithm of a corresponding percentile

²³We also view this analysis as a proof of concept. Often small value medical collections result from charges that exceed the “reasonable and customary” charges that insurers pay or disputes about insurance coverage, whereby insured individuals may incur collections without any knowledge of a missed payment (Brevoort and Kambara, 2015). In contrast, large value medical collections more commonly arise from unpaid emergency room visits or hospital admissions of uninsured individuals. To the extent that the reform provided insurance to the previously uninsured, we would expect a relatively greater impact on the incidence of large value debts.

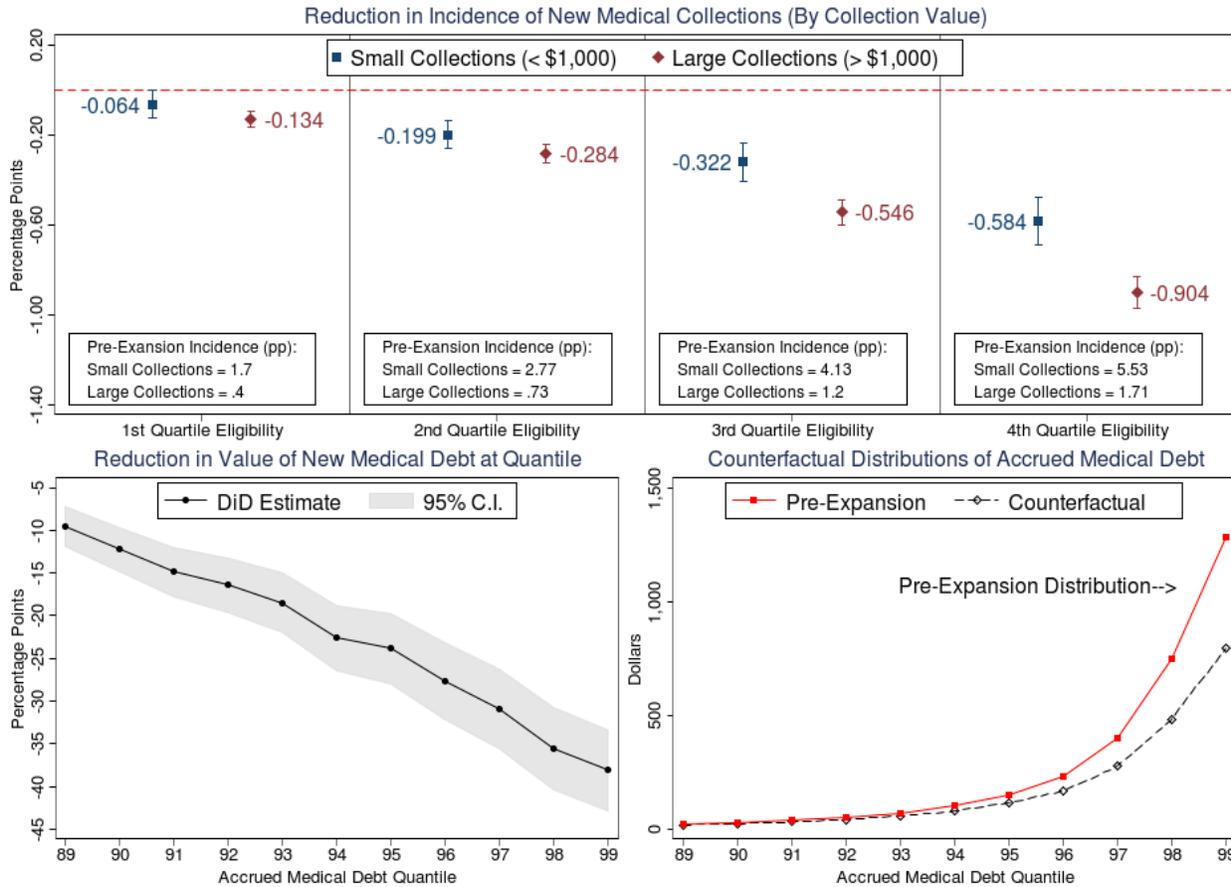


Figure 3: Distributional Effects of Expansion on Medical Debt

Notes: The figure shows distributional effects of the reform on the accrual of medical debt. Data are from CFPB's CCP and span the third quarter of 2011 to the end of 2016. The top panel show treatment effects and 95% confidence intervals on the incidence of large ($\geq \$1,000$) vs. small ($< \$1,000$) collections using Equation 8 separately for each eligibility quartile. The bottom left panel plots treatment effects and point-wise confidence intervals at each quantile of medical debt in tract c and quarter t . Regressions include Tract and quarter-year fixed effects and are weighted by the synthetic control weight. Standard errors (in parentheses) are clustered by tract.

in the distribution of newly-accrued medical debt at the Census tract quarter-year level.²⁴ The bottom right panel then plots the corresponding level effects, where we simply scale the percentage reduction with the pre-reform levels among adopting states.

Our findings suggest that the policy is more effective at eliminating tail-end risk for uninsured individuals. Specifically, the policy induced reduction in new medical debt rises from approximately 10 percent at the 89th quantile to 40 percent at the 99th quantile. We transform these relative effects into levels in the bottom right panel of Figure 3. The estimated

²⁴Although fewer than 5% of consumers receive a medical collection in each quarter on average, this may mask some variation across Census tracts. This is why we can identify effects at the 89th quantile.

10 percent reduction at the 89th quantile, on a base of about \$20, translates to a modest savings of about \$2. However, the savings become quite substantial past the 95th percentile. For the highest quantile, a 40 percent reduction in the accrual of new medical debt, on a base of \$1,450 translates to a more than \$580 quarterly reduction in medical debt accruals.

6 Indirect Effect of Medicaid on Financial Health

6.1 Medical Debt and Distress

Building on the direct effects of the Medicaid expansion on the accrual of medical debt, we now turn to the relationship between medical debt and measures of financial distress. As outlined above, medical collections are reported to credit bureaus and may adversely affect an individual's credit score and subsequently the terms of credit. To provide first evidence on this negative relationship, we investigate the consequences of a new medical collection on credit scores using an event study framework, which we detail in Appendix C. Specifically, we focus on a person's first medical collection and test for a concurrent drop in that person's credit score around the timing of the collection. The event study reveals a sharp and persistent 20 point decline in credit scores among individuals with ex-ante credit scores less than 620 (subprime) following the accrual of the first medical collection. Among individuals with ex-ante credit scores above 620 (prime), we observe a still larger 40 point drop in credit scores (Figure C.1). We further document adverse impacts of a newly accrued medical collection on future delinquencies. Overall, this event study provides first direct evidence of the adverse relationship between medical debt in collection on the one hand and credit scores and delinquency rates on the other. Motivated by this evidence, we now turn to the indirect effects of the Medicaid expansion on repayment delinquency and credit scores.

6.2 Medicaid and Periods of Worsening Distress

We begin our analysis of the indirect effect of insurance by studying the impact of the Medicaid expansion on the incidence of new delinquencies in a tract-quarter-year. Table 3 shows the effects of the reform on new 30 day (top panel) and 90 day (bottom panel) delinquencies on any debt obligation. A new 30 day delinquency indicates a moderate degree of financial distress, whereby a substantial portion of these newly delinquent individuals are able to get back on track with their payments. In contrast, becoming 90 day or more delinquent shows a more serious degree of distress. Few individuals ever more than 90 days delinquent become current on their payments.

Table 3: Decline in New Delinquencies

	By Proportion of Newly Medicaid Eligible Adults				
	All (1)	1 st Quartile (2)	2 nd Quartile (3)	3 rd Quartile (4)	4 th Quartile (5)
<i>Decline in New 30 Day Delinquencies</i>					
1 st Year ($Exp \times Post_1$)	-0.06 (0.04)	-0.04 (0.06)	-0.01 (0.07)	-0.08 (0.07)	-0.13 (0.08)
2 nd Year ($Exp \times Post_2$)	-0.09 (0.04)	-0.07 (0.07)	-0.09 (0.07)	-0.07 (0.07)	-0.13 (0.09)
3 rd Year ($Exp \times Post_3$)	-0.06 (0.04)	-0.06 (0.07)	0.07 (0.07)	-0.07 (0.08)	-0.22 (0.09)
Average	-0.07 (0.03)	-0.06 (0.05)	-0.01 (0.06)	-0.07 (0.06)	-0.16 (0.07)
<i>Pre-Expansion Average</i>	6.03	5.71	6.14	6.35	6.41
Observations	685,747	218,435	178,745	150,524	138,043
<i>Decline in New 90 Day Delinquencies</i>					
1 st Year ($Exp \times Post_1$)	-0.04 (0.02)	0.00 (0.03)	-0.01 (0.03)	-0.06 (0.04)	-0.07 (0.04)
2 nd Year ($Exp \times Post_2$)	-0.07 (0.02)	-0.05 (0.03)	-0.02 (0.03)	-0.09 (0.04)	-0.14 (0.05)
3 rd Year ($Exp \times Post_3$)	-0.03 (0.02)	0.01 (0.03)	0.01 (0.03)	-0.04 (0.04)	-0.12 (0.05)
Average	-0.05 (0.01)	-0.01 (0.02)	-0.01 (0.03)	-0.06 (0.03)	-0.11 (0.03)
<i>Pre-Expansion Average</i>	2.03	1.74	2.07	2.24	2.56
Observations	762,157	228,841	197,773	173,256	162,287

Notes: The table shows effects of the Medicaid expansion on the rate of new delinquencies using Equation 8. Data are from the CFPB's Consumer Credit Panel described in section 4. The unit of observation is a Census-tract in a quarter-year (*ct*). All regressions include Census tract and quarter-year fixed effects and are weighted by the synthetic control weight. Standard errors (in parentheses) are clustered by tract.

During its first three years the expansion reduced new 30 day delinquencies by $\frac{0.07\%}{6.03\%} = 1.2$ percent of the pre-reform mean. Severe, or 90 day, delinquencies declined relatively more at $\frac{0.05\%}{2.03\%} = 2.5$ percent. Like for medical debt (Table 1), the impact of the reform on new delinquencies rose substantially (in magnitude) between the first and the second post-reform years, retreating back in the third year. Further, for both 30 and 90 day delinquencies, the decline in financial distress was highly concentrated in Census tracts with a high proportion of newly eligible adults (Column 4). On both measures, the effect was about twice as large for the highest quartile of newly eligibles as for the population as a whole.

Putting these intent-to-treat estimates into perspective, we again divide the estimates by the average coverage increase. This scaling suggests that transitions into 30 day delinquency among the newly insured declined by $\frac{0.07}{0.044} = 1.59$ percentage points. The size of the effect is large exceeding the differences in the pre-reform averages between Census tracts with the smallest and largest fraction of newly eligible individuals. For severe 90 day delinquencies, we estimate a $\frac{0.05}{0.044} = 1.13$ percentage point reduction, which amounts to nearly 50 percent of the pre-reform mean rate (2.56) in Census tracts with the largest fraction of newly eligible adults (Column 4).

6.3 Medicaid and Credit Scores

Turning to credit scores, Table 4 presents the effects of the Medicaid expansion on the average credit score in the Census tract. Overall we find a statistically significant increase of about 0.5 points in the first post-reform year (Column 1). The effect triples by the third post-reform year.²⁵ In line with previous outcomes, the rise in credit scores is concentrated in communities with a high proportion of newly eligible adults (Column 5). In these communities, the overall effect is more than two times larger than the average. Once again, we scale the overall average effect by the estimated coverage increase and find that treated individuals saw a $\frac{1.04}{0.044} = 23.6$ point increase in credit scores. The effect increases to a $\frac{1.52}{0.044} = 34.5$ points in the third year, which corresponds to almost half of the pre-reform difference in average credit scores between the bottom and the top eligibility quartile.

To better understand these effects, we explore heterogeneity in the reform’s effect on credit scores at different points on the credit score distribution. To this end, we replace the average credit score with the credit score quantile in tract quarter-year ct as the dependent variable in Equation 8. The top panels of Figure 4 show these effects for the first (left), second (middle) and third (right) post-reform year. The points correspond to an average treatment effect at a given quantile, and the shaded region shows the point-wise 95 percent confidence interval for the estimate. The bottom panel shows the pre-treatment distribution of credit scores among expanding states.

²⁵The majority of enrollments occurred in the first year following the reform (Section 3). Consequently, we interpret the growth in our intent-to-treat estimates as evidence for an increasing treatment effect on the treated. Credit scores are calculated using past repayment behavior, usually incorporating a complex weighting of debts and delinquencies dating back 7 or more years. As a result, improvements in credit scores are likely slower to manifest than declines in medical debt accruals and worsening delinquency, both flows. Note also that this is not in contradiction with the seemingly immediate effect of a new medical collection on an individual’s credit scores. This is because, although new collections are immediately incorporated, they continue to influence scores for many quarters after. As shown in Figure C.1 in Appendix Section C, credit score drops are persistent and take a long time to recover from.

Table 4: Rise in Consumer Credit Scores

	By Proportion of Newly Medicaid Eligible Adults				
	All (1)	1 st Quartile (2)	2 nd Quartile (3)	3 rd Quartile (4)	4 th Quartile (5)
1 st Year ($Exp \times Post_1$)	0.53 (0.16)	0.4 (0.31)	0.3 (0.29)	0.41 (0.31)	1.25 (0.4)
2 nd Year ($Exp \times Post_2$)	1.08 (0.2)	1.05 (0.38)	0.56 (0.37)	0.88 (0.38)	2.33 (0.49)
3 rd Year ($Exp \times Post_3$)	1.52 (0.23)	0.96 (0.44)	0.98 (0.41)	1.26 (0.45)	3.5 (0.56)
Average	1.04 (0.17)	0.8 (0.33)	0.61 (0.32)	0.85 (0.33)	2.36 (0.42)
<i>Pre-Expansion Average Score</i>	679.25	701.63	678.94	661.83	627.55
<i>Percent Newly Eligible Adults</i>	0-100	0-10	10-19	19-32	32-100
Observations	707,616	211,488	180,402	162,305	153,421

Notes: The table shows effects of the Medicaid expansion on credit scores using Equation 8. Data are from the CFPB’s Consumer Credit Panel described in section 4. The unit of observation is a Census-tract in a quarter-year (ct). As discussed in section 4.3, all regressions include Tract and quarter-year fixed effects and are weighted by the CCP population, scaled by the synthetic control weight, in a given tract-quarter-year. Standard errors (in parentheses) are clustered by tract.

The graphs illustrate an inverse u-shaped pattern indicating that the reform’s effect are smaller at the very bottom and the top of the credit score distribution. The effects in year 2 and 3 peak at around the 30th percentile, which corresponds to the border between subprime and deep subprime, or a credit score of around 600 as indicated in the bottom graph. In the third post-reform year, we see that individuals in the middle two quartile of baseline credit score distribution benefit from a 2 point increase in their credit score. Individuals in the bottom and top quartile see smaller increases in their credit scores.

6.4 Medicaid and the Terms of Offered Credit

Given the reform’s positive role in improving individuals’ repayment outcomes and subsequently their credit scores, we now turn to its impact on the pricing of credit offered to consumers. Specifically, we look at the four most common types of debt obligations held by consumers: credit cards; personal loans (unsecured installment credit); auto loans; and mortgages. For credit cards and unsecured personal loans, we observe direct offerings; for auto loans and mortgages, we impute interest rates based on observed credit scores and

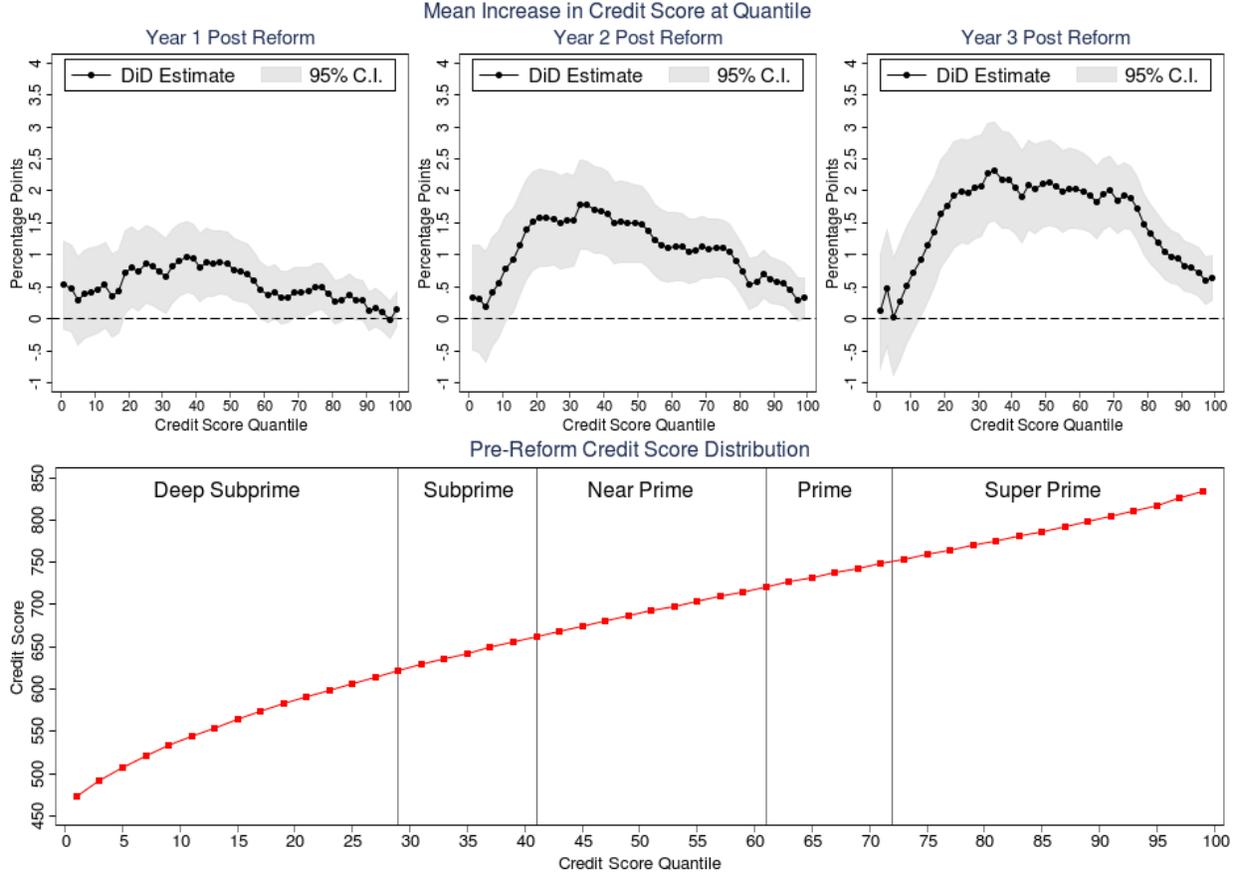


Figure 4: Distributional Effects of Expansion on Credit Scores

Notes: The figure shows distributional effects of the reform on credit scores. The top panels plot treatment effects (Equation 8) and pointwise confidence intervals at each quantile of medical debt in tract c and quarter t for the first, second, and third years post reform, respectively. Regressions are weighted by the synthetic control weight and standard errors are clustered by tract. The bottom panel plots the pre-reform distribution of credit scores.

the credit score interest rate crosswalk provided by aggregated rate sheet data (Appendix Section D).

Leveraging the county of residence information from the Mintel data, we estimate the reform’s effect on interest rates at the county year-quarter level using Equation 8. We also split the sample into subpopulations based on the fraction of newly Medicaid-eligible adults. However, due to the limited sample size, we group counties into categories depending on whether their share of newly eligible adults exceeds the share in the median county.²⁶

Consistent with our earlier findings, Table 5 shows statistically significant reductions in both credit card (columns 1-3) and personal loan (columns 4-6) interest rates. Furthermore,

²⁶Importantly, to maintain consistency in our interpretation of these heterogeneous effects, and for the refinancing exercise that follows, we define the median eligibility rate using the CCP and not Mintel.

the effects are more pronounced in communities with a higher proportion of newly eligible adults (Columns 3 and 6 for credit cards and personal loans, respectively). With respect to

Table 5: Decrease in Offered Rates (Measured in Basis Points)

	Credit Cards			Personal Loans		
		Below	Above		Below	Above
	All	Median	Median	All	Median	Median
	(1)	(2)	(3)	(4)	(5)	(6)
1 st Year ($Exp \times Post_1$)	-0.09 (0.16)	-0.11 (0.17)	-0.05 (0.43)	-0.49 (0.48)	-0.48 (0.54)	-0.50 (0.97)
2 nd Year ($Exp \times Post_2$)	-0.27 (0.25)	-0.27 (0.28)	-0.29 (0.55)	-0.72 (0.47)	-0.76 (0.51)	-0.57 (1.19)
3 rd Year ($Exp \times Post_3$)	-0.78 (0.22)	-0.62 (0.26)	-1.31 (0.40)	-0.33 (0.27)	-0.35 (0.30)	-0.21 (0.61)
Average	-0.42 (0.15)	-0.36 (0.17)	-0.60 (0.36)	-0.50 (0.26)	-0.52 (0.29)	-0.40 (0.52)
<i>Pre-Expansion Mean Offered Rates</i>	15.42	15.40	15.57	9.29	9.21	9.74
<i>Percent Newly Eligible Adults</i>	0-100	0-19	19-100	0-100	0-19	19-100
Observations	35,941	28,657	7,284	5,804	4,737	1,067

Notes: The table shows effects of the Medicaid expansion on credit card and personal loan interest rates offered to consumers. The unit of observation is an individual in a month, and the dependent variable is the average offered rate for each respective product offered to a consumer. All regressions include County and quarter-year fixed effects and are weighted by the synthetic control weight. Standard errors (in parentheses) are clustered by County.

credit cards, we estimate an average intent-to-treat effect of -42 basis points (bps) overall, and -60 bps in communities with more newly eligibles. Moreover, consistent with what we find for delinquencies and credit scores, the impact of the reform on offered credit card rates grew over time. While in its first year rates dropped by an average of -9 bps, by the third year this decline was 9 times larger. Among those living in communities with more newly eligible adults, this pattern is still more evident. While individuals in these communities saw only a 5 bp reduction during the first year following the reform, the average drop in the third year exceeded one percentage point.

For personal loans, the timing of the effects is reversed ranging from a reduction of 49 bps in the first year to a reduction of 33 bps in the third year. Unlike credit cards, which are widely used by prime borrowers, personal loans are a smaller market which largely focuses on subprime customers (Section 4). Therefore, it is conceivable that improvements in credit scores change the composition of borrowers of personal loans over time as these individuals gain access to more alternative forms of credit. These changes may well be reflected in the interest rate changes.

Robustness: Like for the direct effect, we complement our regressions with graphical evidence to substantiate the assumptions underlying our identification strategy. First, we plot the time series for delinquencies, credit scores, and interest rates by type of credit for on-time adoption states and synthetic control states. The patterns suggest that both time series run in parallel until the end of 2013 and start to diverge in the post-reform quarters, see Appendix Section A. Along with the presented differences between more and less affected Census Tracts, these findings corroborate our main conclusions concerning the indirect effects of the Medicaid expansion.

6.5 Dollar Value of Improved Financial Health

To ease the interpretation of our results on pricing, we calculate the potential annual interest savings. To this end, we restrict our population to individuals living in treatment states and consider a refinancing of their pre-form debt level at improved terms of credit. We abstract away from dynamic considerations and use as our preferred estimates the average interest rate savings over the three post-reform years. We also assume that the credit cards and personal loans are amortized over 36 months, that auto loans are refinanced as 5-year loans, and that mortgages are refinanced at 30-year, fixed-rate loans.²⁷ Unlike mortgages and auto loans, credit cards and personal loans are not backed by valuable assets. To provide a conservative estimate on the consumer savings on credit card debt and personal loans, we subtract increased repayments (via lower delinquency) from the estimated interest rate savings. Our simulation exercise is detailed in Appendix Section D.

Column 1 of Table 6 shows the annualized interest savings by type of credit. Overall, we find annual interest savings worth \$14.6 per person, which is about 30% of the per-person reduction in medical debt (Table 1).²⁸ Once again, we divide these intent-to-treat estimates by the fraction of non-elderly treated adults and find interest savings of about $\frac{\$14.60}{0.044} = \332 per treated person and year. Savings on credit card debt dominate the total effect, accounting for 74 percent of the total savings. A potential explanation for the large credit card savings is given by the reform’s negative effect on bankruptcy filing.²⁹ Consistent with evidence from

²⁷The amortization period for auto loans and mortgages is consistent with the interest rates published by FICO.

²⁸We note a small *negative* savings for mortgages among individuals in high eligibility Census tracts of \$0.04. This is because mortgages are generally not given to individuals with low credit scores, who are those more likely to live in these poorer communities. As a result, our regression using imputed credit scores yield a small positive and insignificant estimate of this effect (see Appendix D for details). For transparency we left these small negative effects as is, though we interpret it as effectively a zero.

²⁹We find that the Medicaid expansion reduced the number of personal bankruptcies by 25,000 per year in the on-time expansion states.

Table 6: Simulated Annual Interest Rate Savings

	All (1)	Below Median (2)	Above Median (3)
<i>Credit Cards</i>			
Average Per Person (\$)	10.78	11.18	9.92
Total (\$Millions)	496.74	352.12	144.61
<i>Personal Loans</i>			
Average Per Person (\$)	1.82	2.02	1.38
Total (\$Millions)	83.84	63.77	20.07
<i>Auto Loans</i>			
Average Per Person (\$)	1.20	1.13	1.33
Total (\$Millions)	55.07	35.69	19.37
<i>Mortgages</i>			
Average Per Person (\$)	0.80	1.19	-0.04
Total (\$Millions)	37.04	37.62	-0.57
<i>All</i>			
Average Per Person (\$)	14.60	15.53	12.59
Total (\$Millions)	672.69	489.21	183.48

Notes: The table shows results from simulations of consumer savings using intent-to-treat estimates in Tables 5, D.3, and D.1. The table shows per-person effects and total effects calculated using the CCP population. See Appendix Section D for further details.

the literature (Domowitz and Sartain, 1999; White, 2006), we observe that individuals filing for bankruptcy hold disproportionate amounts of unsecured debt, which can be discharged in the proceedings. Lenders may therefore predominantly reduce interest on credit card debt (and other unsecured debt) following a reduction in bankruptcy filings.

The interest savings differ across communities with a low (Column 2) and high (Column 3) proportion of newly eligible adults. Despite smaller interest rate reductions, we find larger interest savings in communities with a lower proportion of newly eligible adults. There are two reasons for this finding, which can be illustrated in the case of credit card debt. First, consumers in low eligibility communities hold significantly more credit card debt which contributes to higher overall savings. Second, we find larger reductions in the rate of delinquency in high eligibility communities (Table 3), leading to higher loan repayments and consequently lower net savings to consumers.

Discussion: Our calculation of interest savings is predicated on the assumption that individuals can refinance their outstanding debt at the improved terms of credit. This assumption is particularly important for credit card debt, which accounts for the lion share of

interest savings. Evidence in the literature on this topic remains mixed. While [Ponce, Seira and Zamarripa \(2017\)](#) argue that Mexican borrowers forgo arbitrage opportunities based on interest differences between credit cards, [Stango and Zinman \(2015\)](#) find that borrowers in the U.S. seem to minimize financing costs when allocating their debt across credit cards. Interestingly, [Stango and Zinman \(2015\)](#) also report significant differences in paid interest rates between borrowers, which they reconcile by differences in shopping intensities combined with differences in mail-out offers they receive. Our analysis focuses on pre-screened mail out offers, which alleviates concerns that consumers do not take advantage of the improved terms of credit.³⁰ Furthermore, we note that our calculated interest savings abstract away from changes in the optimal borrowing amount providing a lower bound on the benefits from improved terms of credit.

Finally, an important detail is that the CCP provides end-of-quarter snapshots of loan balances for respective individuals. A potential concern is that a portion of reported credit card balances may not constitute credit card borrowing per se. This is because a portion of reported credit card balances may still be held within the 'grace' period, and as a result not incur any financial charges.³¹ Unfortunately, we cannot distinguish interest bearing credit card balances from overall balances in individuals' credit records directly. However, detailed credit card data from the Board of Governors of the Federal Reserve System indicate that 85 percent of outstanding balances is held beyond the grace period and bears interest.³² We expect that the share of 'true' credit card debt in overall outstanding balances may be even higher in our low-income treatment population with lower than average credit scores.³³

7 Medical Bills and Consumer Welfare in Practice

We now return to the theoretical framework outlined in Section 2 as a basis for quantifying and decomposing the effects of insurance on consumer welfare. In particular, we calculate how a decline in the mean and variance of accrued medical bills impact consumer welfare

³⁰Consistent with this assessment, we note that the average interest rate on offered credit card debt in our sample, 15.4 percent see (Table 5), corresponds quite well to estimates of realized credit card rates from large and representative credit card account databases, see for example Figure 1 in ([Alexandrov, Grodzicki and Bedre-Defolie, 2018](#)).

³¹Individuals who pay their balance in full each billing cycle, e.g. while still in their grace period, are commonly called *transactors*. Individuals who carry, or revolve, balances across billing cycles are called *revolvers*. Unlike transactors, revolvers are often charged interest on their balances. Once a balance has been carried across a billing cycle, there is no longer a grace period on any balances until the account is repaid in full.

³²See page 73 in the CFBP's consumer credit card report from 2017, <https://bit.ly/2o0bTJj>, last accessed June 25th, 2018.

³³The Board's credit card data also indicate that individuals with lower credit scores are more likely to accrue interest on their credit card debt.

directly and *indirectly* via the credit channel. We carry out this decomposition using two alternative and conceptually disparate approaches, which we discuss and implement below. We provide details on the derivations in the Appendix Section E.

7.1 Revealed Preference Approach

In the first approach, we calibrate our model by combining estimates from the literature with our direct evidence on medical debt in collection. A complication then arises from the fact that individuals in our model possess preferences over consumption and unpaid medical bills. Assessing the CV and RP effects of lower medical bills requires knowledge of both. While we can borrow preference structures over consumption from the literature, we are not aware of analogous preference structures over unpaid bills. Therefore, we only calibrate preferences over consumption and reveal the disutility over unpaid medical bills from realized "optimal" out-of-pocket payments. The intuition behind this method is that the observed share of the total medical bill that is paid out-of-pocket provides information on the disutility of higher levels of unpaid medical bills. For instance, if the disutility over unpaid medical bills is very small, then we expect consumers to pay only a very small fraction of the overall bill out-of-pocket. We refer to this approach as the *revealed preference approach*.

To quantify the CV, we impose two simplifying assumptions on disutility of unpaid medical bills and the repayment rate of incurred medical costs. First, and motivated by the graphical analysis in Figure 1, we use a linear approximation to the marginal utility function around the proportion of a bill left unpaid, b^* . This simplification means that we only need to characterize the first and the second derivative of the disutility of medical debt. Second, and supported by our data, we assume that the fraction of unpaid bills, $\bar{\tau} = \frac{b^*}{\epsilon_{MB}}$, is constant.³⁴ Using the first order condition, and applying the implicit function theorem, we can then express the disutility of medical debt in terms of the utility over consumption and the fraction of unpaid bills.

We focus on the CV under increasing marginal disutilities in medical debt, $h''(\cdot) > 0$ which provides a lower bound for the case, $h''(\cdot) = 0$. We can then express the CV as a function of the curvature of consumption utility, $\phi(\cdot) = -\frac{g'(\cdot)}{g''(\cdot)}$, the share of unpaid medical bills, $\bar{\tau}$, and the size of the medical bill, ϵ_{MB} as follows

$$CV = -\phi(\cdot) + (1 - \bar{\tau})\epsilon_{MB} + \sqrt{\phi(\cdot)^2 + 2\bar{\tau}\phi(\cdot)\epsilon_{MB} - \bar{\tau}(1 - \bar{\tau})\epsilon_{MB}^2}, \quad (9)$$

³⁴We see larger repayment rates for very small medical debts worth less than \$1,000. Importantly though, repayment rates are very similar for debts between \$1,000-\$10,000 or more than \$10,000. This pattern is inconsistent with quasilinear preferences, $h''(\cdot) = 0$, in which case individuals repay medical bills up to a given amount and borrow the rest in the form of medical debt.

if $\epsilon_{MB} \leq \frac{\phi(\cdot)}{1-\bar{\tau}}$, where $(\cdot) = (Y - (1 - \bar{\tau}) \cdot \epsilon_{MB})$.³⁵

This function possesses a number of intuitive properties. First, the CV is decreasing in the curvature of consumption utility. This is because, holding the repayment rate fixed, the implicit function theorem reconciles less curvature in consumption with less curvature in the disutility of medical debt. Graphically speaking, a decrease in curvature flattens out both marginal utility curves in Figure 1. This reduces the value of borrowing, R , and hence raises the CV.³⁶ Second, provided minimal curvature and sufficiently small medical bills, the CV is decreasing in the share of unpaid medical bills $\bar{\tau}$. This is because a decrease in $\bar{\tau}$ raises out-of-pocket spending and hence the CV. Moreover, a lower $\bar{\tau}$ signals that the credit consequences of unpaid medical bills are particularly costly from the point of view of the patient.³⁷ Lastly, provided minimal curvature in consumption utility, the ratio of CV and the medical bill, $\frac{CV}{\epsilon_{MB}}$, decreases in ϵ_{MB} . In other words, the credit channel is relatively more important for smaller medical bills.

To quantify the RP , we adopt a second order Taylor approximation to the utility function evaluated at the average medical bill $\bar{\epsilon}_{MB}$ again holding the repayment ratio $(1 - \bar{\tau})$ fixed. Let $\bar{r}^* = \bar{\epsilon}_{MB} - \bar{b}^* = (1 - \bar{\tau}) \cdot \bar{\epsilon}_{MB}$ be the average repayment. Building again on the first order condition and the implicit function theorem, we can implicitly express the RP in terms of preferences over consumption and the repayment ratio:

$$g(Y - \bar{r}^*) - g(Y - \bar{r}^* - RP) = -\frac{1}{2} \cdot (1 - \bar{\tau}) \cdot g''(Y - \bar{r}^*) \cdot var(\epsilon_{MB}). \quad (10)$$

Our direct effect benchmark for the RP is derived from a nested version of the model which ignores the impact of unpaid bills on utility: $h(\cdot) = 0$. We refer to this benchmark RP as RP^{oop} . By similar method to Equation 10, we can express RP^{oop} implicitly as

$$g(Y - \bar{r}^*) - g(Y - \bar{r}^* - RP^{oop}) = -\frac{1}{2} \cdot (1 - \bar{\tau})^2 \cdot g''(Y - \bar{r}^*) \cdot var(\epsilon_{MB}), \quad (11)$$

whereby Equations 10 and 11 together imply that

$$RP^{oop} \leq RP \leq \frac{1}{1 - \bar{\tau}} \cdot RP^{oop}, \quad (12)$$

or that the credit channel can increase the risk premium by a factor of $\frac{1}{1-\bar{\tau}}$.

³⁵The condition $\epsilon_{MB} \leq \frac{\phi(\cdot)}{1-\bar{\tau}}$ requires that the extrapolated marginal utility of consumption at $c = Y$ is weakly greater than zero.

³⁶For example, as $g''(\cdot)$ converges to zero, both marginal utility curves become horizontal and the CV converges to ϵ_{MB} .

³⁷An extreme case is $\bar{\tau} = 0$, in which case medical bills are fully repaid, the CV equals ϵ_{MB} .

Finally, we calibrate preferences over consumption, income, mean and variance in medical bills, and repayment decisions, which allows us to quantify the CV and the RP via Equations 9, 10, and 11. We follow (Finkelstein, Hendren and Luttmer, 2015) and consider CRRA utility over consumption with a parameter of relative risk aversion $\theta = 3$. We further normalize income to \$3,800. Following Coughlin (2014), we set the cost of care to \$2,400 and consider that the average patient pays 20 percent, or \$480 of their cost of care out-of-pocket.

As discussed in Section 3, uncompensated care comprises 'charity care', which is not billed to the consumer, and bad debt or unpaid medical bills, which are sent to collection agencies. To isolate the tradeoff between paid and unpaid medical bills, we first net out charity care. Building on our results on the direct effect of the reform in Table 1, we calculate that $\alpha_{charity} = 29\%$ of overall health care utilization is charity care.³⁸ Subtracting charity care from overall utilization reduces average annual health care utilization to $\bar{\epsilon}_{MB} = (1 - 0.29) \times \$2,400 = \$1,704$ and implies that individuals pay $1 - \bar{\tau} = \frac{0.2}{1-0.29} \approx 28\%$ of this net amount out-of-pocket.

Given the above, we can infer the relative importance of the CV to the out-of-pocket benchmark from the ratio of the implied CV (Equation 9) and the corresponding medical bill, $\frac{CV}{\epsilon_{MB}}$, divided by $1 - \bar{\tau}$. This is because the out-of-pocket spending is defined as $(1 - \bar{\tau}) \cdot \epsilon_{MB}$. Mathematically, we have:

$$\frac{CV}{OOP} = \frac{CV}{\epsilon_{MB}} \cdot \frac{1}{1 - \bar{\tau}}. \quad (13)$$

We evaluate Equation 13 at the average overall medical bill $\bar{\epsilon}_{MB} = \$1,704$, which provides a lower bound as $\frac{CV}{\epsilon_{MB}}$ is convex in ϵ_{MB} , see Figure E.1. Combining Equation 13 with Equation 9, we calculate that $\frac{CV}{\epsilon_{MB}} = 0.7$ or that $\frac{CV}{OOP} = \frac{0.7}{0.28} = 2.5$.³⁹ This suggests that the CV exceeds annual out-of-pocket savings by a factor of 2.5. Based on annual out-of-pocket spending of \$450 among the uninsured, we estimate a CV of $\$480 \times 2.5 = \$1,200$ and attribute $\$1,200 - \$480 = \$720$ of the CV to the indirect credit channel. Our results are detailed in Column 1 of Table 7.

Risk averse consumers also benefit from a reduction in risk. We evaluate the RP (Equation 10) around average annual consumption of \$3,300 and consider a standard deviation in consumption of \$768 as in (Finkelstein, Hendren and Luttmer, 2015).⁴⁰ Combining the

³⁸From Section 5, we find reductions in medical debt of $\frac{\$13.5 \times 4}{0.044} = \$1,227$ per treated person and year, which corresponds to 51% of overall health care utilization. For further details of these calculations, see Appendix Section E.5.

³⁹In Appendix Section E.5, we provide further details regarding this calculation, including a full plot of $\frac{CV}{\epsilon_{MB}}$ against ϵ_{MB} and under various assumptions about consumer risk preferences.

⁴⁰The consumption level corresponds to income net of average out-of-pocket spending: $\$3,800 - \$480 \approx \$3,300$.

Table 7: Overall Annual Financial Benefits

	Revealed Preference					Direct
	Baseline	$\theta=2$	$\theta=4$	OOP=15%	OOP=25%	
<i>Mean Effects</i>						
Credit Channel (Indirect)	720	870	620	700	740	332
OOP Spending (Direct)	480	480	480	480	480	480
Compensating Variation (CV)	1,200	1,320	1,100	1,180	1,220	812
Ratio: CV/OOP	2.5	2.75	2.29	2.46	2.54	1.7
<i>Variance Effects</i>						
Risk Premium (RP)	680	538	768	826	576	(680)
OOP Benchmark (RP OOP)	240	170	225	240	240	(240)
Ratio: RP/RP OOP	2.8	3.2	3.4	3.4	2.4	(2.8)
<i>Total Benefit</i>	1,880	1,858	1,868	2,000	1,796	1,492
<i>Total Spending</i>	2,400	2,400	2,400	2,400	2,400	2,400
Ratio: Benefit/Spending	0.78	0.77	0.78	0.83	0.75	0.62

Notes: Column 1 summarizes the consumer welfare gains under the revealed preference approach. This approach considers a parameter of relative risk aversion of $\theta = 3$ and assumes that 20 percent of health care utilization is paid out-of-pocket (OOP=20%). Columns 2-5 revisit the estimates under different parameters of relative risk aversion or different assumptions on out-of-pocket spending as a fraction of overall medical utilization. Column 6 summarizes the consumer welfare gains under the direct approach.

variation in consumption with the observed repayment ratio allows us to construct the variance in medical bills, which enters Equations 10 and 11.⁴¹ Combining these estimates we calculate a RP of \$680, which exceeds the pure OOP benchmark (Equation 11) by a factor of 2.8. This is also shown in Column 1 of Table 7. Combining the CV and RP in the table, we calculate an overall annual financial benefit of \$1,880. In other words, using the revealed preference approach, we calculate that the total value $TV = CV + RP$ amounts to about 78% of overall medical spending and exceeds the out-of-pocket benchmark by a factor of 2.6.

We note that the estimated consumer welfare gains rely on strong functional form assumptions. We revisit some of these assumptions in additional robustness exercises summarized in columns 2-5. Specifically, we consider alternative parameters of relative risk aversion, columns 2-3, as well as different assumptions on out-of-pocket spending as a fraction of overall medical utilization, columns 4-5. We find that the overall financial benefit of health insurance is remarkably robust across these different specifications.

⁴¹The variance in consumption equals $(1 - \bar{\tau})^2 \times var(\epsilon_{MB})$. Note that we have already netted out the role of charity care.

7.2 Direct Approach

Our second approach builds directly on our estimated interest savings from Section 6.5. Specifically, we define

$$CV = \Delta OOP + \Delta Interest \quad (14)$$

where ΔOOP is the change in out-of-pocket costs and $\Delta Interest$ is the value to consumers of better credit options. We view this approach as a conservative lower bound for the CV. This is because it ignores the benefits from increased access to credit, reduced hassle costs from dealing with debt collectors, and other various costs related to consumer bankruptcy. While this approach does not require additional assumptions on the utility function, it also provides a less comprehensive evaluation of the financial benefits from paid and unpaid medical bills. In particular, we cannot use this method to evaluate the consumer welfare gains from a variance reduction in medical bills.

We benchmark our calibration to our direct estimates presented above. The estimated indirect benefits from reduced costs of credit $\Delta Interest = \frac{\$10.78 + \$1.82 + \$1.20 + \$0.80}{0.044} = \frac{14.60}{0.044} = \332 per year (Column 2 of Table 7). Combined with the reduction in out-of-pocket spending, we calculate a compensating variation of \$812, which exceeds the out-of-pocket reduction by 69%.

When compared to the overall reduction in medical debt, the estimated credit channel (indirect) benefit is valued at $\frac{\$14.6}{\$13.5 \times 4} = \$0.27$ per dollar of reduced medical debt. Taking repayments of medical debt on the order of 8 percent into account (Table 2), we find a total financial benefit of a reduction in unpaid medical bills of about $\$0.27 + \$0.08 \approx \$0.35$ per dollar of reduced medical debt.

Unfortunately, our direct approach does not yield an estimate for the risk premium. Therefore, we borrow the corresponding estimates from the revealed preference approach to calculate an overall annual financial benefit of $\frac{\$1,492}{\$2,400} \approx 62\%$ of overall medical spending. This exceeds the out-of-pocket benchmark by a factor of 2.

7.3 Other Insurance Value

The above suggests that, absent changes in health care utilization, individuals may not be willing to buy Medicaid insurance, even when offered at a fair premium.⁴² This may be because of charity care and the option to not pay the medical bill, including the option to file bankruptcy, provide implicit health insurance. Dividing the CV by overall medical spending we find an effective price of only 50 cents per dollar, when considering the revealed

⁴²This finding is consistent with the results in (Finkelstein, Hendren and Luttmer, 2015) and (Finkelstein, Hendren and Shepard, 2017).

preference approach. This suggests that charity care and default options may insure about 50% of health care spending.

We revisit the role of charity care and medical debt in two thought experiments. In the first, we use the revealed preference approach to calculate the benefit-spending ratio in the absence of charity care, holding constant utilization and the proportion of the bill paid out-of-pocket. Specifically, we scale the net medical bill up to the raw bill amount and assume that 28 percent is paid out-of-pocket. Our model implies that $\frac{CV}{\epsilon_{MB}}$ now increases to 63%, or that the out-of-pocket spending would understate the CV by a factor of $\frac{63\%}{28\%} = 2.25$. Out-of-pocket spending increases to $0.28 \times \$2,400 = \672 suggesting a new CV of $2.25 \times \$672 = \$1,512$. The new risk premium equals \$1,046. This leads to a total benefit of $\$1,512 + \$1,046 = \$2,558$, which is now slightly larger than the annual cost of health care utilization.

In the second, we consider one possible mechanism for the net value of unpaid medical bills: the insurance value of bankruptcy protection (Mahoney, 2015). Medical debt can be discharged in bankruptcy proceedings which may explain why patients value a one dollar reduction in medical debt at only 20-28 cents. In a complementary analysis, we find that subprime borrowers discharge on average only \$860 in medical debt per bankruptcy filing. Considering an annual reduction of about 25,000 bankruptcies from the reform, this can account for only about $\$860 \times 25,000 = \21.5 million in medical debt or 1 percent of the overall reduction in medical debt. However, we note that the marginal filers, who were affected by the Medicaid expansion, may hold considerably more medical debt. If so, the \$21.5 million estimate provides a very conservative estimate of the potential insurance value of bankruptcy protection. In all, this suggests that charity care may be more important in explaining low valuations of health insurance than insurance through insolvency.

8 Discussion

Uninsured individuals pay on average only about 20 percent of overall health care utilization out-of-pocket. If the residual 80 percent of utilization is provided as charity care, then out-of-pocket payments provide a good estimate of the financial cost of their health care utilization. In this paper, we argue that unpaid medical bills in fact present a financial strain on consumers. This is because a large fraction of them are submitted to third-party collections and reported to credit bureaus, with dire effects on their credit market outcomes. As a result, substantial indirect financial benefits of health insurance accrue through protection against unpaid bills and resultant improvements in terms of credit. In the context of the Medicaid expansion under the Affordable Care Act, we estimate that insurance provision led to a \$5.89 billion decline in unpaid medical bills sent to collections, to higher credit scores, and to bet-

ter credit terms valued at \$670 million annually. Using our novel conceptual framework, we employ two distinct approaches to show that the financial benefits of Medicaid double when accounting for this indirect credit channel of health insurance.

Our conceptual framework reconciles estimated improvements in terms of credit with reductions in unpaid medical bills. However, it is conceivable that changes in out-of-pocket payments and related borrowing decisions, including credit card or payday loan borrowing (Allen et al., 2017), may also directly affect beneficiaries' future terms of credit. We respond to this in several ways. First, in an event study analysis we present direct evidence of a sudden, sharp, and persistent drop in credit scores after a medical collection is first placed on a record. We interpret this as a corroboration of our proposed mechanism. Second, we note that individuals make efforts to repay a non-trivial portion of these medical collections, at a significant personal expense. As a result, we infer there must be a cost to leaving medical bills unpaid. Otherwise, consumers would not pay any of the overall bill out-of-pocket. Our revealed preference approach formalizes this idea and finds even larger consumer welfare gains than the direct approach. This suggests that our calculated interest savings might understate the benefits from protection against unpaid medical bills. Consistent with this assessment, we see credit card debt declines following the placement of a medical collection on an individual's record. Such declines are in line with findings in Dobkin et al. (2016) and might counteract the negative effects of unpaid medical bills on the terms of credit.

Finally, we emphasize that implications of the credit channel for the value of Medicaid do not depend on whether the credit channel operates exclusively through protection from unpaid medical bills. Our overall assessment of indirect effects via improved credit terms can also be more broadly interpreted as a first attempt to propose and highlight the importance of a credit channel of health insurance.

References

- Abadie, Alberto, and Javier Gardeazabal. 2003. "The economic costs of conflict: A case study of the Basque Country." *American economic review*, 93(1): 113–132.
- Alexandrov, Alexei, Daniel Grodzicki, and Özlem Bedre-Defolie. 2018. "Consumer Demand for Credit Card Services." *CEPR Discussion Paper*, , (12506).
- Allen, Heidi, Ashley Swanson, Jialan Wang, and Tal Gross. 2017. "Early Medicaid Expansion Associated With Reduced Payday Borrowing In California." *Health Affairs*, 36(10): 1769–1776.

- Argys, Laura M, Andrew I Friedson, M Melinda Pitts, and D Sebastian Tello-Trillo.** 2017. “Losing Public Health Insurance: TennCare Disenrollment and Personal Financial Distress.”
- Bachrach, D, P Boozang, and M Lipson.** 2015. “The impact of Medicaid expansion on uncompensated care costs: Early results and policy implications for states.” *Princeton: Robert Wood Johnson Foundation*.
- Brevoort, Kenneth P., and Michelle Kambara.** 2015. “Are all collections equal? The Case of Medical Debt.” *The Journal of Credit Risk*, 11(4).
- Cabral, Marika, and Mark R Cullen.** 2016. “Estimating the Value of Public Insurance Using Complementary Private Insurance.”
- Caswell, Kyle J, and Timothy A Waidmann.** 2017. “The Affordable Care Act Medicaid Expansions and Personal Finance.” *Medical Care Research and Review*, 1077558717725164.
- Coughlin, Teresa A.** 2014. “Uncompensated care for the uninsured in 2013: A detailed examination.”
- Courtemanche, Charles, James Marton, Benjamin Ukert, Aaron Yelowitz, and Daniela Zapata.** 2016. “Impacts of the Affordable Care Act on health insurance coverage in Medicaid expansion and non-expansion states.” National Bureau of Economic Research.
- Dobkin, Carlos, Amy Finkelstein, Raymond Kluender, and Matthew J Notowidigdo.** 2016. “The Economic Consequences of Hospital Admissions.” *Working Paper*.
- Domowitz, Ian, and Robert L. Sartain.** 1999. “Determinants of the Consumer Bankruptcy Decision.” *Journal of Finance*, 54: 403–420.
- Dranove, David, Craig Garthwaite, and Christopher Ody.** 2016. “Uncompensated care decreased at hospitals in Medicaid expansion states but not at hospitals in nonexpansion states.” *Health Affairs*, 35(8): 1471–1479.
- Federal Trade Commission.** 2013. “The structure and practices of the debt buying industry.”
- Finkelstein, Amy.** 2007. “The Aggregate Effects of Health Insurance: Evidence from the Introduction of Medicare.” *The Quarterly Journal of Economics*, 122(1): 1–37.

- Finkelstein, Amy, and Robin McKnight.** 2008. “What did Medicare do? The initial impact of Medicare on mortality and out of pocket medical spending.” *Journal of public economics*, 92(7): 1644–1668.
- Finkelstein, Amy, Nathaniel Hendren, and Erzo FP Luttmer.** 2015. “The value of Medicaid: Interpreting results from the Oregon Health Insurance Experiment.”
- Finkelstein, Amy, Nathaniel Hendren, and Mark Shepard.** 2017. “Subsidizing health insurance for low-income adults: Evidence from Massachusetts.” *NBER Working Paper*.
- Finkelstein, Amy, Sarah Taubman, Bill Wright, Mira Bernstein, Jonathan Gruber, Joseph P Newhouse, Heidi Allen, Katherine Baicker, et al.** 2012. “the Oregon Health Insurance Experiment: Evidence from the First Year.” *Quarterly Journal of Economics*, 127(3).
- Frean, Molly, Jonathan Gruber, and Benjamin D Sommers.** 2016. “Premium subsidies, the mandate, and Medicaid expansion: Coverage effects of the Affordable Care Act.” *NBER Working Paper*.
- Gallagher, Emily, Radhakrishnan Gopalan, and Michal Grinstein-Weiss.** 2018. “The effect of health insurance on home payment delinquency: Evidence from aca marketplace subsidies.”
- Garthwaite, Craig, Tal Gross, and Matthew J Notowidigdo.** 2015. “Hospitals as insurers of last resort.”
- Gross, Tal, and Matthew J Notowidigdo.** 2011. “Health insurance and the consumer bankruptcy decision: Evidence from expansions of Medicaid.” *Journal of Public Economics*, 95(7): 767–778.
- Hamel, Liz, Mira Norton, Karen Pollitz, Larry Levitt, Gary Claxton, and Mollyann Brodie.** 2016. “The burden of medical debt: Results from the Kaiser Family Foundation.” *New York Times medical bills survey*.
- Hu, Luoia, Robert Kaestner, Bhashkar Mazumder, Sarah Miller, and Ashley Wong.** 2016. “The effect of the Patient Protection and Affordable Care Act Medicaid expansions on financial well-being.” *NBER Working Paper*.
- Kaiser Family Foundation.** 2012a. “A Guide to the Supreme Court’s Decision on the ACA’s Medicaid Expansion.” <http://kff.org/health-reform/issue-brief/a-guide-to-the-supreme-courts-decision/>, Accessed: 2015-03-25.

- Kaiser Family Foundation.** 2012*b*. “How Is the Affordable Care Act Leading to Changes in Medicaid Today? State Adoption of Five New Options.” Kaiser Commission on Medicaid and the Uninsured.
- Kowalski, Amanda E.** 2015. “What do longitudinal data on millions of hospital visits tell us about the value of public health insurance as a safety net for the young and privately insured?” *NBER Working Paper*.
- Mahoney, Neale.** 2015. “Bankruptcy as implicit health insurance.” *The American Economic Review*, 105(2): 710–746.
- Mazumder, Bhashkar, and Sarah Miller.** 2016. “The effects of the Massachusetts health reform on household financial distress.” *American Economic Journal: Economics Policy*, 8(3): 284–313.
- Ponce, Alejandro, Enrique Seira, and Guillermo Zamarripa.** 2017. “Borrowing on the Wrong Credit Card? Evidence from Mexico.” *American Economic Review*, 107(4): 1335–61.
- Sojourner, Aaron, and Ezra Golberstein.** 2017. “The ACA’s Medicaid expansion and its effects on financial wellbeing.” *Working Paper*.
- Stango, Victor, and Jonathan Zinman.** 2015. “Borrowing high versus borrowing higher: price dispersion and shopping behavior in the US credit card market.” *The Review of Financial Studies*, 29(4): 979–1006.
- White, Michelle.** 2006. *Bankruptcy and consumer behavior: Theory and US evidence*. Vol. 7, ch.
- Zeckhauser, Richard.** 1970. “Medical Insurance: A case study of the tradeoff between risk spreading and appropriate incentives.” *Journal of Economic Theory*, 2(1): 10–26.

Appendix

A Empirical Appendix

In this appendix we detail the identification strategy underlying our difference-in-difference (DD) framework (Section 4.3). We provide evidence of the validity of our DD approach in two ways. In the first, we present data trends comparing states adopting the policy on January 1st 2014, the initial adopters, to those states that chose not to expand Medicaid, the control. In the second, we illustrate the impact of the reform in a regression context, tracing quarter-year trends in our DD estimates. The appendix closely follows the structure of the main paper, whereby we first show evidence for the *direct effect* of the reform on medical debt followed by the *indirect effect* of the reform on financial distress and credit terms.

A.1 Direct Effect

We begin by presenting graphical evidence in Figure A.1. The left panel of the figure plots data trends in newly-acrued medical collections for initial adoption states and the synthetic control states, respectively. As illustrated in the figure, trends follow relatively well across

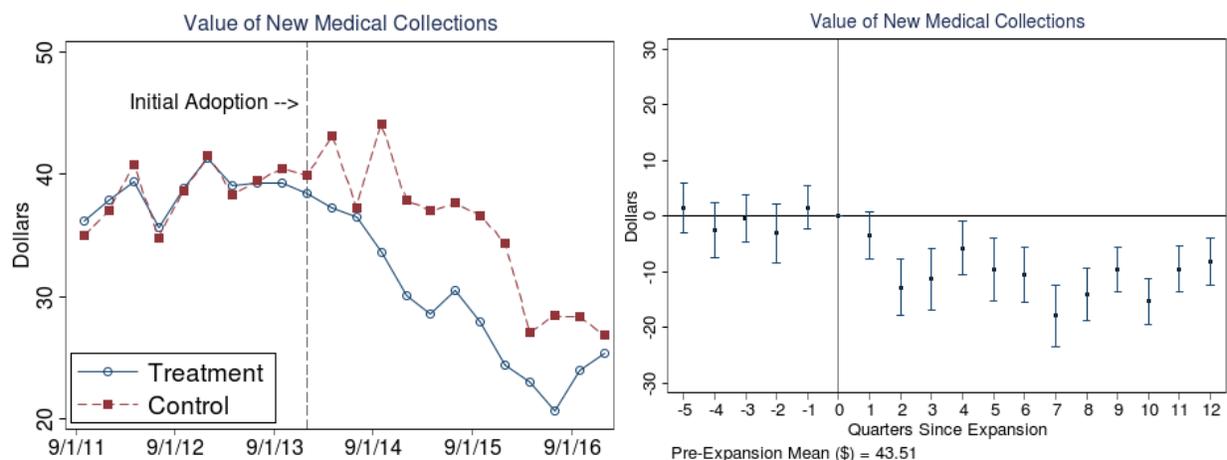


Figure A.1: Trends in Newly Accrued Medical Collections

Notes: The left panel of figure shows trends in the value of newly accrued non-medical collections. The right panel shows regression coefficients from Equation A.1. Data are from the CFPB’s Consumer Credit Panel described in section 4. Trends are quarterly means of newly accrued non-medical collections for initial adopters, treatment, and non adopters, synthetic control, respectively. Trends are weighted by the synthetic control weights. Vertical lines highlight the initial implementation date of the expansion - January 1st, 2014. All regression are weighted by the synthetic control weights, and include tract and quarter-year fixed effects. Standard errors are clustered by tract.

treatment and synthetic control prior to the adoption date. With the exception of a modest early decline in the value of newly accrued collections among treatment states one quarter

prior to the expansion date, there seems to be little systematic difference across these pre-treatment time series. Nevertheless, two-years after the reform, the average value of newly accrued medical debt declined by about 25 percent in adopting states relative to control.

In the right panel of Figure A.1, we revisit the graphical evidence in a regression framework with the following specification:

$$y_{ct}^k = \alpha_c^k + \eta_t^k + \sum_{r=S}^{r=-1} \beta_r^k + \sum_{r=1}^{r=F} \beta_r^k + \epsilon_{ct}^k . \quad (\text{A.1})$$

Here, y_{ict}^k denotes the respective outcome k for Census tract c during year-quarter t . We also include tract fixed effects α_c^k and quarter-year fixed effects η_t^k . The key parameters of interest are the β^r , which are indicators for time, in quarters, relative to the expansion date. Outcomes are normalized to the end of the quarter just prior to expansion. Our analysis extends from 5 quarters before to 12 quarters following an expansion.

As shown in the figure, we find no systematic differences, or pre-trends, between treatment and control states prior to the reform. Congruent with our graphical evidence, the figure also illustrates a sharp decline of medical collections in treatments following the reform. In the two years after the expansion, medical collections in treatment states decline by about 20 percent relative to those in control states.

We corroborate these findings in two robustness checks. First, we check that our findings are not driven by differential openings of private market insurance exchanges in treatment states. Other factors governing medical debt may be associated with the opening of the exchanges and, specifically, platform choice among states. To account for these factors, we subset our sample to include only states, treatment and control, that adopted the federal platform. In other words, for these states, all individuals using the exchanges did so on the same platform.⁴³

Second, we check that the findings are not driven by systematic changes in overall collections activities among adopting states coinciding with the Medicaid expansion. To the extent that reduction in accrued medical collections is driven by higher medical insurance rates, trends in non-medical collections should not differ to greatly in treatment states relative to control following the reform.

Figure A.2 plots trends in medical collections for states with only Federal exchanges (left panel) and for medical collections for all states in our sample (right panel). As shown in the left panel of the figure figure, this restricting to states operating federal platforms does

⁴³These states are: Alabama, Arizona, Florida, Georgia, Hawaii, Kansas, Maine, Mississippi, Missouri, Nebraska, North Carolina, North Dakota, Ohio, Oklahoma, South Carolina, South Dakota, Tennessee, Texas, Utah, Virginia, Wisconsin, Wyoming

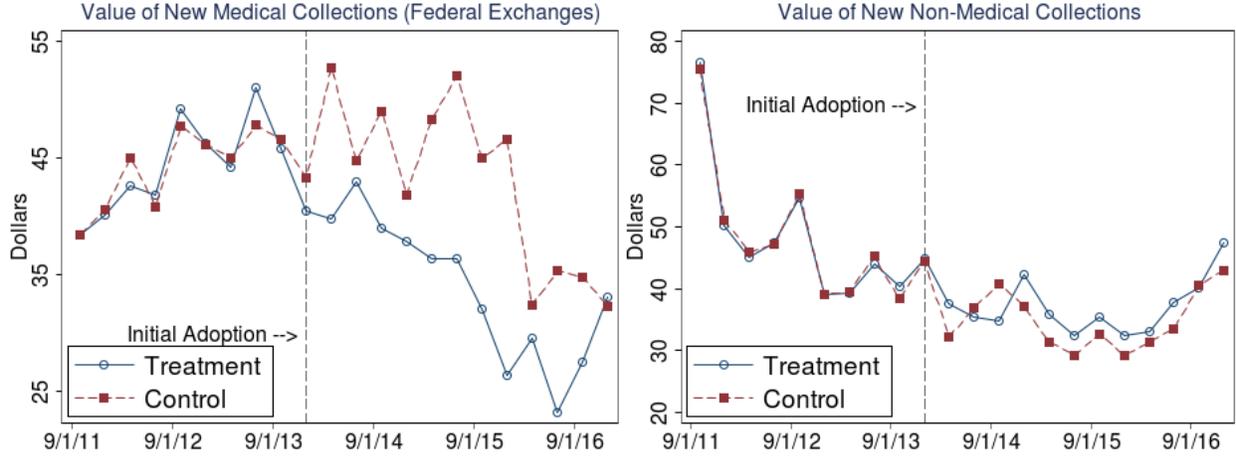


Figure A.2: Federal Exchanges & Non-Medical Collections

Notes: The figure shows new medical collections for states operating a federal platform for private exchanges (left panel) and in non-medical collections for all states in our sample (right panel). Data are from the CFPB’s Consumer Credit Panel described in section 4. The figure displays quarterly means of newly accrued non-medical collections for treatment and synthetic control states, respectively, and are weighted using the synthetic control weights. Vertical lines highlight the initial implementation date of the expansion - January 1st, 2014.

not materially alter the results. Although some noise is now more visible in the trends due to the significantly smaller sample, we note that the accrual of medical collections declines dramatically in propensity and volume within this subset of initial adopters. Moreover, the magnitudes are quite similar when considered alongside the full sample.

Further, we note no evidence of dramatic differences in trends of non-medical collections for adopting states following the initial roll out of the reform, right panel of Figure A.2. We conclude that there was likely no systematic change in overall collections activity driving the reduction in medical debt accruals. Rather, reductions in unpaid medical bills sent to collections are a result of newly-insured households not generating newly-unpaid medical bills following adverse health events.

A.2 Indirect Effect

A.2.1 Financial Distress: Delinquencies and Credit Scores

Next we turn to measures of financial distress, our *indirect effect*. Figure A.3 shows trends in the incidence of new delinquencies, or worsening credit, for Census tracts in initial adopting states (treatment) as compared to those in non adopting states (synthetic control). The left and the middle panel show trends for the 30 day and the 90 day delinquency rate, respectively. The right graphs shows trends in the credit score. While the trends for both

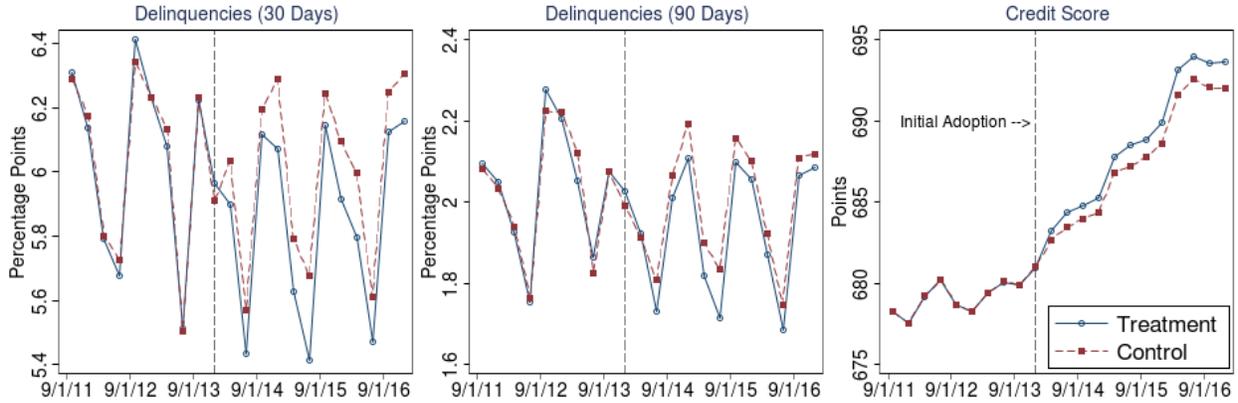


Figure A.3: Trends in Incidence of New Delinquency for Initial Adopting States

Notes: The figure shows quarterly flows into new delinquency, as defined in Section 4, and credit scores using the CFPB’s CCP. The left and the middle panel show trends for the 30 day and the 90 day delinquency rate. The right panel shows trends in credit scores. Delinquencies and credit scores are weighted using the synthetic control weights. Vertical lines highlight the implementation date of the expansion - January 1st, 2014.

groups are similar during the pre-expansion period, delinquency rates trend notably lower after the expansion in states that expanded Medicaid (e.g. treatment states). Similarly, credit scores trend in parallel prior to the expansion. Following the expansion, however, credit scores trend higher in the expansion in states.

Figure A.4 shows the *indirect effects* of the reform on delinquencies (left panel) and credit scores (right panel) using the regression framework in Equation A.1. The figure reiterates previous graphical evidence showing a decline in both measures of distress in the three years following the reform. Moreover, also congruent with the above graphical evidence, as with the regression results in the main paper (Tables 3 and 4), these effects grow between the first and second year after expansion and are seemingly long lasting.

A.2.2 Credit Supply: Pricing and Availability

Finally, we repeat the above analyses focusing on the reform’s impact on the pricing of credit to consumers using the data from *Mintel Comperemedia*. Figure A.5 shows trends in average rates of offered credit cards (left panel) and unsecured personal loans (right panel). Consistent with our findings on delinquency rates and credit scores, we see a relative decline in average credit card interest rates of around 0.5-0.8 percentage points in treatment states in the second and third post-expansion year.

Unlike credit cards, personal loans form part of a smaller and nascent market which largely focuses on highly indebted subprime customers. As a result, the incidence of personal

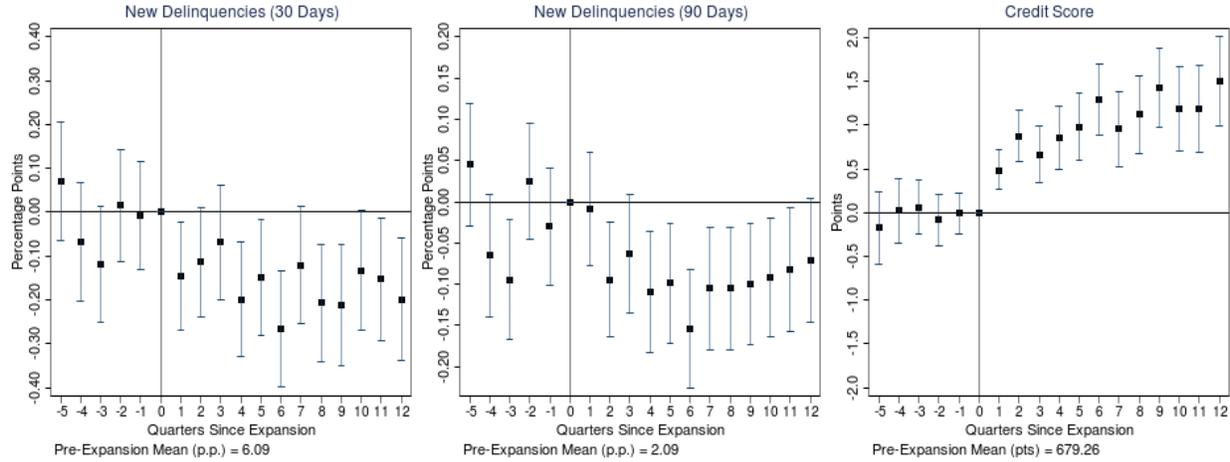


Figure A.4: Indirect Effect of Expansion on Financial Distress: Event Study

Notes: The figure shows changes in new delinquencies and credit scores using Equation A.1, weighted by the synthetic control weights. Data are from the CFPB’s Consumer Credit Panel. Confidence intervals in the figure are calculated using standard errors clustered at the Census tract level.

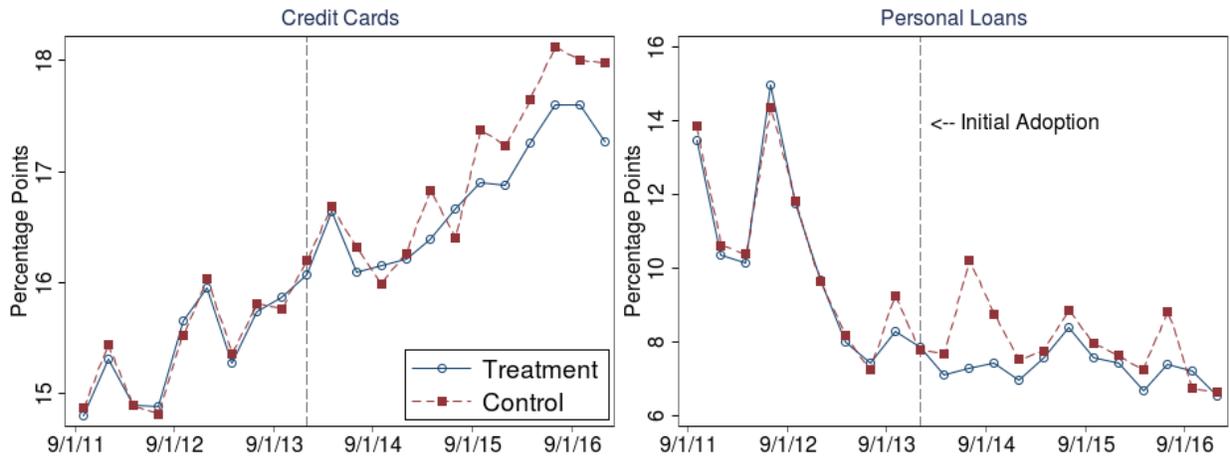


Figure A.5: Trends in Offered Rates for Initial Adopting States

Notes: The figure shows average interest rates for credit cards (left panel) and unsecured personal loans (right panel). The figure shows quarterly averages of rates offered to screened consumers weighted by the synthetic control weights. See Figure A.1 for details.

loan offers in the data is much lower than for credit cards (See Table 5). This smaller sample size on offers leads to noisier trends. Nevertheless, as shown in the right panel of Figure A.5, offered rates on personal loans seem to decline for recipients in expanding states relative to non-expanding states following the reform.

Figure A.6 further shows this effect within the regression framework, Equation A.1. In line with the main specification in the paper, we control for balances in each respective

product to cut out variation driven by differences in demand for borrowing. Moreover, the level of observation for regressions in this table is the individual-month. We find similar

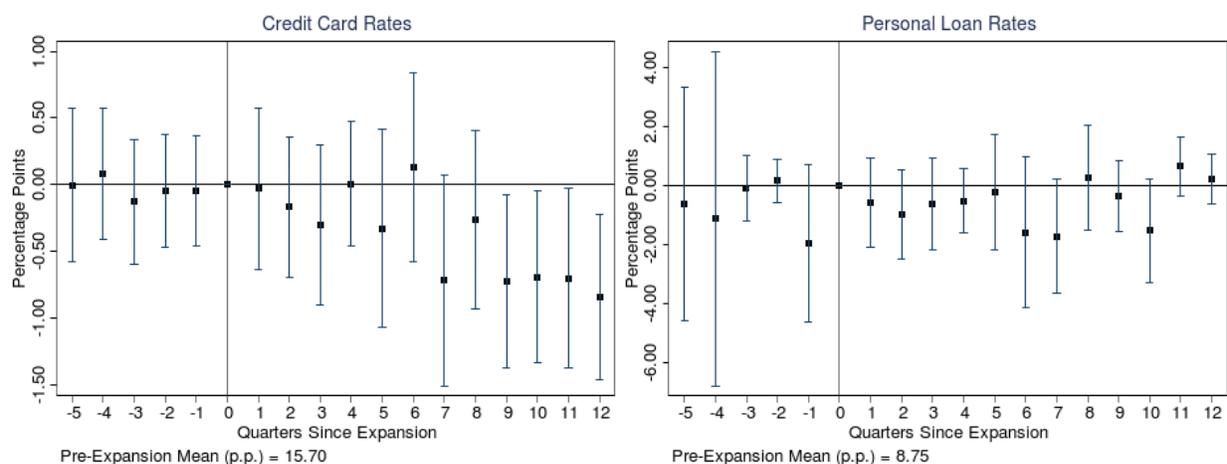


Figure A.6: Indirect Effect of Expansion on Credit Pricing: Event Study

Notes: The figure shows changes in credit card and personal loan rates using Equation A.1. Data are from Mintel Comperemedia. All regressions include county and quarter-year fixed effects and are weighted using the synthetic control weights. Confidence intervals in the figure are calculated using standard errors clustered at the county level.

results using the event study analysis. For credit cards, we see negative point estimates in all post-reform quarters (except for quarter 6). However, it takes 8 quarters until the point estimates become statistically significant. The evidence on personal loan interest rates is qualitatively similar. However, at least in parts because personal loans are offered less frequently, the point estimates are (individually) not statistically significant. Another explanation for differences between credit card and personal.⁴⁴

Overall, like for the case in the direct effect, we believe that collectively these results corroborate the notion that improved financial health and resulting improvements in credit availability and pricing observed since 2014 were driven by the Medicaid expansion.

⁴⁴A careful inspection of the point estimates suggests subtle differences between credit card rates and rates for personal loans. While we find that credit card rates decline even further in the third year following the reform, personal loan rates decline initially but bounce back towards the end of the third post-reform year. These differences likely arise from the fact that personal loans and credit cards are substitutes. Early on, improvements in financial health may have driven lenders to offer products to consumers for the purpose of debt consolidation, whereas sustained improvements in credit worthiness and pay down likely prompted a delayed decline in the offer of better credit card rates to consumers following the reform.

B Details on Synthetic Control Method

In this section, we provide further details on the synthetic control methodology. We construct different control groups for different dependent variables. For a given outcome measure, we first aggregate the data to the state-quarter level and then choose state weights that minimize the difference in the pre-reform outcomes between initial adoption states and non-adopters, following [Abadie and Gardeazabal \(2003\)](#). Table B.1 summarizes the weights for the control states by outcome variable.

Finally, we scale the state weights by the relative census tract population weights in each state to map the state weights into census tract weights.

Table B.1: Synthetic Control Weights by State

	Medical Collections	Non-Medical Collections	30-Day Delinq.	90-Day Delinq.	FICO	Credit Card	Personal Loans	Auto Loans	Mortgages
AL	0	0	0	0	0.046	0	0	0.01	0
FL	0	0	0	0	0	0.438	0.087	0	0
GA	0.148	0.152	0.321	0.222	0	0	0	0.087	0.013
ID	0	0.04	0.072	0	0.138	0	0	0.226	0
KS	0.018	0	0	0.028	0.193	0.012	0	0.021	0
ME	0.133	0.492	0	0.077	0.017	0	0	0	0
MO	0	0	0	0.065	0.104	0.117	0	0	0.194
MS	0.061	0	0	0	0.037	0	0.025	0	0
NC	0	0	0.129	0.107	0.026	0.066	0	0.103	0.298
NE	0	0	0.171	0.023	0	0	0	0.166	0.202
OK	0	0.079	0	0	0.116	0	0	0	0.099
SC	0	0	0	0	0	0.1	0	0	0
SD	0.439	0.017	0.035	0	0.049	0	0	0.034	0.068
TN	0	0.22	0	0.119	0	0	0	0.114	0.034
TX	0	0	0	0	0	0.033	0.33	0	0
UT	0.007	0	0	0	0	0	0	0	0.041
VA	0.123	0	0.224	0.264	0	0.049	0.185	0.238	0
WI	0	0	0.048	0.095	0.274	0.139	0.373	0	0
WY	0.072	0	0	0	0	0.046	0	0	0.052

Notes: The table shows the synthetic control weights by outcome variable.

C Collections, Debts, & Distress: An Event Study

In this section we discuss the direct relationship between medical collections and financial distress. In doing so we provide further motivation and detail on the analysis in Section 6. We investigate this link using a non-parametric event study. Like in Appendix A, our approach closely follows the event study methodology in [Dobkin et al. \(2016\)](#) which tracks how individuals' financial outcomes fare following a hospital admission. As we do not observe hospitalization, we replace the event of hospital admission with reporting of a new medical collection ($> \$100$).

There are several differences between a hospital admission and a medical collection. For example, new collections are generally not reported for up to 180 days following services rendered. Moreover, not all hospital admissions result in patients having their unpaid medical

bills sent to collections (Section 3). However there are also similarities, especially when considering uninsured individuals. Consequently, in addition to illustrating the relationship between collections and distress, we benchmark our event study results to those in [Dobkin et al. \(2016\)](#).

We subset our sample to include those whose first collection valued at more than \$100, or those most likely to result from hospital admissions and/or doctor’s visit. We then follow each of these individuals from six quarters prior to receiving the collection and for eight quarters, or two years, following the event. Our specification is as follows:

$$y_{ict}^k = \alpha_c^k + \eta_t^k + \sum_{r=S}^{r=-1} \beta_r^k + \sum_{r=1}^{r=F} \beta_r^k + \epsilon_{ict}^k \quad (\text{C.1})$$

where y_{ict}^k denotes the respective outcome k for record i in Census tract c during year-quarter t , such as delinquency. We also include Census tract fixed effects α_c^k and quarter-year fixed effects η_t^k . The key parameters of interest are the β^r , which are indicators for time, in quarters, relative to the quarter prior to having a collection placed on the record. Outcomes are normalized to the end of the quarter just prior to a collection being placed on the account. Standard errors are clustered by Census tract.

Figure C.1 plots the raw β_r^k s and their respective 95 percent confidence intervals. The figure plots these for medical collection balances (left panel), serious delinquencies (middle panel), and credit scores (left panel) separately for individuals with base credit score < 620 and > 620 , or subprime and prime borrowers, respectively. As shown in the figure, following a new collections, and by construction, individuals’ collections balances increase substantially. However, not by construction, the rise in medical debt is long lasting. High level of medical collections balances remains on individuals’ accounts for at least 2 years after the first one is reported. For prime individuals, beginning in the second quarter after the event, the medical debt level remains stable, or is paid off slightly. In contrast, subprime individuals’ medical debt balances continue to rise years after the initial event.

The middle panel of Figure C.1 shows that movement into serious delinquency rises substantially following a new medical collection (middle panel). Moreover, this effect is long lasting. However, in contrast to medical debt balances, there is a stronger surge in delinquency for prime borrowers. This is likely because prime borrowers’ base levels of delinquency are low to begin with, whereas subprime borrowers are likely troubled by delinquencies prior to receiving a new medical collection. It follows that a new medical collection also dramatically reduces borrowers’ credit scores (right panel), and that this effect is much greater among prime borrowers. As is shown in the figure, credit scores begin to fall prior to the collection,

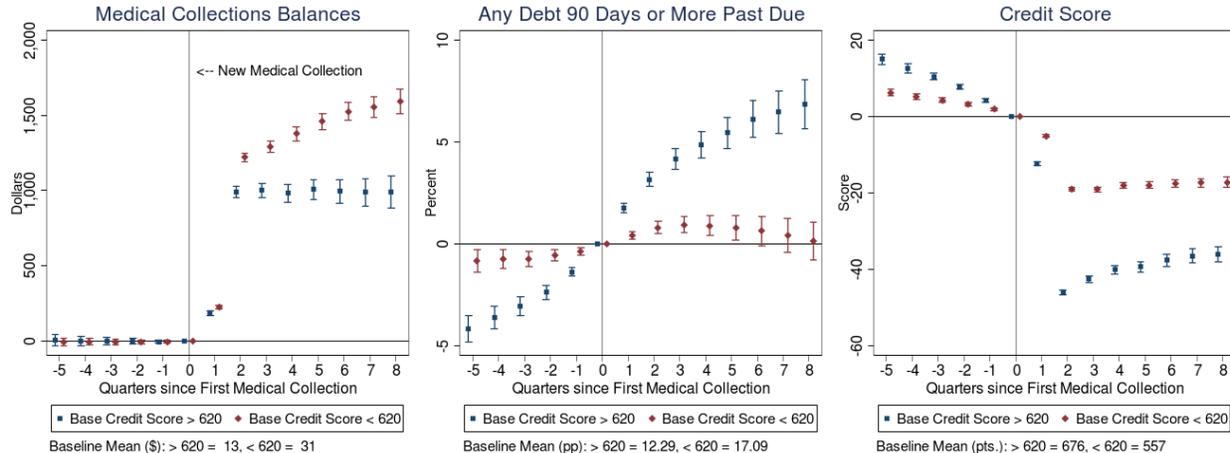


Figure C.1: Event Study: Credit Worthiness (By Risk)

Notes: The figure shows how 'healthy' individuals who receive a medical collection fair in the eight quarters (2 years) following the event. It does so along three dimensions: (1) Overall medical collections balances (left panel) (2) serious (90 day or more) delinquency (middle panel) (3) credit score (right panel). Serious delinquency is defined as the individual ever having become delinquent on a non-medical credit line, or debt, by that quarter. Data are from the CFPB's Consumer Credit Panel. The figure includes only individuals whose first collection value exceeds \$100. Effects are as of the end of each quarter and are normalized to the end of the quarter just prior to the first collection an individual receives on their record (the event). All regressions (Equation C.1) include Census tract and year-quarter fixed effects. Confidence intervals in the figure are calculated using standard errors clustered at the Census tract level.

likely because the actual health event, and distress resulting from it, begin some time before a medical collection is placed on individuals' records. However, there is a substantial drop, especially for prime individuals, just after the first collection is reported. This drop persists, despite a modest rebound, for years following the initial event. This final result is evidence of a direct and mechanical link between a medical collection and credit scores, which are designed to predict future delinquencies.

D Calculations of Simulated Decline in Monthly Bills

As described in Section 6.5, we use (screened) offer data for credit cards and personal loans from Mintel Compmedia and rate sheet pricing data for auto loans and mortgages from MyFico to estimate how the interest rates available to consumers were affected by the Medicaid expansion. In this section, we detail how we convert those interest rate changes into the savings in interest rate expenses that were available to consumers via a simulated refinancing.

First, note that a borrower i residing in Census tract c paying a monthly interest rate r_c (e.g. $\frac{APR}{12}$) with current balance $B_{i,0}$ and amortization period m (e.g. 12, 24 or 36 months)

faces a monthly payment of

$$P_{i,c}(m, r_c, B_{i,0}) = B_{i,0} \cdot \frac{r_c \cdot (1 + r_c)^m}{(1 + r_c)^m - 1}. \quad (\text{D.1})$$

As aforementioned (Section 6.5), our exercise simulates a debt refinancing of an individual’s balances as of the end of 2013Q4, just prior to the expansion. It follows that for each borrower we take B_0 to be their outstanding debt of that loan type as of that date. Moreover, in our calculations we assume fixed-payment loans with fixed interest and loan terms of 5-years for auto loans, 30-years for mortgages, and 3-years for credit cards and personal loans.⁴⁵ Because credit cards are revolving debt, they generally do not have fixed repayment terms or fixed payments. We use 3 years as an admittedly arbitrary estimate of how long it would take consumers to pay off their existing balances. Our results do not vary much if we reduce the payoff period to 1 year.

For unsecured loans, the scheduled monthly payments for a loan can overstate the expected cost to borrowers since some borrowers will fail to repay. A borrower who fails to repay an auto loan or mortgage loses the car or house backing the loan and is deprived of the flow of transportation and housing services those products provide. As a result, any money saved by not making payments will be at least partially offset by the loss of collateral. In contrast, unsecured borrowers do not surrender collateral when they default and are unlikely to face any directly offsetting expenses (though they do incur the costs of dealing with debt collectors and may have to pay higher costs for credit in the future).⁴⁶ For these borrowers, the stream of scheduled monthly payments likely overstates the cost of the loan. We therefore calculate an expected repayment amount for these loans as

$$\bar{P}_{i,c}(m, r_c, B_{i,0}, d_c) = (1 - d_c) \cdot P_c(m, r_c, B_{i,0}) + d_c \cdot 0 = (1 - d_c) \cdot P_c(m, r_c, B_{i,0}) \quad (\text{D.2})$$

where d_c is the monthly default rate in tract c . We measure default d_c as the likelihood of having a new 90-day delinquency or worse during a month for a respective debt type (e.g. credit card or personal loan). Following 90 day delinquencies, the probability of ever repaying a loan is nearly zero. Borrowers who become 30 days or more delinquent are much more likely to return to repayment. We then estimate the effects of the policy on default rates for consumers in each debt category separately using our baseline specification (Equation 8) in

⁴⁵Specifically, mortgage rates are for a 30-year, fixed rate mortgage of \$150,000 on a single-family owner-occupied property with a loan-to-value ratio of 80% and 1 point in origination fees. Auto rates are for a 60-month loan of between \$10,000 and \$20,000 for a new automobile.

⁴⁶While lenders can seek wage garnishments or other ways of compelling payment from unsecured borrowers, these options are not commonly pursued.

which the dependent variable is y_{ct}^k is the proportion of new delinquencies in Census tract c with $k \in \{\text{Credit Card, Personal Loan}\}$.

These estimates are shown in Table D.1. To conform with the analysis in Section 6.4, the regression is run separately for Census tracts below the median in the proportion of newly eligible adults (column 1) from those above the median (column 2). Since our specification

Table D.1: 90 Day Delinquency For Credit Cards and Personal Loans

	Below Median (1)	Above Median (2)
<i>Credit Cards</i>		
DD Coefficient	0.00045 (0.0002)	-0.00042 (0.00022)
90 Day Delinquency Rates	0.01101	0.01189
<i>Unsecured Personal Loans</i>		
DD Coefficient	-0.00002 (0.00005)	-0.00012 (0.00008)
90 Day Delinquency Rates	0.00729	0.01532

Notes: This table shows effects of the Expansion on new 90 day or more delinquencies for credit cards and personal loans. Each regression is estimated using Equation 8. See Section 4.3 for details. Standard errors (in parentheses) are clustered by tract.

provides estimates for quarterly flows into delinquency (q_c), we approximate the monthly default rate as $d_c \approx \frac{q_c}{3}$. Note that under the assumption of independence in delinquency over time we have $\frac{q_c}{3} = \hat{m}(1 - \hat{m})^2$ whereby $\hat{m} < m$ so our simplification in fact modestly understates net savings.

Next, we must define *baseline* ($\bar{r}_{baseline,c}^k$) and *refinanced* ($\bar{r}_{refinanced,c}^k$) rates for each of the four loan categories. One complication in determining baseline rates is that borrowers' individual interest rates on loans are not observed. As a result, we calculate an expected baseline rate by assigning to borrowers residing in Census tract c the average rate in their respective tract. For credit cards and personal loans we used the average of offered rates from Mintel, see Table 5. For auto loans and mortgages, we use aggregated rate sheets from *MyFico*, shown in Table D.2, to assign individuals in the CCP rates they qualify for given their credit score in quarter t .⁴⁷

⁴⁷D.2 shows the *MyFico* aggregated rate sheets for 5-year auto loans and 30-year fixed rate mortgages as of March 19, 2017. Consumers with credit scores below the bottom price tiers, or without a credit score are excluded from calculations, as they are not eligible for a loan.

Table D.2: Rate Sheets for Auto Loans and Mortgages

<i>Auto Loan Pricing Tiers</i>						
Credit Score Bin	500-589	590-619	620-659	660-689	690-720	>720
Auto Loan APR	15.117	13.970	9.653	6.948	4.863	3.514
<i>Mortgages Pricing Tiers</i>						
Credit Score Bin	620-639	640-659	660-679	680-699	700-759	>760
Mortgage APR	5.484	4.938	4.508	4.294	4.117	3.895

Notes: This table shows rate sheets for Mortgages and Auto Loans from the Fair Isaac Corporation's (FICO) MyFico web page (<http://www.myfico.com/credit-education/calculators/loan-savings-calculator>)

Refinanced rates ($r_{refinanced,c}^k$) are determined as the counterfactual rates implied by effects of the reform, e.g. the sum of baseline rates and the difference-in-difference estimates. As we allow for the reform to differentially impact communities with low vs. high proportion of newly eligible adults, we calculate different refinanced rates for each type of Census tract.

For credit cards and unsecured personal loans we use our estimates from Table 5. For auto loans and mortgages we estimate these effects using our imputed prices. Table D.3 shows the impact of the reform on imputed annual percentage rates (APR) for auto loans and mortgages. The table illustrates a modest decline in these annual rates as a result of the

Table D.3: Effects of Reform on Auto Loan and Mortgage Rates

	Auto Loans			Mortgages		
	All	Below	Above	All	Below	Above
	(1)	Median (2)	Median (3)	(4)	Median (5)	Median (6)
<i>1st Year ($Exp \times Post_1$)</i>	-0.0240 (0.0068)	-0.0246 (0.0087)	-0.0233 (0.0107)	0.0001 (0.0010)	-0.0029 (0.0012)	0.0034 (0.0016)
<i>2nd Year ($Exp \times Post_2$)</i>	-0.0424 (0.0084)	-0.0389 (0.0107)	-0.0466 (0.0132)	-0.0003 (0.0011)	-0.0007 (0.0014)	0.0001 (0.0019)
<i>3rd Year ($Exp \times Post_3$)</i>	-0.0537 (0.0095)	-0.0394 (0.0121)	-0.0710 (0.0150)	-0.0033 (0.0012)	-0.0037 (0.0015)	-0.0029 (0.0020)
<i>Average</i>	-0.0401 (0.0072)	-0.0343 (0.0092)	-0.0470 (0.0113)	-0.0012 (0.0009)	-0.0024 (0.0011)	0.0002 (0.0015)
<i>Pre-Expansion Average Rate</i>	7.4854	7.0007	8.6721	4.2091	4.1872	4.2742
<i>Percent Newly Eligible Adults</i>	0-100	0-19	19-100	0-100	0-19	19-100
<i>Observations</i>	729,898	405,283	324,615	705,398	392,629	312,769

Notes: This table shows effects of the Expansion on imputed auto loan and mortgage rates. Each regression is estimated using Equation 8. See Section 4.3 for details. All regression include Census tract and quarter-year fixed effects and are weighted by the population in a tract, scaled by the synthetic control weight. Standard errors (in parentheses) are clustered by tract.

reform. Moreover, consistent with previous evidence, these effects are larger in communities with a higher proportion of eligible adults. Note all rates are annual. As a result, we divide our estimates by 12 to transform APRs into monthly rates, which approximates monthly compounding. As aforementioned, delinquencies are directly observed in the CCP. Consequently, we set as the baseline the delinquency rate in each tract

$$\bar{d}_{baseline,c}^k = \frac{1}{3} \cdot \bar{d}_{baseline,c}^k \quad (D.3)$$

for $k \in \{CC, PL\}$.

More specifically, to determine *refinanced* rates and delinquencies, we predict counterfactuals of each using the difference in difference estimates (Table D.1, Table 5 and Table D.3) as follows

$$\bar{r}_{refinanced,c}^k = \bar{r}_{baseline,c}^k + \beta_c^k \quad (D.4)$$

for $k \in \{CC, PL, AUT, MTG\}$. Again, β_c^k is the key difference-in-difference coefficient from Equation 8. We then divide both baseline and refinanced APR by 12 to transform our estimated reduction into a monthly interest rate decline. Similarly for delinquencies, we calculate

$$\bar{d}_{refinanced,c}^k = \bar{d}_{baseline,c}^k + \frac{1}{3} \times \beta_c^k \quad (D.5)$$

for $k \in \{CC, PL\}$. Finally we define expected annual savings (*SV*) to be the sum of expected monthly savings as follows

$$SV_{i,c} = 12 \times [\bar{P}_{i,c}(m, \frac{\bar{r}_{baseline,c}}{12}, B_{i,0}, \bar{d}_{baseline,c}^k) - \bar{P}_{i,c}(m, \frac{\bar{r}_{refinanced,c}}{12}, B_{i,0}, \bar{d}_{refinanced,c}^k)] \quad (D.6)$$

for each of $\{CC, RP, AUT, MTG\}$.

In our simulations we calculate an average per-person annual savings. As aforementioned, these intent-to-treat effects on rate savings are generated using slightly different methods for the secured and unsecured loans. For our estimates on secured products, we use the entire sample. Our estimates for the unsecured products, however, were estimated conditional on receiving a credit offer. We have no information on the correlation between receiving an offer and Medicaid eligibility. Absent this information, we assume independence between these receiving an offer and Medicaid enrollment and treat our estimates as intent-to-treat similar to the case for secured loans.

There is another interpretation of this approach. Suppose there is non-zero correlation between Medicaid enrollment and the propensity to receive credit offers. Nevertheless, all individuals with improved credit scores still qualify for new loans at an equally lower rate, were they to seek them out. This interpretation assumes zero correlation between Medicaid

enrollment and eligibility for lower rates, which is a weaker and quite plausible condition. Finally, we simulate aggregate potential savings by multiplying our per person effects with the CCP Population in at the end of 2013, similar to the exercise presented in Table 6 in the main text.

E Details on the Consumer Welfare Analysis

E.1 Model Details

Before turning to the CV and the RP, we first show how we can use the first order condition and the implicit function theorem to express the disutility over medical debt in terms of preferences over consumption.

As indicated in the main text, we assume that consumers have existing medical debt \bar{D} and decide on the optimal amount of new medical bills $0 \leq b \leq \epsilon_{MB}$ that go unpaid, trading off utility from consumption and disutility from medical debt. Conditional on a realized medical bill, ϵ_{MB} , consumers maximize:

$$\max_{0 \leq b \leq \epsilon_{MB}} g(Y - (\epsilon_{MB} - b)) - h(\bar{D} + b)$$

where in optimality

$$F(\epsilon_{MB}, b) = g'(Y - (\epsilon_{MB} - b^*)) - h'(\bar{D} + b^*) = 0 . \quad (\text{E.1})$$

Applying the implicit function theorem, it follows that

$$\begin{aligned} \frac{\partial F(\epsilon_{MB}, b)}{\partial \epsilon_{MB}} \Delta \epsilon_{MB} + \frac{\partial F(\epsilon_{MB}, b)}{\partial b} \Delta b &= -g'' \Delta \epsilon_{MB} + [g'' - h''] \Delta b = 0 \\ \iff \frac{\Delta b}{\Delta \epsilon_{MB}} &= \frac{g''(Y - \epsilon_{MB} + b^*)}{g''(Y - \epsilon_{MB} + b^*) - h''(\bar{D} + b^*)} \in [0, 1] \quad (\text{E.2}) \end{aligned}$$

where we normalize $b^*(\epsilon_{MB} = 0) = 0$. It follows that a fraction $\tau(\epsilon_{MB}) \in [0, 1]$ of new medical bills remains unpaid and becomes medical debt with

$$b^* = \tau(\epsilon_{MB}) \cdot \epsilon_{MB} \Rightarrow \frac{\Delta b}{\Delta \epsilon_{MB}} = \tau'(\epsilon_{MB}) \epsilon_{MB} + \tau(\epsilon_{MB}). \quad (\text{E.3})$$

Equations E.1, E.2, and E.3 allow us to express (locally) the first and second derivative of $h(D)$ in terms of $g'(c)$, $g''(c)$, and $\tau(\epsilon_{MB})$. We return to this observation below.

E.2 Details on Compensating Variation

To gauge the transfer gain from insurance, in dollars, we quantify the compensating variation (CV). As outlined in the main text, we assume that the demand for medical care is price inelastic. Then, if consumers do not have the option to leave bills unpaid (e.g. borrow), we trivially have

$$CV = e(p_0, u_0) - e(p_1, u_0) = e(\epsilon_{MB}, u_0) - e(0, u_0) = Y - (Y - \epsilon_{MB}) = \epsilon_{MB}$$

where $e(\cdot)$ denotes the expenditure function. If consumers can leave bills unpaid, then we have to take the substitution patterns between consumption and unpaid bills into account. The compensating variation is implicitly defined by

$$\begin{aligned} u_0 &= g(Y - (1 - \tau(\epsilon_{MB}))\epsilon_{MB}) - h(\bar{D} + \tau(\epsilon_{MB})\epsilon_{MB}) \\ u_0 &= g(Y - dc) - h(\bar{D} - dd) \end{aligned} \quad (\text{E.4})$$

with

$$CV = dc - dd \geq [1 - \tau(\epsilon_{MB})] \cdot \epsilon_{MB}.$$

It follows that dc and dd correspond to the optimal reductions in consumption and unpaid bills (medical debt) if the income is reduced by CV. Under the assumption that consumers cannot take out medical debt to finance consumption, absent a new medical bill, we also have that $dd \geq 0$. The first order condition combined with, $g''(\cdot) < 0$, and $h''(\cdot) > 0$ imply that $g'(Y - dc) - h'(\bar{D}) > 0$ if $dc \geq (1 - \tau(\epsilon_{MB}))\epsilon_{MB}$. Therefore, individuals will not be willing reduce consumption in exchange for fewer unpaid bills. Hence, they optimally choose $dd = 0$, $dc = CV$. Consequently, we can rewrite Equation E.4 as

$$\int_{Y-CV}^{Y-(1-\tau(\epsilon_{MB}))\epsilon_{MB}} g'(x)dx = \int_{\bar{D}}^{\bar{D}+\tau(\epsilon_{MB})\epsilon_{MB}} h'(x)dx . \quad (\text{E.5})$$

In the context of Figure 1, $Y - CV$ corresponds to the point on the horizontal axis such that the corresponding area underneath MU_C bounded by $Y - CV$ from the left and $Y - (\epsilon_{MB} - b^*)$ from the right equals the blue area (I). It is evident from here that the CV is bounded from below by $(1 - \tau(\epsilon_{MB}))\epsilon_{MB}$ and by the entire bill ϵ_{MB} from above.⁴⁸

⁴⁸The lower bound is achieved if the right hand side of Equation E.5 equals zero. The upper bound is achieved if $-\int_{\bar{D}}^{\bar{D}+\tau(\epsilon_{MB})\epsilon_{MB}} h'(x)dx \geq \int_{Y-\epsilon_{MB}}^{Y-(1-\tau(\epsilon_{MB}))\epsilon_{MB}} g'(x)dx$.

E.3 Proposition 1

An advantage of the revealed preference approach is that we can also consider the comparative statics of the CV with respect to the underlying bill amount, the repayment rate, and the curvature in utility as stated in the following proposition

Proposition 1 *If $g'(\cdot) > 0$, $g''(\cdot) < 0$ and $h(\cdot) > 0$, $h''(\cdot) > 0$ and $b^* = \bar{\tau}\epsilon_{MB}$, then the linear approximation to the marginal utility function around b^* can be characterized as follows*

1. *The CV is given by:*

$$CV = -\phi(\cdot) + (1 - \bar{\tau})\epsilon_{MB} + \sqrt{\phi(\cdot)^2 + 2\bar{\tau}\phi(\cdot)\epsilon_{MB} - \bar{\tau}(1 - \bar{\tau})\epsilon_{MB}^2},$$

where $\phi(\cdot) = -\frac{g'(\cdot)}{g''(\cdot)}$ and $\cdot = Y - (1 - \bar{\tau})\epsilon_{MB}$ if $\epsilon_{MB} \leq \frac{\phi(\cdot)}{1 - \bar{\tau}}$.

2. *The CV is increasing in $\phi(\cdot)$*

3. *The CV is decreasing in $\bar{\tau}$ if $\frac{g'''(\cdot)g'(\cdot)}{g''(\cdot)^2} \leq 2$ and $\epsilon_{MB} < \min\{\frac{\phi(\cdot)}{\bar{\tau} + \frac{1}{8}}, 4\phi(\cdot)\}$*

4. *CV over ϵ_{MB} is decreasing in the medical bill amount if $\frac{g'''(\cdot)g'(\cdot)}{g''(\cdot)^2} \leq 1 + \frac{\phi(\cdot)}{1 - \bar{\tau}}$.*

The specific value of CV depends on the shape of both marginal utility functions. Unfortunately, it is difficult to calibrate $h'(\cdot)$ directly. However, we can combine the first order condition and the result from the implicit function theorem with observed out-of-pocket payments to approximate the marginal disutility of medical debt in terms of the marginal utility of consumption. We start with the case $h''(\cdot) > 0$ and turn to the case $h''(\cdot) = 0$ below. Specifically, we propose a local linear approximation of the marginal disutility of debt around the optimal borrowing decision:

$$\begin{aligned} h'(\bar{D} + x) &= h'(\bar{D} + \tau(\epsilon_{MB})\epsilon_{MB}) + h''(\bar{D} + \tau(\epsilon_{MB})\epsilon_{MB}) * (x - \tau(\epsilon_{MB})\epsilon_{MB}) \\ &= g'(Y - (1 - \tau(\epsilon_{MB}))\epsilon_{MB}) - \frac{1 - \tau'(\epsilon_{MB})\epsilon_{MB} - \tau(\epsilon_{MB})}{\tau'(\epsilon_{MB})\epsilon_{MB} + \tau(\epsilon_{MB})} \\ &\quad * g''(Y - (1 - \tau(\epsilon_{MB}))\epsilon_{MB}) * (x - \tau(\epsilon_{MB})\epsilon_{MB}), \end{aligned}$$

where the second equality uses the first order condition and the implicit function theorem. Similarly, using a local linear approximation around $g'(\cdot)$ and assuming that locally a constant

fraction of medical bills is unpaid $\tau(\epsilon_{MB}) = \bar{\tau}$, we can rewrite Equation E.5 as:

$$\begin{aligned}
& g'(Y - (1 - \bar{\tau})\epsilon_{MB}) \left[CV - (1 - \bar{\tau})\epsilon_{MB} \right] \\
+ & g''(Y - (1 - \bar{\tau})\epsilon_{MB}) \int_{Y-CV}^{Y-(1-\bar{\tau})\epsilon_{MB}} (x - (Y - (1 - \bar{\tau})\epsilon_{MB})) dx \\
= & g'(Y - (1 - \bar{\tau})\epsilon_{MB}) * \bar{\tau}\epsilon_{MB} - \frac{1 - \bar{\tau}}{\bar{\tau}} * g''(Y - (1 - \bar{\tau})\epsilon_{MB}) \int_{\bar{D}}^{\bar{D} + \bar{\tau}\epsilon_{MB}} (x - (\bar{D} + \bar{\tau}\epsilon_{MB})) dx .
\end{aligned}$$

Simplifying terms, we have

$$\begin{aligned}
& g'(Y - (1 - \bar{\tau})\epsilon_{MB}) \left[CV - (1 - \bar{\tau})\epsilon_{MB} \right] - g''(Y - (1 - \bar{\tau})\epsilon_{MB}) \int_0^{CV - (1 - \bar{\tau})\epsilon_{MB}} x dx \\
= & g'(Y - (1 - \bar{\tau})\epsilon_{MB}) * \bar{\tau}\epsilon_{MB} + \frac{1 - \bar{\tau}}{\bar{\tau}} * g''(Y - (1 - \bar{\tau})\epsilon_{MB}) \int_0^{\bar{\tau}\epsilon_{MB}} x dx .
\end{aligned}$$

and

$$\begin{aligned}
& g'(Y - (1 - \bar{\tau})\epsilon_{MB}) \left[CV - \epsilon_{MB} \right] - \frac{1}{2} g''(Y - (1 - \bar{\tau})\epsilon_{MB}) \left[CV - (1 - \bar{\tau})\epsilon_{MB} \right]^2 \\
= & \frac{1 - \bar{\tau}}{2 * \bar{\tau}} * g''(Y - (1 - \bar{\tau})\epsilon_{MB}) \left[\bar{\tau}\epsilon_{MB} \right]^2 .
\end{aligned}$$

Finally, we have

$$CV = \frac{\left[-g'(\cdot) - (1 - \bar{\tau})\epsilon_{MB}g''(\cdot) \right] + \sqrt{g'(\cdot)^2 - 2\bar{\tau}g'(\cdot)g''(\cdot)\epsilon_{MB} - \bar{\tau}g''(\cdot)^2\epsilon_{MB}^2(1 - \bar{\tau})}}{-g''(\cdot)} .$$

Let $\phi(\cdot) = -\frac{g'(\cdot)}{g''(\cdot)}$, then we have

$$CV = -\phi(\cdot) + (1 - \bar{\tau})\epsilon_{MB} + \sqrt{\phi(\cdot)^2 + 2\bar{\tau}\phi(\cdot)\epsilon_{MB} - \bar{\tau}(1 - \bar{\tau})\epsilon_{MB}^2} ,$$

which establishes the first part of the proposition.

Case $h''(\cdot) = 0$: Before we turn to the comparative statics, we establish that the CV discussed above provides a lower bound for the case $h''(\cdot) = 0$. Simplifying the former derivation we now have,

$$g'(Y - (1 - \bar{\tau})\epsilon_{MB}) \left[CV - \epsilon_{MB} \right] - \frac{1}{2} g''(Y - (1 - \bar{\tau})\epsilon_{MB}) \left[CV - (1 - \bar{\tau})\epsilon_{MB} \right]^2 = 0 .$$

This implies the following compensating variation:

$$\begin{aligned} CV^* &= -\phi(\cdot) + (1 - \bar{\tau})\epsilon_{MB} + \sqrt{\phi(\cdot)^2 + 2\bar{\tau}\phi(\cdot)\epsilon_{MB}} \\ &\geq -\phi(\cdot) + (1 - \bar{\tau})\epsilon_{MB} + \sqrt{\phi(\cdot)^2 + 2\bar{\tau}\phi(\cdot)\epsilon_{MB} - \bar{\tau}(1 - \bar{\tau})\epsilon_{MB}^2}, \end{aligned}$$

where the second row replicates the CV derived above.

Comparative statics: We now turn to the comparative statics. We first show that $\frac{dCV}{d\phi(\cdot)} > 0$. Taking the first derivative, we have

$$\frac{dCV}{d\phi(\cdot)} = -1 + \frac{\phi + \bar{\tau}\epsilon_{MB}}{\sqrt{\cdot}}.$$

Now we show that $[\phi + \bar{\tau}\epsilon_{MB}]^2 > (\sqrt{\cdot})^2$. So we have

$$\begin{aligned} [\phi + \bar{\tau}\epsilon_{MB}]^2 &> (\sqrt{\cdot})^2 \\ \phi(\cdot)^2 + 2\bar{\tau}\epsilon_{MB}\phi(\cdot) + \bar{\tau}^2\epsilon_{MB}^2 &> \phi(\cdot)^2 + 2\bar{\tau}\phi(\cdot)\epsilon_{MB} - \bar{\tau}(1 - \bar{\tau})\epsilon_{MB}^2 \\ \Leftrightarrow 0 &> -\bar{\tau}\epsilon_{MB}^2, \end{aligned}$$

which establishes the second part of the proposition.

Next we show that $\frac{dCV}{d\bar{\tau}} < 0$. Taking the first derivative, we have

$$\begin{aligned} \frac{dCV}{d\bar{\tau}} &= -\epsilon_{MB} + \frac{1}{2 * \sqrt{\cdot}} [2\phi(\cdot)\epsilon_{MB} - \epsilon_{MB}^2 + 2\bar{\tau}\epsilon_{MB}^2] - \frac{d\phi(\cdot)}{d\bar{\tau}} + \frac{1}{2 * \sqrt{\cdot}} [2\phi(\cdot)\frac{d\phi(\cdot)}{d\bar{\tau}} + 2\epsilon_{MB}\bar{\tau}\frac{d\phi(\cdot)}{d\bar{\tau}}] \\ &= -\epsilon_{MB} \underbrace{\left[1 - \frac{\sqrt{[\phi + (\bar{\tau} - \frac{1}{2})\epsilon_{MB}]^2}}{\sqrt{\phi(\cdot)^2 + 2\bar{\tau}\phi(\cdot)\epsilon_{MB} - \bar{\tau}(1 - \bar{\tau})\epsilon_{MB}^2}} \right]}_A \\ &\quad - \underbrace{\frac{d\phi(\cdot)}{d\bar{\tau}} \left[1 - \frac{\sqrt{(\phi(\cdot) + \bar{\tau} * \epsilon_{MB})^2}}{\sqrt{\phi(\cdot)^2 + 2\bar{\tau}\phi(\cdot)\epsilon_{MB} - \bar{\tau}(1 - \bar{\tau})\epsilon_{MB}^2}} \right]}_B \end{aligned}$$

First, we note that $\sqrt{\left[\phi + (\bar{\tau} - \frac{1}{2})\epsilon_{MB}\right]^2} < \sqrt{\phi(\cdot)^2 + 2\bar{\tau}\phi(\cdot)\epsilon_{MB} - \bar{\tau}(1 - \bar{\tau})\epsilon_{MB}^2}$, which implies that term A is greater than 0. Hence, we have

$$\begin{aligned} \left[\phi + (\bar{\tau} - \frac{1}{2})\epsilon_{MB}\right]^2 &< \left(\sqrt{\cdot}\right)^2 \\ \phi(\cdot)^2 + 2(\bar{\tau} - \frac{1}{2})\epsilon_{MB}\phi(\cdot) + (\bar{\tau} - \frac{1}{2})^2\epsilon_{MB}^2 &< \phi(\cdot)^2 + 2\bar{\tau}\phi(\cdot)\epsilon_{MB} - \bar{\tau}(1 - \bar{\tau})\epsilon_{MB}^2 \\ \Leftrightarrow -\phi(\cdot)\epsilon_{MB} + [\bar{\tau}^2 - \bar{\tau} + \frac{1}{4}]\epsilon_{MB}^2 &< [\bar{\tau}^2 - \bar{\tau}]\epsilon_{MB}^2 \\ \Leftrightarrow \frac{\epsilon_{MB}^2}{4} &< \phi(\cdot)\epsilon_{MB} \\ \Leftrightarrow \epsilon_{MB} &< 4\phi(\cdot) . \end{aligned}$$

which is true if $\epsilon_{MB} < \min\{\frac{\phi(\cdot)}{\bar{\tau} + \frac{1}{8}}, 4\phi(\cdot)\}$ as required in the proposition.

Second, we have that $\sqrt{(\phi(\cdot) + \bar{\tau} * \epsilon_{MB})^2} \geq \sqrt{\phi(\cdot)^2 + 2\bar{\tau}\phi(\cdot)\epsilon_{MB} - \bar{\tau}(1 - \bar{\tau})\epsilon_{MB}^2}$, which implies that $sign(B) = sign(-\frac{d\phi(\cdot)}{d\bar{\tau}})$. Here, we have

$$\frac{d\phi(\cdot)}{d\bar{\tau}} = -\frac{d\frac{g'(\cdot)}{g''(\cdot)}}{d\bar{\tau}} = -\frac{g''(\cdot)^2\epsilon_{MB} - \epsilon_{MB}g'''(\cdot)g'(\cdot)}{g''(\cdot)^2}.$$

If $\frac{g'''(\cdot)g'(\cdot)}{g''(\cdot)^2} \leq 2$ then $\frac{d\phi(\cdot)}{d\bar{\tau}} \leq \epsilon_{MB}$. Then we have

$$\begin{aligned} \frac{dCV}{d\bar{\tau}} &\geq -\epsilon_{MB} * \left[2 - \frac{\sqrt{\left[\phi + (\bar{\tau} - \frac{1}{2})\epsilon_{MB}\right]^2} + \sqrt{\left[\phi(\cdot) + \bar{\tau} * \epsilon_{MB}\right]^2}}{\sqrt{\phi(\cdot)^2 + 2\bar{\tau}\phi(\cdot)\epsilon_{MB} - \bar{\tau}(1 - \bar{\tau})\epsilon_{MB}^2}} \right] \\ &= -\epsilon_{MB} * \left[2 - \frac{\phi + (\bar{\tau} - \frac{1}{2})\epsilon_{MB} + \phi(\cdot) + \bar{\tau} * \epsilon_{MB}}{\sqrt{\phi(\cdot)^2 + 2\bar{\tau}\phi(\cdot)\epsilon_{MB} - \bar{\tau}(1 - \bar{\tau})\epsilon_{MB}^2}} \right] \\ &= -\epsilon_{MB} * \left[2 - 2\frac{\sqrt{(\phi(\cdot) + (\bar{\tau} - \frac{1}{4})\epsilon_{MB})^2}}{\sqrt{\phi(\cdot)^2 + 2\bar{\tau}\phi(\cdot)\epsilon_{MB} - \bar{\tau}(1 - \bar{\tau})\epsilon_{MB}^2}} \right]. \end{aligned}$$

Finally, we show that

$$\begin{aligned}
& (\phi(\cdot) + (\bar{\tau} - \frac{1}{4})\epsilon_{MB})^2 < \phi(\cdot)^2 + 2\bar{\tau}\phi(\cdot)\epsilon_{MB} - \bar{\tau}(1 - \bar{\tau})\epsilon_{MB}^2 \\
\leftrightarrow & \phi(\cdot)^2 + 2\phi(\cdot)(\bar{\tau} - \frac{1}{4})\epsilon_{MB} + (\bar{\tau} - \frac{1}{4})^2\epsilon_{MB}^2 < \phi(\cdot)^2 + 2\bar{\tau}\phi(\cdot)\epsilon_{MB} - \bar{\tau}(1 - \bar{\tau})\epsilon_{MB}^2 \\
\leftrightarrow & -\frac{1}{2}\phi(\cdot)\epsilon_{MB} + \frac{1}{2}\bar{\tau}\epsilon_{MB}^2 + \frac{1}{16}\epsilon_{MB}^2 < 0 \\
\leftrightarrow & \phi(\cdot) > (\bar{\tau} + \frac{1}{8})\epsilon_{MB} \\
\leftrightarrow & \epsilon_{MB} < \frac{\phi(\cdot)}{\bar{\tau} + \frac{1}{8}}
\end{aligned}$$

which is true if $\epsilon_{MB} < \min\{\frac{\phi(\cdot)}{\bar{\tau} + \frac{1}{8}}, 4\phi(\cdot)\}$ as required in the proposition. This establishes the third part of the proposition.

Finally, we turn to

$$\frac{CV}{\epsilon} = -\frac{\phi(\cdot)}{\epsilon_{MB}} + (1 - \bar{\tau}) + \sqrt{\frac{\phi(\cdot)^2}{\epsilon_{MB}^2} + 2\frac{\bar{\tau}\phi(\cdot)}{\epsilon_{MB}} - \bar{\tau}(1 - \bar{\tau})}$$

Here we have

$$\begin{aligned}
d\frac{CV}{\epsilon_{MB}}/d\epsilon_{MB} &= -\left[\frac{\frac{d\phi(\cdot)}{d\epsilon_{MB}} - \phi(\cdot)}{\epsilon_{MB}^2}\right] + 2 * \frac{\left[\frac{\phi(\cdot)}{\epsilon_{MB}} + \bar{\tau}\right]}{2 * \sqrt{\cdot}} \left[\frac{\frac{d\phi(\cdot)}{d\epsilon_{MB}} - \phi(\cdot)}{\epsilon_{MB}^2}\right] \\
&= -\left[\frac{\frac{d\phi(\cdot)}{d\epsilon_{MB}} - \phi(\cdot)}{\epsilon_{MB}^2}\right] * \left[1 - \frac{\sqrt{\left[\frac{\phi(\cdot)}{\epsilon_{MB}} + \bar{\tau}\right]^2}}{\sqrt{\cdot}}\right].
\end{aligned}$$

Since $\sqrt{\left[\frac{\phi(\cdot)}{\epsilon_{MB}} + \bar{\tau}\right]^2} \geq \sqrt{\cdot}$, the second factor is smaller than zero. Hence the sign of the effect equals the sign of $\left[\frac{\frac{d\phi(\cdot)}{d\epsilon_{MB}} - \phi(\cdot)}{\epsilon_{MB}^2}\right]$.

We have

$$\begin{aligned}
\frac{d\phi(\cdot)}{d\epsilon_{MB}} - \phi(\cdot) &= -(1 - \bar{\tau})\left[\frac{g''(\cdot)^2 - g'''(\cdot)g'(\cdot)}{g''(\cdot)^2}\right] - \phi(\cdot) \\
&< -(1 - \bar{\tau}) + \phi + (1 - \bar{\tau}) - \phi(\cdot) = 0,
\end{aligned}$$

where the second line uses $\frac{g'''(\cdot)g'(\cdot)}{g''(\cdot)^2} \leq 1 + \frac{\phi(\cdot)}{1 - \bar{\tau}}$. This establishes the last part of the proposition.

E.4 Details on Effects of Variance Reduction

The second order Taylor expansion of utility yields:

$$\begin{aligned}
U(\epsilon_{MB}, \epsilon_{\bar{MB}}) &= g(Y - (1 - \bar{\tau}) * \epsilon_{\bar{MB}}) - h(\bar{D} + \bar{\tau} * \epsilon_{\bar{MB}}) \\
&- \left[(1 - \bar{\tau}) * g'(Y - (1 - \bar{\tau}) * \epsilon_{\bar{MB}}) + \bar{\tau} * h'(\bar{D} + \bar{\tau} * \epsilon_{\bar{MB}}) \right] (\epsilon_{MB} - \epsilon_{\bar{MB}}) \\
&+ \frac{1}{2} \left[(1 - \bar{\tau})^2 g''(Y - (1 - \bar{\tau}) * \epsilon_{\bar{MB}}) - \bar{\tau}^2 * h''(\bar{D} + \bar{\tau} * \epsilon_{\bar{MB}}) \right] (\epsilon_{MB} - \epsilon_{\bar{MB}})^2 .
\end{aligned}$$

The first order condition and the condition from the implicit function theorem allow us to replace the derivatives of $h(\cdot)$ with derivatives of $g(\cdot)$ as follows:

$$\begin{aligned}
U(\epsilon_{MB}, \bar{\epsilon}) &= g(Y - (1 - \bar{\tau}) * \epsilon_{\bar{MB}}) - h(\bar{D} + \bar{\tau} * \epsilon_{\bar{MB}}) \\
&- g'(Y - (1 - \bar{\tau}) * \epsilon_{\bar{MB}}) (\epsilon_{MB} - \epsilon_{\bar{MB}}) \\
&+ \frac{1}{2} * (1 - \bar{\tau}) * g''(Y - (1 - \bar{\tau}) * \epsilon_{\bar{MB}}) (\epsilon_{MB} - \epsilon_{\bar{MB}})^2 .
\end{aligned}$$

Finally, expected utility is given by:

$$EU = \int U(\epsilon, \epsilon_{\bar{MB}}) dG$$

and the risk premium, RP , is implicitly given by:

$$EU = g(Y - (1 - \tau) * \epsilon_{\bar{MB}} - RP) - h(\bar{D} + \tau * \epsilon_{\bar{MB}}) .$$

Hence we have

$$\begin{aligned}
&g(Y - (1 - \bar{\tau}) * \epsilon_{\bar{MB}}) - g(Y - (1 - \bar{\tau}) * \epsilon_{\bar{MB}} - RP) \\
&= -\frac{1}{2} * (1 - \bar{\tau}) * g''(Y - (1 - \bar{\tau}) * \epsilon_{\bar{MB}}) \int (\epsilon_{MB} - \epsilon_{\bar{MB}})^2 dG \\
&= -\frac{1}{2} * (1 - \bar{\tau}) * g''(Y - (1 - \bar{\tau}) * \epsilon_{\bar{MB}}) * var(\epsilon_{MB}) .
\end{aligned}$$

Conversely, had we ignored the impact of unpaid medical bills, we could have applied a second order Taylor approximation around $U^{oop} = g(Y - (1 - \bar{\tau}) * \bar{\epsilon}_{MB})$. This would deliver:

$$\begin{aligned}
U^{oop}(\epsilon_{MB}, \bar{\epsilon}_{MB}) &= g(Y - (1 - \bar{\tau}) * \bar{\epsilon}_{MB}) \\
&- (1 - \bar{\tau}) * g'(Y - (1 - \bar{\tau}) * \bar{\epsilon}_{MB}) (\epsilon_{MB} - \bar{\epsilon}_{MB}) \\
&+ \frac{1}{2} (1 - \bar{\tau})^2 g''(Y - (1 - \bar{\tau}) * \bar{\epsilon}_{MB}) (\epsilon_{MB} - \bar{\epsilon}_{MB})^2 .
\end{aligned}$$

Compared to the case also considering unpaid medical bills, the first and the second order term are now each smaller by a factor of $\frac{1}{1-\bar{\tau}}$. The implied risk premium ignoring the impact of unpaid medical bills RP^{oop} is then

$$\begin{aligned} & g(Y - (1 - \bar{\tau}) \cdot \bar{\epsilon}_{MB}) - g(Y - (1 - \bar{\tau}) \cdot \bar{\epsilon}_{MB} - RP) \\ = & \frac{1}{1 - \bar{\tau}} \cdot \left[g(Y - (1 - \bar{\tau}) \cdot \bar{\epsilon}_{MB}) - g(Y - (1 - \bar{\tau}) \cdot \bar{\epsilon}_{MB} - RP^{oop}) \right]. \end{aligned}$$

It follows that

$$RP^{oop} < RP < \frac{1}{1 - \bar{\tau}} \cdot RP^{oop} .$$

As with the mean reduction, this suggests that considering unpaid medical bills can increase the risk premium by factor of $\frac{1}{1-\bar{\tau}}$. We quantify the risk premium in a numerical example in Section 7.

E.5 Revealed Preference Approach: Calibration Details

As mentioned in the main text, we consider CRRA utilities with parameters of relative risk aversion ranging between 2 and 4. Following (Finkelstein, Hendren and Luttmer, 2015) we normalize income to 3,800. We assume that patients pay 20% of the original medical bill out-of-pocket. The direct evidence indicates average reductions in medical debt of $\frac{\$13.5 \times 4}{0.044} = \$1,227$ per treated person and year, which corresponds to 51% of overall health care utilization. This suggests that the residual 29% of medical bills go as charity care. In other words, individuals are only held responsible $(1 - \alpha_{charity}) \times \$2,400 = \$1,704$ of medical bills.

In Figure E.1, we plot the ratio of the implied compensating variation (CV) over the corresponding medical bill ($\frac{CV}{Medical\ Bill}$) (vertical axis) against the underlying medical bill (horizontal axis). The medical bill amount is net of charity care. As implied by the model, this ratio decreases from a maximum of 100% for small bills to $1 - \bar{\tau} = \frac{0.2}{1-0.29} = 0.28$ for large bills. Moreover, $\frac{CV}{Medical\ Bill}$ is convex in the underlying medical bill amount suggesting that evaluating the ratio at the average medical bill amount would understate the expected $\frac{CV}{Medical\ Bill}$ when considering the full distribution in medical bills. Evaluated at $\theta = 3$ and the average overall net medical utilization of \$1,704 this ratios equals about 70%. The calibration thus implies that restricting consideration to reductions in out-of-pocket payments may understate the effects on consumer welfare by a factor of $\frac{70\%}{28\%} = 2.5$.

Using the above calibrated factor of 2.5, an associated parameter of risk aversion of 3, and considering overall annual health care spending of \$2,400 per uninsured non-elderly person (see Section 3.2), we calculate out-of-pocket spending and implied compensating variation of

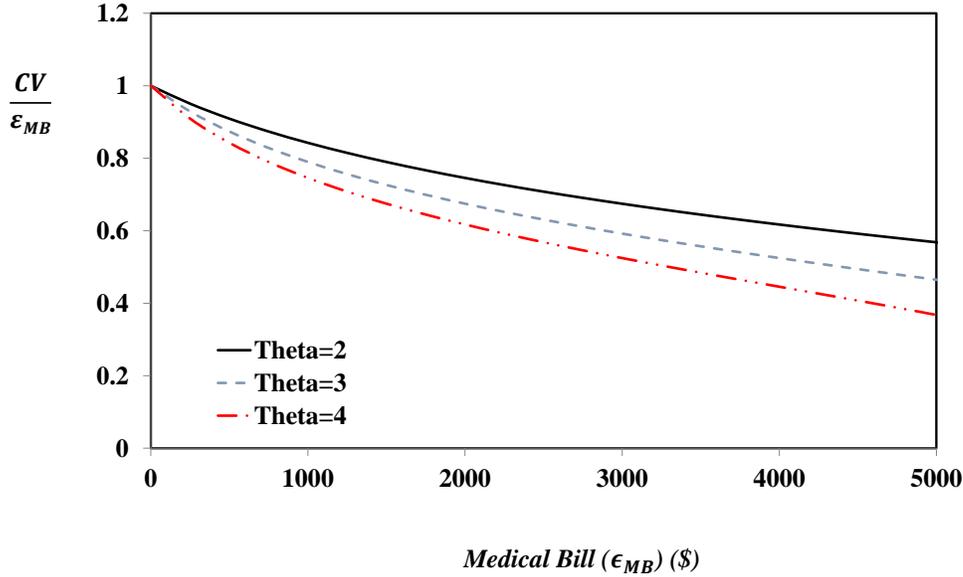


Figure E.1: Compensating Variation/Medical Bill by Medical Bill

\$480 (20%) and $\$480 \times 2.5 = \$1,200$, respectively. This suggests an indirect benefit through the credit channel of $\$1,200 - \$480 = \$720$.

Risk averse consumers also benefit from a reduction in risk. We evaluate the risk premium defined in Equation 10 around average annual consumption of \$3,300 and consider a standard deviation in consumption of \$768 as in (Finkelstein, Hendren and Luttmer, 2015).⁴⁹ To quantify the variance in non-charity medical care, we build on the observation that only $\frac{0.2}{0.71}$ of the non-charity care amount is paid out-of-pocket. Specifically, the variance in consumption then equals $(\frac{0.2}{0.71})^2 \times var(\epsilon_{MB})$. Solving for RP in the revised Equation 10, we find a risk premium of \$680, which exceeds the pure OOP benchmark, building on a revised Equation 11, by a factor of 2.8 (column 1 of Table 7). Combining the estimates, we find an overall annual financial benefit of \$1,880, about 78% of overall medical spending, which exceeds the out-of-pocket benchmark by a factor of 2.6.

⁴⁹The consumption level corresponds to income net of average out-of-pocket spending: $\$3,800 - \$480 \approx \$3,300$.