The Credit Consequences of Unpaid Medical Bills*

Kenneth Brevoort † Daniel Grodzicki ‡ Martin B. Hackmann §

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 a large fraction of unpaid medical bills is 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. We find that 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 Leemore Dafny, 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.

†Federal Reserve Board, ken@brevoort.com.

‡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.

1 Introduction

A fundamental goal of health insurance is to provide financial risk protection against large and unforseen 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% of their overall health care utilization out-of-pocket (Coughlin et al., 2014). This suggests that the financial benefits of health insurance to beneficiaries may be relatively small (Finkelstein, Hendren and Luttmer, 2019). 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 new 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, 2019), 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 dis-utility from leaving medical bills unpaid. They then choose what portion of their medical expenses to leave unpaid, trading off greater consumption with the dis-utility 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 dis-utility from fewer unpaid bills, which operates through the *indirect* credit channel.

We then 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 our empirical analysis.

Our empirical analysis starts out by highlighting an important mechanism underlying the credit channel. Specifically, we argue that health insurance reduces unpaid medical bills sent to third party collectors and reported as medical debt in collections. This decline in the accrual of medical debt raises creditworthiness and leads to higher credit scores among the newly insured, thereby improving the terms of credit available to them.

To quantify the effects of the Medicaid expansion on beneficiaries' accrual of medical debt and subsequent creditworthiness, we combine state adoption decisions and census tract-level variation in eligibility from the Medicaid expansion with administrative data from the Consumer Financial Protection Bureau's (CFPB) Consumer Credit Panel (CCP), a nationally representative panel of over 5 million de-identified credit records. Unlike other credit panels, the CCP contains information on individual credit obligations (trade-lines), the source of each obligation, the date it was credited, and quarterly changes in its repayment condition (e.g. delinquency status). In particular, the CCP allows us to separate medical from non-medical collections and to identify the date on which the collection was placed on an individuals' record. In addition, the CCP provides quarterly updates of individuals' credit score.

We find that the Medicaid expansion reduced newly-accrued medical debt by $54 per person annually, or $1,231 per treated person. 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 35% of overall utilization and 40% of unpaid medical bills (uncompensated care) of the uninsured go into collection. The large reduction in medical debt had a positive effect on beneficiaries creditworthiness, substantially reducing the likelihood of severe delinquencies and raising credit scores.

Using novel data on direct-mail credit offers from Mintel Comperemedia (Mintel) and aggregated lender rate sheets collected by the Fair Isaac Corporation (FICO), we then assess the extent to which improvements in beneficiaries' creditworthiness translated into better terms of credit available to them. We document significant declines in the offered interest rates on credit cards and personal loans, as well as in available interest rates on auto loans and mortgages, due to the reform.

Building on this evidence, we quantify the relative importance of the indirect credit channel of insurance on consumer welfare. To this end, we develop a revealed preference approach, which is predicated on the assumption that individuals make optimal out-of-pocket payment decisions when confronted with a medical bill. Given standard assumptions on consumption utility, this approach allows us to infer the dis-utility of leaving medical bills unpaid from observed out-of-pocket payments. Intuitively, out-of-pocket spending would fall to zero in the absence of dis-utility from unpaid medical bills. Yet, this stands in contrast to observed out-of-pocket payments among the uninsured. Using micro data from the Medical Expenditure Panel Survey (MEPS), we document that annual out-of-pocket spending increases to \$2,300 out of an annual income of only \$4,400 when individuals are confronted with large cumulative medical bills exceeding \$20,000.

Combining the observed repayment decisions from the MEPS with direct evidence on the reform's impact on medical debt in collection, the revealed preference approach implies that the overall financial benefits of Medicaid insurance increase substantially when incorporating the indirect credit channel. Considering Medicaid gross expenditures of \$3,600 per beneficiary and year, we estimate a financial benefit of about \$0.7 per Medicaid dollar. In contrast, we find a benefit of only \$0.2 per Medicaid dollar when focusing only on the benefits arising from a mean and variance reduction in out-of-pocket spending, excluding the credit channel.

We revisit the quantitative importance of the indirect channel in several robustness exercises. We find that the financial benefit per Medicaid dollar remains largely unchanged at \$0.7 when varying the baseline income or the calibrated parameter of relative risk aversion. Our baseline estimate of the financial benefit per Medicaid dollar increases to \$0.8, when we assume that 60% as opposed to only 40% of unpaid medical bills are sent to collection agencies. Conversely, our benefit estimate decreases to \$0.6, when we assume that only 20% of unpaid medical bills are sent to collection agencies. Our estimate falls to \$0.4, when we assume that unpaid bills in collection are bounded from above by the beneficiaries' annual income net of out-of-pocket spending.

Finally, we relate our credit channel estimates to the estimated interest rate reductions discussed above. To this end, we provide a ballpark estimate on the potential interest rate savings using a stylized back-of-the-envelope calculation. Scaling the estimated interest rate reductions by the observed pre-reform debt, we find potential annual aggregate savings of \$18.66 per person annually, or \$424 per treated person, which come predominantly from credit card and unsecured personal loan debt. Adding the interest savings to the reductions in out-of-pocket spending, we find that the financial benefits of a mean reduction in medical bills increases by 90% when considering the indirect credit channel. Unfortunately, this calculation does not provide us with an estimate on value of risk protection. Borrowing the corresponding estimate from our baseline approach, we find a benefit of about \$0.5 per Medicaid dollar.

Overall, we estimate a financial benefit of \$0.4 to \$0.8 per Medicaid dollar across a variety of specifications. In contrast, we find a benefit of only \$0.2 per Medicaid dollar across these specifications when focusing only on the benefits arising from a mean and variance reduction in out-of-pocket spending. This suggests that the financial benefit of Medicaid insurance at least doubles when adding the indirect credit channel to the benefits from a mean and variance reduction in out-of-pocket spending.

Our paper contributes to three main literatures. First, our analysis complements recent studies on the value of Medicaid (Finkelstein, Hendren and Luttmer, 2019) and the value of public insurance more generally (Kowalski, 2015; Cabral and Cullen, 2019; Finkelstein and McKnight, 2008). These studies investigate the overall consumer benefit of public insurance, taking financial and health related benefits into account. In the context of Medicaid, Finkelstein, Hendren and Luttmer (2019) find that beneficiaries value the program by only \$0.2 to \$0.5 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 (2019) 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., 2018; Sojourner and Golberstein, 2017; Caswell and Waidmann, 2017; Gallagher, Gopalan and Grinstein-Weiss, 2019; Allen et al., 2017; Argys et al., 2017). We contribute to this literature by developing and implementing a method to quantify the consumer welfare implications of reductions in unpaid medical bills. We also provide novel evidence on the reform's effect on credit scores and available interest rates on credit cards, personal loans, mortgages, and automobile loans.

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 et al. (2014) and Dranove, Garthwaite and Ody (2016)). Notably, Garthwaite, Gross and Notowidigdo (2018) 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, see Finkelstein, Hendren and Shepard (2019). We contribute to these studies by shedding new light on the incidence of uncompensated care. Our findings suggest that the incidence of uncompensated care falls at least partially on the low-income uninsured patients themselves, through the 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 and Section 4 describes the data. In Section 5, we develop and document an important mechanism for the credit channel of health insurance by tying the Medicaid expansion to measures of creditworthiness and the terms of credit. Building on this evidence, we return to the conceptual framework in Section 6, and discuss the implications for consumer welfare. Finally, Section 7 concludes.

2 The Credit Channel of Health Insurance

To assess the role of the indirect credit channel, we begin by extending the textbook model of insurance to examine the role of unpaid medical bills in consumer welfare. Our conceptual framework focuses on the financial benefits of health insurance. To this end, we treat healthcare 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 a static environment in which an uninsured individual derives positive utility from consumption, c, and faces a utility loss from unpaid bills, b. Utility losses from unpaid bills captures the cost of worsening credit options, and also hassles from dealing with debt collectors and legal complications related to unpaid bills. Let an individual's utility be of the form

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

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 dis-utility of leaving medical bills unpaid is weakly increasing. We return to these properties in Section 6.

An individual earns income Y and is exposed to a random medical bill $\epsilon_{MB} \sim F$, where F denotes the underlying distribution function. Upon incurring a medical bill, she decides how much of her medical bill $0 \leq b \leq \epsilon_{MB}$ to leave unpaid. This decision triggers an inherent trade off in utility from greater consumption and dis-utility from leaving bills unpaid:

$$\max_{0 \leq b \leq \epsilon_{MB}} g(Y - (\epsilon_{MB} - b)) - h(b) \ . \tag{2}$$

The optimal choice of b^* is determined by the first order condition, which equates the marginal utility of consumption and the marginal dis-utility of unpaid medical bills:

$$g'(Y - (\epsilon_{MB} - b^*)) = h'(b^*). \tag{3}$$

We illustrate this trade-off graphically in Figure 1, which plots consumption on the horizontal axis and marginal (dis)utility on the vertical axis. The downward sloping line

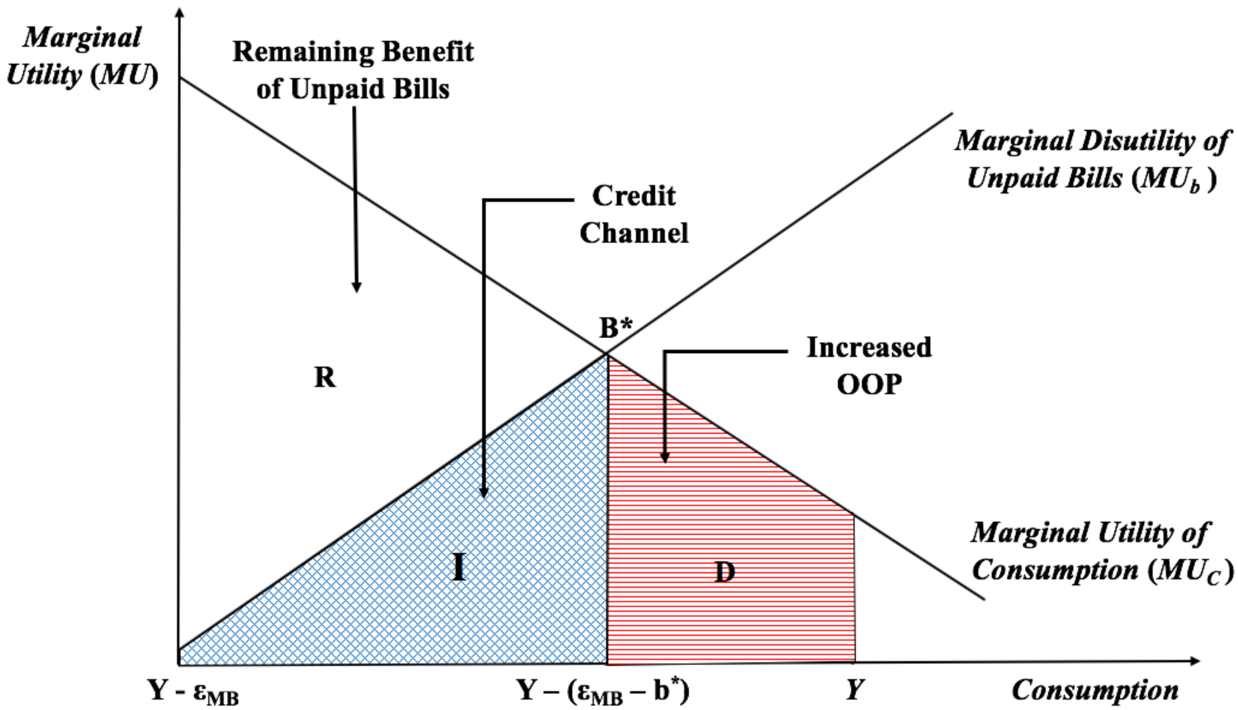

Figure 1: Consumer Welfare Effects of a Medical Bill: Example

represents the marginal utility of consumption (MU_C), and the upward sloping line is the marginal dis-utility of unpaid medical bills (MU_b). Absent any medical expenses, an individual consumes her income Y. When facing a medical bill of size ϵ_{MB}, she pays $\epsilon_{MB} - b^*$ out-of-pocket and consumes $Y - (\epsilon_{MB} - b^*)$, as shown in equation 3. The amount consumed out of income is shown in point B^* and depends on the underlying medical bill amount. We can decompose the consumer's 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, 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 the area **I**, bounded by the marginal dis-utility of unpaid medical bills, final consumption, $Y - (\epsilon_{MB} - b^*)$, and income net of the entire medical bill $Y - \epsilon_{MB}$. As described above, this *credit channel* highlights the potentially adverse consequences of unpaid bills on access to and the price of credit in addition to other costs associated with not paying bills. The sum of areas **I**+**D** captures the overall utility loss from receiving the medical bill shock ϵ_{MB}. Finally, the area **R** captures any remaining net benefit from leaving a medical bill unpaid. This is because, were the individual to pay the entire bill amount (ϵ_{MB}) out-of-pocket, the utility loss would be the entire area underneath the marginal utility of consumption curve = **R**+**I**+**D**.

Building on the tradeoff between consumption utility and the dis-utility of unpaid medical bills, we now turn to the consumer welfare implications of a mean and variance reduction in paid and unpaid medical bills as provided by Medicaid insurance. We define γ as the amount of consumption that the individual would have to give up under Medicaid insurance in order to maintain the same expected utility without Medicaid insurance:

$$\mathbb{E}_{\epsilon_{MB}}\Big[g(Y - (\epsilon_{MB} - b^*(\epsilon_{MB}))) - h(b^*(\epsilon_{MB}))\Big] = g(Y - \gamma) \ . \tag{4}$$

Here, the expectation is taken over the medical bill amount ϵ_{MB}. For simplicity, we abstract away from uncertainty in income. γ is measured from the point of view of the Medicaid beneficiary and can be decomposed into a transfer component, γ^T, and risk-protection component, γ^{RP}. γ^T is implicitly defined by

$$g(Y - \mathbb{E}_{\epsilon_{MB}}[\epsilon_{MB} - b^*(\epsilon_{MB})]) - h(\mathbb{E}_{\epsilon_{MB}}[b^*(\epsilon_{MB})]) = g(Y - \gamma^T). \tag{5}$$

The pure risk-protection component γ^{RP} is then

$$\gamma^{RP} = \gamma - \gamma^T \ . \tag{6}$$

2.1 Revealed Preference Approach

To quantify the transfer and the risk protection component, we specify the utility function over consumption and trace out the dis-utility over unpaid medical bill from observed "optimal" out-of-pocket payment decisions.

Following Finkelstein, Hendren and Luttmer (2019), we assume a constant relative risk aversion (CRRA) utility function in consumption:

$$g(c) = \frac{c^{1-\sigma}}{1-\sigma} \ \ , \tag{7}$$

where σ denotes the coefficient of relative risk aversion. Combining observed medical bill amounts and "optimal" out-of-pocket payment decisions from the MEPS data with the first order condition (equation 3), we can non-parametrically recover the marginal dis-utility of unpaid medical bills $h'(b^*(\epsilon_{MB}))$. Taking an integral of $h'(\cdot)$, we can recover the overall dis-utility of unpaid medical bills $h(\cdot)$ and quantify γ, γ^{RP}, and γ^{T} using equations 4, 5, and 6.

We then benchmark the estimated transfer and risk protection components to those derived under a standard utility framework that ignores the cost of leaving medical bills unpaid. In this out-of-pocket framework, we set $h(b) = 0 \ \forall \ b$ and refer to the analogous transfer and risk protection components as γ^{T}_{OOP} and γ^{RP}_{OOP}, respectively.

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 FPL. 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.[1]

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.[2] Another 7 states expanded coverage at various points after January 1, 2014.[3] 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.

[1] The court case is known as *National Federation of Independent Business v. Sebelius*, 567 U.S. 519 (2012), see Kaiser Family Foundation (2012) for more details.

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

[3] Alaska expanded on September 1, 2015. Indiana expanded on February 1, 2015. Louisiana expanded on July 1, 2016. Michigan expanded on April 1, 2014. Montana expanded on January 1, 2016. New Hampshire expanded on August 15, 2014. Pennsylvania expanded on January 1, 2015.

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).

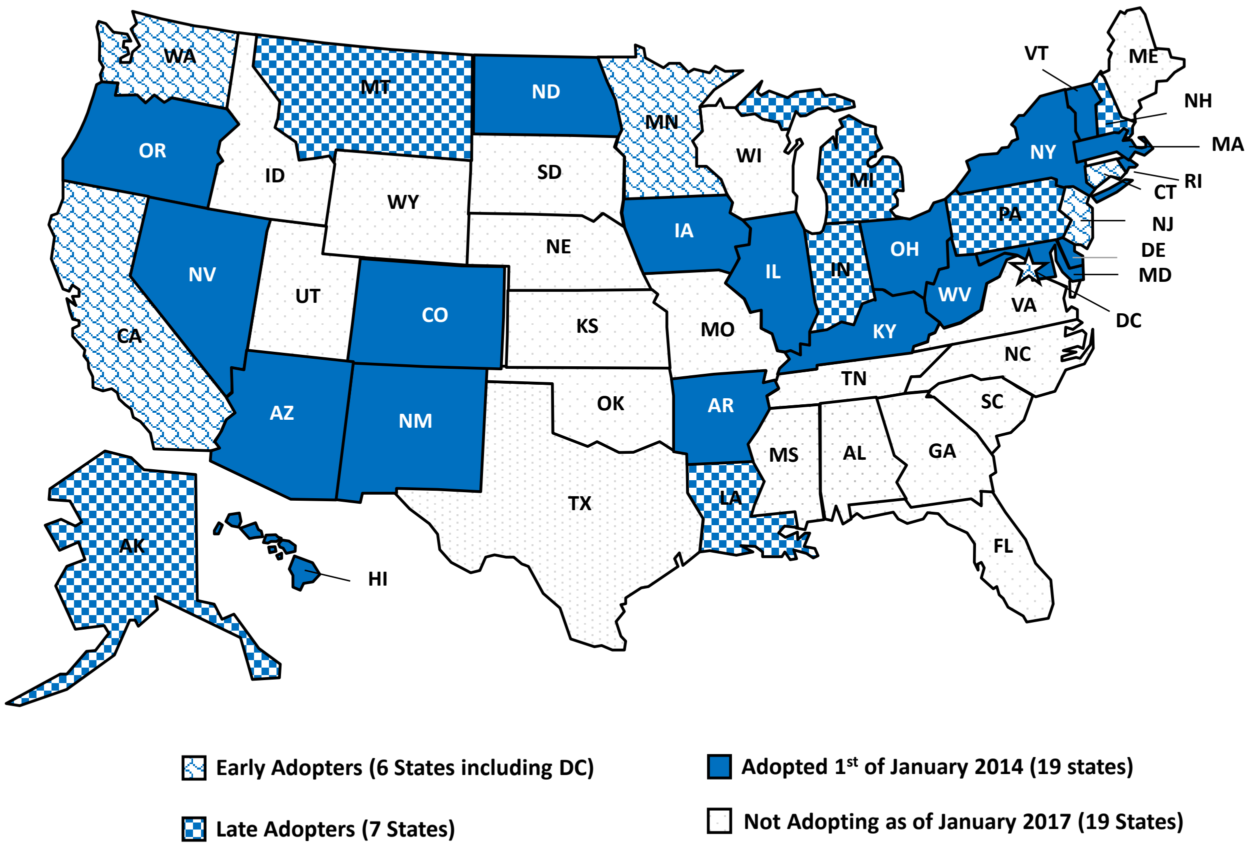

Figure 2: Medicaid Adoption Across States

Source: *Kaiser Family Foundation. For further details see http://kff.org/health-reform/slide/current-status-of-the-medicaid-expansion-decision/.*

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%.[4] Hence, we attribute a differential Medicaid enrollment increase of 3.3 million, about 4.1% of the non-elderly population, to the Medicaid expansion, which is

[4]See 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 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.

roughly consistent with estimates from the literature.[5] Repeating the analysis for December 2016, we find a cumulative coverage 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.

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 only up to 20% of medical bills out of pocket (Finkelstein, 2007). The remaining cost is left as uncompensated care (Coughlin et al., 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

[5] 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 (2017) 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.

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 et al., 2014).

4 Data

4.1 Consumer Credit Panel

Data on medical debt, delinquencies, and credit scores used in this study come from the Consumer Financial Protection Bureau's (CFPB) 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). We focus 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 and late-adopting states to simplify the empirical analysis.[6] 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 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.

Each credit record in the sample includes information about every individual debt obligation, or trade-line, reported on that record. This includes a trade-line'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. The key variables of interest for our analysis include medical debts in collection, the 90-day delinquency rate, and the credit score, which we describe in further detail below.

Medical Debt: We define medical debts as those reported directly by a medical provider or by third-party debt collectors as unpaid medical bills.[7] 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

[6] We 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.

[7] Our 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).

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. Compared with a stock measure of outstanding medical debt, which is provided in other credit panels, the flow measure offers two advantages. First, we can leverage the precise timing of newly accrued medical bills in an event study approach that sheds new light on the relationship between medical debt and measures of financial distress. Second, tracking medical debt since their origination provides a comprehensive assessment of the overall medical debt amount in collection. In contrast, about 38% of debts are no longer owed after one year indicating that stock measures provide a noisier and incomprehensive assessment of the overall debt amount in collection.[8] Related to the timing, we also note that there are oftentimes lags between when debts are acquired and when they are reported to the NCRC. The reporting delay does not affect the reported trade line's opening date, so we can still assign later-reported medical debts to the quarters in which they were incurred. To account for this lag, we identify new tradelines opened in a calendar quarter by looking at the CCP archive from the following quarter. This lag was selected because it represented the best balance between allowing time for new debt collections to be reported, but not so long that they began to be removed as mentioned beforehand.

We note two potential limitations in the measurement of medical debt. First, an amendment to the ACA contained provisions restricting the collection activities of not-for-profit hospitals.[9] These provisions became effective with the publication of the implementing regulations by the Internal Revenue Service (IRS) in December 2014 and may have reduced the amount of medical debt in collection from not-for-profit hospitals. These provisions apply nationwide and we intend to control for these changes via year-quarter fixed effects. Nevertheless, we revisit asymmetric impacts between expansion and non-expansion states in Robustness Section 6.1.

Second, we may understate the amount of medical debt in collection if some of these collections are not reported to one or more credit bureaus. To assess the fraction of medical collections that are not reported to any credit bureau, we first compare our estimates to

[8] Collection accounts that are reported by a collection agency may be removed if that collection agency resells the debt to another buyer, or the third-party debt collector's contract to collect the debt with the owner of the original debt expires, or a consumer or their insurance company repays the debt.

[9] Specifically, not-for-profit hospitals are required to determine whether patients qualify for the hospital's financial assistance or charity policies before initiating 'extraordinary collection actions', which include selling the debt to a third-party debt collector or reporting information about the debt to a credit bureau. These restrictions are described in 26 CFR § 1.501(r), see https://bit.ly/2FznUgU for details, last accessed January 7, 2020.

the results from the 'Survey of Consumer Views on Debt', which was conducted by the CFPB between December 2014 and March 2015.[10] The survey asked respondents about their interactions with debt collectors. 32% of respondents reported being contacted in the past year by a debt collector. Of these, 59% reported being contacted about a medical debt. Together, this implies that about 32% · 59%=19% of consumers reported about a medical debt, which is very close to the estimated 19.5% of consumers with a medical collection account on their credit record in our sample. Second, to explore reporting differences between credit bureaus, we compare our estimates to the results reported by the Federal Reserve Board (FRB), which uses data from a different credit bureau. Based on a random sample of credit records from 2003 for which credit scores can be created, the FRB reports that 16.2% of consumers have a medical collection account on their credit record.[11] Restricting our sample to records for which credit scores can be created, we find that 17% of consumers have a medical collection account on their credit record.[12] The similarity in the estimates alleviates concerns over underreporting of medical debt in collection. Nevertheless, we return to the potential implications for our main estimates in Robustness Section 6.1.

Measures of Financial Distress: Turning to our measures of financial distress, we focus on 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 trade-line 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. We focus our empirical analysis on transitions into 90-day, or serious, delinquency. More often than not, accounts that become 30-day delinquent are returned to current status. However, upon becoming 90+ days delinquent, borrowers almost never become current on their debt. As a result, we interpret transitions into 90-day delinquency as a proxy for stopping payments on a loan altogether.

Credit Score: Creditworthiness is often summarized via a credit score. Each quarterly archive of CCP data includes a widely-used, commercially available credit score that was produced for that record. Although the payment history is generally the most important

[10]See https://bit.ly/2QDPlgc for details, last accessed January 7, 2020.

[11]See https://bit.ly/2NavCT2 for details, last accessed January 7, 2020.

[12]To match the FRB's 2003 estimate as closely as possible, our specific estimate relies on data from the third quarter of 2011 only, which is the first quarter in our sample period. Furthermore, and following the FRB sample construction, we only consider credit records that have existed for at least 18 months. We also note that our estimate is close to the estimates of collection accounts from Avery et al. (2003), who find that just over 30% of credit records have a collection (medical or non-medical) and that slightly over half of collections accounts are medical using data from a different credit bureau than ours and a different time period than the FRB.

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.

4.2 American Community Survey

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 KFF.[13] 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.

4.3 Loan Offers and Pricing (Mintel and MyFico)

Data on credit offers and pricing are from Mintel Comperemedia (Mintel), for credit cards and personal loans, and aggregated rate-sheet data from Fair Isaac Corporation's *MyFico* pricing tool, 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 consumers' credit options.

The Mintel 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. Participants are asked to provide Mintel with all mail solicitations they received during the month. These include offers of new credit from all lenders in the marketplace. 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. Mintel combines offer information from the mailings with the demographic profiles of respondents, including the

[13] Specifically, we use eligibility criteria by state for childless adults as of January 1, 2013.

county in which they reside. Because the Mintel data provide extensive information on the supply decisions of nearly all lenders in the marketplace, as well as demographic information on recipients, they are uniquely suited for exploring changes in the supply of credit to consumers following the Medicaid expansion.

An advantage of the Mintel data is that it allows us to isolate offers made to potential customers whose credit record, and score, have been previously checked by the lender. These offers are commonly referred to as pre-screened offers.[14] We focus on pre-screened offers to provide a tight link between changes in an individual's credit record and the pricing of credit. 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 (Argyle, Nadauld and Palmer, 2017). These 'Rate sheets', which are often set by each lender statewide or nationally, translate credit score ranges into the interest rates available from a lender. The Fair Isaac Corporation uses aggregated rate-sheet information from Informa Data Services 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 for their educational *MyFico* web pricing tool. In our analysis, we use these rate sheets to assign to each consumer the auto loan and mortgage interest rate for which they would have qualified in that quarter based on their credit score. Specifically, we impute rates for 5-year auto loans and 30-year fixed rate mortgages, the most common product within each respective loan category, see Table C.2 in Appendix C for details. We then estimate the impacts of the reform on available auto loan and mortgage rates using these imputed values.[15]

4.4 Medical Expenditure Panel Survey (MEPS)

We complement the credit panel and pricing data with data from the MEPS, which allows us to identify out-of-pocket payments and the medical charge amounts. We closely follow the methodology in Coughlin et al. (2014) and use data from the household and medical provider component. We complement their analysis of the 2008-2010 waves with more recent survey data from 2011-2013 and project health care utilization in 2013, the last pre-expansion

[14] All pre-screened offers must, by federal law, include an opt-out disclosure, which is flagged in the data. This option allows consumers to remove their name from prescreened lists, whereby credit bureaus can no longer sell their credit record information to credit card companies.

[15] Consumers with credit scores below the bottom price tiers are excluded from calculations concerning automobile loans and mortgages, as they likely do not qualify for a loan.

year in the expansion states.[16] To construct the annual medical utilization amount for the uninsured, we follow Coughlin et al. (2014) and multiply the observed annual charge amount (for the uninsured) by the observed aggregate payment-to-charge ratio for the privately insured population. A a result, we are able to match their reported average annual out-of-pocket spending and overall utilization estimates for the non-elderly uninsured population very closely.[17]

To relate the MEPS sample population as closely as possible to the Medicaid expansion context, we focus on the non-elderly adults aged 18-64. We restrict our sample to the uninsured population with incomes below 80% of the FPL, in order to match the average income among new Medicaid beneficiaries of 39% of the FPL ($4,400 per year) reported in Miller et al. (2018). This yields a sample population of 5,430 uninsured individuals with average annual out-of-pocket spending of $480, about 14% of overall annual utilization ($3,530). For comparison, individuals in the control group of the Oregon experiment spend on average $569 of the $2,721 in overall health care utilization out-of-pocket (Finkelstein et al., 2012).

5 Mechanism of the Credit Channel

Our empirical analysis starts with an assessment of the mechanism underlying the indirect effect, or credit channel, of Medicaid insurance. Specifically, we investigate whether Medicaid leads to a decline in the accrual of medical debts, and how this affects the creditworthiness of the newly insured. We then explore the extent to which these changes led to improved credit options for these individuals.

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' adoption decisions in a simple difference-in-differences framework. Our primary specification is as follows:

$$y^k_{c,t} = \alpha^k_c + \eta^k_t + \sum_{j=1}^{3} \beta^k_j \times Exp_c \times Post_j + \epsilon^k_{c,t} \ , \tag{8}$$

where, y^k_{ct} 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^k_c and η^k_t denote census tract and year-quarter fixed effects, respectively.

[16] We use changes from 2011 and 2012 to 2013 in overall population growth and per capita personal health care expenditures from Table 5 and Table 1 of the National Health Expenditure Projections, respectively. See Centers for Medicare and Medicaid Services (2011) for details.

[17] We estimate annual out-of-pocket spending of $470 and overall annual utilization of $2,420. For comparison, Coughlin et al. (2014) report $500 and $2,443 in Figure 1, 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, $\beta_j,\ j = 1, 2, 3$, which map out the full dynamic effects of the policy during its first three years. We interpret the coefficients $\beta_j,\ j = 1, 2, 3$ 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 in treatment states differ 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 are 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.[18]

Second, we exploit within-state variation across Census tracts that differ in the fraction of newly-eligible adults, i.e. the intensity of treatment. 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 sub-sample allowing for a non-parametric comparison of the reform effects between Census tracts.

[18]Our unit of analysis is a census tract by year-quarter. To construct the synthetic control weights, we first aggregate observations at the state and year-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 the 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. In a previous version of the paper, we conducted the empirical analysis without synthetic control weights, i.e. unweighted. While the method of synthetic controls raises our level of confidence in the findings, we note that our main results remain largely unchanged by its inclusion. The previous version of the paper is available upon request.

5.1 Medicaid and Medical Debt

Panel A of 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 and year-quarter (ct). The table shows the full dynamic effects of the reform in each of its first three years ($Exp \times Post_j$) in addition to the average impact over the analysis period ($Exp \times Post$). The first column of the table shows results for the entire sample, while columns 2–5 show these effects by quartile of newly-eligible individuals.

As shown in the panel, the Medicaid expansion led to a decline in the value of newly accrued medical debt on the order of \$13.54 per quarter, or 35%, 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 relative to low eligibility tracts, see the fifth and the second column, respectively. 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, we find that the expansion prevented \$5.89 billion in medical collections from being debited to households balance sheets over the first three post-reform years, see Table A.1 for details.

We also explore the distributional effects of the Medicaid expansion. As expected, 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% at the 89^{th} quantile to 40% at the 99^{th} quantile, see Figure A.3 in Appendix Section A.1.2 for details.

To put our estimated mean 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 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, see Table A.1) we find an annual debt reduction per treated 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

Table 1: New Medical Debt, Delinquencies, and Changes in Credit Scores

		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)
A. New Medical Debt in Collection					
1^{st} Year ($Exp \times Post_1$)	-10.07	-2.99	-8.25	-10.58	-18.94
	(1.31)	(1.35)	(2.11)	(2.55)	(4)
2^{nd} Year ($Exp \times Post_2$)	-17.17	-4.51	-10.87	-21.12	-33.06
	(1.52)	(1.39)	(1.96)	(3.78)	(3.97)
3^{rd} Year ($Exp \times Post_3$)	-13.37	-6.42	-10 .8	-13.54	-23.59
	(1.14)	(1.15)	(1.83)	(2.35)	(3.3)
Average Effect ($Exp \times Post$)	-13.54	-4.64	-9.97	-15.08	-25.19
	(1.01)	(0.9)	(1.58)	(2.18)	(2.83)
Pre-Reform Mean Value (in Expansion States)	38.54	18.91	35.41	56.97	80.16
B. New 90 Day Delinquencies					
1^{st} Year ($Exp \times Post_1$)	-0.04	0.00	-0.01	-0.06	-0.07
	(0.02)	(0.03)	(0.03)	(0.04)	(0.04)
2^{nd} Year ($Exp \times Post_2$)	-0.07	-0.05	-0.02	-0.09	-0.14
	(0.02)	(0.03)	(0.03)	(0.04)	(0.05)
3^{rd} Year ($Exp \times Post_3$)	-0.03	0.01	0.01	-0.04	-0.12
	(0.02)	(0.03)	(0.03)	(0.04)	(0.05)
Average Effect ($Exp \times Post$)	-0.05	-0.01	-0.01	-0.06	-0.11
	(0.01)	(0.02)	(0.03)	(0.03)	(0.03)
Pre-Reform Mean Rate (in Expansion States)	2.03	1.74	2.07	2.24	2.56
C. Credit Score					
1^{st} Year ($Exp \times Post_1$)	0.53	0.4	0.3	0.41	1.25
	(0.16)	(0.31)	(0.29)	(0.31)	(0.4)
2^{nd} Year ($Exp \times Post_2$)	1.08	1.05	0.56	0.88	2.33
	(0.2)	(0.38)	(0.37)	(0.38)	(0.49)
3^{rd} Year ($Exp \times Post_3$)	1.52	0.96	0.98	1.26	3.5
	(0.23)	(0.44)	(0.41)	(0.45)	(0.56)
Average Effect ($Exp \times Post$)	1.04	0.8	0.61	0.85	2.36
	(0.17)	(0.33)	(0.32)	(0.33)	(0.42)
Pre-Reform Mean Score (in Expansion States)	679.25	701.63	678.94	661.83	627.55
Percent Newly Eligible Adults	0-100	0-10	10-19	19-32	32-100

Notes: *The table shows effects of the Medicaid expansion on newly accrued medical debt in collection (panel A), the rate of new delinquencies (panel B), and the credit score (panel C) at the year-quarter level using equation 8. Standard errors (in parentheses) are clustered by census tract.*

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. Using the MEPS data, we find that the uninsured in our sample population pay about 14% of overall health care utilization, worth \$3,530 per year, out-of-pocket. This suggests that uncompensated care equals about \$3,036 per uninsured person and year. 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 $\frac{\$13.54\times4}{0.044}$=\$1,231 per treated person in the first year, which is about 35% of overall health care utilization or about 40% of uncompensated care.

Robustness: We consider various robustness checks to support our main findings, detailed in Appendix Section A. 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 measured effects, see the right graph of Figure A.2.

5.2 Medical Debt and Distress

Medical debt is frequently reported to credit bureaus and has a direct and adverse effect on individuals' credit scores. This is because, as it is indicative of distress and future delinquency, medical debt directly enters the credit scoring algorithm (Brevoort and Kambara, 2015). In Appendix Section A.2.1, we document this direct relationship using an event study approach, which explores the precise timing of newly-accrued medical debt in collection. Specifically, we quantify a positive and a negative effect of new medical collections on new delinquencies and credit scores, respectively. Building on this evidence, we next turn to the indirect effect of the Medicaid expansion on rates of new delinquencies and credit scores.

5.2.1 Medicaid and Periods of Worsening Distress

Panel B of Table 1 shows the reform's effects on the rate of new delinquencies. The dependent variable is the likelihood of a new 90-day, or serious, delinquency on any debt obligation in census tract c during year-quarter t. As shown in the panel, during its first three years the expansion reduced transitions into 90-day delinquencies by $\frac{0.05\%}{2.03\%} = 2.5\%$ of the pre-reform mean. As with medical debt (Panel A), the impact of the reform on new delinquencies rose substantially (in magnitude) between the first and the second post-reform year, though it retreated in the third year. Further, the decline in financial distress was highly concentrated in Census tracts with a high proportion of newly-eligible adults, see the fifth column. The effect was about twice as large for the highest quartile of newly eligibles as for the population as a whole.[19]

Putting these intent-to-treat estimates into perspective, we again divide the estimates by the average coverage increase. This suggests that transitions into 90-day delinquency among the newly insured declined by $\frac{0.05}{0.044} = 1.14$ percentage points. The size of the effect is large and amounts to nearly 50% of the pre-reform mean rate (2.56) in Census tracts with the largest fraction of newly eligible adults, see the fifth column.

5.2.2 Medicaid and Credit Scores

Turning to credit scores, Panel C of Table 1 presents results from equation 8 in which the dependent variable is the average credit score in census tract c during year-quarter t. As shown in the panel, we find a statistically significant increase of about 0.5 points in the first post-reform year, see the first column. Moreover, the effect triples by the third post-reform year. The rise in credit scores is concentrated in communities with a high proportion of

[19] We also explore the effects on new 30-day delinquency, which indicates a less severe degree of financial distress, see Figure A.5 for details. Our findings are qualitatively similar to those on 90-day delinquency.

newly-eligible adults, see the fifth column. In these communities, the overall effect is more than two times larger than the average. Once again, scaling the overall average effect by the estimated coverage increase, we find that treated individuals saw a $\frac{1.04}{0.044}= 23.6$ point increase in credit scores, which corresponds to 30% of the pre-reform difference in average credit scores between the bottom and the top eligibility quartile.

We also explore heterogeneity in the reform's effect on credit scores at different points along the credit score distribution. We find evidence for 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, see Figure A.7 for details.

5.3 Medicaid and the Terms of Offered Credit

Having established the reform's positive role in improving individuals' repayment outcomes, and subsequently their credit scores, we now turn to its impact on the price of credit offered to consumers. Specifically, we look at the four most commonly held debt obligations: credit card debt, personal loans (unsecured installment loans), auto loans, and mortgages.

Building on the Mintel data structure, we estimate the reform's effect on the mean interest rates on credit cards and personal loans offered to a household in a given month. Leveraging the county of residence information in the Mintel data, we adapt the structure of equation 8 and control for county fixed effects as well year-quarter fixed effects. As done previously with medical debt, repayment, and credit scores we assess heterogeneous effects of the reform by estimating its effects separately for communities with a high and low fraction of newly-eligible adults. Due to the limited sample size, we group counties into one of two categories depending on whether their share of newly eligible adults exceeds or falls short of the share in the median county. To maintain consistency in our interpretation of these heterogeneous effects, we define the median eligibility rate using the CCP and not Mintel.

Consistent with our earlier findings, the first column of Panel A of Table 2 shows significant reductions in credit card interest rates. We estimate an average intent-to-treat effect of -42 basis points (bps) overall, and -60 bps in communities with a higher proportion of newly eligibles, see the third column. Consistent with our findings on delinquencies and credit scores, the reform's effects on offered credit card rates grew over time. While during the reform's first year rates changed by an average of -9 bps, by the third year this decline was 9 times larger. Among those living in communities with a larger fraction of newly eligible adults, this pattern is still more evident. While individuals in these communities saw a change of only -5 bp during the first year following the reform, the average drop in the third year exceeded one percentage point.

Table 2: Decrease in Offered (Available) Interest Rates on New Credit

	Share of Newly Eligibles			Share of Newly Eligibles		
	All	Below Median	Above Median	All	Below Median	Above Median
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Credit Cards and Personal Loans	Credit Cards			Personal Loans		
1^{st} Year ($Exp \times Post_1$)	-0.09	-0.11	-0.05	-0.49	-0.48	-0.50
	(0.16)	(0.17)	(0.43)	(0.48)	(0.54)	(0.97)
2^{nd} Year ($Exp \times Post_2$)	-0.27	-0.27	-0.29	-0.72	-0.76	-0.57
	(0.25)	(0.28)	(0.55)	(0.47)	(0.51)	(1.19)
3^{rd} Year ($Exp \times Post_3$)	-0.78	-0.62	-1.31	-0.33	-0.35	-0.21
	(0.22)	(0.26)	(0.40)	(0.27)	(0.30)	(0.61)
Average Effect ($Exp \times Post$)	-0.42	-0.36	-0.60	-0.50	-0.52	-0.40
	(0.15)	(0.17)	(0.36)	(0.26)	(0.29)	(0.52)
Pre-Reform Mean Offered Rate in %	15.42	15.40	15.57	9.29	9.21	9.74
Pre-Reform Mean Balance (in Expansion States)	\$4,323			\$614		
Panel B: Auto Loans and Mortgages	Auto Loans			Mortgages		
1^{st} Year ($Exp \times Post_1$)	-0.024	-0.025	-0.023	0.000	-0.003	0.003
	(0.007)	(0.009)	(0.011)	(0.001)	(0.001)	(0.002)
2^{nd} Year ($Exp \times Post_2$)	-0.042	-0.039	-0.047	-0.000	-0.001	0.000
	(0.008)	(0.011)	(0.013)	(0.001)	(0.001)	(0.002)
3^{rd} Year ($Exp \times Post_3$)	-0.054	-0.039	-0.071	-0.003	-0.004	-0.003
	(0.010)	(0.012)	(0.015)	(0.001)	(0.002)	(0.002)
Average Effect ($Exp \times Post$)	-0.040	-0.034	-0.047	-0.001	-0.002	0.000
	(0.007)	(0.009)	(0.011)	(0.001)	(0.001)	(0.002)
Pre-Reform Mean Available Rate in %	7.485	7.001	8.672	4.209	4.187	4.274
Pre-Reform Mean Debt (in Expansion States)	\$4,845			\$56,459		
Percent Newly Eligible Adults	0-100	0-19	19-100	0-100	0-19	19-100

Notes: *Panel A shows effects of the Medicaid expansion on credit card and personal loan interest rates offered to consumers from Mintel Comperemedia, as implied by equation 8. The unit of observation is a household in a month. The dependent variable is the mean rate for each respective product offered to an individual in a given month. Panel B shows the effects of the Medicaid expansion on available rates for auto loans and mortgages, also using equation 8. The unit of observation is an individual in a quarter, and the dependent variable is the imputed, or 'available', rate for an individual in a given quarter. Columns 1 and 4 show the effects overall. Columns 2,5 and 3,6 report estimates for counties whose share of newly-eligible adults falls below and above the sample median, respectively. Standard errors, in parentheses, are clustered by county.*

For personal loans, we also find a significant reduction in the interest rate, see the fourth column. In contrast to credit card debt, however, the timing of the effects is reversed and ranges from a reduction of 49 bps in the first year to a decline of 33 bps in the third year. We also find no evidence for larger interest rate reductions in communities with a higher proportion of newly eligibles. Unlike credit cards, which are widely used by prime borrowers, personal loans are a smaller market which largely focuses on subprime customers. 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 effects.

Panel B of Table 2 shows the effect of the Medicaid expansion on the imputed interest rates for automobile loans and mortgages. Consistent with the evidence on credit scores, we observe a statistically significant reduction in the rates for automobile loans, see the first column. The effects on automobile loans are more pronounced in counties with a larger fraction of newly-eligible Medicaid beneficiaries and grow in magnitude over time. For mortgages, we find a significant reduction in interest rates in the third post-reform year, see the fourth column. The pooled post-reform effect is negative but not statistically significant. In contrast to our findings on automobile loans, we find no evidence for larger mortgage interest rate reductions in communities with a higher proportion of newly eligibles. This is not entirely surprising, as counties with a larger fraction of newly eligible beneficiaries have fewer homeowners and a larger fraction of individuals who would never qualify for a mortgage.

Robustness: As with 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 Medical Bills and Consumer Welfare in Practice

Finally, we return to the conceptual framework outlined in Section 2 to quantify the effects of a mean and variance reduction in paid and unpaid medical bills on consumer welfare. To implement the proposed revealed preference approach, we must quantify the parameters of individuals' consumption utility, their annual income, the distribution of newly-accrued medical bills, and their (non)-repayment. We must also account for charity care in our treatment of (non)-repayment, as not all unpaid medical bills are sent to collections.

For consumption utility, we follow Finkelstein, Hendren and Luttmer (2019) and assume a baseline coefficient of relative risk aversion of $\sigma = 3$. Following Miller et al. (2018), we consider an annual income $Y = \$4,400$ (39% of FPL), which corresponds to the average income in our baseline MEPS sample. To account for the role of charity care, we use our findings from Section 5.1 which suggest that 40% (60%) of unpaid bills are sent to third party collections (treated as charity by the provider). We further assume that the rate of charity care ($\theta = 60\%$) does not vary by the size of the medical bill. For example, if a fraction

$\tau(\epsilon_{MB})$ of a new medical bill is paid out-of-pocket, then fraction $(1-\theta)\cdot(1-\tau(\epsilon_{MB}))$ of newly accrued medical bills is sold to third party collections.

We estimate the distribution of annually accrued medical bills (F), and individuals' (non)-repayment decisions (b^*), using data from the MEPS. Given the modest sample size, we discretize the support of the annual medical bill amount into 5 bins as outlined at the top of Table 3: $\leq \$2.5k, \$2.5k - \$5k, \$5k - \$10k, \$10k - \$20k, > \$20k$.

Table 3: Revealed Preference Calculation

	Medical Bill Category (m)				
	<\$2.5k	\$2.5k - \$5k	\$5k - \$10k	\$10k - \$20k	>\$20k
	(1)	(2)	(3)	(4)	(5)
Probability mass function: $f(\epsilon_{MB})$	0.8358	0.0571	0.0463	0.0349	0.0359
Average Bill Amount: $\bar{\epsilon}_{MB}$	273.56	3,621.70	7,030.97	14,351.20	63,354.30
Average OOP Spending: $O\bar{O}P$	218.99	981.93	1,839.30	2,174.84	2,304.95
Average OOP Share: $\bar{\tau} = \frac{O\bar{O}P}{\bar{\epsilon}_{MB}}$	0.7976	0.2711	0.2616	0.1515	0.0364
Average New Medical Debt: $\bar{b} = (1-\bar{\tau})\cdot(1-\theta)\cdot\bar{\epsilon}_{MB}$	22.15	1,055.91	2,076.67	4,870.54	24,419.74
Marginal Dis-utility $\left[\frac{h'(\bar{b})}{g'(Y)}\right]$	1.16	2.13	5.07	7.73	9.26
Medical Debt Category	<\$0.2k	\$0.2k - \$1.46k	\$1.46k - \$2.41k	\$2.41k - \$5.8k	>\$5.8k
Dis-utility of Unpaid Bills $\left[\frac{h(b)}{g'(Y)\times 10k}\right]$	0.003	0.206	0.760	2.832	20.808
Baseline Income (Y)	4,400	4,400	4,400	4,400	4,400
Fraction of Charity Care in Uncompensated Care (θ)	0.60	0.60	0.60	0.60	0.60
Relative Risk Aversion (σ)	3	3	3	3	3

Notes: *The first four rows show the distribution of new medical bills and the out-of-pocket payments for uninsured individuals in the MEPS (Section 4). Columns 1–5 show the outcomes separately for different cumulative annual medical bill bins. The fifth and the seventh row combine the bill amount with the out-of-pocket payments and our baseline estimate of the role of charity care (θ), to quantify the average and the range of new medical debt amounts in the relevant medical bill bin. The sixth and the eighth row present estimates of the dis-utility of leaving bills unpaid, $h(\cdot)$ and $h'(\cdot)$, for each medical bill bin, given our assumptions on baseline income (Y), the importance of charity care (θ), and risk preference parameter (σ).*

The first row depicts the probability mass function for annual cumulative medical bills. About 84% of the uninsured individuals accrue medical bills of less than \$2,500. Among these individuals, the average medical bill amount is \$274, see the second row. About 80% or \$219 of the overall bill amount is paid out-of-pocket, see the third and the fourth row, respectively. For larger bill amounts (columns 2-5), we see a gradual decline in the share paid out-of-pocket. At the same time, we observe an increase in annual out-of-pocket spending reaching \$2,175 for cumulative bills worth $\$10k-\$20k$. This corresponds to about 15% of the underlying medical bill amount. Interestingly, we find that overall out-of-pocket spending increases only to \$2,300 as we consider bill amounts exceeding \$20k. This is consistent with a potentially binding liquidity constraint for individuals with annual incomes of only \$4,400.

Leveraging the observed repayment decisions, we apply the first order condition (equation 3) to recover the marginal dis-utility of medical debt. For example, consider individuals in the first category described in the first column of Table 3. At their average bill ($274), the first order condition implies that

$$\begin{aligned} g'(c) &= (Y - \bar{\tau}(\bar{\epsilon}_{MB}) \cdot \bar{\epsilon}_{MB})^{-\sigma} = (4400 - 0.8 \cdot 274)^{-3} \\ &= h'((1-\theta) \cdot (1 - \bar{\tau}(\bar{\epsilon}_{MB})) \cdot \bar{\epsilon}_{MB}) = h'(0.4 \cdot 0.2 \cdot 274) \ . \end{aligned}$$

Hence, combining the repayment decision with the first order condition allows us to infer $h'(x)$ at $x = 0.4 \cdot 0.2 \cdot \$274 = \$22$ from $g'(c)$ at $c = \$4,400 - 0.8 \cdot \$274 = \$4,181$. To facilitate the interpretation of the marginal dis-utility of medical debt, we divide $h'(\cdot)$ by the marginal utility of consumption allowing us to express the dis-utility in dollars. We find $h'(22)/g'(Y) = \$1.16$, see the sixth row, which is only slightly larger than the marginal utility of consumption at $c = Y$. Repeating this exercise for the different medical bill bins, we find that the marginal dis-utility increases in the underlying medical debt amount. For medical amounts exceeding $24k we find a marginal dis-utility of more than $9. Consistent with our assumptions in Section 2, this points to a convex dis-utility function over medical debt in collection.

Finally, we take the integral of $h'(\cdot)$ to recover the dis-utility of medical debt $h(x)$. To this end, we assume that $h'(\cdot)$ is constant within the respective medical bill bins that correspond to analogous medical debt bins. Specifically, we assume that the upper bound of the corresponding medical debt interval m, $b_{UB(m)}$, is given by the medical debt formula, outlined in the fifth row of Table 3, when evaluated at the corresponding upper bound on medical bills , $\epsilon_{UB(m)}$:

$$b_{UB(m)} = (1 - \bar{\tau}(m)) \cdot (1 - \theta) \cdot \epsilon_{UB(m)} \ .$$

The corresponding intervals, displayed in the seventh row, are then simply $[b_{UB(m-1)}, b_{UB(m)}]$ with $b_{UB(0)} = 0$ and $b_{UB(5)} = \infty$. Finally, taking the integral over the step-function $h'(\cdot)$, we have

$$h(x) = \sum_{m<k} \Big(h'(m) \cdot (b_{UB(m)} - b_{UB(m-1)}) \Big) + h'(k) \cdot (x - b_{UB(k-1)}) \ if \ b_{UB(k-1)} < x \leq b_{UB(k)} \ .$$

Of course, the role of the interval bounds decreases as we increase the number medical bill categories and decrease the length of the medical bill interval.[20] The eighth row of Table 3 presents the normalized dis-utility estimates at the corresponding medical debt amount in \$10,000. Building on the estimated probability mass function, the estimates imply an expected dis-utility over medical debt in collection of up to \$8,900. We note that this estimate is based on the marginal utility of consumption at $c = Y$. However, the marginal utility of consumption may increase substantially when $c < Y$ suggesting that the implications for consumer welfare are less drastic, which we turn to next.

Specifically, we recover the transfer (γ^T) and risk protection (γ^{RP}) components of insurance from equations 4, 5, and 6. As shown in the first column of Table 4, accounting for the dis-utility of leaving medical bills unpaid, the indirect credit channel, we calculate a transfer component of $\gamma^T = \$1,532$. This estimate exceeds the out-of-pocket benchmark, which captures expected out-of-pocket spending, by a factor of 3.2. We calculate a risk-protection component $\gamma^{RP} = \$965$, which exceeds the out-of-pocket benchmark by a factor of 3.8. Together we recover an overall value $\gamma = \gamma^T + \gamma^{RP} = \$2,498$. This exceeds the out-of-pocket benchmark by a factor of 3.4. However, it falls short of overall utilization by 30% as indicated in the last row.

6.1 Robustness

Baseline income: Our baseline calculation builds on the observed repayment decisions of the uninsured population in the MEPS data earning less than 80% of the FPL. For our first robustness check, we extend our sample population to the uninsured who earn less than 138% of the FPL, the expansion eligibility cutoff. The average income for these individuals is \$8,400, or 73% of the FPL. As shown in Column 2 of Table 4, we find qualitatively similar estimates for this broader population. The consumer welfare gains from insurance are slightly smaller because incomes are higher and expected out-of-pocket spending falls from \$480 to \$440. This effect is offset by a smaller overall spending estimate.

Relative Risk Aversion: In columns 3 and 4, we consider different parameters of relative risk aversion. A smaller parameter of relative risk aversion lowers the risk-protection component of health insurance but disproportionately so under the out-of-pocket benchmark. Put differently, the ratio of the risk protection components falls from 5.3 for $\sigma = 2$ to 2.8 under $\sigma = 4$. Similarly, the ratio of the overall value of insurance falls to 3 when $\sigma = 4$. This

[20] In the limit, as the number of medical bill categories increases to infinity and the length of the medical bill interval converges to zero, the length of the medical debt interval also converges to zero and $h'(\cdot)$ converges to a continuously differentiable function.

Table 4: Overall Annual Financial Benefits

	Baseline	$\leq$ 138% FPL Y=$8,400	$\sigma = 2$	$\sigma = 4$	$\theta = 0.4$	$\theta = 0.8$	$\epsilon \leq$ $4,400	Interest Savings
	(1)	(2)	(3)	(3)	(5)	(6)	(7)	(8)
Transfer Component: γ^T								
Full Model	1,532	1,487	1,574	1,499	1,817	1,127	827	904
OOP Benchmark	480	440	480	480	480	480	480	480
Ratio: γ^T/γ^T_{OOP}	3.2	3.4	3.3	3.1	3.8	2.3	1.7	1.9
Risk-Protection Component: γ^{RP}								
Full Model	965	833	789	1,060	955	885	608	965
OOP Benchmark	254	190	147	380	254	254	180	254
Ratio: $\gamma^{RP}/\gamma^{RP}_{OOP}$	3.8	4.4	5.3	2.8	3.8	3.5	3.4	3.8
Total Benefit: γ	2,498	2,319	2,362	2,559	2,772	2,012	1,435	1,869
Total Utilization	3,531	3,223	3,531	3,531	3,531	3,531	3,531	3,531
Ratio: $\gamma/Spending$	0.7	0.7	0.7	0.7	0.8	0.6	0.4	0.5

Notes: *Column 1 summarizes the consumer welfare gains in our baseline specification. This approach considers a parameter of relative risk aversion of $\sigma = 3$, an annual income $Y = \$4,400$, and that fraction $\theta = 0.6$ of unpaid bills comprises charity care. Column 2 revisits the estimates in the sub-sample of individuals who earn less than 138% of the FPL who earn an average income of $Y = \$8,400$. Columns 3-4 consider different parameters of relative risk aversion. Columns 5-6 consider different shares of charity care and column 7 considers the case when charity covers every dollar on a medical bill exceeding $4,400. Column 8 summarizes the consumer welfare gains when considering the potential benefits from lower interest rates.*

relationship is, however, not monotone. For example, we find a ratio of 3.5 for $\sigma = 0.01$ as individuals are close to risk-neutral, which falls short of the ratio of 3.8 observed at $\sigma = 2$. Overall, we conclude that the value of insurance more than doubles when considering the indirect credit channel under plausible parameters of relative risk aversion.

Charity Care: In columns 5-7 we consider different degrees of charity care. As mentioned in Data Section 4, potential underreporting of collection activities to credit bureaus may bias our measurement of medical debt in collection downward. Hence, this may lead us to overstate the role of charity care. In the fifth column, we reduce the share of unpaid bills that is covered by charity care from 60% to 40%. As a result we find a slightly larger value of health insurance. The financial benefit per Medicaid dollar increases from \$0.7 to \$0.8.

Conversely, if the ACA restrictions on billing and collection activities, see again Data Section 4, disproportionately affect not-for-profit hospitals located in expansion states, then we might overstate the reduction in medical debt in collection and consequently understate the role of charity care. In Column 6, we increase the share of charity care to 80%. Now we find a slightly smaller value of health insurance. The financial benefit per Medicaid dollar decreases to \$0.6. Overall, we conclude that our financial benefit estimates are quite robust to varying the share of charity care in uncompensated care between 40% and 80%.

This may change, however, when the role of charity care increases in the underlying billing amount. This increases the implicit value of insurance provided by charity care. In column 7 we consider a drastic scenario in which charity care covers any additional dollar exceeding the individuals income of \$4,400. That means that there is no charity care for small bills below \$2,500. For larger bills, $\geq \$4,400$, medical debt is reduced by the difference between \$4,400 and out-of-pocket payments, as the residual is covered by charity care. In this extreme scenario, we find a substantially smaller risk-protection component and a smaller overall value of $\gamma = \$1,435$. We note however that this scenario is inconsistent with our empirical evidence that points to disproportionately larger reduction in larger medical debt amounts, see Figure A.3. This suggests that the charity care share may even be negatively correlated with the underlying billing amount. Regardless, even in this exercise we still find that the value of insurance still exceeds the out-of-pocket benchmark by a factor of 2.2.

Interest Savings: Finally, we revisit the financial benefits accruing through the indirect credit channel using the estimated interest rate reductions displayed in Table 2. Specifically, we provide a ballpark estimate for the potential interest savings using a simplistic back-of-the-envelope calculation. We focus our calculation on individuals living in treatment states and simply scale the outstanding debt amounts in the pre-reform period, see Table 2, by our estimated average interest rate reduction over the three post-reform years ($Exp \times Post$). This calculation assumes that the average debt amounts remain unchanged over time and that consumers can fully refinance existing debt as well as new debt at the improved terms. We return to these assumptions below. Cross-multiplying the debt amounts (for credit cards, personal loans, auto loans, and mortgages) by the average interest reductions, we find annual interest savings of $\$4,323 \cdot 0.42\% + \$614 \cdot 0.5\% + \$4,845 \cdot 0.04\% + \$56,459 \cdot 0.001\%$=\$23.73 per person.

We consider two refinements to this coarse calculation. First, a portion of reported credit card balances is held within the 'grace' period. This portion does not incur any finance (interest) charges and is not considered debt.[21] Unfortunately, the CCP does not allow us to distinguish between balances that are debt from balances that are not. However, using a large panel of credit card accounts, Grodzicki and Koulayev (2019) find that 82% of credit card balances are not covered by a grace period and are charged interest.[22] As such, we scale the average pre-reform credit card balance by 82%. Second, some accounts become delinquent and borrowers may stop paying interest on these accounts. This is particularly relevant for credit cards and personal loans, which are not secured by an underlying asset, such as a car or a house. To provide a more conservative estimate on interest savings, we net out the fraction of credit card and personal loan accounts that have become delinquent. Based on the 'Quarterly Report on Household Debt and Credit' from the Federal Reserve Bank of New York, we assume that about 10% of credit card and personal loan debt is more than 90 days delinquent.[23] Hence, we scale credit card and personal loan debt by 1-10%=90% in our calculation.[24]

Accounting for these adjustments, we find slightly smaller annual interest savings of \$18.66 per person, which is about 34% of the annualized per-person reduction in medical

[21] Individuals 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.

[22] This fraction increases to 96% among subprime borrowers that are more likely to represent our low-income treatment population.

[23] We assume that the personal loan delinquency rate, which is not separately reported, equals the credit card delinquency rate, see https://nyfed.org/35BhgBC for details, last accessed January 7, 2020.

[24] This calculation abstracts away from changes in delinquency rates. While there are overall reductions in the 90-day delinquency rates on any credit account, see Table 1, the evidence on credit cards and personal loans is not conclusive. For credit cards, we estimate a pooled effect of 0.00045 and -0.00042 in counties whose share of newly-eligible adults falls below and above the sample median, respectively. Hence, the average effect of 0.00045-0.00042=0.00003 is very small relative to the baseline share of new delinquencies and if anything positive, see Table C.1 for details. The delinquency effects for personal loans are statistically insignificant, see also Table C.1 for details.

debt (Table 1).[25] Once again, we divide this intent-to-treat estimate by the fraction of non-elderly treated adults and find interest savings of $\frac{\$18.66}{0.044}$=\$424 per treated person and year. Savings on credit card debt dominate the total effect, accounting for 72% of the total savings. A potential explanation for the large credit card savings is given by the fact that these products are unsecured and easily discharged, and as such their pricing is more sensitive to changes in underlying credit worthiness. Combining the estimated interest savings with the reduction in out-of-pocket spending, we calculate a transfer component of \$904, which exceeds the out-of-pocket reduction by 88%, see column 8 of Table 4. Unfortunately, our interest savings calculation does not yield an estimate for the risk-protection component of insurance. Therefore, we borrow the corresponding estimates from our baseline approach to calculate an overall annual financial benefit of \$1,869. This estimate exceeds the out-of-pocket benchmark by a factor of 2.5.

Our stylized calculation of interest savings is subject to several important caveats. We illustrate these limitations in the context of credit card debt, which is quantitatively most important for our calculation. First, we assume that consumers finance their overall debt at the improved terms of credit. At the same time, a growing body of evidence suggests that at least some consumers make suboptimal credit card borrowing and repayment decisions. For example, Ponce, Seira and Zamarripa (2017) find that Mexican borrowers misallocate 24% of their debt by borrowing on high-interest credit cards. This suggests that our baseline estimates overstate the actual interest savings particularly so in the short run. Second, we only observe interest rates on credit card mail offers, which may of course differ from the actual interest rates on the consumer's credit accounts. For example, the actual interest rate on a trade depends the consumer's credit score and market conditions at the time of origination, which we do not take into account in this calculation, and may be quite different from the timing of the mail offers. To alleviate this concern, at least partly, our calculation uses the estimated average rate reduction over the three post-reform years, which may mitigate the role of differences in timing between loan take-up and mail offers. Furthermore, we note that our analysis focuses on pre-screened mail-out offers, which means that the consumers will most likely not be rejected by the offering bank should they decide take that offer.[26] Third, we note that we do not take the value of improved access to credit markets into account.

[25] Specifically, we calculate the interest savings per person as follows: $90\% \cdot 82\% \cdot \$4,323 \cdot 0.42\% + 90\% \cdot \$614 \cdot 0.5\% + \$4,834 \cdot 0.04\% + \$56,459 \cdot 0.001\%$=\$18.66.

[26] To alleviate this concern further, we note that the average interest rate on offered credit card debt in our sample, 15.4% see Table 2, 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).

In contrast to the former two caveats, this may lead us to understate the consumer welfare gains.

Given these limitations, we interpret the calculated interest savings as a broader ballpark estimate regarding the significance of the indirect credit channel. Consistent with our baseline results, we find that the financial benefits roughly double when taking the indirect credit channel into account.

6.2 Discussion

Overall, our findings suggest that incorporating the implicit costs of leaving medical bills unpaid may raise the financial benefits from health insurance substantially. To relate our estimated benefits more directly to the estimates from the literature, we divide the total benefit by the gross cost of Medicaid per beneficiary. To facilitate the comparison, we follow Finkelstein, Hendren and Luttmer (2019) and assume an annual gross cost per beneficiary of $G = \$3,600$. This estimate only considers the medical expenditures and abstracts away from administrative costs.

Across the different approaches outlined in Table 4, we find that our estimates on the financial benefit of insurance, γ, fall short of G. Our estimates of the ratio γ/G range between 0.4 and 0.8. For comparison, the estimated analogues ratio in Finkelstein, Hendren and Luttmer (2019) ranges between 0.2 and 0.5. Hence, by including the costs of leaving medical bills unpaid our estimates exceed these by roughly a factor of 2 even though we do not take potential patient benefits from increased health care utilization into account.

Of course, all our results are sensitive to the assumptions underlying the specific approaches. While we do consider a variety of robustness exercises, we note that there remain several strong modeling assumptions that we do not relax explicitly. Perhaps most significantly, we consider a static model even though the credit consequences of unpaid medical bills have clearly dynamic implications on future terms and access to credit. Incorporating the dynamic tradeoffs into the analysis would be an interesting subject for future research. In addition, we assume that the financial benefits from reduced medical debt in collection are entirely borne by new beneficiaries themselves. Alternatively, some of these benefits may accrue to non-beneficiaries through spillover effects if, for instance, health care providers reduce the amount of bills sent to collection agencies more generally. Therefore, this approach might overstate the benefits accrued by the beneficiary. That being said, recent work by Miller et al. (2018) documents large reductions in medical debt among newly-eligible Medicaid beneficiaries indicating that at least a large fraction of our estimated benefits accrue to beneficiaries.

7 Conclusion

Uninsured individuals pay on average only up to 20% of overall health care utilization out-of-pocket. If the residual 80% 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 the 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 better credit terms. Using a novel conceptual framework, we find that the financial benefits of Medicaid double when accounting for this indirect 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.

Argyle, Bronson, Taylor Nadauld, and Christopher Palmer. 2017. "Real Effects of Search Frictions in Consumer Credit Markets."

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."

Avery, Robert B, Paul S Calem, Glenn B Canner, and Raphael W Bostic. 2003. "An overview of consumer data and credit reporting." *Fed. Res. Bull.*, 89: 47.

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. 2019. "Estimating the Value of Public Insurance Using Complementary Private Insurance." 11(3): 88–129.

Caswell, Kyle J, and Timothy A Waidmann. 2017. "The Affordable Care Act Medicaid Expansions and Personal Finance." *Medical Care Research and Review*, 1077558717725164.

Centers for Medicare and Medicaid Services. 2011. "National Health Expenditure Projections 2012–2022: Centers for Medicare and Medicaid Services."

Coughlin, Teresa A, John Holahan, Kyle Caswell, and Megan McGrath. 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.

Dranove, David, Craig Garthwaite, and Christopher Ody. 2016. "Uncompensated Care Decreased at Hospitals in Medicaid Expansion States but not at Hospitals in Non-expansion 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. 2019. "The value of medicaid: Interpreting results from the oregon health insurance experiment." *Journal of Political Economy*, 127(6): 2836–2874.

Finkelstein, Amy, Nathaniel Hendren, and Mark Shepard. 2019. "Subsidizing health insurance for low-income adults: Evidence from Massachusetts." *American Economic Review*, 109(4): 1530–67.

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. 2017. "Premium subsidies, the mandate, and Medicaid expansion: Coverage effects of the Affordable Care Act." *Journal of Health Economics*, 53: 72–86.

Gallagher, Emily A, Radhakrishnan Gopalan, and Michal Grinstein-Weiss. 2019. "The effect of health insurance on home payment delinquency: Evidence from ACA Marketplace subsidies." *Journal of Public Economics*, 172: 67–83.

Garthwaite, Craig, Tal Gross, and Matthew J Notowidigdo. 2018. "Hospitals as insurers of last resort." *American Economic Journal: Applied Economics*, 10(1): 1–39.

Grodzicki, Daniel, and Sergei Koulayev. 2019. "Sustained Credit Card Borrowing." *Available at SSRN 3403045*.

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, Luojia, Robert Kaestner, Bhashkar Mazumder, Sarah Miller, and Ashley Wong. 2018. "The effect of the affordable care act Medicaid expansions on financial wellbeing." *Journal of public economics*, 163: 99–112.

Kaiser Family Foundation. 2012. "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.

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*.

Mazumder, Bhashkar, and Sarah Miller. 2016. "The Effects of the Massachusetts Health Reform on Household Financial Distress." *American Economic Journal: Economic Policy*, 8(3): 284–313.

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

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 onfFinancial wellbeing." *Working Paper*.

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 section, we corroborate our baseline difference-in-differences estimation strategy in two ways. First, we present data trends comparing states adopting the policy on January 1^{st} 2014, treatment states, to those states that chose not to expand Medicaid by the end of 2016, control states. Second, we illustrate the impact of the reform in a regression context, tracing year-quarter trends in our difference-in-differences estimates. The appendix section 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 on Medical Debt in Collection

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

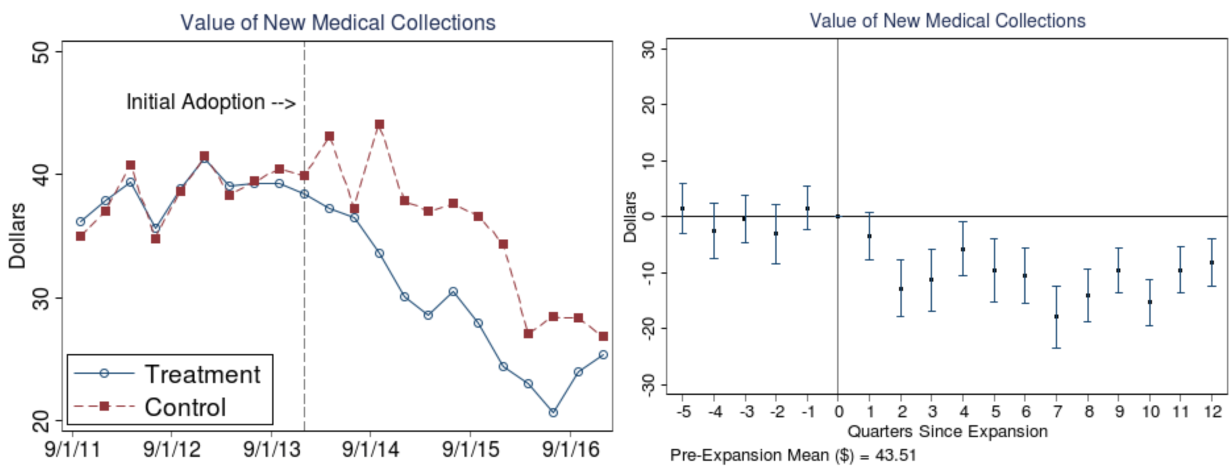

Figure A.1: Trends in Newly Accrued Medical Collections

Notes: *The left panel shows trends in the value of newly accrued medical collections. The right panel shows regression coefficients from equation A.1. Outcomes are weighted by the synthetic control weights. Vertical lines highlight the initial implementation date of the expansion - January* 1^{st}*, 2014. Standard errors are clustered at the census tract level.*

treatment and control states 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 difference across these pre-treatment time series. Nevertheless, over the three post-reform years, we see a gap in the average value of newly accrued medical debt between treatment and control states of about $10 per quarter, which is consistent with the evidence from Table 1.

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

$$y_{ct}^{k} = \alpha_{c}^{k} + \eta_{t}^{k} + \sum_{r=S}^{r=-1} \beta_{r}^{k} + \sum_{r=1}^{r=F} \beta_{r}^{k} + \epsilon_{ct}^{k} \ . \tag{A.1}$$

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}^{k}, which denote the coefficients on a series of indicator variables for the time relative to the expansion date, measured in quarters. 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 right panel of Figure A.1, we find no systematic differences, or pre-trends, between treatment and control states prior to the reform. Congruent with the evidence from the left graph, the figure also illustrates a sharp decline of medical collections in treatment states following the reform. In the post-reform period, the coefficients fluctuate around -$10, which is again consistent with the evidence from Table 1.

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. 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.[27] The left graph of Figure A.2 plots trends in medical collections for states with only Federal exchanges, which does 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.

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 greatly between treatment and control states following the reform. As evidenced by the right graph of Figure A.2, we find no

[27]These 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

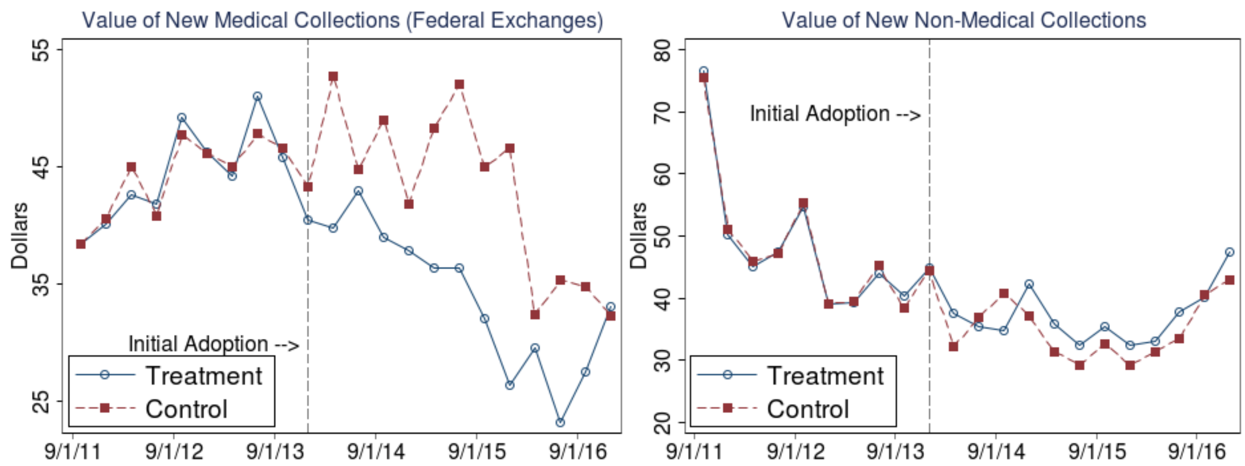

Figure A.2: Federal Exchanges & Non-Medical Collections

Notes: *The left figure shows quarterly means of newly accrued medical collections in states operating a federal platform for private exchanges. The right figure shows quarterly means of newly accrued non-medical collections for all states in our sample. Outcomes are weighted using the synthetic control weights. Vertical lines highlight the initial implementation date of the expansion - January* 1^{st}*, 2014.*

evidence for differential changes in non-medical collections in adopting states in the post-reform period. 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.1.1 Aggregate Savings and Repayment Decisions

In this subsection, we present estimates on the aggregate reduction in medical debt in collection as well as repayment decisions and the associated consumer savings. Column 1 in the top panel of Table A.1 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.

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

Table A.1: Reduction and Repayment of Medical Debt

	All	1^{st} *Quartile*	2^{nd} *Quartile*	3^{rd} *Quartile*	4^{th} *Quartile*
	(1)	(2)	(3)	(4)	(5)
Annual Aggregate Decrease in Accrued Medical Collections ($ Millions)					
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
Proportion of New Medical Collections Repaid (p.p)					
After One Year					
Repaid	7.53	9.67	8.79	8.16	5.46
Repaid or Removed	38	35.85	37.34	38.55	38.69
After Two Years					
Repaid	8.57	10.59	10	9.36	6.3
Repaid or Removed	53.02	50.22	52.53	53.87	53.58
Lower Bound of Medical Collection Repayments ($ Millions)					
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 annual reduction in medical debt are calculated by multiplying estimates from Table 1 by the CCP population in respective census tracts during each respective year-quarter. Repayment rates are calculated directly from the data. Percent repaid is the proportion of new medical collections accrued in a given year-quarter 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 and year-quarter by 48, the inverse of the population sampling rate (Section 4)*

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 A.1, 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 A.1 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.[28]

A.1.2 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.[29] 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 post-reform period.

The top panel of Figure A.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 census tract and year-quarter. 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

[28] Of 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. 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 6.

[29] We 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.

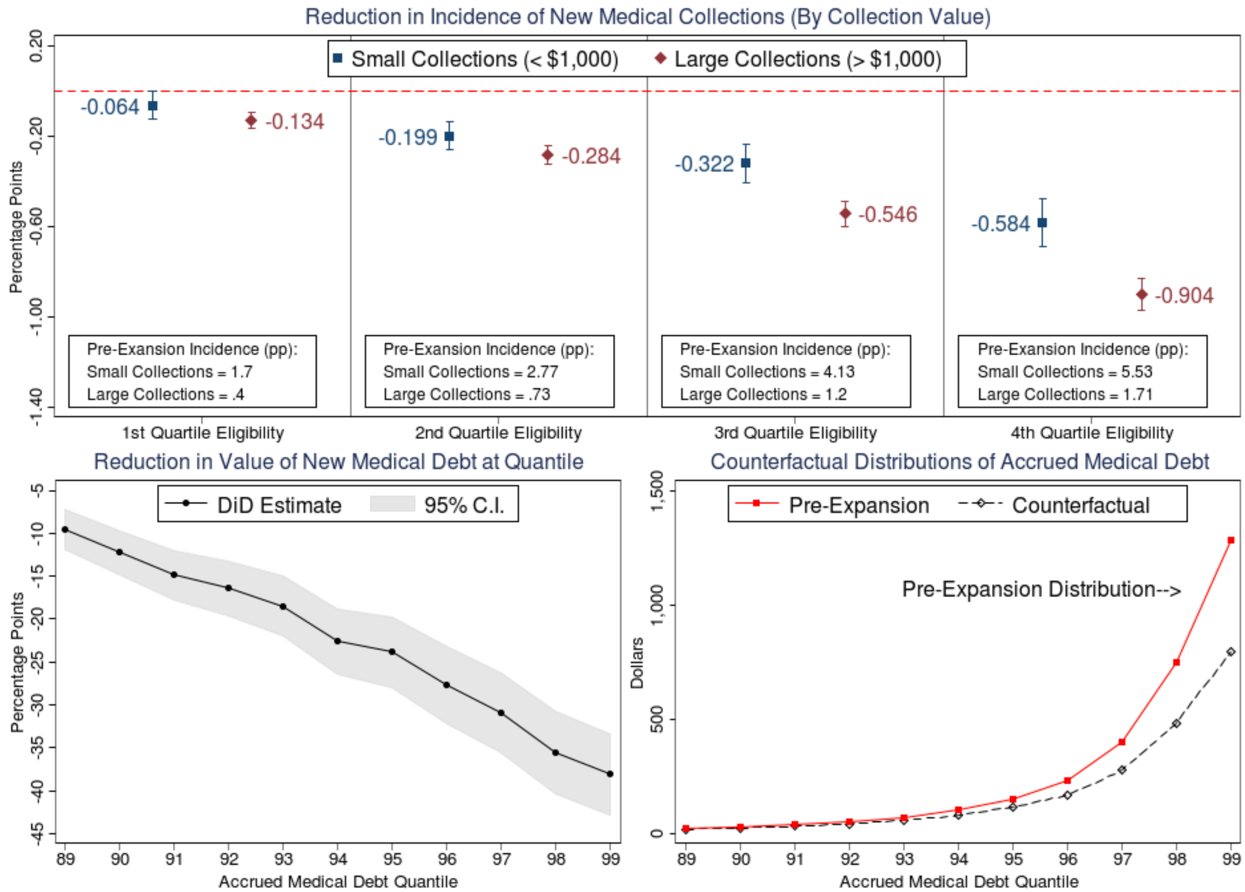

Figure A.3: Distributional Effects of Expansion on Medical Debt

Notes: *The figure shows distributional effects of the reform on the accrual of medical debt. The top panel show treatment effects and 95% confidence intervals on the incidence of large ($\geq \$1,000$) and 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 at the census tract and year-quarter level. Regressions include census tract and quarter-year fixed effects and are weighted by the synthetic control weight. Standard errors (in parentheses) are clustered by census tract.*

eligibility, the incidence of a large medical collection declines by approximately 34%, as compared to 4% for a small collection. In communities with high rates of newly eligible adults, the incidence of small debts declines by about 11%. 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 A.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 in the distribution of newly-accrued medical debt at the census tract year-quarter level. 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% at the 89^{th} quantile to 40% at the 99^{th} quantile. We transform these relative effects into levels in the bottom right panel of Figure A.3. The estimated 10% reduction at the 89^{th} quantile, on a base of about \$20, translates to a modest savings of about \$2. However, the savings become quite substantial past the 95^{th} percentile. For the highest quantile, a 40% 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.

A.2 Indirect Effect

A.2.1 Medical Debt and Financial Distress: An Event Study Approach

We start exploring the relationship between medical debt in collection and measures of financial distress using an event study approach. To this end, we focus on the sub-sample of individuals who receive their first medical collection valued at more than \$100 during our period of analysis. We explore the precise timing of medical collection and follow each of these individuals from six quarters prior to receiving the collection to two years after. Our specification is as follows:

$$y_{ict}^k = \alpha_c^k + \eta_t^k + \sum_{r=-6}^{r=-1} \beta_r^k + \sum_{r=1}^{r=8} \beta_r^k + \epsilon_{ict}^k \tag{A.2}$$

where y_{ict}^k denotes the respective outcome k for record i in census tract c during year-quarter t. We control for census tract (α_c^k) and calendar year-quarter (η_t^k) fixed effects. The key parameters of interest are β_r^k, which denote the coefficients on a series of indicator variables for the time relative to the event, the accrual of the first medical collection. Consequently, we normalize the end of the quarter just prior to a collection being placed on the account to zero.

Figure A.4 plots the estimated β_r^k coefficients along with their 95% confidence intervals for medical debt balances (left panel), serious delinquencies (middle panel), and credit scores (right panel). As a proof of concept, we see that following the accrual of a new medical collection, individual medical debt balances rise sharply and persist for at least two years (left panel). The evidence from the middle panel suggests that a rise in medical debt is followed by movement into serious delinquency. Most importantly for our analysis, we observe a sharp

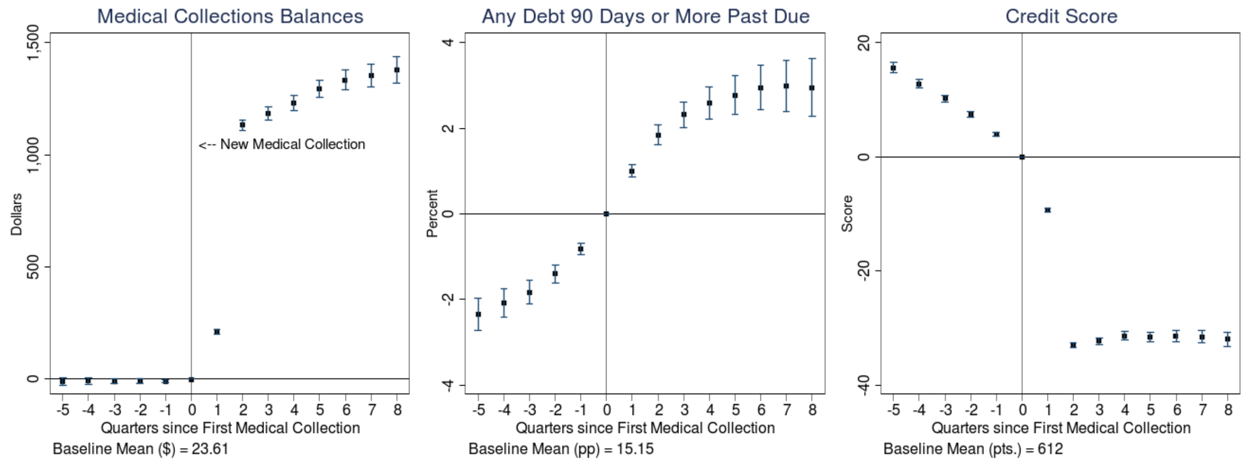

Figure A.4: Event Study: Medical Collections, Credit Worthiness, and Credit Scores

Notes: *The figure shows medical debt, 90-day delinquency, and the credit score from 5 quarters before to 8 quarters after receiving a consumer's 'first' medical collection using equation A.2. Effects are normalized to the end of the quarter just prior to the collection. Standard errors clustered at the Census tract level.*

and persistent drop in credit scores following the accrual of a new medical collection, see the right panel.[30]

A.2.2 Financial Distress: Delinquencies and Credit Scores

Next we turn to the reform effects on measures of financial distress, our *indirect effect*. Figure A.5 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 graph shows trends in the credit score. While the trends for both groups are similar during the pre-expansion period, delinquency rates trend notably lower after the expansion in 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.6 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, and congruent with the above graphical evidence and the regression results in Table 1, these effects grow between the first and second post-expansion year and are seemingly long lasting.

[30] Though credit scores begin to fall prior to the collection, likely due to other borrowing related to the health event and the health event itself, the clearest and sharpest drop is just after the collection is reported.

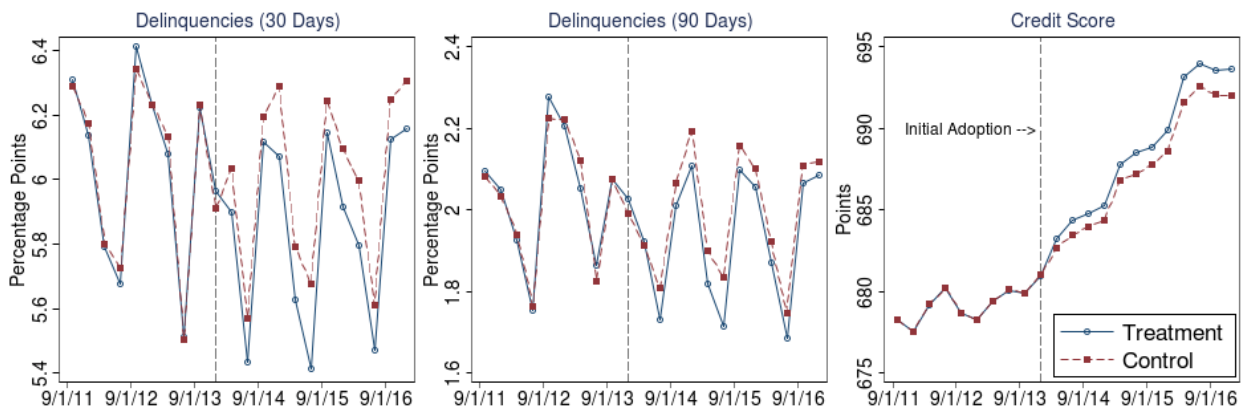

Figure A.5: Trends in Incidence of New Delinquency for Initial Adopting States

Notes: *The figure shows quarterly flows into new delinquency, as defined in Section 4, and credit scores. 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* 1^{st}*, 2014.*

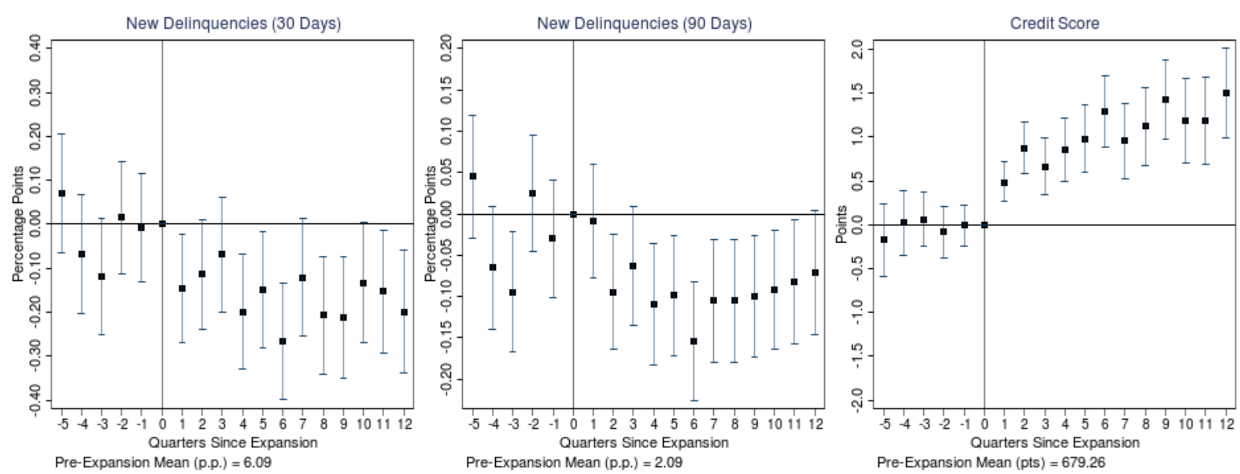

Figure A.6: Indirect Effect of Expansion on Financial Distress: Event Study

Notes: *The figure shows changes in new delinquencies and credit scores using equation A.1, weighted by the synthetic control weights. Confidence intervals in the figure are calculated using standard errors clustered at the census tract level.*

A.2.3 Heterogenous Effects on Credit Scores

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 each census tract and year-quarter as the dependent variable in equation 8. The top panels of Figure A.7 show these effects for the first (left), second (middle) and third (right) post-reform year. The points correspond to an

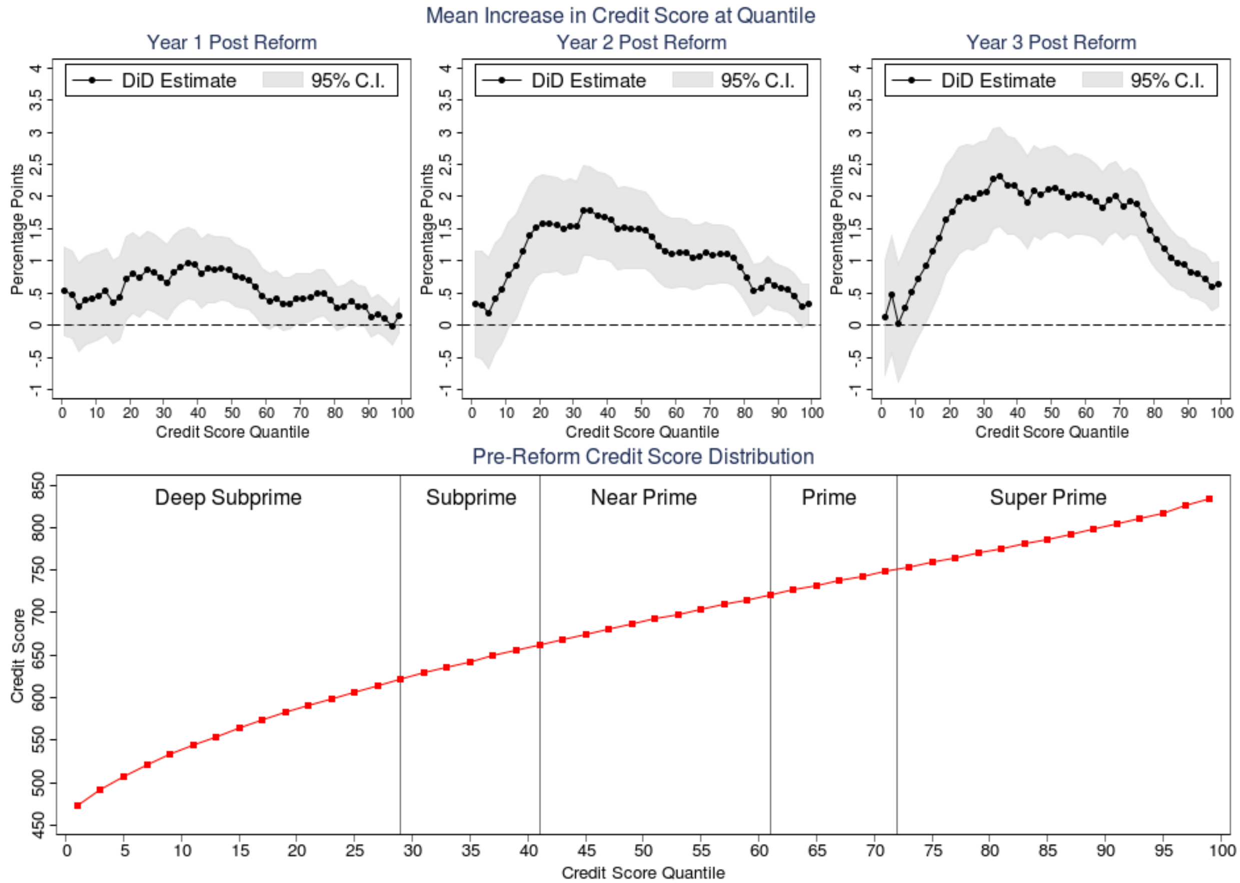

Figure A.7: Distributional Effects of Expansion on Credit Scores

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 credit score quantile for the first, second, and third post-reform year, 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.*

average treatment effect at a given quantile, and the shaded region shows the point-wise 95% confidence interval for the estimate. The bottom panel shows the pre-treatment distribution of credit scores among expanding states.

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.

A.2.4 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.8 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

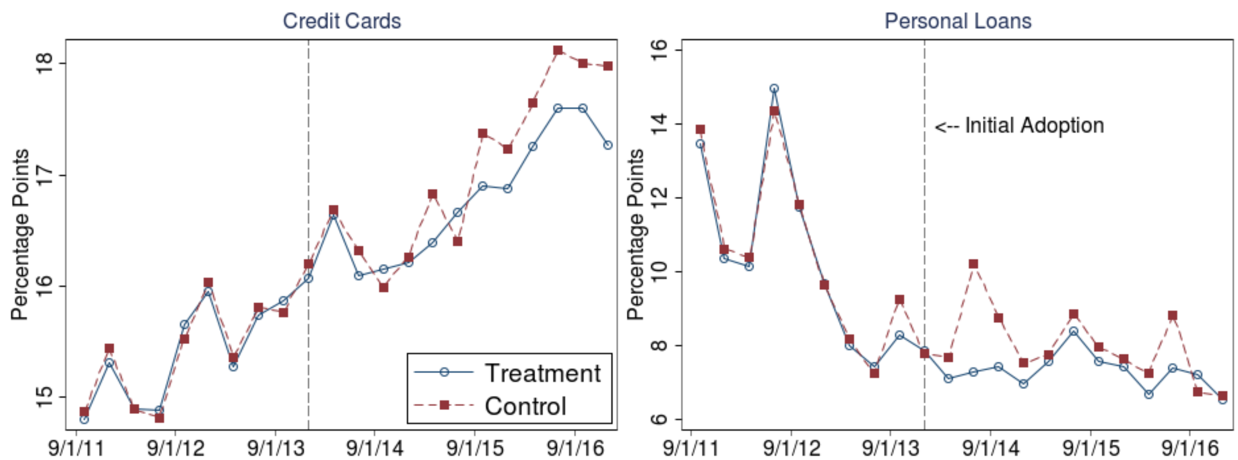

Figure A.8: Trends in Offered Rates for Initial Adopting States

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.*

in average credit card interest rates of around 0.5 to 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 loan offers in the data is much lower than for credit cards (Table 2). This smaller sample size on offers leads to noisier trends. Nevertheless, as shown in the right panel of Figure A.8, offered rates on personal loans seem to decline for recipients in expanding states relative to non-expanding states following the reform.

Figure A.9 further shows this effect within the regression framework, equation A.1. We find similar 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.[31]

[31] A 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

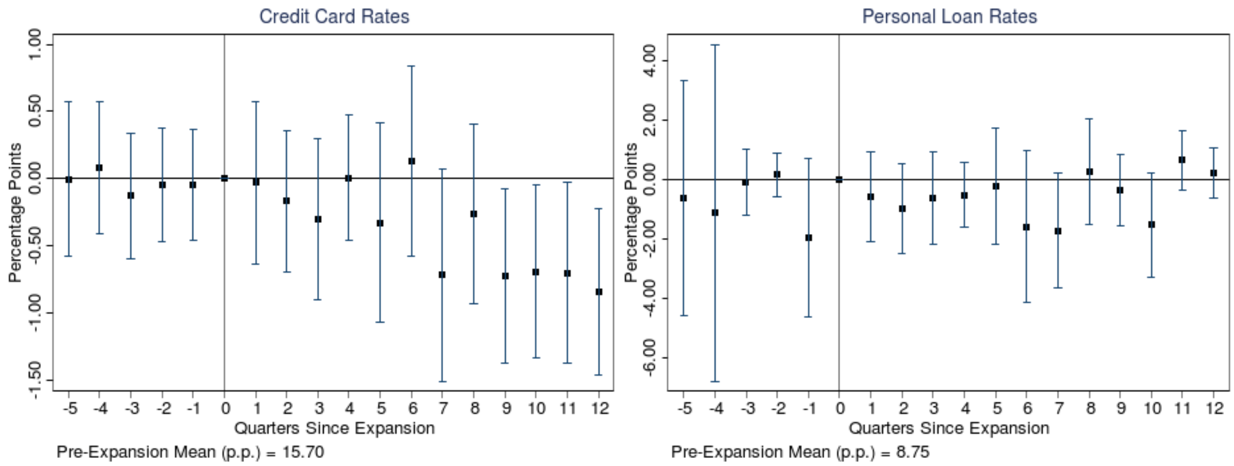

Figure A.9: Indirect Effect of Expansion on Credit Pricing: Event Study

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.*

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 and year-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.

C Details on Calculated Interest Savings

Delinquency: Table C.1 presents the reform effects on the quarterly 90-day delinquency rate for credit cards (top panel) and personal loans (bottom panel). To conform with the analysis in Section 5.3, the regression is run separately for census tracts, whose proportion

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 consumes following the reform.

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	Mortages
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.*

of newly eligible adults falls below and above the median proportion, see the first and the second column, respectively.

Table C.1: 90+ Delinquency For Credit Cards and Personal Loans

	Below Median (1)	Above Median (2)
Credit Cards		
DD Coefficient	0.00045	-0.00042
	(0.0002)	(0.00022)
90 Day Delinquency Rates	0.01101	0.01189
Unsecured Personal Loans		
DD Coefficient	-0.00002	-0.00012
	(0.00005)	(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 5 for details. Standard errors (in parentheses) are clustered by Census tract.*

Rate Sheets: One limitation of our analysis is that we do not see interest rates on mortgages and auto loans directly. To fill this data gap, we leverage the rate sheet information, displayed in Table C.2, which shows the *MyFico* aggregated rate sheets for 5-year auto loans and 30-year fixed rate mortgages as of March 19, 2017. This information allows us to map observable credit scores into an average interest rate by credit score bin (separately for mortgages and auto loans). Specifically for auto loans and mortgages, we assign individuals in the CCP rates they qualify for given their credit score in quarter t. Consumers with credit scores below the bottom price tiers, or without a credit score are excluded from calculations, as they are most likely not eligible for a loan.

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

Auto Loan Pricing Tiers						
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
Mortgages Pricing Tiers						
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*