Hospital Mergers, Health Provider Concentration, and Medical Debt*

Sérgio Pinto† Alaa Abdelfattah‡

Marshall Steinbaum§ Justin C. Wiltshire¶

October 8, 2026

We leverage consumer credit panel data covering 2011–2019 to examine how hospital consolidation impacts individual medical debt. In our preferred event study specifications, exposure to hospital mergers increases the likelihood of having medical debt in collections, the flow of new debt generated over the previous twelve months, and the total outstanding stock. Five years after treatment, the effects on our outcomes of interest range from 11% to 20%. We also find evidence that additional balances are carried on credit cards. In a placebo test, we find no effect of merger exposure on analogous outcomes for non-medical debt in collections. Using comprehensive county-level healthcare spending and utilization data from the Institute for Health Metrics and Evaluation's Disease Expenditure Project, we investigate the mechanisms underlying these effects. We identify three complementary channels: hospital mergers raise the price of commercially insured care, expand the set of healthcare encounters for which patients bear out-of-pocket financial responsibility, and shift the local payer mix away from Medicare and toward commercial insurance. Together, these channels increase both the cost and the breadth of patients' financial exposure to healthcare.

JEL Codes: I11, I13, G51, L41, L42.

*We thank participants at the University of Victoria Department of Economics brown bag, the University of Utah Economics Department CAKE seminar, the UC Davis Department of Economics Alumni conference, the APPAM (2025) conference, the ASSA (2026) conference, and the CEA (2026) conference for helpful comments. We thank the California Policy Lab for hosting, documenting, and facilitating access to the University of California Consumer Credit Panel. Zack Cooper and Stuart Craig generously shared their panel dataset of hospital ownership changes. Miriam Straus of the Center for Health Law and Policy at Georgetown University Law Center provided crucial background information about the use of anti-steering provisions in hospital-insurer contracts, which Blake Eliason built on to inform this paper.

†University of Liverpool and ISCTE-IUL, DINAMIA'CET, Lisbon, Portugal; s.pinto@liverpool.ac.uk.

‡Occidental College; abdelfattah@oxy.edu.

§Corresponding author. University of Utah Department of Economics; marshall.steinbaum@utah.edu.

¶University of Victoria Department of Economics and UC Berkeley CWED@IRLE; wiltshire@uvic.ca.

I Introduction

In a multi-payer healthcare system such as prevails in the United States, health providers perform services under the assumption that they will be reimbursed from multiple sources: patients themselves, the patient's commercial insurer, a government insurer, or even other patients. Moreover, the price for a given healthcare service depends on both the type and the identity of the payer (Cooper et al., 2019). Most commercial insurance plans build patient participation into ongoing care and use that mechanism to steer patients to different healthcare services and to preferred providers as a means of cost-sharing (Aron-Dine, Einav and Finkelstein, 2013; Gowrisankaran, Nevo and Town, 2015).

Medical debt is the logical result: the provider seeks payment from the patient for delivered-but-uncompensated care (Messac, 2024), and patients may be surprised that services they thought were covered are nonetheless being billed (Cooper, Scott Morton and Shekita, 2020). Moreover, healthcare needs are inherently risky (hence the need for insurance in the first place), so necessary care may be delivered to both uninsured and under-insured patients with high-deductible health plans (Tolbert et al., 2025; Brot et al., 2017), leading patients to turn to crowd-funding when their insurance coverage (and their own resources) are insufficient (Kenworthy, 2024).

A large literature now demonstrates that a lack of competition in the healthcare industry is responsible, to a considerable degree, for the fact that healthcare prices are higher and actual healthcare provided is lower in the United States, relative to peer countries. A subset of that literature has identified horizontal concentration in the hospital sector as responsible for both higher prices for the care that is delivered and follow-on effects, such as a high marginal cost of employing an additional worker, in the form of inflated premiums for employer-provided health insurance (Brot et al., 2024). However, one outcome of hospital mergers that has not been studied to the same degree is medical debt, likely because individual-level credit records are difficult to obtain, let alone link to individual-level healthcare data, such as insurance claims.[1]

[1] An exception is Andre et al. (2023), who study hospital closures as opposed to mergers. The paper finds that closures bring about increased hospital market concentration and medical debt outcomes broadly similar to ours.

In this paper, we link together these two strands in the literature: merger-induced horizontal consolidation among hospitals, and medical debt. Using a staggered difference-in-differences and event study research design with treatment assigned by geography, we show that mergers predicted to increase local area hospital concentration cause substantial increases in medical debt in the population of that local area. In our preferred specification, five years after treatment, hospital mergers increase the likelihood of having medical debt by 11%, the amount of medical debt accumulated over the previous 12 months by 14% (a point estimate that is not statistically significant), and the total stock of outstanding medical debt by 20%. Moreover, we find that hospital mergers lead to further increases in credit card balances, suggesting that some patients carry medical expenses on their credit cards, perhaps because providers withhold access to subsequent care if patients have unpaid balances.

We then investigate the mechanisms through which hospital mergers generate medical debt, using comprehensive county-level healthcare spending and utilization data (sourced from insurance claims as well as provider records) from the Institute for Health Metrics and Evaluation's Disease Expenditure Project (Dieleman et al., 2025). We identify three complementary channels. First, mergers raise the price of commercially-insured care: private insurance spending per enrollee and per healthcare encounter increase by approximately 3% five years after treatment, while utilization remains stable, indicating that the spending increase is driven by higher unit prices rather than a greater volume of care. Second, mergers expand the set of healthcare encounters for which patients bear direct out-of-pocket financial responsibility. The number of out-of-pocket encounters rises approximately 4%, while out-of-pocket spending per encounter, if anything, declines. In other words, the greater out-of-pocket burden on patients is driven entirely by more encounters involving patient payments, not by higher charges per encounter—a novel finding in the hospital mergers literature, to our knowledge. Third, mergers shift the local payer mix away from Medicare and toward commercial insurance. Medicare enrollment declines by about 0.6% and the Medicare spending share declines by about 1.2 percentage points, so that a larger share of healthcare utilization is subject to the higher prices and broader cost-sharing obligations documented in the first two

channels.

Beyond these three quantitative channels, we identify a fourth mechanism that, while not directly testable in our data, provides the institutional foundation for the price and cost-sharing effects we document: anticompetitive provisions in hospital–insurer contracts. Dominant hospital systems contractually prevent insurers from directing patients to lower-cost competitors or placing high-cost hospitals in less-preferred network tiers. These provisions are the hospital-sector analog of the anti-steering restraints at issue in recent antitrust litigation, e.g. Ohio v. American Express Co. (2018): they neutralize downstream price competition by prohibiting the intermediary from channeling demand toward lower-cost alternatives. When a merged hospital system imposes anti-steering language on insurers, it forecloses the primary mechanism through which insurers could discipline the price increases that mergers enable. One of the few remaining instruments available to insurers in this situation is to control costs by offering employer-sponsored plans with greater patient cost-sharing, thereby shifting the burden of deterring utilization of high-cost providers onto patients themselves. This institutional channel, which has been documented in federal and state antitrust enforcement actions against Atrium Health, Sutter Health, and New York-Presbyterian, among others, helps explain both the price increases and the expansion of patient financial exposure that we observe in the IHME data (United States v. Charlotte-Mecklenburg Hospital Authority, 2019; UFCW & Employers Benefit Trust v. Sutter Health, 2021; United States v. New York Presbyterian Hospital, 2026).

A recent literature has found that forgiving medical debt already in collections has little measurable effect on recipients' financial wellbeing, mental health, or healthcare utilization (Kluender et al., 2025).[2] This finding might seem to suggest that medical debt is not itself a meaningful source of harm to patients or consumers, even if it does result from hospital mergers as we show. But our findings suggest a re-interpretation. The medical debt observable on credit reports represents only the final, residual layer of a much larger transfer of healthcare costs onto patients when healthcare markets become less competitive. Before a medical bill reaches collections, patients

[2] Duarte et al. (2026) estimate a null effect on credit reporting outcomes from the non-reportability of medical debt on credit reports after 2023, which is short of forgiving it outright but supports the conclusion of Kluender et al. (2025).

may pay it outright, finance it with a credit card, negotiate a reduced amount, or forgo care altogether to avoid incurring the expense (Brot et al., 2017). The fact that any debt that ends up in collections can be forgiven at pennies on the dollar with little welfare consequence may indicate that the economically significant harm occurs upstream: in the higher prices patients face, the broader set of encounters on which they bear out-of-pocket responsibility, and the less effective insurance coverage that results from provider consolidation. Our contribution is to document those upstream, direct consequences of reduced provider competition and their downstream residuum. The medical debt increases we estimate should be understood not as the totality of the consumer harm from hospital mergers, but as a downstream indicator of a broader shift in the financial burden of healthcare from insurers and providers to patients.

Medical debt can lead to adverse credit consequences for patients and their families (also demonstrated by Kluender et al. 2025), which is why policy surrounding its reportability on credit records has fluctuated considerably (Gibbs et al., 2025). Along with the advent of the pandemic, this is the reason why our study concludes in 2019, after which the medical debt observable in our credit panel dwindles to the point of unreliability. But beyond the immediate consequences of the debt that is our main outcome of interest, our findings are informative about access to healthcare in a context in which it is rationed. The mechanisms we study indicate that mergers increase patients' financial exposure to healthcare costs along both intensive and extensive margins: each episode of commercially insured care becomes more expensive, and more episodes generate out-of-pocket obligations. These findings imply that, in the wake of hospital mergers, commercial health insurance plans become less effective at insulating patients from the rising cost of care, even absent any explicit change in plan design. But our findings are also consistent with plans being redesigned to include greater patient participation, e.g., through higher deductibles, and with employers gravitating toward high-deductible plans to contain their own cost of providing employer-sponsored insurance.

This paper proceeds as follows: Section II explains how our analysis sample of hospital mergers, credit records, and aggregated claims data is constructed. Section III explains our empirical

approach, including summary statistics and our main estimating equation for the event study. Section IV reports our main results as well as a placebo test. Section V uses healthcare spending and utilization data derived from insurance claims to investigate the mechanisms through which hospital mergers increase medical debt. Section VI concludes.

II Data

II.1 Data Sources

This study relies on three primary datasets, supplemented from other sources for controls and ancillary analysis. The first and most important is the University of California's Consumer Credit Panel (UCCCP), an individual- and tradeline (loan)-level panel commencing in 2004. UCCCP provides anonymized credit records for a nationally-representative, replenishing 2 percent sample of U.S. adults with credit histories. It includes detailed demographic and geographic information, which allows us to identify the location of each individual. It also includes tradeline-level data on every loan or credit record obtained by each individual and every debt that has gone into collections. Tradeline-level covariates include the date the tradeline was opened, the initial amount, the current outstanding balance, and various indicators of the tradeline's status, such as whether it is or ever has been delinquent. Crucially for this paper, starting in 2011, UCCCP includes comprehensive tradeline-level data on medical debt in collections. Collections tradelines are conceptually different from standard tradelines on credit reports in that they are not credit that is applied for and extended, but rather unpaid bills. Usually a medical bill is put into collections when the healthcare provider sells or assigns the debt to a debt collection agency, which then utilizes the apparatus of credit reporting to secure payment. Sometimes providers themselves put an outstanding bill into collections and reports it to a credit bureau, although this is comparatively uncommon, and they have other mechanisms for collecting on unpaid bills short of utilizing debt collections and credit reporting, as discussed further below.

To avoid potential confounding caused by the pandemic and possible biases introduced by

changes in the reporting of medical debt by credit bureaus, we restrict the dataset to the 2011–2019 period.[3] Although UCCCP archives update at a quarterly frequency, we use only the Q4 archives because our indicator for hospital mergers is defined on an annual basis. Our analysis using UCCCP is conducted at the individual level.

Appendix A describes our process for cleaning the credit panel data in detail. However, there is a particularity of credit reporting data that is especially relevant in this context: approximately 18% of medical debt in collections tradelines in UCCCP do not appear on credit reports in the year when they are first created. For example, a collections tradeline for medical debt might report a creation date of January 2012 but only appear in UCCCP data for the first time in 2014, even though the individual it corresponds to has been in the UCCCP sample the entire period. This is likely due to the lag between the date when healthcare was provided and the charge incurred and its sale to a debt collector (or placement into collections directly by the provider). Given the impact that this could have on the levels and trends of the key outcomes of interest that we define in Section II.2 and in Appendix A, we take the tradeline creation date as accurate, create observations for the missing years if the individual is sampled, and estimate the corresponding balances through linear interpolation. For our outcomes of interest, we construct both a "raw/unadjusted" measure where we take the data in UCCCP at face value, and an "adjusted" measure that corrects missing balances. All of our main results concern the adjusted measure; we report the same results with the raw/unadjusted measure in Appendix D.

Our second data source is hospital-level ownership data and associated covariates from Cooper et al. (2024), which is itself sourced from the American Hospital Association's annual survey of hospitals. This covers the period from 2000 to 2024 but, due to the UCCCP data restrictions mentioned above, we focus on mergers through 2019. These data include the universe of hospitals, their locations, their number of beds, and the hospital systems to which they belong. The hospital-level ownership data allows us to identify, for each year and commuting zone, the mergers predicted

[3] For example, starting in 2023, medical debt in collections of less than $500 stopped being included in credit reports. The waiting period for credit reporting was also extended from six months to one year, and some states banned it entirely.

to increase local hospital market concentration relative to the previous year. As we outline in Section III.1, this is central to our methodology. Hence, the hospital-level data is actually analyzed at the level of the geographic market, which in our case is the commuting zone.[4] We geolocate individuals in UCCCP to commuting zones using their zip code.[5]

Our third main data source is the Disease Expenditure estimates produced by the Institute for Health Metrics and Evaluation (IHME) and described in Dieleman et al. (2025). The IHME project combines more than 40 billion insurance claims with 836 million facility records to produce county-level estimates of healthcare spending, utilization, and enrollment for 148 health conditions across four payer categories—private insurance, Medicare (including Medicare Advantage), Medicaid, and out-of-pocket (OOP)—and seven types of care: ambulatory, inpatient, emergency department, nursing facility, home health, dental, and pharmacy. All estimates are age- and sex-standardized to a common reference population. The data span the years 2010–2019, covering 3,110 counties that account for 97.3 percent of the U.S. population. The claims cover 76.6 percent of personal healthcare expenditures.[6] Appendix B gives a more complete account of the construction of the IHME dataset.

We aggregate the county-level IHME estimates to commuting zones (CZs). County-to-CZ crosswalks follow the 1990 CZ definitions used throughout the paper. CZ population and other demographic covariates come from the Intercensal Population Estimates.

We supplement our primary datasets with information on the timing of the ACA Medicaid

[4] The commuting zone is coarser than the market definitions typically used in the hospital-pricing literature: Hospital Service Areas, fixed-radius markets (Cooper et al., 2019), or diversion-based markets (Gowrisankaran, Nevo and Town, 2015), so a CZ-level HHI may understate the concentration increase in the submarket where a given merger actually bites. At the same time, Dranove, Kim and Vatter (2026) show that mergers between hospitals in geographically distinct markets also raise prices, and that the effect runs through common customers: a system that spans the markets in which a large employer's workers live becomes a must-have provider for any insurer bidding for that employer's business, and can exploit that leverage in both markets. The same logic implies that a within-CZ merger can raise prices in other CZs whenever insurers serve employers whose workforces span them. For our design this means that CZ-level concentration understates treatment intensity where the relevant submarket is smaller, which biases our estimates toward zero.

[5] We first use the yearly US Department of Housing and Urban Development's zip-code-to-county crosswalks to assign each individual to a county (U.S. Department of Housing and Urban Development, 2019); then, we use Dorn's county-to-commuting-zone (CZ) crosswalk to assign each individual to a CZ (Autor and Dorn, 2013).

[6] Exclusions include spending on medical equipment and transportation, as well as spending that is not covered by insurance claims, including Veterans Affairs, Tricare (for serving military personnel), and the Indian Health Services.

expansion across states (Kaiser Family Foundation, 2025), the count of uninsured people at the commuting zone level (from the Small Area Health Insurance Estimates program), and zip code rurality (US Census Bureau, 2023). We also use Quarterly Workforce Indicators to measure the impact of hospital mergers on local labor markets.

II.2 Outcomes of Interest

This subsection introduces the commuting zone- and individual-level outcomes of interest that we use for the remainder of the paper.

HHI For a given commuting zone-year cell, the Herfindahl-Hirschman Index is computed as the sum of squared market shares for each hospital system. The market share of a system in commuting zone z at year t is the share of hospital beds owned by that system in z and t. Independent hospitals that are not part of a hospital system are counted for the purposes of this computation as a 'system' composed of only one hospital.

As described in Subsection III.1, our treatment is defined as a merger predicted to increase HHI in a commuting zone, which therefore means both parties had positive market share in the year prior to the merger. We also use HHI as an outcome variable, testing whether mergers predicted to increase HHI actually did and whether the effect is transitory or permanent, as well as whether treated and control CZs are on parallel trends with respect to hospital market concentration prior to treatment.

Medical Debt For medical debt, we focus on three individual-level outcomes of interest: (i) whether an individual has medical debt in collections; (ii) the flow of new medical debt in collections generated over the previous 12 months; and (iii) the total amount of outstanding medical debt in collections. Medical debt represents unpaid bills, i.e. credit involuntarily extended by the merchant, a healthcare provider.[7] We aggregate all of a consumer's tradelines reflecting medical

[7] Appendix A.1 reports how we identify medical debt in collections in the credit panel to which we have access. That is not necessarily the full universe of medical debt either on credit records or in general, because some tradelines that may in fact be medical collections are not identified as such (and this may differ across bureaus), and some medical

debt in collections in a given credit archive for the purpose of computing these outcomes.

Utility and Cable/Cellular Debt We use utility and cable/cellular debt (hereafter "utility debt") as a placebo test, because it functions similarly to medical debt on credit reports (i.e. it represents unpaid bills as opposed to credit extended in advance), but there is no direct mechanism by which we expect hospital mergers to affect the amount of utility debt consumers carry. Our utility debt outcomes are defined equivalently to our medical debt outcomes.

Credit Card Debt Credit card tradelines differ from medical and utility debt tradelines on credit reports in a fundamental respect: consumers apply for and are granted an unsecured line of credit, as opposed to credit being involuntarily extended by the merchant (healthcare provider or utility). Hence, the outcomes we use to measure hospital mergers' effect on individual-level credit card usage differ slightly from those we use for medical debt. Two outcomes are constructed analogously to the medical debt outcomes: an indicator for having a positive credit card balance and the total outstanding balance. As credit card borrowing often takes place on existing accounts, a flow measure based on newly opened tradelines, as used for medical debt, would miss most new borrowing; therefore, we instead use the year-on-year change in the total balance. We then add a fourth outcome: the share of the total balance that is delinquent. For medical and utility debt in collections, there is no coherent concept of delinquency; or, put differently, all of it is definitionally delinquent.

Healthcare Spending and Utilization We use the county-level aggregated insurance claims data from the IHME Disease Expenditure Project to compute the following healthcare spending, insurance coverage, and utilization outcomes. The IHME data is divided into types of care and payer. We aggregate five types of care: Ambulatory, Emergency, Inpatient, Dental, and Nursing. We exclude Pharmacy and Home Health care types because encounter counts for those two care types show parallel trends violations between our treated and control CZs, indicating that underlying

collections are not reportable to credit bureaus at all. Hence, our findings with respect to medical debt should be understood to refer to the medical debt that we observe on credit reports. See Gibbs et al. (2025) for additional details.

population trends and not hospital mergers or market concentration drive utilization of those two types of care. Appendix B.3 gives further background behind the decision to exclude Pharmacy and Home Health.[8] The four payers are Private (i.e., Commercial) Insurance, Medicare, Medicaid, and Out-of-Pocket.

- **Spending by payer**: Total spending across all five included care types, by payer.
- **Spending shares**: Each payer's spending as a fraction of total spending in the CZ–year cell.
- **Beneficiary count**: Total enrollment in Medicare, Medicaid, and Private Insurance. These categories are not mutually exclusive for an individual (i.e., the same patient can be enrolled in more than one type of insurance within a given year), and there is no concept of enrollment for out-of-pocket spending.
- **Encounters**: Total number of healthcare encounters, by payer. An encounter is defined in the IHME data as treatment for a given medical condition. If the same patient is observed receiving treatment for the same condition over successive days within a year, the entire episode is considered one encounter. But if the same patient is observed receiving treatment for the same condition on non-successive days within a year, each episode is considered a separate encounter. Treatment for two different conditions is two encounters, even if that treatment happens on the same day.
- **Spending per beneficiary**: divides total spending by a payer across all included types of care by the number of individuals enrolled in that payer category, yielding average annual spending per beneficiary, a measure that combines both utilization intensity and the price of care consumed. This outcome again excludes the OOP payer.
- **Spending per encounter** replaces enrollment with encounters in the denominator. It captures the average payer outlay per healthcare encounter and is available for all four payers.

[8] The implication of this omission is that our findings related to the effect of hospital mergers on both spending and utilization are restricted to the five care types we aggregate and do not pertain to the two we omit.

Note that the same episode of receiving healthcare for the same condition from the patient's perspective counts as multiple encounters if different payers pay for it. So, for example, an outpatient doctor's appointment covered by commercial insurance with a co-pay will be two encounters, one privately-insured, the other out-of-pocket.

Because we include population controls in our commuting-zone-level regressions, we specify all commuting-zone-level outcomes of interest in absolute rather than per-capita terms, except spending per beneficiary and per encounter, for which the denominator plausibly varies independently of population trends (and we also include those denominators as separate outcomes of interest).

II.3 Time Trends for Debt in Collections

Figure I plots the levels of the outcomes of interest for medical debt in collections and for non-medical debt in collections for the 2011–2019 period, showing both the "raw" and "adjusted" measures in each case. For medical debt in collections, we observe the trend previously documented by Kluender et al. (2021): increasing levels until 2013/2014 and large, sustained declines thereafter, which they demonstrate to be at least partly due to the ACA Medicaid expansion adopted in 2014 and subsequent years by a large number of states. This motivates our inclusion of a Medicaid-expansion covariate as a control variable in our estimating equations.

The total amount of debt in collections in our credit panel differs significantly from Kluender et al. (2021). There are plausible explanations for this, including the fact that Kluender et al. (2021) rely on data from a different credit reporting agency than ours and, as they point out, the measures of medical debt in collections may vary across agencies. Likewise, the specific criteria based on which a tradeline is deemed to correspond to medical debt in collections (e.g., because the original creditor is or isn't identified as a healthcare provider), as well as the variables present in the data made available by each credit reporting agency, are not necessarily identical.

Figure I also shows that the "raw" and "adjusted" values for the outcomes of interest regarding medical debt in collections, in addition to having similar levels, also exhibit similar trends. There-

fore, unless explicitly stated otherwise, the remainder of this paper will consider only the outcomes based on the adjusted amounts, our preferred definition for the outcomes of interest.

III Empirical approach

III.1 Empirical specification

We examine the impact of hospital consolidation and merger exposure on individual- and commuting zone-level outcomes using a dynamic difference-in-differences ("event study") design. This approach relies on the differential exposure to hospital mergers and variations in their timing to identify their dynamic effects on our outcomes of interest.

We take commuting zones (CZs) to be the relevant geographic market, and therefore treatment is defined at the CZ level. A CZ is treated in year t if it experiences a hospital merger predicted to increase concentration in year t, which in turn requires that the parties to the merger were each present with positive market share in the CZ in $t-1$. Hereafter we refer to these simply as "hospital mergers."[9]

Equation 1 is our estimating equation for CZ-level outcomes of interest:

$$y_{z,t} = \sum_{e=-5}^{\bar{E}} \beta_e D_z \mathbb{1}\{t - c_z = e\} + \gamma_z + \lambda_t + \phi x_{z,t} + \varepsilon_{z,t} \tag{1}$$

$y_{z,t}$ represents the outcome of interest for commuting zone z in year t. D_z indicates that commuting zone z experienced a hospital merger during the 2014–2018 treatment window. c_z is the year z was treated. If z is treated more than once in our analysis period 2014-2018, we assign it to the earliest cohort in which it was treated. γ_z and λ_t are CZ and year fixed effects, respectively. $x_{z,t}$ represents time-varying CZ-level control variables: i. a binary variable indicating whether the ACA Medicaid

[9] Following the DOJ and FTC Merger Guidelines (U.S. Department of Justice and Federal Trade Commission, 2023) we define treatment as a "change in HHI", $2ab$, resulting from a merger between two incumbent local firms with pre-merger shares a and b. This definition means the actual level of HHI can decrease in the year following a merger predicted to increase HHI due to actual variation in entry and market shares between $t-1$ and t. Nonetheless, the HHI-increasing mergers that we consider as treatments induce durable post-merger increases in concentration. Figure II shows that consolidation amplifies in the years after the merger.

expansion was enacted in z at t,[10] ii. the log of commuting zone population, and iii. the log of commuting zone population aged 65 and over.

Equation 2 is our estimating equation for individual-level outcomes of interest:

$$y_{i,t} = \sum_{e=-5}^{\bar{E}} \beta_e D_{z(i,t)} \mathbb{1}\{t - c_{z(i,t)} = e\} + \alpha_i + \lambda_t + \gamma_{z(i,t)} + \phi x_{z(i,t),t} + \varepsilon_{i,t} \tag{2}$$

$y_{i,t}$ represents the outcome of interest for individual i in year t. $z(i,t)$ denotes the commuting zone in which individual i resides in year t; D_z and c_z are defined as in Equation 1. Treatment status is therefore assigned by current residence: an individual contributes to β_e in year t if they reside in a treated CZ e years from its merger, and to the never-treated comparison group if they reside in a never-treated CZ. For individuals who never move, $z(i,t) = z_i$ and Equation 2 reduces to a standard staggered-adoption event study with a fixed cohort assignment. α_i, γ_z, and λ_t are individual, CZ, and year fixed effects, respectively. $x_{z,t}$ represents time-varying CZ-level control variables. In Equation 2, the only time-varying control is the Medicaid expansion binary.

The control group consists of never-treated CZs, i.e., those that were not exposed to a merger that was predicted to increase HHI during 2014-2019. Our study period covers a time of intense merger activity in the hospital sector; however, another such period occurred in the 1990s to early 2000s, so our control CZs are, by and large, "already" concentrated either due to small population and hence few hospitals, or due to earlier mergers.

In recent years, a large body of literature has developed on the biases that can potentially result from two-way fixed effects estimation when treatment is staggered (e.g., Goodman-Bacon 2021). To address this, we estimate Equations 1 and 2 using the interaction-weighted estimator of Sun and Abraham (2021), which accounts for heterogeneity in treatment timing and dynamic treatment effects (see also Callaway and Sant'Anna 2021; Borusyak, Jaravel and Spiess 2024).[11][12]

[10] For the purpose of assigning the Medicaid expansion binary variable, which varies at the state level, to commuting zones that cross state lines, we assign the binary associated with the state where the majority of the commuting zone population resides.

[11] Imputation-based alternatives (e.g., Borusyak, Jaravel and Spiess 2024) may offer efficiency gains over Sun and Abraham (2021) but are more exposed to model misspecification because they extrapolate a fixed-effects model to treated observations.

[12] Because Equation 2 assigns treatment by current residence, an individual's treatment status can switch on or off as

We do not consider treatment cohorts before 2014 to ensure that every cohort has at least three years in their pre-treatment period; to prevent large imbalances in the pre-treatment period, we restrict it to a maximum of five years for every treated cohort. As unpaid medical debt takes at least six months to appear on credit records, we impose that every treatment cohort must have a minimum of one year in their post-treatment period, resulting in 2018 being the last cohort whose treatment effect we estimate. The 20 CZs that, during the 2014-2019 period, were first exposed to a merger in 2019 are dropped from our sample. Consequently, there are 137 treated CZs, with the 2014 cohort being by far the largest (64 CZs) thanks to our assignment of multiply-treated CZs to their earliest treatment date, and 556 control CZs.

As is typical of staggered designs, the post-treatment period varies by cohort, from a minimum of one year (for the 2018 cohort) to a maximum of five years (for the 2014 cohort).[13] β_e are our key parameters of interest for each relative time period and measure the effect of merger exposure on our outcomes of interest.

The empirical specification described above and embodied in Equations 1 and 2 relies on the traditional identification assumptions for difference-in-differences approaches, particularly that there is no anticipation of treatment and that there are common trends in potential outcomes between the treated and control cohorts.

While we cannot test the common trends assumption directly, the results in Sections IV and V are consistent with parallel trends for the outcomes of interest described in Section II.2 in the pre-treatment period, making it plausible to assume that such trends would have continued in the absence of the treatment. We also find no evidence suggesting that the treatment (a hospital merger that is predicted to increase HHI in the CZ where patients live) is anticipated. This point is further supported in Appendix C.

---

they move across CZs. Treatment is therefore non-absorbing at the individual level and the baseline specification falls outside the staggered-adoption setting of Sun and Abraham (2021), in which each unit belongs to a single treatment cohort. In the never-movers specification (column (4) of Tables E.I–E.III), residence is fixed, each individual belongs to exactly one cohort, and the setting matches that of Sun and Abraham (2021) exactly; Section IV shows that the two sets of estimates are similar.

[13] $\bar{E}$ therefore takes a maximum value of 5 and is defined as $\bar{E} = min\{5, 2019 - c_z\}$, with $c_z \in \{2014, 2015, 2016, 2017, 2018\}$.

III.2 Summary statistics

Table I provides an overview of the basic demographic characteristics, the key outcomes of interest related to medical debt in collections, and indicators related to other types of debt, across the treatment and control cohorts in the individual-level credit panel. All the values in the table refer to 2013, before the treatment period for each cohort.

The gender composition of individuals in treated and control CZs is very similar, while those in control CZs are slightly older than those in treated CZs. It also shows that the control CZs have a substantially higher ex-ante HHI level than the treated cohorts. Likewise, control CZs typically had a higher share of individuals with medical debt in collections, new medical debt in collections generated over the previous 12 months, and total outstanding medical debt in collections, than most or all of the treated CZs. As stated earlier, control CZs were "treated" by the earlier wave of hospital mergers in the 1990s-2000s, and are generally smaller/more rural. This is why we include CZ-level population and 65+ population as time-varying controls when estimating Equation 1. Table I also shows that the 2013 levels of the "raw" and "adjusted" measures for the outcomes of interest were very similar for each cohort.

We emphasize that difference-in-differences and event study approaches do not require identical baseline levels between the treatment and control groups. Instead, the key identification assumptions are those previously discussed in Section III.1, i.e., no anticipation of treatment and common trends in potential outcomes. If those hold, as we have argued is the case, the control cohort provides an appropriate counterfactual and allows us to assess the causal impact of hospital merger exposure on medical debt in collections, even with differences in baseline characteristics and levels.

Table II reports summary statistics for commuting-zone-level outcomes from the IHME Disease Expenditure data, as well as commuting zone population and count per treatment cohort. As with Table I, all values are for 2013.

IV Results

IV.1 The effect of hospital mergers on market concentration

Our research design relies on the assumption that the CZs treated by mergers predicted to increase HHI (due to ex-ante market overlap by the parties) are not subject to other changes, such as the entry of new hospitals, that would offset the competitive effect of the mergers. We test this assumption by estimating Equation 1 using the log-transformed CZ-level HHI as the outcome of interest. Figure II displays the resulting event study plot. It shows a large and persistent increase in HHI immediately after the treatment takes place. The figure therefore shows no evidence that subsequent trends counteracted the concentrating effect of the merger and brought the HHI of treated CZs back down. In fact, it implies the opposite: the treatment effect amplifies over time such that, five years post-treatment, the treated CZs have an HHI that is 24.5% higher than if they had not been exposed to a hospital merger.[14] It is also consistent with the parallel trends assumption for CZ-level hospital HHI, i.e. the mergers in our sample are not selected into CZs in which concentration is already increasing (or decreasing).

IV.2 Baseline results: medical debt in collections

Figure III shows the event study estimates for the three outcomes of interest related to medical debt discussed in Section III.1. Here, we focus only on our preferred specification that (i) uses the full sample, (ii) includes year, CZ, and individual fixed effects, and (iii) includes a binary control that identifies if and when states expanded Medicaid.[15]

As mentioned in Section III.1, the two key identifying assumptions are the common trends

[14] This percentage increase is based on the fact that the point estimate for $t = 5$ is 0.219 and $exp(0.219) - 1 \approx 24.5\%$. It is important to highlight that the coefficient at $t = 5$ is estimated solely from the 2014 cohort, but the increasing trend in the point estimates is present in every preceding year, when it is estimated from a larger number of cohorts.

[15] In Appendix D, we re-estimate the baseline specification using the "raw/unadjusted" version of the outcomes for medical debt in collections. Figure D.I shows that the resulting estimates are generally very similar, with a slight difference arising in the outcome for new medical debt in collections generated over the previous 12 months. This is to be expected, given that, as outlined briefly in Section II and in more detail in Appendix A.1, the data corrections we implement have a larger effect on both the amounts and the timing of this outcome variable.

in potential outcomes and no anticipation of treatment. Figure III illustrates that, across all panels, the pre-treatment estimates are typically tightly estimated around zero. This implies that we observe parallel trends between treatment and control CZs in the pre-treatment period, making it plausible that both groups have common trends in potential outcomes such that, in the absence of treatment, we would expect treated and control CZs to have remained on identical paths regarding the outcomes of interest.

We address the no-anticipation assumption by conducting a test where we assign a placebo treatment year to each cohort that precedes the true treatment date by one year (i.e., the 2014 cohort is assigned a placebo treatment date of 2013, the 2015 cohort is assigned 2014, and so on) and then re-estimate our baseline specifications. Appendix C shows the results of this procedure: for every outcome in (placebo) event time $t = 0$, corresponding to the year the (placebo) treatment took place, the resulting event study point estimate is a tightly estimated zero, in line with what would be expected in the absence of any systematic anticipation of merger events.

Panels III.1 through III.3 show a large and significant effect of hospital mergers on medical debt that does not attenuate over time. Hospital mergers cause an 11% increase in the likelihood of having medical debt in collections five years after treatment, and an increase of 20% in the total amount outstanding.[16] The point estimate, at $t = 5$, for the effect on the amount of new medical debt generated over the previous 12 months amounts to a 14% increase relative to the year preceding treatment, but that coefficient estimate is not statistically significant.

The upward trend in Panels III.1-III.3 indicate that the merger effect on medical debt outcomes amplifies over time. However, the estimates five years post-treatment are based solely on the 2014 cohort—the largest and earliest-treated one. It is possible that the dynamic amplification we find is true for that cohort but not later-treated ones. Nevertheless, Panels III.1 through III.3 show that, even just two years post-treatment, the effects were already positive and significant for the likelihood of having medical debt (a 2.3% increase) and the total amount of outstanding medical

[16] We computed these by dividing, for each outcome of interest, the corresponding point estimates at $t = 5$ by the baseline pre-treatment value for the 2014 cohort shown in Table I. The point estimates are detailed in the event study regression tables in Appendix E; the preferred specification illustrated by Figure III corresponds to the results in column (3).

debt (a 6.0% increase); for the new medical debt generated over the previous 12 months, the point estimate implies an increase of 8.1% but is again non-significant.[17]

Appendix E shows the event study regression tables for each of the three outcomes depicted in Figure III, in Tables E.I, E.II, and E.III, respectively. All three tables follow the same format: columns (1) to (3) use the full sample and sequentially add fixed effects and a binary control for the state-level Medicaid expansion that occurred from 2014 onward; column (3) corresponds to our preferred specification, plotted in Figure III. Column (4) restricts the sample to never movers, and column (5) further restricts it to never movers under the age of 65. Column (6) returns to the full sample as the starting point and then restricts the treatment group to CZs where the merger-induced HHI increase exceeded 150; column (7) further restricts the treatment group to CZs where the post-merger HHI exceeded 1,800, and column (8) increases the post-merger HHI threshold to 2,500.

Overall, Tables E.I-E.III show that the baseline results are robust to various sample and treatment group definitions. Column (5) shows that restricting the sample to never-movers under 65 results in similar point estimates for all three outcomes. Restricting the treatment group based on merger-induced HHI increases and post-merger HHI levels results in larger and more significant point estimates for the outcomes related to the flow of new medical debt over the previous 12 months and the total outstanding amount. This occurs simultaneously with lower magnitudes and significance regarding the effect on the likelihood of having medical debt, suggesting an effect along the intensive margin for medical debtors in CZs where the hospital mergers result in higher levels of concentration.

The baseline results are robust to using an alternative control group defined as the set of CZs that did not have any merger predicted to increase HHI between 2011 and 2019 (i.e., making our definition of "never-treated" more onerous, to exclude control CZs that experienced HHI-increasing mergers during our pre-period). This reduces the number of control CZs from 556 to

[17] We followed an analogous procedure to the one described in footnote 16. In this case, at $t = 2$, the estimate is based on the 2014, 2015, 2016, and 2017 cohorts; therefore, the denominator used the average of the baseline values for each cohort in 2013, as detailed in Table I, weighted by the number of individuals in the cohort.

507. Appendix F shows that the resulting estimates are very similar, though with marginally higher point estimates relative to those in Figure III.

Another potential concern with our baseline results is that they may suffer from dynamic bias as a result of our baseline model (Equation 2) disregarding dynamic feedback in the outcome variable (Klosin, 2024). However, the inclusion of lagged outcomes to address this concern may alternatively induce Nickell bias (Nickell, 1981; de Chaisemartin and D'Haultfoeuille, 2023). In Appendix G we replicate our baseline event study estimates but include the lagged dependent variable as an additional control in every specification. Across all outcomes, the estimates remain similar in magnitude, well within the confidence interval of the baseline results, indicating no evidence of either type of bias in the relevant estimates.[18]

We also explore several dimensions of heterogeneity. Splitting the sample by gender yields nearly identical estimates for female and male individuals (Appendix Figure H.I). Estimating separately for states that did and did not expand Medicaid during the study period, we find no notable difference in the effects across the two groups; if anything, the effects are driven by individuals in expansion states (Appendix Figure H.II). Finally, restricting the sample to individuals who already had medical debt in collections in 2013 produces estimates very similar to the full-sample baseline, suggesting that the effect of merger exposure operates similarly across both the intensive and extensive margins of medical indebtedness (Appendix Figure H.III).

IV.3 Results for other outcomes

IV.3.1 Credit card debt

One potential limitation of the approach followed so far for estimating the effect of hospital mergers on medical debt, also highlighted by Kluender et al. (2021), is the possibility that patients sometimes pay for out-of-pocket medical expenses with a credit card. That type of medical debt would not be captured by our medical-debt-in-collections-related outcomes reported so far. Con-

[18] We note that this comparison would be uninformative if both biases were similar in magnitude and direction, though we view this as unlikely given their distinct sources.

sequently, the baseline estimates might underestimate the total effect of mergers on medical debt outcomes. To address this, we adapt our specification in this section to include new outcomes of interest based on credit card debt.

One important difference between credit card debt and medical debt concerns how each appears on credit reports: credit card tradelines are opened with the express purpose of obtaining access to credit, whereas medical debt is, in effect, an unpaid bill that has been sent to collections. Hence, for our analysis of credit card debt, we do not stipulate that the tradelines in question need to be in collections.

Figure IV displays the resulting event study plots, suggesting that exposure to hospital mergers affects credit card outcomes as well. As detailed in section II, Panels IV.1 and IV.3 are constructed analogously to the baseline outcomes for medical debt, while Panel IV.2 reports the year-on-year change in the total credit card balance. In contrast to the baseline results, the effects on credit card debt operate almost entirely through the intensive margin. Panel IV.1 shows little-to-no effect of merger exposure on the probability of having credit card debt, relative to individuals in non-exposed CZs. On the other hand, Panels IV.2 and IV.3 show that merger exposure leads to frequently positive changes in outstanding credit card balances relative to the previous year, which translates to a gradually increasing stock of total credit card debt.[19]

The intensive-margin pattern for credit cards—more debt conditional on having any, but no change in the probability of holding credit card debt—is consistent with a specific institutional feature of medical billing. Hospitals routinely condition access to non-emergency care on the resolution of outstanding balances, effectively requiring patients to pay prior debts before receiving further treatment (Kliff and Silver-Greenberg, 2023). Patients facing such demands are likely to turn to credit cards as a means of financing those payments, particularly when the alternative is being denied care. This channel would generate the pattern we observe: merger-induced increases

[19] The estimates in Panels IV.2 and IV.3 are not always based on the same cohorts. Since the year-on-year change is not defined in 2011, $t = -5$ in Panel IV.2 excludes the 2016 cohort but $t = -5$ in Panel IV.3 does not. Because our data end in 2019, only the 2014 cohort contributes to $t = 5$, while both the 2014 and 2015 cohorts contribute to $t = 4$. As a result, the rise in Panel IV.3 between $t = 4$ and $t = 5$ may partly reflect the 2015 cohort dropping out, and does not necessarily need to match the estimate at $t = 5$ in Panel IV.2.

in medical costs (and in the likelihood those costs are borne by patients themselves, as we show in Section V) raise the amount that existing credit card holders charge, without drawing new individuals into credit card borrowing. It also implies that the medical debt in collections captured by our baseline estimates understates the full financial burden that hospital mergers impose on patients, since some merger-induced medical expenses are absorbed into credit card balances rather than appearing as medical collections tradelines.

The effect of merger exposure on delinquent credit card debt would arguably provide a better measure of the kind of financial distress that could eventually lead to having medical debt in collections. Panel IV.4 shows mergers correspond to a *decline* in the share of credit card debt that is delinquent. It is important to emphasize again here that all medical debt on credit reports is in some sense delinquent—it represents unpaid bills—whereas only some credit card balances are delinquent. The event study specifications suggest that severe pre-treatment trends are present for this outcome: individuals in CZs that experienced hospital mergers were already experiencing a decline in credit card delinquency relative to control CZs. This violation of a key condition for identification makes a causal interpretation of the resulting estimates impossible.

IV.3.2 Placebo Test: Non-medical debt in collections

A potential concern regarding the research design we use in this paper is that there is an underlying phenomenon that broadly overlaps, across space and time, with hospital mergers, which could be causing an increase in debt in collections across the board, rather than just medical debt. While it is never possible to fully disprove this alternative, conducting placebo tests in which we apply the baseline specification to outcomes that are plausibly unaffected by the treatment (i.e., by hospital merger exposure) should help allay concerns that the results we have found so far result from spurious correlation between hospital mergers and overall deteriorating credit conditions.

In this sub-section, we construct outcomes based on other types of non-medical debt in collections. Specifically, we focus on categories of non-medical debt in collections that meet the following three criteria: First, they can reasonably be assumed not to be directly affected by the treatment in question. Second, they are prevalent enough in the UCCCP sample that they are ge-

ographically widespread and the resulting estimates will not be overly noisy and underpowered. We satisfy this by relying on categories that represent at least 5% of the total number of tradelines in collections (for comparison, medical debt represents 58% of the number of collections tradelines in the UCCCP data). Third, like medical debt, they must represent unpaid bills as opposed to credit that is extended ex-ante, so that the outcomes of interest can be constructed analogously from credit reporting data.

Applying these criteria, we focus on two categories of non-medical debt in collections: utilities and cable/cellular. Like medical debt, utility debt represents an unpaid bill for which the provider is to some degree restrained from discontinuing service for non-payment, and so those providers are granted access to consumers' credit reports as a second-best means of enforcing payment.[20] We pool these tradelines together and define the three resulting outcomes analogously to the baseline outcomes for medical debt. Ex-ante, we expect the treatment not to directly affect these categories of debt in collections; this should translate into point estimates that are not statistically significant and are centered around zero.

Figure V displays the resulting event study plots, which are broadly in line with this expectation. Panels V.1 and V.3 show largely non-significant estimates that oscillate around zero both before and after the treatment. Panel V.2 (new debt in collections) shows a slight increase in point estimates after the treatment year, but these are largely non-significant and very small in magnitude, such that they never translate into a positive effect on total utility/cable/cellular debt in collections, as shown by Panel V.3.

V Testing Mechanisms using Aggregated Claims Data

The preceding section establishes that hospital mergers cause economically meaningful increases in individual medical debt. A natural next question is *why*? Medical debt arises when patients face out-of-pocket costs they cannot (or choose not to) pay, so the causal chain from a hospital merger to higher medical debt must run through some combination of higher prices for care, reduced

[20] The alternative for secured loans is repossession.

insurance coverage for a given quantity of care received, changes in the volume or intensity of services delivered, reduced patient ability to pay thanks to reductions in income, or reduced patient willingness to pay due to dissatisfaction and/or surprise. In this section, we use the county-level IHME Disease Expenditure data described in Section II to investigate which of those mechanisms is in operation.

V.1 Hospital mergers lead to higher prices for commercially-insured patients

A large body of evidence documents that hospital mergers raise the prices negotiated between hospitals and commercial insurers (Cooper et al., 2019; Gowrisankaran, Nevo and Town, 2015). Higher negotiated prices translate mechanically into higher out-of-pocket costs for commercially insured patients through co-insurance, co-pays, and deductibles: even if the cost-sharing percentage remains unchanged, the dollar amount the patient owes rises proportionally with the price of care. This channel provides perhaps the most direct causal pathway from hospital mergers to medical debt.

We test for this pathway by estimating Equation 1 using spending and utilization outcomes from the IHME data. The results are presented in Figure VI. Panel VI.1 shows the merger impact on log private insurance spending: point estimates rise steadily to approximately 2.8% at $t = 5$, with no evidence of differential pre-trends.[21] This establishes that hospital mergers raise the overall cost of healthcare in affected commuting zones. Panel VI.2 establishes that the increased spending is not due to increased utilization: the log of privately-insured healthcare encounters shows no merger impact.

To further decompose the private insurance spending increase, Panels VI.3 and VI.4 present the effect of mergers on log spending per private insurance beneficiary and log spending per privately insured encounter, respectively. Both outcomes show clean pre-treatment periods and clear post-

[21] Throughout this section, percentage interpretations of log-point coefficients are computed using the exact transformation $\exp(\hat{\beta}) - 1$.

treatment increases that grow over time. Spending per beneficiary increases 2.6% at $t = 5$, while spending per encounter reaches approximately 2.9% at $t = 5$. Because the number of encounters per privately-insured enrollee is stable, this result indicates that the cost per episode of care has risen, not that more care is being delivered.

Taken together, these results provide strong evidence for the price channel. Hospital mergers raise the amount that commercial insurers—and by extension, their beneficiaries—pay each time they receive healthcare and overall, with no effect on the amount of healthcare they receive. Holding constant the cost-sharing provisions of commercial insurance plans, these higher prices translate directly into higher out-of-pocket obligations for patients. Given the well-documented sensitivity of medical debt to out-of-pocket costs, this channel provides a clear and direct explanation for the increases in medical debt documented in Section IV.

V.2 Hospital mergers lead to broader patient participation in healthcare costs

The price channel documented above operates along the intensive margin: each episode of commercially insured care becomes more expensive, and patients bear some portion of those costs through co-insurance. However, hospital mergers could also increase medical debt along the extensive margin of out-of-pocket care by expanding the set of care episodes for which patients bear direct financial responsibility. If mergers lead to restructured cost-sharing arrangements, for instance, by introducing co-pays on services previously covered without them, or by shifting patients into plans with higher deductibles, then more healthcare encounters would generate out-of-pocket obligations, even if the overall cost of each encounter remained unchanged. This extensive-margin channel would increase both the probability that any given patient faces an unexpected medical bill and the total burden of out-of-pocket spending across the population.

We investigate this pathway using four out-of-pocket (OOP) spending and utilization outcomes, estimated via Equation 1. The results are presented in Figure VII. Panel VII.1 shows the effect on log total OOP spending. There is a clear and immediate increase in OOP spending at the time of

treatment ($t = 0$) and the point estimate reaches 2.1% at $t = 5$, albeit not significant at the 95% confidence threshold. While the earliest pre-treatment estimate (at $t = -5$) is slightly negative, the estimates from $t = -4$ to $t = -2$ are close to zero, and the sharp upward break at $t = 0$ is evident.

To determine whether this increase reflects higher charges per encounter or a broader set of encounters involving patient payments, we turn to the remaining three panels. Panel VII.2 shows the effect on log OOP encounters, the count of encounters for which patients made direct out-of-pocket payments. The point estimates are consistently positive after the merger, rising to a 3.7% increase at $t = 5$. Although the estimates are somewhat imprecise, particularly at longer horizons, the direction of the shift is clear and statistically significant: mergers increase the number of care episodes that involve out-of-pocket payments.

Panel VII.3 presents the complementary result. It shows the effect on log OOP spending per encounter, which captures the average patient payment conditional on an encounter occurring. In contrast to the total spending effect in Panel VII.1, this outcome shows no post-treatment increase. If anything, per-encounter OOP spending drifts slightly downward at longer horizons. The combination of Panels VII.2 and VII.3 implies that the mechanism operates along the extensive margin rather than the intensive margin of out-of-pocket care: mergers increase OOP spending not because each encounter becomes more expensive for the patient, but because a larger proportion of each patient's healthcare episodes generate out-of-pocket obligations.

Finally, Panel VII.4 shows the effect on the OOP share of total healthcare spending. The estimates are tight and show a small positive effect, on the order of 0.4 percentage points (against an OOP baseline share of $\approx 13\%$ of total spending).

This pattern—more OOP encounters without higher per-encounter charges—is consistent with several possible mechanisms at the plan-design level. Merged hospital systems may negotiate contracts that shift more costs to patients, insurers responding to higher hospital prices may redesign plans with broader cost-sharing requirements which become more attractive to employers as a means of containing their own health insurance costs, or the elimination of competing providers may reduce patients' ability to seek care at facilities where their coverage provides more generous

terms. Whatever the specific plan-level mechanism, the aggregate result is that mergers expand the set of care episodes on which patients have direct financial exposure—a novel finding in the hospital mergers literature, to our knowledge. This extensive-margin channel complements the price channel: the price increases documented in the preceding subsection raise the cost of each episode where patients have financial exposure, while the cost-sharing channel documented here expands how many episodes create such exposure. Further, our results are consistent with overall commercial health insurance plans becoming less generous to members relative to the total (rising) cost of their healthcare.

V.3 Hospital mergers shift the payer mix away from Medicare and toward commercially-insured and out-of-pocket care

The two channels documented above, higher commercial prices and broader patient cost-sharing, both increase the financial exposure of commercially insured patients. A third channel can amplify these effects: if hospital mergers shift the local payer mix away from public insurance and toward commercial coverage, then a larger share of healthcare utilization is subject to the price and cost-sharing increases documented in the preceding subsections.

We investigate this channel using Medicare-related outcomes from the IHME data. Medicare spending in the IHME data includes Medicare Advantage (MA) plans alongside traditional fee-for-service (FFS) Medicare. Consequently, a decline in IHME Medicare enrollment cannot be attributed to a shift from FFS to MA within the Medicare program.

Figure VIII presents the results. Panel VIII.1 shows that treated CZs experience a decline in Medicare enrollment that reaches approximately -0.0064 at $t = 5$. Relative to the 2013 baseline for the 2014 cohort (244,155, from Table II), this corresponds to a decline of roughly 1,550 individuals. Note that the number of privately-insured beneficiaries shows no merger effect (Figure VIII.2).

Panel VIII.3 shows the corresponding effect on Medicare spending share. The point estimates turn negative immediately after treatment and remain so throughout the post-treatment period,

reaching approximately -1.2 percentage points at $t = 5$ (against a pre-treatment baseline of $\approx 25 - 30\%$). Figure VIII.4 shows the countervailing rise in spending share for privately-insured patients. Together, these panels document a re-composition of the local payer mix away from Medicare and toward commercial insurance. The private insurance share increase exceeds the Medicare share decline. Note also that figure VII.4 shows an approximately 0.4 percentage point increase in the OOP spending share.

The decline in Medicare enrollment and spending share could reflect several underlying mechanisms. Merged hospital systems, which derive higher margins from commercially insured patients (Cooper et al., 2019), may steer capacity toward commercial admissions at the expense of Medicare-covered ones. Service line consolidation or facility closures following mergers may disproportionately reduce access to the types of care most heavily utilized by Medicare-covered patients. Alternatively, merged systems may reclassify or restructure episodes in ways that shift billing away from Medicare and toward patients' supplemental or commercial coverage. We cannot distinguish among these explanations with the IHME data, but the aggregate pattern is clear: mergers reduce the share of care covered by Medicare and shift the payer mix toward commercial insurance. To the extent that care previously billed to Medicare is instead billed to commercial plans or directly to patients, the individuals affected face the higher prices and broader cost-sharing obligations documented in the preceding subsections. This coverage channel thus acts as a multiplier on the price and cost-sharing effects: not only are commercial prices higher and out-of-pocket episodes more frequent, but a larger share of local healthcare utilization is subject to commercial pricing.

This recomposition of the payer mix operates along the intensive rather than the extensive margin of insurance coverage. As we show in Appendix J.2, merger exposure does not increase the number of uninsured residents in treated commuting zones (Appendix Figure J.I); if anything, the point estimates are negative, though imprecisely estimated. The shift away from Medicare thus reflects a reallocation of covered care toward commercial payers, not patients losing insurance coverage altogether.

V.4 Anticompetitive hospital–insurer contract provisions as a mechanism linking mergers to higher prices and increased cost-sharing

The three quantitative channels documented in subsections V.1-V.3 describe *what* changes in the healthcare market following hospital mergers. An important remaining question is *how* merged systems bring about these changes, given that commercial insurers have strong incentives to resist price increases and steer their members toward lower-cost providers. In this subsection, we describe provisions in hospital–insurer contracts that can provide the institutional foundation for the price and cost-sharing effects we document. While we cannot directly test this channel with our data, the mechanism is well-documented in antitrust enforcement actions and is consistent with the patterns in our preceding IHME results.

V.4.1 Anti-steering contractual provisions and their economic logic

In a competitive hospital market, insurers steer plan members toward lower-cost providers in response to provider price increases, disciplining the exercise of provider market power. Steering typically takes the form of network tiers for which members face lower co-pays or coinsurance rates if they patronize providers on the preferred tier. In the extreme, higher-cost providers can be excluded from the network entirely. Tiering creates demand elasticity facing individual providers: a provider that raises its prices risks losing patient volume as the insurer redirects demand to cheaper competitors. This is the central pro-competitive role commercial insurance plays in healthcare markets, if it plays any such role (Prager, 2020).

Dominant hospital systems, including, critically, systems that have grown through the type of horizontal mergers studied in this paper, have developed contractual instruments that we refer to collectively as *anti-steering provisions*, which neutralize insurers' ability to channel members away from the contracting system.[22] The most important categories are:

Anti-tiering provisions, which require either that the hospital system be placed in the insurer's

[22] These provisions are the subject of a 2026 report from the Council of Economic Advisors finding that banning them would reduce provider prices by 18%, although that estimate is not based on a policy impact evaluation (Council of Economic Advisors, 2026).

most preferred network tier regardless of its cost or quality, or more generally that the insurance plan not use financial incentives to direct members to other providers. Either formulation prevents the insurer from signaling to members that cheaper alternatives exist and incentivizing them to switch, suppressing the sources of price elasticity that would otherwise exist for insured patient demand.

All-or-nothing provisions, which require that, if an insurer includes any facility in a health system's network in an insurance plan, it must include all of them. This forecloses the insurer's ability to selectively contract with, for example, a lower-cost suburban hospital in a system while excluding the higher-cost flagship—effectively extending the system's bargaining leverage across its entire portfolio of facilities (and incentivizing further horizontal and vertical consolidation among providers).

Gag clauses and price secrecy provisions, which prevent insurers from disclosing negotiated prices to members, employers, or competing providers, thereby suppressing the price transparency that is necessary for market-based steering to operate, even in the absence of explicit anti-tiering or all-or-nothing language.[23]

These provisions are the hospital-sector analog of the Most-Favored-Nations (MFN) clauses that have been the subject of antitrust litigation in platform markets. In Ohio v. American Express Co. (2018), the Supreme Court considered whether a credit card network could contractually prohibit merchants from steering customers to lower-cost payment methods. Like hospital anti-steering clauses, the provisions at issue in that case operate by preventing the intermediary from channeling demand toward cheaper alternatives. The economic logic is identical: a firm with sufficient bilateral bargaining power to insist on vertical contractual restraints to foreclose downstream price competition insulates itself from the demand elasticity that would otherwise discipline its pricing, further increasing its power to price in excess of cost.

[23] Price secrecy provisions were prohibited federally in 2020 (Kabrick and Keith, 2026).

V.4.2 How mergers amplify anticompetitive contracting

The relationship between hospital mergers and anti-steering provisions is mutually reinforcing. First, mergers increase a system's bargaining leverage over insurers by making the system more essential for network adequacy: an insurer that cannot include the dominant local system in its provider network cannot credibly sell coverage in that market. This leverage makes it easier to demand and enforce anticompetitive contract terms at each contract renewal. Second, large national hospital chains responsible for much of the acquisition activity in our sample period maintain standardized contracting templates that systematically include these provisions. When such a system acquires an independent hospital or smaller local system, it may import these contracting practices to a market where the acquired hospital would not previously have had the leverage to demand them. Third, post-merger consolidation reduces the number of credible alternative providers available to insurers, making the threat to exclude the merged system from a network less credible and thereby reducing the insurer's ability to resist anticompetitive contractual provisions.

V.4.3 Documentary evidence: Atrium Health, Sutter Health, and New York-Presbyterian

Although contracts between hospitals and insurers are generally proprietary, antitrust enforcement actions have produced detailed public records documenting these practices.

Atrium Health (Charlotte, NC). In 2016, the U.S. Department of Justice filed a complaint against Carolinas Healthcare System (now Atrium Health), alleging that the system used anti-steering provisions in its contracts with every major commercial insurer in the Charlotte metropolitan area (United States v. Charlotte-Mecklenburg Hospital Authority, 2019). Atrium's market position had grown substantially through a series of acquisitions over the preceding decade, and the DOJ's complaint alleged that anti-steering provisions were imposed as that leverage increased, directly linking the provisions to acquired market power. The complaint documented that Atrium's own senior leadership identified insurer steering as the single greatest threat to the system's revenues, and that Atrium responded by contractually prohibiting insurers from steering and from disclosing price information that would facilitate market-based provider choice.

The consent decree made public the actual contractual language in Atrium's insurer agreements. For example, Atrium's contract with UnitedHealthcare stated: "Plan shall not directly or indirectly steer business away from Hospital." Its contract with Cigna provided that "CIGNA will not steer business away from MHR Participating Providers" and further required that "All MHR entities...will be represented in the most preferred benefit level for any and all CIGNA products." The contract with Aetna permitted termination if Aetna introduced products that had "the effect of placing Participating PHO Providers in a non-preferred position." These provisions spanned every major commercial insurer in the Charlotte area. The case settled with a consent decree prohibiting Atrium from using anti-steering provisions for ten years.

Sutter Health (Northern California). The Sutter Health litigation provides an even more direct illustration of the link between hospital mergers and anticompetitive contracting. In December 1999, Sutter merged its Alta Bates hospital with Summit Medical Center in the East Bay market. The merger was challenged by the State of California (California v. Sutter Health System, 2000), but the district court declined to grant a preliminary injunction to block it on the grounds that insurer steering would prevent Sutter from exercising market power. The ruling relied in part on Sutter's own expert testimony that "the primary mechanism by which MCOs and IPAs keep prices low is through the 'steering' of patients" and that insurers facing price increases "can attempt to steer patients to lower cost health care providers and away from the hospital imposing a price increase, thereby pressuring the hospital to eliminate the price increase."

Having secured merger approval on this basis, Sutter then imposed contract provisions that prohibited the steering the court had relied upon to conclude the merger would not be anticompetitive. The California Attorney General's subsequent complaint alleged that beginning no later than 2002, Sutter compelled commercial insurers across Northern California to accept all-or-nothing terms, anti-tiering provisions, and price secrecy clauses (UFCW & Employers Benefit Trust v. Sutter Health, 2021). The complaint documented that Sutter used punitive pricing to enforce these terms: if an insurer resisted including a high-priced newly-acquired facility, Sutter would increase rates

at its existing facilities, coercing the insurer into accepting the system-wide terms. An FTC study found that post-merger price increases at Summit ranged from approximately 23 to 50 percent depending on the insurer, among the highest in California (Tenn, 2011). The case settled in 2021 for $575 million, with behavioral remedies requiring Sutter to cease all-or-nothing contracting, anti-tiering, and price secrecy practices.

New York-Presbyterian (New York City). The most recent enforcement action targets New York-Presbyterian (NYP), the largest hospital system in New York City (United States v. New York Presbyterian Hospital, 2026).[24] The DOJ's complaint alleges that NYP used its market power to impose plan restrictions on every major commercial insurer operating in the city, contractually preventing payers from offering narrow-network plans that excluded NYP, tiered plans that placed NYP in a less-preferred benefit tier, or any plan design that offered patients lower cost-sharing for choosing a lower-priced rival. NYP further required payers to contract on an all-or-nothing basis, so that any insurer wishing to include a single NYP facility had to include all of them. The complaint alleges that NYP's own internal analyses calculated that the introduction of tiered plans alone would reduce its profits by hundreds of millions of dollars, confirming that the restrictions served to insulate NYP from the price competition that budget-conscious plan designs would otherwise create.

As in the Sutter litigation, the public enforcement action was catalyzed by a labor union's direct experience with anticompetitive contracting. Beginning in 2022, the health fund of SEIU 32-BJ excluded NYP from its provider network after concluding that NYP's prices for routine procedures were far higher than those of comparable New York City hospitals (Goldstein, 2025). When the union subsequently attempted to contract with Aetna as its plan administrator while maintaining NYP's exclusion, it discovered that Aetna's contract with NYP required the hospital system's permission before NYP could be omitted from any client's network. NYP responded by presenting the union with a $25 million bill for allegedly unpaid past claims, effectively conditioning its

[24] DOJ also sued OhioHealth around the same time for similar practices in the Ohio market; a proposed settlement banning the use and enforcement of anti-steering contractual language is pending in that case as of this writing (United States and State of Ohio v. OhioHealth Corporation, 2026).

consent to exclusion on payment of a sum roughly equal to what the union had spent at NYP in an entire year (Mathews, 2024). The union's 2024 memo to the DOJ requesting a civil investigation detailed these obstacles and provided claims data documenting NYP's price premium; that memo prompted the investigation that culminated in the March 2026 complaint (Goldstein, 2025).[25] The complaint states that "NYP's plan restriction generally forbid payers from offering plans that either exclude NYP or offer more generous benefits (e.g., lower copays) when patients choose a rival provider" and "payers cannot resist this exercise of market power because they cannot viably do business in New York City without at least one plan that includes access to NYP's hospitals."

The NYP case extends the pattern established by Atrium and Sutter in two respects. First, whereas Sutter illustrates how a single merger event can trigger anticompetitive contracting, like Atrium, NYP's dominance reflects a sustained program of consolidation, beginning with the 1998 merger of New York Hospital and Presbyterian Hospital and continuing with the subsequent absorption of community hospitals across the boroughs and Westchester, which has made NYP a system no commercial insurer can credibly exclude from its provider network. Second, NYP illustrates the mechanism through which anti-steering provisions generate the extensive-margin cost-sharing effects documented in Subsection V.2: when insurers are contractually barred from steering patients away from a high-priced system, one of their few remaining instruments is to redesign plans with broader patient cost-sharing, shifting the burden of utilization management, which is the insurer's ostensible reason for being, onto patients themselves.

V.4.4 Consistency with the IHME results and connection to patient cost-sharing

The anti-steering mechanism helps explain two features of our IHME results that are otherwise puzzling. First, the private insurance spending increase documented in Subsection V.1 is sustained and apparently amplifying over time, reaching approximately 2.6% per beneficiary five years after treatment. In the absence of anti-steering provisions, we would expect insurer steering to partially offset post-merger price increases over time as insurers renegotiate contracts and redirect members.

[25] Two other unions sued NYP over similar allegations; those cases have since been consolidated (UFCW Local 1500 Welfare Fund v. The New York and Presbyterian Hospital, 2025; Cement and Concrete Workers DC Benefit Fund v. The New York and Presbyterian Hospital, 2025).

The fact that the effect grows rather than attenuates is consistent with merged systems locking in contractual protections against steering that prevent competitive responses from eroding the price increase.

Second, the extensive-margin cost-sharing result documented in Subsection V.2 and Figure VII—more healthcare encounters generating out-of-pocket obligations, without higher per-encounter charges—can be explained by the interaction between anti-steering provisions and insurer plan design. When anti-steering clauses prevent insurers from channeling members to lower-cost providers through network design, one of the few remaining instruments available to insurers is to redesign their benefit plans to include broader patient cost-sharing, which may be attractive to employers as an alternative to increasing their own insurance premiums. By increasing deductibles or imposing co-pays on a wider range of services, insurers can shift the burden of utilization management onto patients themselves: if the insurer cannot deter members from using a high-cost hospital by placing it on a non-preferred tier, it can at least ensure that members internalize more of the cost of each encounter and thereby, hopefully, steer themselves. This institutional logic produces the pattern we observe in the IHME data: a broader set of encounters with patient financial exposure (the extensive margin) without higher per-encounter out-of-pocket charges (the intensive margin), because the plan redesign operates by expanding cost-sharing obligations rather than increasing the rate charged per encounter, which anti-steering language prohibits.

We emphasize that we cannot directly test whether the hospitals in our treatment sample used anti-steering provisions. Hospital–insurer contracts are proprietary and not available at the scale necessary to construct a panel that would enable estimating Equation 1 with anti-steering contractual provisions as the outcome of interest. However, the enforcement record establishes that these provisions are a documented, empirically consequential practice among consolidated hospital systems that achieve market power in local markets by means of mergers. The mechanism is consistent with our quantitative findings and provides a coherent institutional account of how merged systems translate acquired bargaining leverage into sustained price increases and expanded patient financial exposure.

The three quantitative channels documented above (higher prices for commercially insured care, broader patient cost-sharing, and a shift in the payer mix away from Medicare) do not exhaust the set of plausible mechanisms. In Appendix K, we investigate whether an indirect labor market channel contributes to the medical debt effects, namely that hospital mergers weaken local labor markets and thereby reduce patient incomes and ability to pay for out-of-pocket medical expenses. We find no evidence that hospital mergers depress CZ-wide earnings despite providing the upstream condition (higher commercial insurance costs) that Brot et al. (2024) argue should do so. Appendix L addresses a potential threat to identification: hospitals in financial distress are both more likely to be acquired in a merger and more likely to utilize bill-collection practices that generate credit records, meaning that the merger effect we estimate in section IV is in fact the effect of financial distress. Our analysis shows that distressed hospitals are indeed more likely to be acquired, but post-acquisition their financial position improves such that hospital financial distress is unlikely to explain our baseline medical debt findings. Finally, mergers may transmit the acquiring system's debt collection norms (such as more frequent use of lawsuits, wage garnishment, or credit reporting) to newly-acquired facilities, increasing observed medical debt on credit reports even if out-of-pocket costs and the number and amount of unpaid bills are unchanged. In Appendix M we investigate whether this is the case and conclude that it is not operative: in fact, system hospitals collect *less* aggressively than independents, and mergers tend to make previously-independent target hospitals operate more similarly to system hospitals in this respect. Therefore, if anything hospital-level variation in collections practices pushes against the medical debt findings in section IV.

VI Discussion and Conclusion

This paper demonstrates that hospital mergers cause economically meaningful increases in medical debt. Using a staggered difference-in-differences design that exploits variation in the timing and geography of mergers predicted to raise local hospital market concentration, we find that merger exposure increases the likelihood of having medical debt in collections by 11 percent, the flow of

new medical debt by 14 percent (a point estimate that is not statistically significant), and the total outstanding stock by 20 percent, five years after treatment. These effects are robust to alternative sample definitions and treatment thresholds, and are not driven by pre-existing trends. A placebo test using utility and cable/cellular debt in collections yields null results, reinforcing the interpretation that the medical debt effects are specific to the healthcare sector rather than reflecting a broader deterioration in credit conditions in merger-affected areas.

Beyond the medical debt that appears on credit reports as collections tradelines, we find that hospital mergers increase credit card balances along the intensive margin: individuals in merger-exposed areas carry higher credit card balances with almost no change in the probability of having a positive balance. This pattern is consistent with patients financing medical expenses on credit cards, particularly when hospitals condition access to non-emergency care on the resolution of outstanding balances (Kliff and Silver-Greenberg, 2023). The implication is that the medical collections tradelines we study as our primary outcome represent a lower bound on the total financial burden that hospital mergers impose on patients.

Using comprehensive county-level healthcare spending and utilization data from the IHME Disease Expenditure Project (Dieleman et al., 2025), we identify three complementary mechanisms linking hospital mergers to individual medical debt. First, mergers raise the price of commercially insured care. Private insurance spending per beneficiary increases by approximately 2.6% five years after treatment and spending per healthcare encounter increases by 2.9%, while utilization remains stable. This finding is consistent with the large existing literature documenting that hospital mergers confer pricing power over commercial insurers (Cooper et al., 2019; Gowrisankaran, Nevo and Town, 2015), and it confirms that the standard cost-sharing provisions of commercial insurance plans translate those higher negotiated prices into higher out-of-pocket obligations for patients.

Second, mergers expand the set of healthcare encounters for which patients bear direct out-of-pocket financial responsibility. Total out-of-pocket spending rises by roughly 2%, but this increase is driven entirely by a greater number of encounters involving patient payments rather than by

higher charges per encounter. This extensive-margin result is, to our knowledge, novel in the hospital mergers literature. It implies that mergers do not merely make each episode of care more expensive for patients; they also broaden the scope of care on which patients have direct financial exposure, whether through restructured cost-sharing arrangements, the elimination of competing providers where patients had more favorable coverage terms, or other changes to the insurance-provider contracting environment that follow from consolidation.

Third, mergers shift the local payer mix away from Medicare and toward commercial insurance. Medicare enrollment declines by about 0.6% and the Medicare spending share by 1.2 percentage points, while the commercial insurance spending share increases by 1.8 percentage points. Because commercially insured care is subject to the higher prices and broader cost-sharing documented in the first two channels, this re-composition of the payer mix acts as a multiplier: not only are commercial prices higher and out-of-pocket episodes more frequent, but a larger share of local healthcare utilization is subject to commercial pricing in the first place.

These three channels operate in concert. The price channel raises the cost of each episode where patients have financial exposure. The cost-sharing channel expands how many episodes create such exposure. And the payer mix channel ensures that a growing share of care is subject to both forces. Together, they imply that commercial health insurance plans effectively become less generous relative to the total cost of care in the wake of hospital mergers, even absent any explicit change in plan design. This interpretation is consistent with the credit card findings: as patients face both higher and more frequent out-of-pocket obligations, some absorb those costs through credit card borrowing rather than allowing them to proceed to collections, while others cannot pay at all and accumulate medical debt.

A fourth mechanism, anticompetitive provisions in hospital–insurer contracts, provides the institutional foundation for these effects. Anti-steering clauses contractually prevent insurers from directing members to lower-cost competitors, thereby eliminating the demand elasticity that would otherwise discipline post-merger price increases. When merged systems bar insurer steering by contract, one of the few instruments remaining to insurers is to redesign their plans to include

broader patient cost-sharing, shifting the burden of utilization management onto patients and generating the extensive-margin out-of-pocket spending increases documented in our IHME analysis. While we cannot directly test the anti-steering channel with our data, the enforcement record, combined with the consistency between the mechanism and our quantitative findings, strongly suggests that anticompetitive contracting is a central institutional channel through which hospital mergers translate into higher medical debt. One area for future research is whether legal settlements that bar the use of such language by dominant provider networks, state-level bans on their use, and the federal ban on price secrecy enacted in 2020 are effective at reversing their effects (Corlette, Knapp and Pogue, 2026).

Several limitations of our analysis merit discussion. First, our study period ends before the COVID-19 pandemic. The pandemic brought an unprecedented combination of hospital financial distress, emergency federal relief, and further consolidation activity, all of which may have altered the relationship between hospital market structure and medical debt in ways our data cannot capture. Second, we cannot link individual credit records to individual insurance claims, which prevents us from observing the specific insurance plans, cost-sharing terms, and medical services underlying each debt. The IHME data allow us to characterize the aggregate spending and utilization channels at the commuting zone level, but individual-level variation in these mechanisms remains unobserved.

Our findings carry implications for antitrust enforcement and healthcare policy. The prevailing framework for evaluating hospital mergers focuses on the price effects of consolidation—whether mergers raise the prices that hospitals negotiate with commercial insurers. Our results indicate that the consequences of hospital mergers extend well beyond negotiated prices to encompass the broader financial relationship between patients and the healthcare system. The extensive-margin cost-sharing channel and the payer mix shift we document would not be captured by an enforcement framework that focuses narrowly on negotiated prices, yet they contribute materially to the financial burden that patients bear. To the extent that merger review aims to protect consumer welfare, the downstream effects on medical debt and household financial distress warrant consid-

eration alongside the direct effect on negotiated prices.

More broadly, our findings speak to the interaction between market structure and the design of the U.S. healthcare financing system. Medical debt arises in the gap between the cost of care and what insurance covers, and that gap is shaped by the bargaining power of hospitals relative to insurers and patients. Hospital mergers widen this gap from both sides: they raise the cost of care through higher commercial prices, and they erode the effective generosity of insurance coverage by expanding the set of encounters on which patients have financial exposure and by shifting care toward payers with less administered price regulation. The result is a transfer of financial risk from hospitals and insurers to patients. Addressing medical debt as a policy matter therefore requires attention not only to the credit reporting rules and collection practices that determine how medical debt appears on consumers' records, but also to the upstream market structure conditions that generate the debt in the first place.

References

Allegretto, Sylvia A., and Dave Graham-Squire. 2023. "Monopsony in Professional Labor Markets: Hospital System Concentration and Nurse Wages." Institute for New Economic Thinking Working Paper 197.

Andre, Jennifer, et al. 2023. "The Impact of Hospital Closures on Medical Debt in Collections: Analysis Using Consumer Credit Bureau Data."

Andreyeva, Elena, et al. 2024. "The Corporatization of Independent Hospitals." *Journal of Political Economy Microeconomics*, 2(3): 602–663.

Aron-Dine, Aviva, Liran Einav, and Amy Finkelstein. 2013. "The RAND Health Insurance Experiment, Three Decades Later." *Journal of Economic Perspectives*, 27(1): 197–222.

Autor, David H, and David Dorn. 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review*, 103(5): 1553–1597.

Bernhardt, Annette, et al. 2009. "Broken Laws, Unprotected Workers: Violations of Employment and Labor Laws in America's Cities." National Employment Law Project.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess. 2024. "Revisiting Event Study Designs: Robust and Efficient Estimation." *Review of Economic Studies*, 91(6): 3253–3285.

Brot, Zarek C., et al. 2017. "What does a Deductible Do? The Impact of Cost-Sharing on Health Care Prices, Quantities, and Spending Dynamics." *The Quarterly Journal of Economics*, 132(3): 1261–1318.

Brot, Zarek, et al. 2024. "Who Pays for Rising Health Care Prices? Evidence from Hospital Mergers." National Bureau of Economic Research Working Paper 32613.

California v. Sutter Health System. 2000. 84 F. Supp.2d 1057. aff'd, 217 F.3d 846 (9th Cir. 2000), amended by 130 F. Supp. 2d 1109 (N.D. Cal. 2001).

Callaway, Brantly, and Pedro H. C. Sant'Anna. 2021. "Difference-in-Differences with multiple time periods." *Journal of Econometrics*, 225(2): 200–230.

Cement and Concrete Workers DC Benefit Fund v. The New York and Presbyterian Hospital. 2025. No. 1:25-cv-05571 (E.D.N.Y.); complaint filed Oct. 3, 2025 (Cement and Concrete Workers DC Benefit Fund, putative class action); consolidated with UFCW Local 1500 Welfare Fund v. The New York and Presbyterian Hospital, No. 2:25-cv-05023 (E.D.N.Y.), May 8, 2026; consolidated complaint filed June 5, 2026; motion to dismiss filed June 22, 2026; statement of interest filed by the United States July 6, 2026; case pending.

Cooper, Zack, Fiona Scott Morton, and Nathan Shekita. 2020. "Surprise! Out-of-Network Billing for Emergency Care in the United States." *Journal of Political Economy*, 128(9): 3626–3677.

Cooper, Zack, et al. 2019. "The Price Ain't Right? Hospital Prices and Health Spending on the Privately Insured." *The Quarterly Journal of Economics*, 134(1): 51–107.

Cooper, Zack, et al. 2024. "Measuring Hospital Ownership Transitions in the United States." Dataset.

Corlette, Sabrina, Abigail Knapp, and Stacey Pogue. 2026. "State Spotlight: Texas Confronts Anti-Competitive Tactics That Drive Up Health Care Costs."

Council of Economic Advisors. 2026. "Effects of Banning Anti-Competitive Hospital Contracts."

de Chaisemartin, Clément, and Xavier D'Haultfoeuille. 2023. "Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: a survey." *The Econometrics Journal*, 26(3): C1–C30.

Dieleman, Joseph L., et al. 2025. "Tracking US Health Care Spending by Health Condition and County." *Journal of the American Medical Association*, 333(12): 1051–1061.

Dorn, David. 2021. "FIPS County Code Changes." Dataset.

Dranove, David, Sungjin Kim, and Benjamin Vatter. 2026. "Mergers with Intermediation: the Price Effect of Distant Hospital Mergers."

Duarte, Victor, et al. 2026. "The Effects of Deleting Medical Debt from Consumer Credit Reports."

Gibbs, Christa, et al. 2025. "Consumer Credit Reporting Data." *Journal of Economic Literature*, 63(2): 598–636.

Goldstein, Joseph. 2025. "U.S. Opens Antitrust Investigation Into NewYork-Presbyterian." *The New York Times*.

Goodman-Bacon, Andrew. 2021. "Differences-in-Differences with Variation in Treatment Timing." *Journal of Econometrics*, 225(2): 254–277.

Gowrisankaran, Gautam, Aviv Nevo, and Robert Town. 2015. "Mergers When Prices Are Negotiated: Evidence from the Hospital Industry." *American Economic Review*, 105(1): 172–203.

Kabrick, Isaac, and Katie Keith. 2026. "Challenges To Anti-Competitive Hospital Contracting Practices: The New York And Presbyterian Hospital."

Kaiser Family Foundation. 2025. "Status of State Medicaid Expansion Decisions."

Kenworthy, Nora. 2024. *Crowded Out: The True Costs of Crowdfunding Healthcare*. Cambridge, Massachusetts:The MIT Press.

Kliff, Sarah, and Jessica Silver-Greenberg. 2023. "Nonprofit Health System Ends Practice of Denying Care to Patients in Debt." *The New York Times*.

Klosin, Sylvia. 2024. "Dynamic Biases of Static Panel Data Estimators." *Working paper*.

Kluender, Raymond, et al. 2021. "Medical Debt in the US, 2009-2020." *JAMA*, 326(3): 250–256.

Kluender, Raymond, et al. 2025. "The Effects of Medical Debt Relief: Evidence from Two Randomized Experiments." *The Quarterly Journal of Economics*, 140(2): 1187–1241.

Lazear, Edward P., and Kristin McCue. 2018. "What Causes Labor Turnover To Vary?" Na-

tional Bureau of Economic Research Working Paper 24873.

Lown Institute. 2026. "Hospital financial assistance and collections policies." Lown Institute Hospitals Index.

Mathews, Anna Wilde. 2024. "Hospital to Union: Pay Up or You're Stuck With Us in Your Health Plan." *Wall Street Journal*.

Messac, Luke. 2024. *Your Money or Your Life: Debt Collection in American Medicine.* New York:Oxford University Press.

Nickell, Stephen. 1981. "Biases in Dynamic Models with Fixed Effects." *Econometrica*, 49(6): 1417–1426.

Ohio v. American Express Co. 2018. 585 U.S. 529 (2018).

Prager, Elena. 2020. "Healthcare Demand under Simple Prices: Evidence from Tiered Hospital Networks." *American Economic Journal: Applied Economics*, 12(4): 196–223.

Prager, Elena, and Matt Schmitt. 2021. "Employer Consolidation and Wages: Evidence from Hospitals." *American Economic Review*, 111(2): 397–427.

Sun, Liyang, and Sarah Abraham. 2021. "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects." *Journal of Econometrics*, 225(2): 175–199.

Tenn, Steven. 2011. "The Price Effects of Hospital Mergers: A Case Study of the Sutter-Summit Transaction." *International Journal of the Economics of Business*, 18(1): 65–82.

Tolbert, Jennifer, et al. 2025. "The Uninsured Population and Health Coverage."

UFCW & Employers Benefit Trust v. Sutter Health. 2021. No. CGC-14-538451 (Cal. Super. Ct. 2021), consolidated with California ex rel. Becerra v. Sutter Health, No. 18-565398; AG complaint filed Mar. 2018; settlement filed Dec. 19, 2019; final judgment entered Aug. 27, 2021 (Judge Anne-Christine Massullo).

UFCW Local 1500 Welfare Fund v. The New York and Presbyterian Hospital. 2025. No. 2:25-cv-05023 (E.D.N.Y.); complaint filed Sept. 8, 2025 (UFCW Local 1500 Welfare Fund, putative class action); consolidated with Cement and Concrete Workers DC Benefit Fund v. The New York and Presbyterian Hospital, No. 1:25-cv-05571 (E.D.N.Y.), May 8, 2026; consolidated complaint filed June 5, 2026; motion to dismiss filed June 22, 2026; statement of interest filed by the United States July 6, 2026; case pending.

United States and State of Ohio v. OhioHealth Corporation. 2026. No. 2:26-cv-207 (S.D.O.H.), also known as United States and State of Ohio v. OhioHealth; complaint filed Feb. 20, 2026; settled.

United States v. Charlotte-Mecklenburg Hospital Authority. 2019. No. 3:16-cv-00311-RJC-DCK (W.D.N.C. 2019), also known as United States v. Atrium Health; complaint filed June 9, 2016 (DOJ and State of North Carolina); motion to dismiss denied Mar. 30, 2017; settlement announced Nov. 15, 2018; final judgment entered Apr. 24, 2019.

United States v. New York Presbyterian Hospital. 2026. No. 1:26-cv-02480 (S.D.N.Y.), also

known as United States v. NewYork-Presbyterian; complaint filed Mar. 26, 2026 (DOJ Antitrust Division and U.S. Attorney's Office for the S.D.N.Y.); case pending.

US Census Bureau. 2021. "Substantial Changes to Counties and County Equivalent Entities: 1970-Present." Dataset.

US Census Bureau. 2023. "Urban and Rural."

U.S. Department of Housing and Urban Development. 2019. "HUD-USPS ZIP Code Crosswalk Files." Dataset.

U.S. Department of Justice, and Federal Trade Commission. 2023. "Merger Guidelines."

Weil, David. 2014. *The Fissured Workplace: Why Work Became So Bad for So Many and What Can Be Done to Improve It.* Cambridge, MA:Harvard University Press.

Tables and Figures

Table I. SUMMARY STATISTICS FOR INDIVIDUAL-LEVEL DATA, 2013

	2014 cohort	2015 cohort	2016 cohort	2017 cohort	2018 cohort	Control CZs
Age	47.3	47.7	48.3	47.5	47.9	48.5
Female	44%	44%	45%	44%	44%	45%
Male	44%	44%	44%	44%	45%	45%
Has medical debt (adj)	18%	22%	23%	20%	26%	23%
Has medical debt (raw)	17%	21%	22%	19%	25%	22%
New medical debt, last 12m (adj)	161	182	216	176	214	219
New medical debt, last 12m (raw)	152	179	222	179	202	208
Total medical debt (adj)	477	558	693	507	804	708
Total medical debt (raw)	424	499	638	457	735	628
Has non-med debt (adj)	31%	33%	35%	31%	37%	34%
Has non-med debt (raw)	30%	33%	34%	31%	36%	33%
New non-med debt, last 12m (adj)	384	388	474	386	420	407
New non-med debt, last 12m (raw)	376	389	467	386	413	402
Total non-med debt (adj)	1,313	1,291	1,665	1,251	1,575	1,407
Total non-med debt (raw)	1,231	1,207	1,555	1,162	1,498	1,304
Has utility debt (adj)	13%	14%	14%	13%	16%	14%
Has utility debt (raw)	12%	13%	13%	12%	14%	13%
New utility debt, last 12m (adj)	48	52	56	45	55	47
New utility debt, last 12m (raw)	45	48	52	41	49	43
Total utility debt (adj)	108	113	120	109	120	103
Total utility debt (raw)	98	104	106	98	109	94
Has credit card debt (raw)	66%	65%	66%	66%	61%	62%
Δ-credit card debt, last 12m (raw)	-18.8	2.4	-25	-0.2	18.6	-5.8
Total credit card debt (raw)	3,949	3,942	4,069	3,978	3,601	3,376
Observations	1,734,941	659,227	267,041	311,076	53,674	1,398,797

Notes: This table displays demographic characteristics and outcomes of interest across treatment and control cohorts in 2013, one year before the treatment start of the treatment period. Female and Male shares do not sum to 100% because gender is missing for some of the individuals in the sample. For the outcomes related to medical debt in collections, non-medical debt in collections, and utility/cable/cellular debt in collections, both the raw and adjusted amounts are shown. For credit card debt, no such adjustments were made, and therefore only the raw amounts are presented.

Table II. SUMMARY STATISTICS FOR COMMUTING ZONE-LEVEL DATA, 2013

	2014 cohort	2015 cohort	2016 cohort	2017 cohort	2018 cohort	Control CZs
Panel A: Spending by payer ($ millions)						
Total spending	9,365.5	8,970.8	7,178.2	7,161.3	1,746.9	862.6
Private ins. spending	4,024.7	4,037.5	3,049.3	3,073.6	658.4	337.4
Medicare spending	2,502.2	2,278.4	1,996.6	1,811.3	510.6	252.2
Medicaid spending	1,626.5	1,453.8	1,147.8	1,244.9	378.1	164.5
OOP spending	1,212.1	1,201.1	984.5	1,031.4	199.8	108.5
Panel B: Spending shares						
Private ins. spending share	38.3%	39.6%	39.8%	39.1%	35.9%	36.8%
Medicare spending share	30.2%	29.7%	30.0%	28.0%	30.6%	31.2%
Medicaid spending share	19.2%	18.0%	17.5%	20.5%	22.8%	19.8%
OOP spending share	12.3%	12.7%	12.7%	12.4%	10.8%	12.2%
Panel C: Enrollment						
Private ins. enrollment	1,051,611	1,061,665	780,005	873,681	223,672	106,867
Medicare enrollment	244,155	237,057	196,845	186,669	56,400	28,994
Medicaid enrollment	318,699	268,702	211,063	224,062	70,019	33,081
Panel D: Encounters (millions)						
Private ins. encounters	10.37	8.57	8.06	7.69	1.69	0.82
Medicare encounters	23.71	17.91	17.22	16.03	5.00	2.27
Medicaid encounters	5.04	3.64	3.46	3.60	1.19	0.53
OOP encounters	11.37	8.45	8.88	8.67	2.07	0.96
Panel E: Spending per enrollee ($)						
Private ins. spending per enrollee	3,189	3,353	3,598	3,237	2,817	2,961
Medicare spending per enrollee	9,111	9,333	9,790	9,011	8,981	8,048
Medicaid spending per enrollee	5,212	5,302	5,219	5,397	5,611	4,941
Panel F: Spending per encounter ($)						
Private ins. spending per encounter	381	406	388	387	387	414
Medicare spending per encounter	112	123	117	108	109	120
Medicaid spending per encounter	295	369	312	316	316	317
OOP spending per encounter	107	122	113	108	95	112
Number of CZs	64	34	12	17	10	556
Average HHI per CZ	2,096	2,250	2,236	2,331	3,725	5,173
Total cohort population	110,481,434	44,579,791	14,918,122	20,974,148	3,662,797	95,918,268
Average population per CZ	1,726,272	1,311,170	1,243,177	1,233,773	366,280	172,515

Notes: This table displays unweighted CZ-level means of healthcare spending variables from the IHME Disease Expenditure Project (Dieleman et al., 2025), by treatment cohort and for control CZs, in 2013, as well as HHI. Spending, enrollment, and encounter counts are CZ-level totals; spending is measured in 2019 US dollars and is age- and sex-standardized. Spending and encounter variables are CZ-level totals computed across five types of care (ambulatory, dental, emergency department, inpatient, and nursing facility), excluding home health and pharmacy; see appendix B.3 for discussion. Enrollment counts reflect payer enrollment (the number of individuals enrolled in each payer category), derived from the ratio of total spending to spending per enrolled individual on ambulatory care reported by the IHME. Enrollment by payer is not mutually exclusive. Medicare is inclusive of Medicare Advantage. Population data are from the U.S. Census Bureau intercensal county population estimates.

Figure I. MEDICAL AND NON-MEDICAL COLLECTIONS OUTCOMES, 2011–2019

I.1 Share with debt in collections

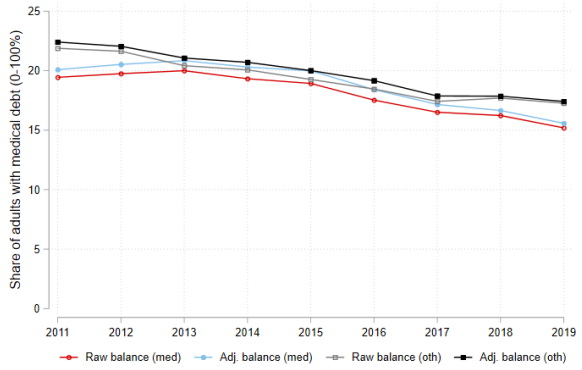

I.2 New debt in collections (last 12 months)

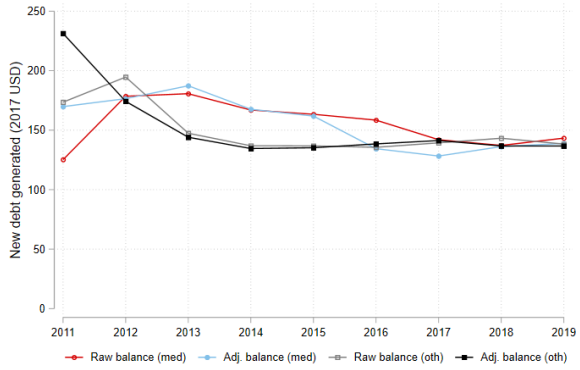

I.3 Total debt in collections

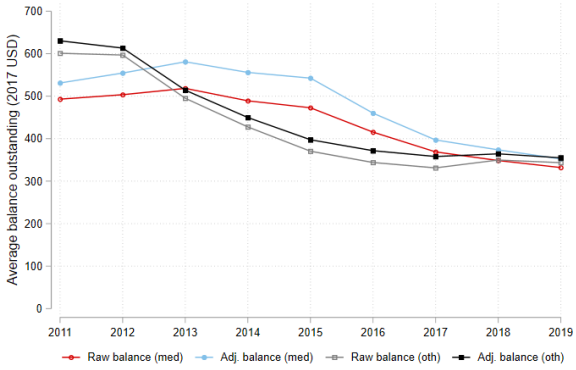

This figure shows the evolution of each of the three key outcomes for the 2011–2019 period, for both medical and non-medical debt in collections. Each panel shows both the 'raw' and 'adjusted' estimates for both medical and non-medical debt. Panel I.1 shows the share of individuals that have medical and non-medical debt in collections, respectively. Panel I.2 shows the total amount of new medical debt in collections and new non-medical debt in collections generated in the previous 12 months, measured in constant 2017 USD, regardless of whether it was subsequently paid. Panel I.3 shows the total outstanding amount of medical debt in collections and non-medical debt in collections, again measured in constant 2017 USD.

Figure II. EVENT STUDY PLOTS FOR THE EFFECT OF HOSPITAL MERGER EXPOSURE ON CZ-LEVEL HHI

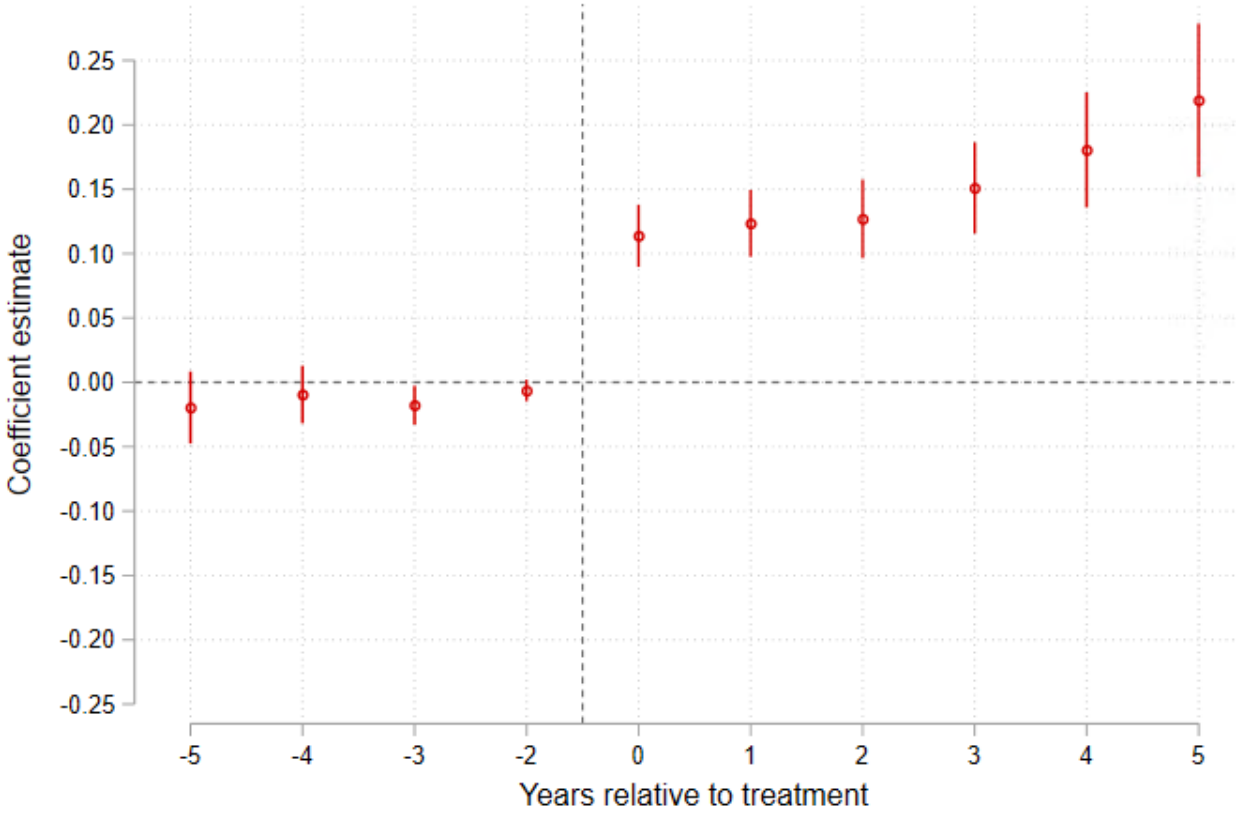

This figure shows the point estimates and 95% confidence intervals for $\hat{\beta}_e$, obtained from estimating Equation 1 when the outcome of interest is the log of commuting-zone-level HHI. The estimates include year and CZ fixed effects, as well as a binary control that accounts for the timing of state-level Medicaid expansion, and the logs of the total population and the population aged 65+, as outlined in Section III.1; standard errors are clustered at the CZ level. The horizontal axis denotes years relative to the merger event; event time $t = -1$ is omitted as the reference period, with the dashed vertical line indicating the period just before the merger took place.

Figure III. EVENT STUDY PLOTS FOR THE EFFECT OF HOSPITAL MERGER EXPOSURE ON MEDICAL DEBT

III.1 Likelihood of having medical debt

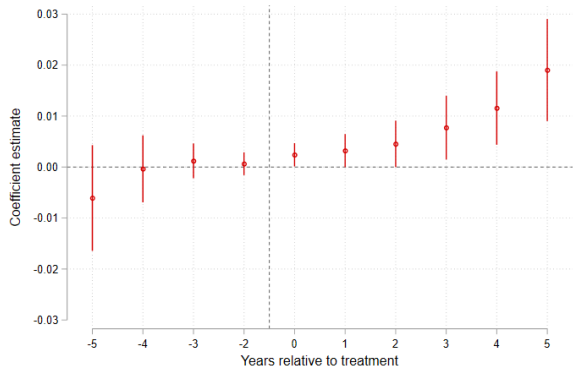

III.2 New medical debt (last 12 months)

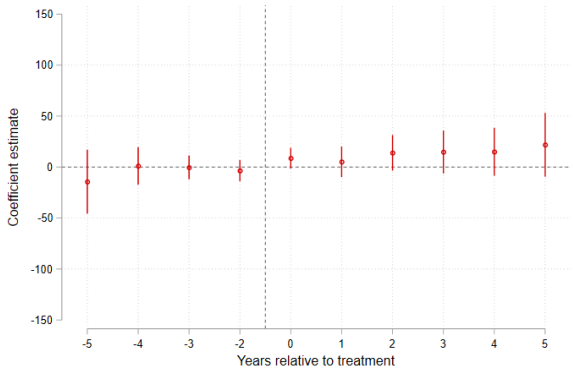

III.3 Total medical debt

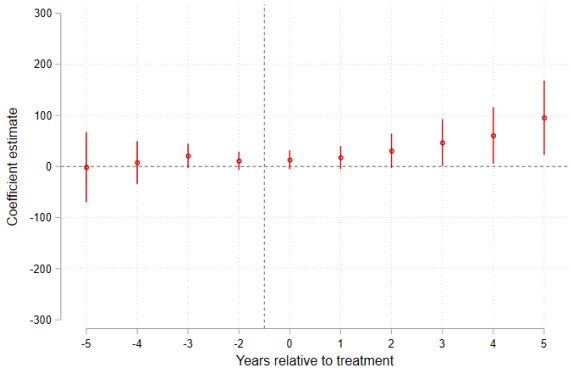

Each panel shows the point estimates and 95% confidence intervals for $\hat{\beta}_e$, obtained from estimating Equation 2, and illustrates the effect of exposure to hospital mergers on medical debt outcomes. The estimates included year, CZ, and individual fixed effects, as well as a binary control that accounts for the timing of state-level Medicaid expansion, as outlined in Section III.1; standard errors are clustered at the CZ level. The outcome for Panel III.1 is a binary variable that takes a value of 1 if an individual has medical debt in collections, and 0 otherwise. The outcome for Panel III.2 is the total amount of new medical debt in collections (measured in constant 2017 USD) generated in the previous 12 months, regardless of whether it was subsequently paid. The outcome for Panel III.3 is the total outstanding amount of medical debt in collections (measured in constant 2017 USD). The horizontal axis denotes years relative to the merger event; event time $t = -1$ is omitted as the reference period, with the dashed vertical line indicating the period just before the merger took place. The full regression results for each panel are shown in column (3) of Tables E.I, E.II, and E.III, respectively.

Figure IV. EVENT STUDY PLOTS FOR THE EFFECT OF HOSPITAL MERGER EXPOSURE ON CREDIT CARD DEBT

IV.1 P(credit card balance > 0)

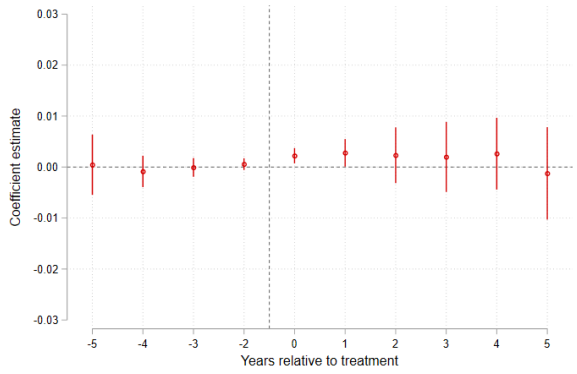

IV.2 Δ-credit card balance (last 12 months)

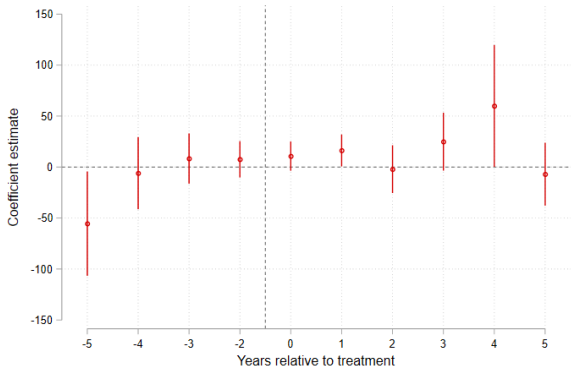

IV.3 Total credit card balance

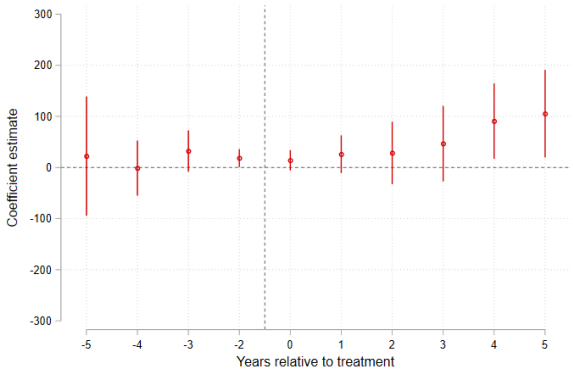

IV.4 Share of delinquent credit card balance

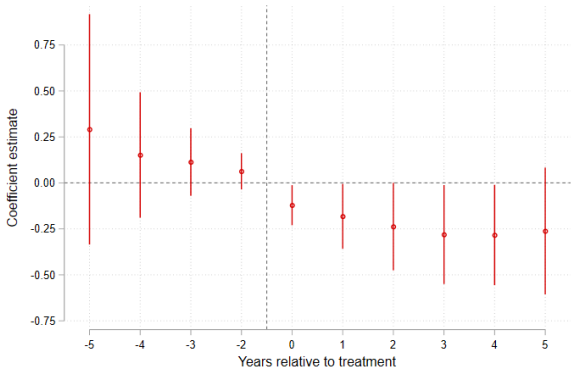

Each panel shows the point estimates and 95% confidence intervals for $\hat{\beta}_e$, obtained from implementing Equation 2, and illustrates the effect of exposure to hospital mergers on credit card debt outcomes. The estimates included year, CZ, and individual fixed effects, as well as a binary control that accounts for the timing of state-level Medicaid expansion, as outlined in Section III.1; standard errors are clustered at the CZ level. The outcome for Panel IV.1 is a binary variable that takes a value of 1 if an individual has a positive credit card balance, and 0 otherwise. The outcome for Panel IV.2 is the change in an individual's credit card balance over the last 12 months, i.e., between the end of the current year and the end of the previous year, measured in constant 2017 USD. It is constructed using a first difference, so it not defined in the first year an individual is observed or in the year right after a missing year in the individual's panel; consequently, event time $t = -5$ is identified only from the 2017 and 2018 cohorts, since the year-on-year change is not defined for 2011. The outcome for Panel IV.3 is the total outstanding credit card balance (also measured in constant 2017 USD). The outcome for Panel IV.4 is defined as the delinquent credit card balance divided by the total outstanding credit card balance. The horizontal axis denotes years relative to the merger event; event time $t = -1$ is omitted as the reference period, with the dashed vertical line indicating the period just before the merger took place.

Figure V. EVENT STUDY PLOTS FOR THE EFFECT OF HOSPITAL MERGER EXPOSURE ON UTILITY, CABLE, AND CELLULAR DEBT IN COLLECTIONS

V.1 Likelihood of having utility debt in collections

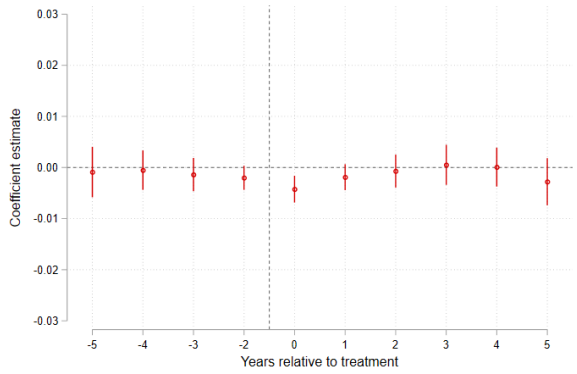

V.2 New utility debt in collections (last 12 months)

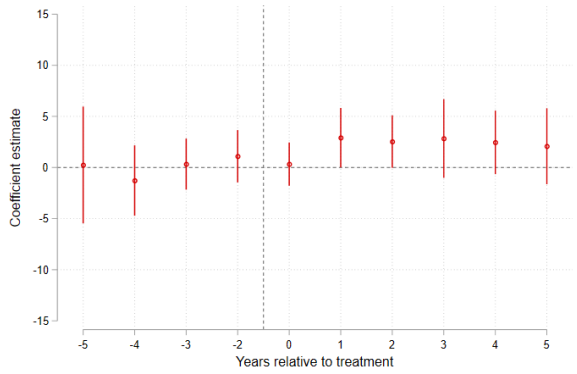

V.3 Total utility debt in collections

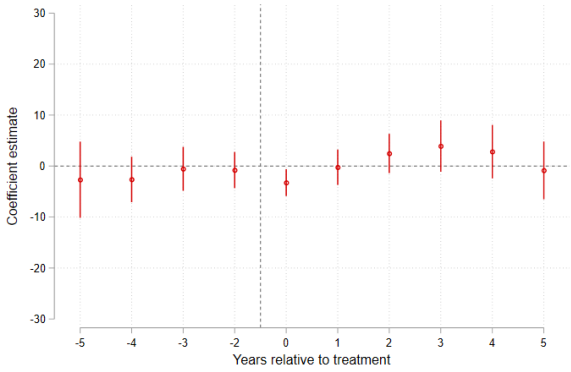

Each panel shows the point estimates and 95% confidence intervals for $\hat{\beta}_e$, obtained from implementing Equation 2, and illustrates the effect of exposure to hospital mergers on outcomes related to utility and cable/cellular debt in collections. The estimates included year, CZ, and individual fixed effects, as well as a binary control that accounts for the timing of state-level Medicaid expansion, as outlined in Section III.1; standard errors are clustered at the CZ level. The outcome for Panel V.1 is a binary variable that takes a value of 1 if an individual has utility and cable/cellular debt in collections, and 0 otherwise. The outcome for Panel V.2 is the total amount of new utility and cable/cellular debt in collections (measured in constant 2017 USD) generated in the previous 12 months, regardless of whether it was subsequently paid. The outcome for Panel V.3 is the total outstanding amount of utility and cable/cellular debt in collections (measured in constant 2017 USD). The horizontal axis denotes years relative to the merger event; event time $t = -1$ is omitted as the reference period, with the dashed vertical line indicating the period just before the merger took place.

Figure VI. EVENT STUDY PLOTS FOR THE EFFECT OF HOSPITAL MERGER EXPOSURE ON PRIVATELY INSURED HEALTHCARE SPENDING AND UTILIZATION

VI.1 Log private insurance spending

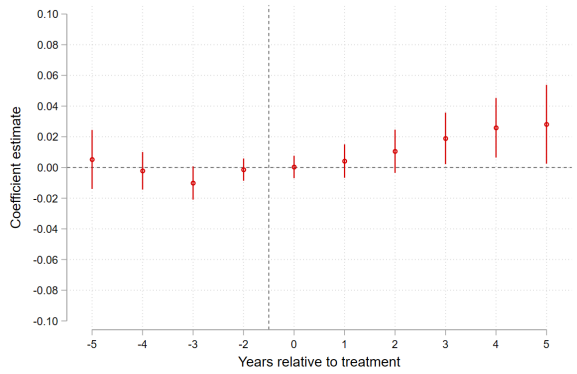

VI.2 Log private insurance encounters

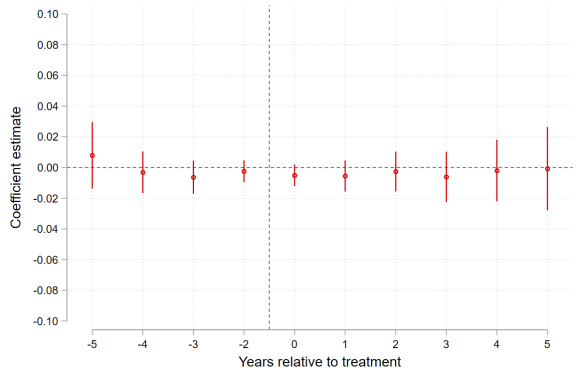

VI.3 Log private ins. spending per beneficiary

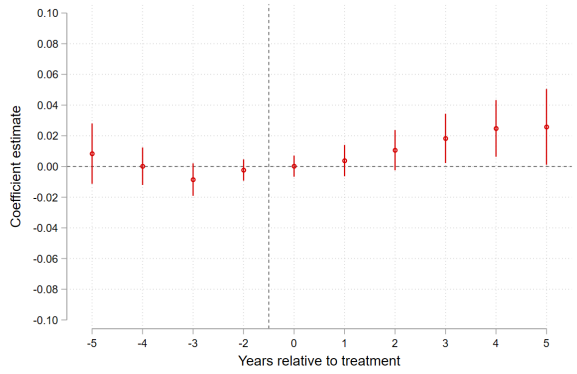

VI.4 Log private ins. spending per encounter

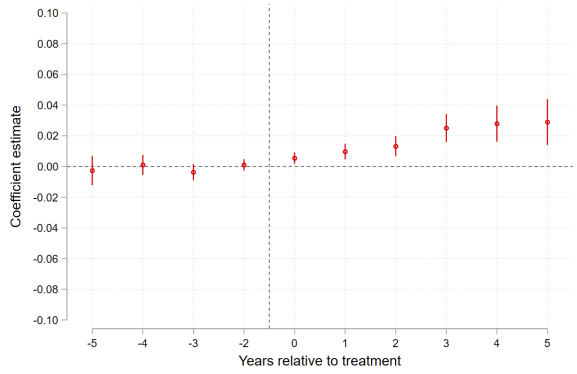

Each panel shows the point estimates and 95% confidence intervals for $\hat{\beta}_e$, obtained from estimating Equation 1 using the IHME Disease Expenditure data. All outcomes are log-transformed. Panel VI.1 uses total private insurance spending as y_{zt}. Panel VI.2 uses the number of privately-insured healthcare encounters. Panel VI.3 reports spending per private insurance beneficiary. Panel VI.4 uses spending per privately-insured healthcare encounter. All regressions include CZ and year fixed effects, a binary control for state-level Medicaid expansion, and time-varying controls for the log of CZ population and population aged 65+. Standard errors are clustered at the CZ level. The horizontal axis denotes years relative to the merger event; event time $t = -1$ is omitted as the reference period, with the dashed vertical line indicating the period just before the merger took place. Regression results are reported in Table J.I.

Figure VII. EVENT STUDY PLOTS FOR THE EFFECT OF HOSPITAL MERGER EXPOSURE ON OUT-OF-POCKET HEALTHCARE SPENDING AND UTILIZATION.

VII.1 Log OOP spending

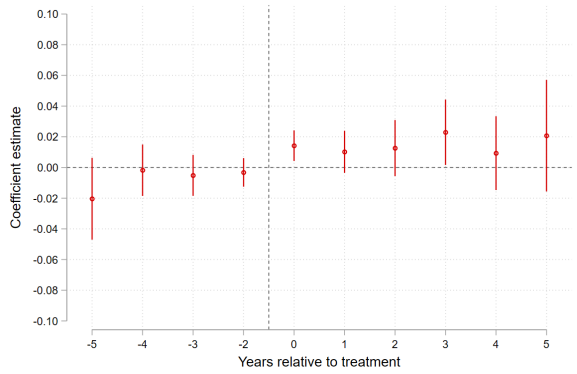

VII.2 Log OOP encounters

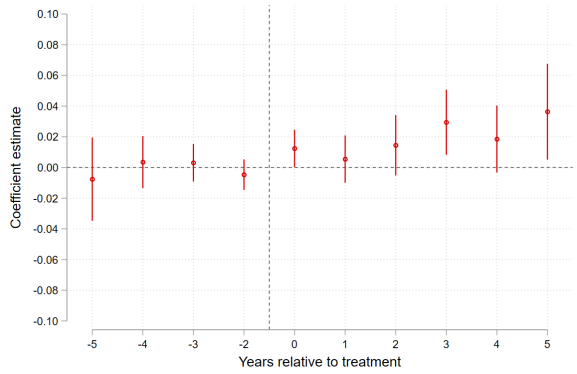

VII.3 Log OOP spending per encounter

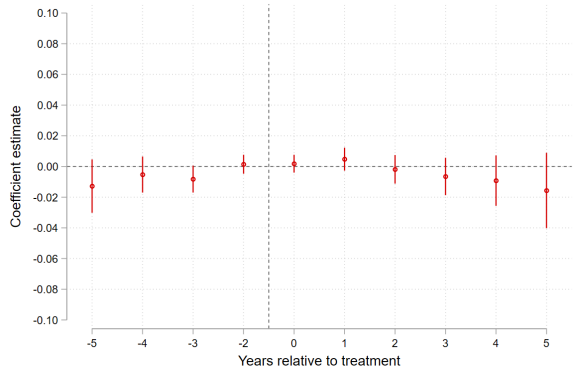

VII.4 OOP spending share

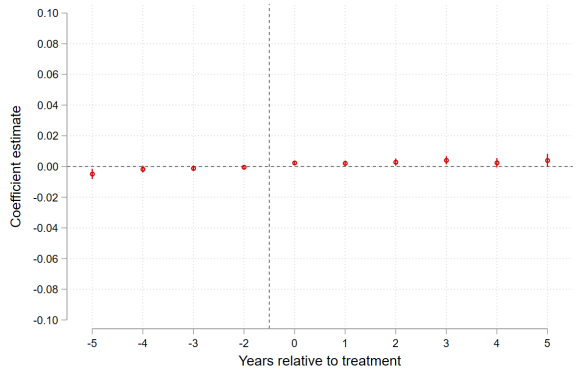

Each panel shows the point estimates and 95% confidence intervals for $\hat{\beta}_e$, obtained from estimating Equation 1 using the IHME Disease Expenditure data. Panels VII.1 and VII.3 use log-transformed spending outcomes; Panel VII.2 uses the log of the count of OOP encounters; Panel VII.4 uses the OOP share of total healthcare spending. All regressions include CZ and year fixed effects and a binary control for state-level Medicaid expansion as well as the log of CZ population and population aged 65+; standard errors are clustered at the CZ level. The horizontal axis denotes years relative to the merger event; event time $t = -1$ is omitted as the reference period, with the dashed vertical line indicating the period just before the merger took place. Regression results are reported in Table J.II.

Figure VIII. EVENT STUDY PLOTS FOR THE EFFECT OF HOSPITAL MERGER EXPOSURE ON MEDICARE AND PRIVATE INSURANCE COVERAGE.

VIII.1 Log Medicare enrollment

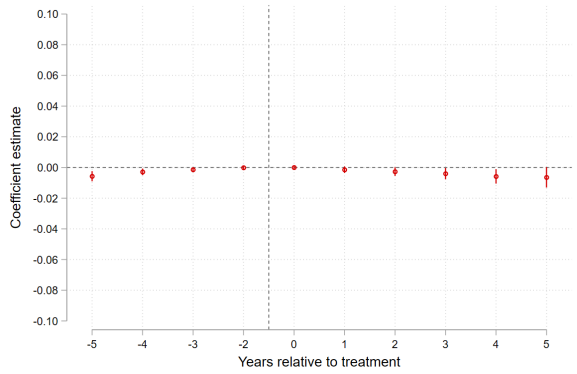

VIII.2 Log Private Insurance Enrollment

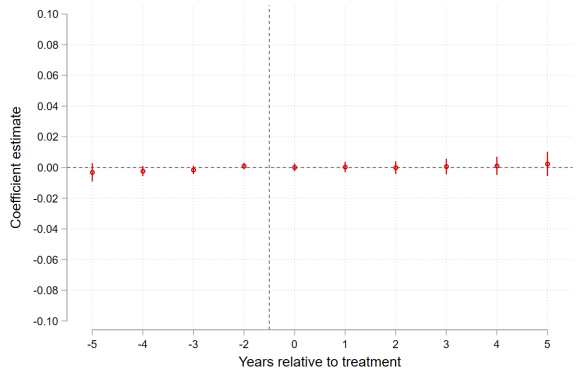

VIII.3 Medicare spending share

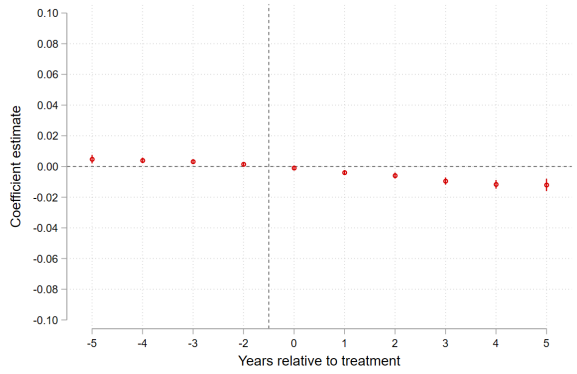

VIII.4 Private insurance spending share

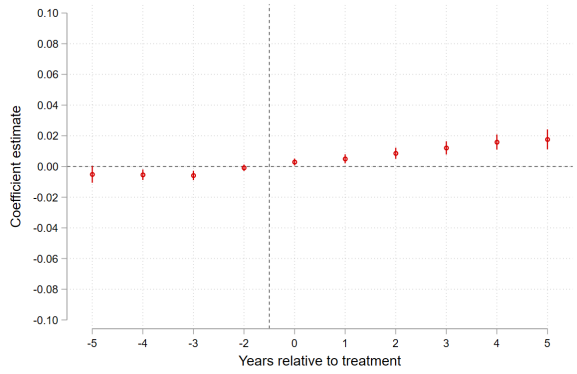

Each panel shows the point estimates and 95% confidence intervals for $\hat{\beta}_e$, obtained from estimating Equation 1 using the IHME Disease Expenditure data. Panel VIII.1 uses the log of Medicare enrollment as y_{zt}; Panel VIII.2 uses log private insurance enrollment; Panels VIII.3 and VIII.4 use the Medicare and private insurance shares of total healthcare spending, respectively. Medicare is inclusive of Medicare Advantage. All regressions include CZ and year fixed effects and a binary control for state-level Medicaid expansion as well as the log of CZ population and population aged 65+. The horizontal axis denotes years relative to the merger event; event time $t = -1$ is omitted as the reference period, with the dashed vertical line indicating the period just before the merger took place. Regression results are reported in Table J.III.

Appendices

A Data extraction and cleaning **A3**

A.1 Data extraction A3

A.2 Clean attribute files A4

A.3 Clean tradeline files for medical debt in collections A8

A.4 Clean tradeline files for non-medical debt in collections and for utility, cable, and cellular debt in collections A11

A.5 Clean tradeline files for credit card debt A11

B The IHME Disease Expenditure Data **A13**

B.1 Background and methodology A13

B.2 Aggregation to commuting zones A15

B.3 Selection of types of care A15

C Event study plots for medical debt outcomes with placebo treatment dates **A20**

D Event study plots for "raw/unadjusted" medical debt outcomes **A21**

E Event study regression tables for medical debt, baseline specifications **A22**

F Alternative control group **A25**

G Event study regression tables for medical debt, specifications with lagged dependent variable **A27**

H Heterogeneity dimensions **A30**

H.1 Gender A30

H.2 Medicaid expansion states A32

H.3 Subset with medical debt in collections in 2013 . . . . . A34

J Event study regression tables and supplemental results for the IHME mechanisms analysis **A35**

J.1 Event study regression tables A35

J.2 Hospital mergers and the uninsured population A39

K The Labor Market Channel **A41**

L The Financial Distress Channel **A47**

M Hospital Debt-Collection Practices **A51**

M.1 Data and linkage A51

M.2 Empirical approach A53

M.3 Results A53

M.4 Independent hospitals collect unpaid bills differently than system hospitals A56

M.5 Collection policy is set at the system level A58

M.6 Interpretation A59

A. Data extraction and cleaning

A.1 Data extraction

A.1.i Attribute files

- Extracted Q4 for each year from 2011 to 2024. Even though the last period used is 2019Q4, the additional time periods are useful when making some demographic corrections (e.g., gender or race).

- The extraction is limited to "Primary Records", i.e., the observations where the variable "source_cd_pr" is defined as "PR".

A.1.ii Tradeline files

- Extracted Q4 from 2011 to 2024, although only the 2011-2019 period is used the paper. The reason for the additional years is that having a longer coverage, especially post-2019, allows us to better implement the adjustments detailed in section A.3. We focus on the last quarter of every year because the hospital data from which we obtain merger information is defined at the yearly level.

- **Medical debt in collections:** these tradelines are categorized as being in "Collection" or with a "Debt buyer", according to the "loan_type" variable, created by UCCCP and based on the classification provided by variables that identify the account type ("enhanced_acct_type_cd") and that identify the creditor's business code ("kob_cd"). In addition, the following conditions need to be met:

  - The variable for the original creditor classification code ("orig_creditor_cls_cd") is coded as "02" ("Medical/Health Care").

  - The variable "enhanced_acct_type_cd" is coded as "90" ("Medical debt").

  - The variable "kob_cd" is set to "MG" ("Medical group"), "MH" ("Hospitals and Clinics"), or "MZ" ("Medical and related health - nonspecific").

- **Utility, cable, and cellular debt in collections:** these tradelines are categorized as being in "Collection" or with a "Debt buyer", according to the "loan_type" variable. In addition, the following conditions need to be met:
  - The variable for the original creditor classification code ("orig_creditor_cls_cd") is coded as "10" ("Utilities") or "11" ("Cable/Cellular").
  - The variable for the enhanced account type code ("enhanced_acct_type_cd") is coded as "92" ("Utility company") or "4D" ("Telecommunications/Cellular").
  - The variable identifying the creditor's business code ("kob_cd") is set to "UA" ("Water utilities/Bottled water"), "UC" ("Cable TV provider"), "UD" ("Garbage and rubbish disposal companies"), "UE" ("Electric light and power companies"), "UF" ("Fuel oil distributors"), "UG" ("Gas companies, natural and bottled"), "UH" ("Coal and wood suppliers"), "UL" ("Long-distance telephone service"), "UO" ("Online/internet access"), "UP" ("Cellular and paging company"), "UR" ("Waste recycling/Waste handlers"), "US" ("Satellite TV/Direct broadcast"), "UT" ("Local telephone service provider"), "UV" ("Home security company"), "UW" ("Wireless telephone service"), or "UZ" ("Utilities and fuel - nonspecific").
- **Non-medical debt in collections:** these correspond to the set of tradelines that are categorized as being in "Collection" or with a "Debt buyer", according to the "loan_type" variable, but identified as not being medical debt, based on the criteria defined earlier in this section.
- **Credit card debt:** these tradelines are identified based on the variable "loan_type". Specifically, we consider all the tradelines identified as "Credit card" or "Retail credit card" by that variable.

A.2 Clean attribute files

We adopt the following procedures to improve the consistency and reliability of the key demographic variables included in the attribute files, as well as to address the issues caused by frag-

mented files, i.e., partial files that should not be in the data, such as cases where two people are present with different unique identifiers because the credit bureau is unable to link the two records (Gibbs et al., 2025):

- **Birth date:** initially missing for 21% of observations. We correct this in two main ways. First, for each individual, we assign them their modal birth date. Second, in cases where only the birth date is missing but the birth year is present, we assign January as the birth month. With this procedure, we reduce the share of observations with missing birth date to 9%.

- **Age:** this is a variable created by UCCCP, but it is not clear how it is being created, i.e., how age is being imputed when the birth date variable is missing. We correct this variable (regardless of whether it was previously missing or not) based on the updated birth date variable and the reporting date associated with the corresponding attribute file. This results in edits to 20% of the observations and reduces the share of missing observations from 10% to 9%.

- **Drop fragmented files:** this is done in 3 steps, following UCCCP's procedure (in total, these drop about 17% of the observations):
  - Drop observations whose credit score is below 300.
  - Drop observations with a missing birth date.
  - Drop observations where the individual is listed as deceased by the "deceased_in" variable.

- **Drop individuals inconsistently tagged as deceased:** this is defined as being identified as not deceased in year t after being identified as deceased in year $t-1$. This procedure results in dropping approximately 0.9% of the individuals in the sample.

- **Gender:** initially missing for about 23% of observations. To reduce this share as well as the cases of inconsistent gender coding within individuals, we assign each individual their modal gender during the period. In cases where there is more than one modal gender, we set it to missing. With this procedure, we reduce the share of observations with missing gender to 13%.

- **Race:** this variable is coded in UCCCP as a probability of being non-Hispanic white, non-Hispanic Black, Hispanic, non-Hispanic Asian, non-Hispanic American Indian or Alaska Native, or non-Hispanic multiracial or another category. It is initially missing for 87% of observations, as it is only assigned in 2023 and 2024. For each individual where this probability is non-missing, we take the average of the 2023 and 2024 values and assign the race based on the highest probability value. With this procedure, we reduce the share of observations with missing race to 23%.

- **Location:**

  - UCCCP includes several geographic location variables, including information on the state, county, and zip code where each individual is located. We take the zip code variable as the most reliable. Ultimately, we want to assign each individual to a commuting zone.

  - We use the zip-to-county crosswalks available at the US HUD's Office of Policy Development and Research for Q4 of each year to assign each zip to a county. For about 75% of all US zip codes, they are fully contained within a county, such that the match is exact. In the remaining cases, we assign each zip to the county that has the highest share of total addresses from that zip; in the cases where there is a tie, we drop those zip codes. We use HUD's zip-to-county crosswalk to assign individuals in the UCCCP data to a "fips_code." As this crosswalk covers the near universe of observations in UCCCP, this approach allows us to reduce the share of observations with a missing

fips_code to less than 0.1%. To address those remaining observations without a county identifier, we take the following steps:

  * For each individual that (i) has a missing fips_code in year t but non-missing in $t-n$ or $t+n$, and that (ii) has not changed zip code between those years: assign the fips_code from year $t-n$ or $t+n$ to year t.
  * In the cases where a given zipcode is always assigned to the same county in the UCCCP data but the fips_code is missing for that observation (e.g., because it is a zip code that no longer exists in a given year t and therefore does not have a match in the yearly zip-to-county crosswalk, we assign the modal fips_code that corresponds to that zip.
  * We generate an alternative county identifier variable by concatenating the variables "state_abb" and "fips_county_cd" that are present in the UCCCP raw data. We then take, for each zip, the corresponding modal county identifiers that result both from the "fips_code" variable obtained from the zip-to-county crosswalks and from this alternative county identifier. In the cases where the modal value from both approach is identical, we assign that county.
  * We then assign the alternative county identifier.
  * We again repeat the procedure described above where, for each individual that (i) has a missing "fips_code" in year t but non-missing in $t-n$ or $t+n$, and that (ii) has not changed zip code between those years, we assign the "fips_code" from year $t-n$ or $t+n$ to year t.

– **Run two geographic crosswalks.** The first, based on county change records from Dorn (2021) and US Census Bureau (2021), accounts for county FIPS code changes over time and maps them into consistent 1990 county FIPS codes. The second uses David Dorn's crosswalk (Autor and Dorn, 2013) to map the county FIPS codes to 1990 commuting zones (CZs), assigning each individual to the 1990 CZ classification.

* The majority of the counties that do not match with Dorn's county-to-CZ crosswalk are in AK or in US territories. The remaining cases seem to result from fips_codes that do not exist (e.g., because of credit bureau data entry errors in either the state or the county FIPS). In those cases, the fips_code is replaced with fips_code_hud.

– Drop remaining observations with missing fips_code (<0.05%).

A.3 Clean tradeline files for medical debt in collections

- **Drop individuals when a tradeline is assigned to more than one person.** This drops approximately 0.1% of individuals in the sample (i.e., the set of individuals who show up in UCCCP as having medical debt in collections sometime between 2011 and 2019).

- **Assign the modal opening date:** for cases where acct_open_dt is not constant within the tradeline, assign the modal acct_open_dt; if there is more than one mode, the oldest date is assigned. This affects less than 0.3% of the observations.

- **Assign the modal original amount:** first, set to missing in cases where it is recorded as zero (<0.01% of observations). Then, for the cases where the original amount in collections is not constant within the tradeline, assign the modal amount; if there is more than one mode, the smaller mode is assigned. This affects approximately 0.3% of the observations.

- **Linearly interpolate the outstanding balance:** if the balance for a given tradeline is zero in year t and non-zero in years $t-n$ and $t+n$, replace the zero(s) with the linear interpolation between the two non-zero values. This affects less than 0.01% of the observations.

- **Set the outstanding balance to zero in tradelines deemed to have been paid or discharged:** we define this type of tradeline based on the "acct_cond_cd" and "enhanced_spcl_cmmt" variables. Specifically, when the former takes the values "A2" ("Paid account/zero balance"), "A3" ("Closed account"), "21" ("Consumer reported as deceased"), "67" ("Debt included or

discharged through Chapter 7, 11, or 12 bankruptcy"), "68" ("Account legally paid in full for less than the full balance"), and "69" ("Debt included or discharged through Chapter 13 bankruptcy"), or when the latter takes the values "18" ("Account closed at credit grantor's request"), "19" ("Account closed at consumer's request"), "25" ("Paid by cosigner"), "28" ("Account closed due to transfer or refinance"), "50" ("Account closed due to refinance"), "51" ("Account closed due to transfer"), "52" ("Account legally paid in full for less than the full balance"), "73" ("Paid through insurance"). This affects only about 0.05% of.

- **For accounts deemed to have been paid or discharged in year t, set any subsequent balance amount to zero:** this affects about 0.01% of tradelines.
- **Set the outstanding balance to zero if the debt is disputed by the consumer:** this follows the approach from Kluender et al. (2021) and is based on the "enhanced_spcl_cmnt_cd" variable taking the values "78" ("Account information disputed by consumer") or "83" ("Account closed at consumer's request and in dispute under FCRA"). This affects approximately 2% of tradelines. In this case, the tradeline can take non-zero values in subsequent periods, if the code associated with it changes.
- **Drop tradelines where the ouststanding amount is always zero**.
- **Cases where a tradeline is absent in year t but present in both $t-n$ and $t+n$:** add those observations to the data, linearly interpolate the outstanding balance amounts, and carry forward the location variables. This affects approximately 0.8% of tradelines.
- **Cases where a tradeline does not appear in the credit bureau data in the year when it is first created:** approximately 18% of tradelines take more than 12 months to show up in the credit reporting data, relative to the date when they were created (given by the variable "acct_open_dt"). For these cases, the corresponding observations are created and added to the dataset. The values for the key variables are then imputed, for that set of observations,

based on the following rules:

  - Location/demographic variables: first, they are assigned the individual's modal value for that year. If the modal value is not defined, the values are assigned based on the first/earliest observation available for the individual.
  - Original tradeline amount: imputed based on the modal value for the tradeline.
  - Account balance amount: linearly interpolated, taking the starting balance as the original tradeline amount and the final balance as the balance amount that is observed when that tradeline is first recorded in the UCCCP data.

- **Define the key outcome variables of interest.** For the three outcomes of interest, two alternative definitions were constructed:

  1. An "adjusted" measure that took into account (i) the changes that set the balance amounts outstanding to zero (because they were deemed to have been paid or because they were disputed by the debtors) and (ii) the various linear interpolations implemented to address the issue of missing tradelines, as described above.
  2. A "raw" measure that did not do either of those adjustments.

  The three outcomes of interest were defined as follows:

  - **Total medical debt in collections:** at the individual level and for each year, this corresponds to the sum of the outstanding balance amounts for each tradeline that is categorized as medical debt in collections. The individuals who are present in the attributes files but not in the tradeline files for medical debt in collections for a given year (or for all years) are considered to not have medical debt in collections and are assigned a value of zero.
  - **New medical debt generated over the previous 12 months:** this variable seeks to measure the flow of medical debt that occurs each year and complements the "stock"

measure that considers the total amount of debt outstanding.

For the "adjusted" measure, this was defined based on the "acct_open_dt" variable. For the observations where the opening year is the same as the reporting year, then the respective original loan amount for this tradeline is defined as new medical debt in collections generated over the previous 12 months.

For the "raw" measure, it corresponds to the original loan amount in the reporting year that a tradeline is first observed.

For both types of measures, the tradeline amount is then aggregated at the individual-year level.

- **Has medical debt:** binary variable that takes the value of 1 if the "total medical debt in collections" has a positive value.

- **Balances are winsorized at the 99.99th percentile.** This is done at the individual - rather than tradeline - level and based on the distribution for each year, so the winsorization threshold varies accordingly.

A.4 Clean tradeline files for non-medical debt in collections and for utility, cable, and cellular debt in collections

The cleaning of these types of tradelines followed the same rules as the cleaning of tradelines related to medical debt in collections, outlined in section A.3.

A.5 Clean tradeline files for credit card debt

The cleaning of these types of tradelines followed the same rules as the cleaning of tradelines related to medical debt in collections, outlined in section A.3, except that no adjusted amounts are computed and only the raw amounts are used. Credit card borrowing mostly takes place on existing accounts, whose balance recurrently fluctuates, and does so in a way that makes the potential imputation procedures described for medical debt no suitable. For a similar reason, the outcome

that corresponds to a “flow” measure in medical debt is not used for credit card debt; instead, the change in credit card balances over the previous 12 months is used.

B. The IHME Disease Expenditure Data

This appendix provides background on the Institute for Health Metrics and Evaluation's Disease Expenditure Project, the source of the county-level healthcare spending and utilization estimates introduced in Section II and analyzed in Section V. Subsection B.1 describes the underlying data and IHME's estimation methodology. Subsection B.2 describes our aggregation of the county-level estimates to commuting zones. Subsection B.3 presents the diagnostic evidence underlying our decision to aggregate five of the seven types of care reported in the data and exclude the remaining two, home health and pharmacy, from our outcome definitions.

B.1 Background and methodology

The Disease Expenditure Project (DEX), maintained by the Institute for Health Metrics and Evaluation (IHME) at the University of Washington and described in Dieleman et al. (2025), produces annual estimates of personal healthcare spending and utilization for each of 3,110 U.S. counties over the period 2010–2019. The estimates are disaggregated along four dimensions: 148 health conditions (based on the cause list of the Global Burden of Disease study), 38 age–sex groups, seven types of care (ambulatory, dental, emergency department, home health, inpatient, nursing facility, and prescribed retail pharmaceuticals), and four payer categories (Medicare, Medicaid, private insurance, and out-of-pocket). Medicare is inclusive of Medicare Advantage. The out-of-pocket category comprises both spending by the uninsured and cost-sharing payments by the insured, such as deductibles, co-pays, and co-insurance. The estimates capture 76.6 percent of U.S. personal healthcare expenditures and reflect the spending of 97.3 percent of the population.[26]

The input data consist of more than 40 billion insurance claims and 836 million facility records. Medicare and Medicaid claims are drawn from the respective federal administrative sources; commercial claims come from MarketScan, Kythera, the Health Care Cost Institute, and Fluent Dental Strategies; hospital administrative records come from the Agency for Healthcare Research and

[26] Omitted categories include durable and non-durable medical equipment, over-the-counter drugs, medical transportation, spending by payers other than the four listed (Veterans Affairs, Tricare, and the Indian Health Service), and non-personal health spending such as public health programs and research and development.

Quality's Healthcare Cost and Utilization Project; and household survey data come from the Medical Expenditure Panel Survey. Each claim or encounter record is assigned to a county based on the patient's residence, to an age–sex group, to a type of care, and to a payer category.

IHME's estimation proceeds in four steps, described in detail in the supplementary material to Dieleman et al. (2025). First, the input data are harmonized: claims are grouped into healthcare encounters, with all claims sharing the same person, type of care, and date of service (plus or minus one day for hospital admissions) treated as a single encounter. Second, each encounter is assigned to one of the 148 health conditions on the basis of its *International Classification of Diseases* (ICD-9 or ICD-10) diagnosis codes or National Drug Codes, with non-specific codes probabilistically redistributed to precise conditions. Third, the data are adjusted for known imperfections: sources reporting only facility charges or only billed (rather than paid) amounts are scaled using ratios estimated from more complete sources; a portion of spending is reallocated from primary diagnoses to comorbidities that systematically raise the cost of treatment; and a small-area model with structured random effects is used to borrow strength across counties, age groups, and years where the underlying data for low-population counties are sparse. Fourth, the resulting estimates are constrained to be internally consistent across the county, state, and national levels and to align with the official government spending totals in the State Health Expenditure Accounts.

The published estimates report, for each county–year–payer–type-of-care cell, total spending, spending per capita, spending per beneficiary (for the three insurance payers), spending per encounter, encounters per capita, and encounters per beneficiary. All dollar amounts are expressed in constant 2019 dollars, and IHME's published rates are age- and sex-standardized to the 2019 U.S. population. The unit of utilization—the encounter—is defined by episode of treatment for a given condition, as described in Section II.2, and its physical interpretation varies by type of care: an admission for inpatient and nursing facility care, a visit for ambulatory, emergency department, home health, and dental care, and a prescription fill for pharmacy. Payer enrollment (the beneficiary count) is not reported directly; we derive it as the ratio of total spending to spending per

beneficiary within each county–payer–year cell.[27]

B.2 Aggregation to commuting zones

Our unit of analysis is the commuting zone, so we aggregate the county-level IHME estimates to the 1990 commuting zone definitions used throughout the paper, using the same county-to-CZ crosswalk applied to the credit panel (Autor and Dorn, 2013). Spending totals, encounter counts, and enrollment counts are summed across the counties belonging to each CZ; ratio measures (spending per beneficiary and spending per encounter) are then recomputed from the aggregated numerators and denominators rather than averaged across counties. CZ population denominators and demographic covariates come from the Census Bureau's Intercensal Population Estimates. The result is a balanced panel with one observation per CZ–year–payer combination for each outcome defined in Section II.2.

B.3 Selection of types of care

As stated in Section II.2, every IHME-based outcome in this paper aggregates five of the seven types of care: ambulatory, dental, emergency department, inpatient, and nursing facility, excluding home health and pharmacy. This subsection presents the diagnostic evidence behind that decision.

Figures B.I and B.II plot privately-insured encounter counts by type of care for the 2014 treatment cohort, the earliest-treated and most populous cohort and hence the one with the longest post-treatment window. For Figure B.I, we sum private insurance encounter counts across all CZs in the 2014 cohort, separately by type of care and year, and index each type-of-care series to its value in 2013, the last year before treatment. For Figure B.II, we construct the same annual sums for the never-treated control CZs, compute the difference between the treated and control totals in each year, and normalize that difference by its 2013 value, so that each series equals one in the baseline year by construction. The first figure therefore displays the raw utilization trends within treated CZs, while the second shows whether those trends are differenced out by comparison with

[27] The raw county-level data report slight variations in this derived enrollment count across types of care within a county–payer–year cell, which we interpret as an artifact of the age–sex reweighting. We use the enrollment count associated with ambulatory care; this choice is arbitrary and does not affect our results.

the control group—the descriptive analog of the parallel trends requirement underlying Equation 1. In both figures, the dashed vertical line marks the 2014 treatment year.

Two patterns stand out. First, the five types of care we retain move closely together: their indexed series cluster tightly around one in the pre-treatment period, in both the treated cohort alone and the treated-minus-control comparison, with no monotone drift before treatment. Second, home health and pharmacy behave altogether differently. Home health encounters rise steeply and monotonically over the entire decade, from roughly 20 percent below their 2013 level in 2010 to roughly 35 percent above it by 2019; pharmacy encounters decline monotonically over the same period. Critically, Figure B.II shows that these movements do not difference out against the control group: the treated-minus-control gap in home health encounters roughly triples between 2010 and 2019, growing steadily on both sides of the treatment year, and the pharmacy gap likewise drifts downward throughout the pre-period. Utilization of these two types of care was thus evolving differentially between treated and control CZs well before any merger in our sample occurred. This violation of the parallel trends assumption would mechanically contaminate any event study estimated on outcomes that include them.

The pattern is also readily interpretable. Home health and pharmacy utilization are governed by secular, economy-wide changes in healthcare delivery over this period—the sustained shift of post-acute and long-term care into the home, and changes in prescription drug utilization and filling behavior—interacting with local demographic composition. Neither is plausibly responsive to hospital market structure in the way that hospital-based and hospital-adjacent care is, and the smooth, monotone trajectories visible in Figure B.I, which begin years before treatment and continue undisturbed through it, are the signature of such secular forces rather than of a treatment effect. Retaining these care types would therefore add trending noise, and potential bias, to outcomes designed to measure the causal effect of mergers.

These diagnostics motivate the two scope decisions applied consistently throughout Section V. First, home health and pharmacy are excluded from every IHME-based outcome. Applying the same exclusion to numerators and denominators preserves the price $\times$ quantity logic of the

decomposition in Section V: spending per encounter would not be interpretable as a price measure if its numerator and denominator covered different sets of care. The corollary, noted in Section II.2, is that our findings pertain to the five included care types and are silent on the two excluded ones. Second, the five retained types of care are aggregated rather than analyzed separately. Hospital systems can shift care across settings, for example, from inpatient to ambulatory or emergency department settings, so aggregation captures the total effect of mergers on spending and utilization without misattributing reallocation across settings to changes in overall volume. The common trends evident in Figures B.I and B.II indicate that this aggregation is innocuous: although an admission, a visit, and a dental appointment are different physical units, our event studies use log outcomes with CZ fixed effects, so identification comes from within-CZ changes over time rather than from level comparisons, and the five included components exhibit the common pre-treatment evolution that pooling them presumes.

Figure B.I. PRIVATELY-INSURED ENCOUNTERS BY TYPE OF CARE: 2014 TREATMENT COHORT, INDEXED TO 2013

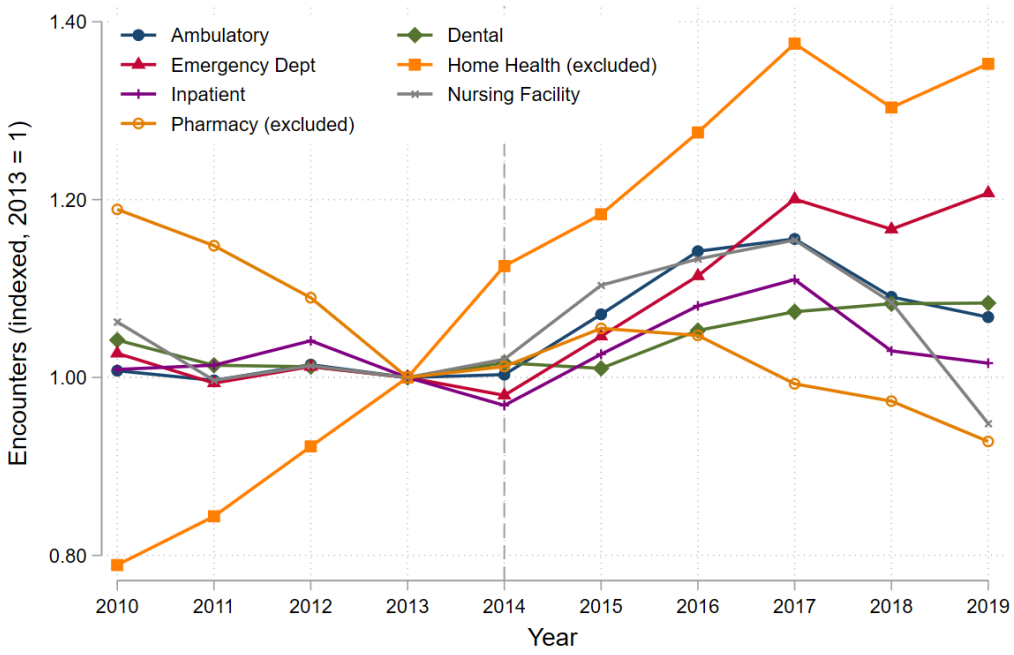

Each series plots total privately-insured encounter counts from the IHME Disease Expenditure data, summed across all commuting zones in the 2014 treatment cohort, separately by type of care, and indexed to its 2013 value (the last pre-treatment year). The dashed vertical line marks the 2014 treatment year. The five types of care we aggregate in our outcome definitions (ambulatory, dental, emergency department, inpatient, and nursing facility) track one another closely throughout the period, while the two we exclude (home health and pharmacy) display steep, monotone secular trends that predate treatment.

Figure B.II. PRIVATELY-INSURED ENCOUNTERS BY TYPE OF CARE: 2014 TREATMENT COHORT RELATIVE TO CONTROL CZS, INDEXED TO 2013

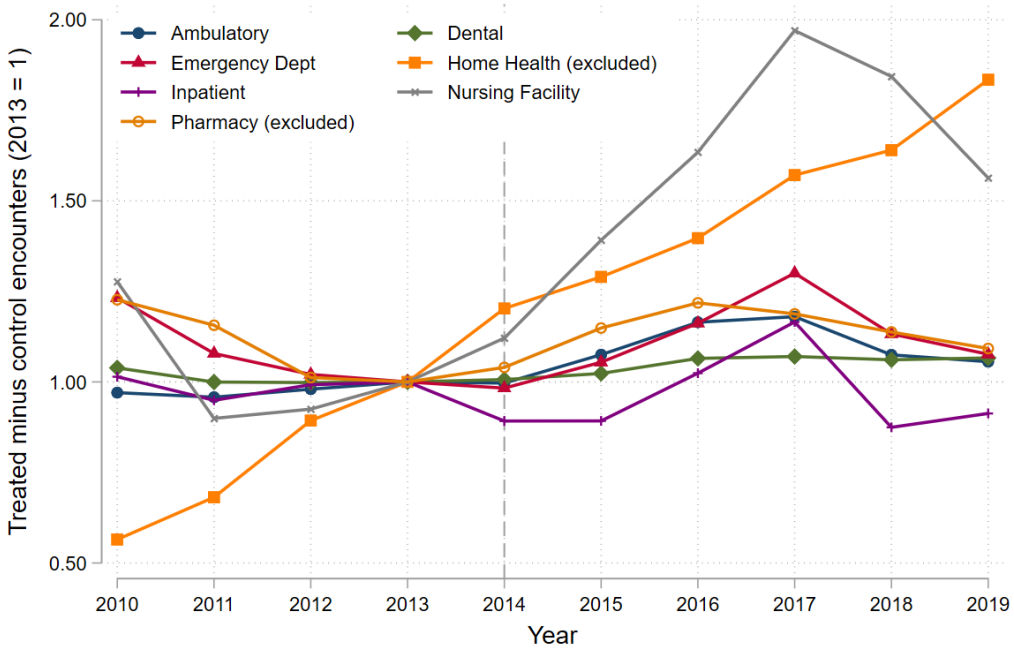

Each series plots the difference between total privately-insured encounter counts in the 2014 treatment cohort and in the never-treated control CZs, by type of care and year, normalized by the 2013 value of that difference so that every series equals one in the baseline year. Encounter counts are from the IHME Disease Expenditure data and are summed across the commuting zones in each group before differencing. The dashed vertical line marks the 2014 treatment year. For the five types of care we aggregate, the treated-minus-control series fluctuate around one in the pre-treatment period; for home health and pharmacy, the series drift monotonically throughout the pre-period, indicating differential trends between treated and control CZs that predate any merger in our sample.

C. Event study plots for medical debt outcomes with placebo treatment dates

Figure C.I. EVENT STUDY PLOTS FOR THE EFFECT OF HOSPITAL MERGER EXPOSURE ON MEDICAL DEBT, PLACEBO TREATMENT DATE

C.I.1 Likelihood of having medical debt

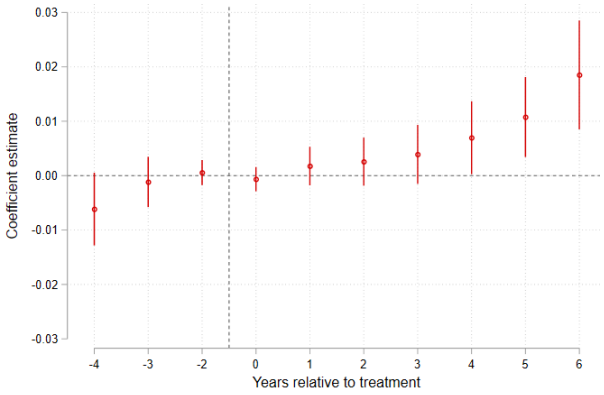

C.I.2 New medical debt (last 12 months)

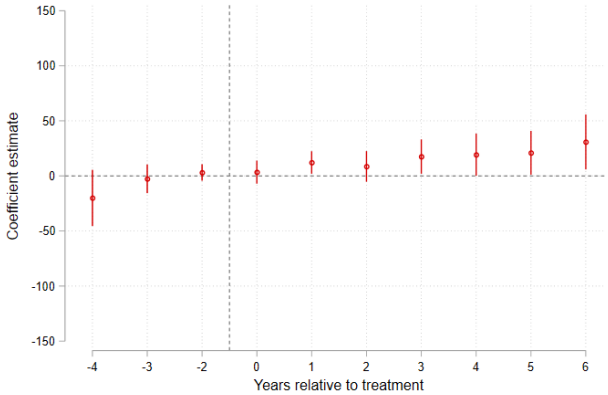

C.I.3 Total medical debt

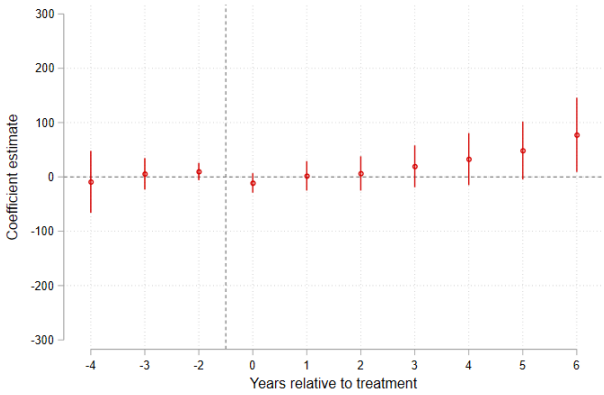

This figure follows the same structure as Figure III in the main draft, but assigning to each cohort a placebo treatment date that precedes the true treatment date by one year (i.e., the 2014 cohort is assigned a placebo treatment date of 2013, the 2015 cohort is assigned a placebo treatment date of 2014, and so on). Each panel shows the point estimates and 95% confidence intervals for β_e, obtained from implementing Equation 2, and illustrates the effect of exposure to hospital mergers on medical debt outcomes. The panels, outcomes, specification fixed effects and controls, and standard errors are otherwise as defined in Figure III. The horizontal axis denotes years relative to the (placebo) merger event; event time $t = -1$ is omitted as the reference period, with the dashed vertical line indicating the period just before the (placebo) merger took place. For every outcome, the point estimate at event time $t = 0$ is never statistically different from zero, suggesting that there was no systematic anticipation of the merger events.

D. Event study plots for "raw/unadjusted" medical debt outcomes

Figure D.I. EVENT STUDY PLOTS FOR THE EFFECT OF HOSPITAL MERGER EXPOSURE ON MEDICAL DEBT ("RAW/UNADJUSTED" OUTCOMES)

D.I.1 Likelihood of having medical debt

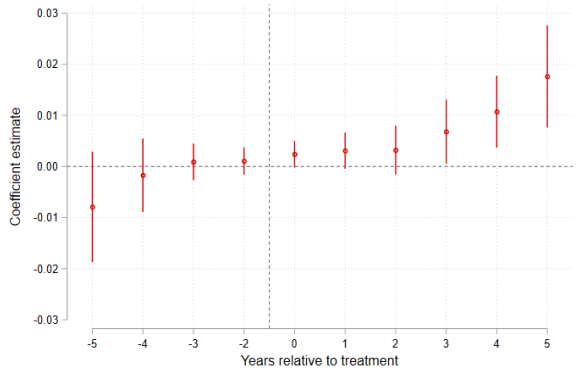

D.I.2 New medical debt (last 12 months)

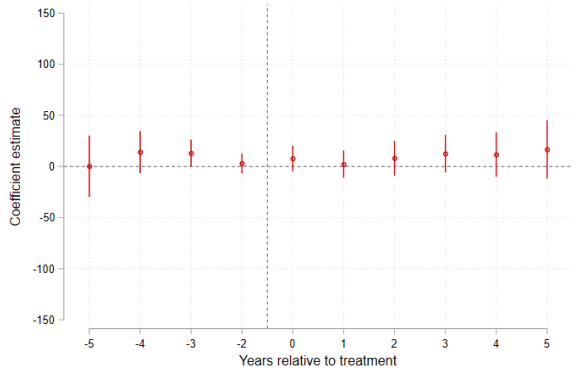

D.I.3 Total medical debt

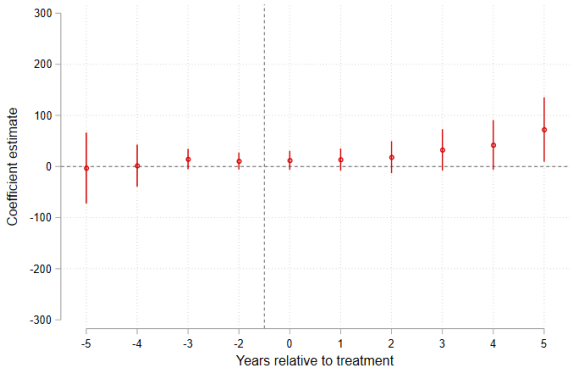

This figure replicates Figure III in the main draft, but using the "raw/unadjusted" outcomes instead of the adjusted ones. Each panel shows the point estimates and 95% confidence intervals for β_e, obtained from implementing Equation 2, and illustrates the effect of exposure to hospital mergers on "raw/unadjusted" medical debt outcomes. The panels, outcomes, specification fixed effects and controls, and standard errors are otherwise as defined in Figure III. The horizontal axis denotes years relative to the merger event; event time $t = -1$ is omitted as the reference period, with the dashed vertical line indicating the period just before the merger took place.

E. Event study regression tables for medical debt, baseline specifications

Table E.I. Effect of merger exposure on likelihood of having medical debt (adjusted) in collections (2011-2019).

	(1) Has medical debt (adj) Full sample	(2) Has medical debt (adj) Full sample	(3) Has medical debt (adj) Full sample	(4) Has medical debt (adj) Never movers	(5) Has medical debt (adj) Never movers Under 65	(6) Has medical debt (adj) HHI change >150	(7) Has medical debt (adj) HHI change >150 level>1800	(8) Has medical debt (adj) HHI change >150 level>2500
-5	-0.0056	-0.0060	-0.0061	-0.0063	-0.0086	-0.0083*	-0.0068	0.0000
	(0.0056)	(0.0053)	(0.0053)	(0.0053)	(0.0060)	(0.0048)	(0.0056)	(0.0073)
-4	0.0002	-0.0004	-0.0003	-0.0001	-0.0005	0.0002	0.0015	0.0060
	(0.0035)	(0.0034)	(0.0034)	(0.0033)	(0.0039)	(0.0044)	(0.0051)	(0.0070)
-3	0.0024	0.0012	0.0012	0.0011	0.0012	-0.0016	-0.0008	0.0013
	(0.0018)	(0.0017)	(0.0017)	(0.0018)	(0.0020)	(0.0022)	(0.0028)	(0.0043)
-2	0.0011	0.0006	0.0006	0.0007	0.0008	-0.0010	-0.0006	-0.0020
	(0.0012)	(0.0011)	(0.0011)	(0.0011)	(0.0013)	(0.0016)	(0.0018)	(0.0021)
0	0.0023**	0.0024**	0.0024**	0.0022*	0.0034**	0.0000	-0.0023	-0.0080**
	(0.0011)	(0.0011)	(0.0012)	(0.0012)	(0.0014)	(0.0022)	(0.0027)	(0.0035)
1	0.0021	0.0032**	0.0032*	0.0030*	0.0046**	0.0006	-0.0019	-0.0037
	(0.0016)	(0.0016)	(0.0017)	(0.0017)	(0.0020)	(0.0026)	(0.0026)	(0.0024)
2	0.0025	0.0045*	0.0045*	0.0037	0.0058**	0.0049	0.0024	-0.0024
	(0.0024)	(0.0023)	(0.0023)	(0.0023)	(0.0028)	(0.0031)	(0.0034)	(0.0039)
3	0.0054	0.0077**	0.0077**	0.0063**	0.0090**	0.0032	-0.0007	0.0028
	(0.0034)	(0.0032)	(0.0032)	(0.0032)	(0.0037)	(0.0043)	(0.0052)	(0.0057)
4	0.0088**	0.0115***	0.0116***	0.0100***	0.0131***	0.0061	0.0060	0.0091
	(0.0039)	(0.0036)	(0.0037)	(0.0037)	(0.0043)	(0.0048)	(0.0067)	(0.0076)
5	0.0164***	0.0190***	0.0190***	0.0170***	0.0214***	0.0081	0.0110	0.0033
	(0.0058)	(0.0051)	(0.0051)	(0.0050)	(0.0060)	(0.0064)	(0.0084)	(0.0092)
Observations	40,811,971	40,506,377	40,506,377	31,641,694	24,642,705	17,781,787	16,384,023	14,368,359
Year FEs	Y	Y	Y	Y	Y	Y	Y	Y
CZ FEs	Y	Y	Y	Y	Y	Y	Y	Y
Individual FEs	N	Y	Y	Y	Y	Y	Y	Y
Medicaid exp.	N	N	Y	Y	Y	Y	Y	Y

*** p<0.01; ** p<0.05; * p<0.10

Notes: Robust standard errors clustered at the CZ level in parenthesis. The dependent variable is binary and takes the value of 1 if the individual has outstanding medical debts in collection. Columns (1) to (3) use the full sample and sequentially add fixed effects and a binary control for the state-level Medicaid expansion from 2014 onward. Column (4) restricts the sample to never movers and column (5) further restricts it to never mover under the age of 65. Column (6) goes back to the full sample as the starting point and then restricts the treatment group to CZs where the merger-induced HHI increase exceeded 150; column (7) further restricts the treatment group to CZs where the post-merger HHI exceeded 1800 and column (8) increases the post-merger HHI threshold to 2500.

Table E.II. Effect of merger exposure on new medical debt in collections (adjusted) created over the previous 12 months (2011-2019).

	(1) New med debt last 12m adj) Full sample	(2) New med debt last 12m adj Full sample	(3) New med debt last 12m adj Full sample	(4) New med debt last 12m adj Never movers	(5) New med debt last 12m adj Never movers Under 65	(6) New med debt last 12m adj HHI change >150	(7) New med debt last 12m adj HHI change >150 level>1800	(8) New med debt last 12m adj HHI change >150 level>2500
-5	-12.5	-12.6	-14.3	-13.0	-16.6	-30.1	-14.5	4.1
	(15.8)	(16.3)	(16.0)	(14.3)	(16.8)	(20.4)	(14.1)	(19.2)
-4	-0.4	-0.9	1.1	2.6	1.7	-17.3	-8.3	9.6
	(9.1)	(9.4)	(9.4)	(8.4)	(9.8)	(16.0)	(15.4)	(11.7)
-3	0.9	-0.8	-0.4	-1.8	-1.8	-2.9	2.2	-0.2
	(5.9)	(6.0)	(5.9)	(5.4)	(6.6)	(7.0)	(7.4)	(8.6)
-2	-3.9	-4.7	-3.5	-3.8	-4.0	-4.7	-0.2	-1.8
	(5.0)	(5.2)	(5.4)	(4.9)	(5.9)	(7.2)	(8.5)	(10.0)
0	2.9	2.8	8.7*	4.1	6.6	-5.7	-7.5	-3.1
	(4.7)	(5.0)	(5.2)	(4.7)	(5.8)	(6.4)	(7.3)	(12.9)
1	-1.7	-0.7	5.2	-1.8	-0.7	-4.5	-6.4	11.7
	(7.7)	(8.4)	(7.7)	(7.1)	(8.7)	(9.9)	(11.3)	(16.6)
2	5.4	7.7	14.0	6.1	10.0	4.0	5.2	19.7
	(8.6)	(9.3)	(8.9)	(8.2)	(10.1)	(10.9)	(13.6)	(21.1)
3	5.3	8.4	14.8	5.3	8.7	-6.0	1.7	26.7
	(10.4)	(11.3)	(10.8)	(10.0)	(12.3)	(12.1)	(16.1)	(22.7)
4	3.6	7.8	15.0	4.3	6.5	9.6	17.9	47.8
	(11.8)	(12.6)	(12.0)	(11.1)	(13.7)	(15.2)	(19.9)	(31.2)
5	9.5	14.5	21.9	11.6	16.0	35.5	52.0*	77.1
	(15.9)	(16.9)	(16.0)	(14.7)	(17.9)	(25.1)	(27.0)	(50.1)
Observations	40,811,971	40,506,377	40,506,377	31,641,694	24,642,705	17,781,787	16,384,023	14,368,359
Year FEs	Y	Y	Y	Y	Y	Y	Y	Y
CZ FEs	Y	Y	Y	Y	Y	Y	Y	Y
Individual FEs	N	Y	Y	Y	Y	Y	Y	Y
Medicaid expansion	N	N	Y	Y	Y	Y	Y	Y

*** p<0.01; ** p<0.05; * p<0.10

Notes: Robust standard errors clustered at the CZ level in parenthesis. The dependent variable is the flow of new medical debt in collections over the past 12 months, measured in 2017 USD. Columns (1) to (3) use the full sample and sequentially add fixed effects and a binary control for the state-level Medicaid expansion from 2014 onward. Column (4) restricts the sample to never movers and column (5) further restricts it to never mover under the age of 65. Column (6) goes back to the full sample as the starting point and then restricts the treatment group to CZs where the merger-induced HHI increase exceeded 150; column (7) further restricts the treatment group to CZs where the post-merger HHI exceeded 1800 and column (8) increases the post-merger HHI threshold to 2500.

Table E.III. Effect of merger exposure on the total amount (adjusted) of medical debt in collections (2011-2019).

	(1) Total (adj) balance	(2) Total (adj) balance	(3) Total (adj) balance	(4) Total (adj) balance	(5) Total (adj) balance	(6) Total (adj) balance	(7) Total (adj) balance	(8) Total (adj) balance
	Full sample	Full sample	Full sample	Never movers	Never movers Under 65	HHI change >150	HHI change >150 level>1800	HHI change >150 level>2500
-5	-0.1	0.6	-1.2	-1.4	0.4	-55.5	-30.4	3.2
	(35.9)	(34.2)	(35.0)	(32.0)	(39.0)	(36.3)	(37.3)	(65.2)
-4	5.8	5.8	7.9	9.1	12.1	-6.9	7.1	37.1
	(21.9)	(21.2)	(21.3)	(19.7)	(24.2)	(25.1)	(26.2)	(41.3)
-3	23.2*	20.5*	21.0*	16.1	21.5	-1.6	3.9	15.5
	(12.1)	(12.1)	(12.0)	(11.0)	(13.4)	(14.4)	(15.8)	(26.9)
-2	10.6	9.8	11.1	9.6	12.5	-6.2	-0.2	-0.5
	(8.6)	(8.9)	(9.1)	(8.4)	(10.1)	(12.4)	(12.4)	(20.1)
0	8.6	6.6	13.1	6.9	10.6	1.1	-9.4	-25.9
	(8.6)	(8.8)	(9.5)	(8.6)	(10.8)	(13.0)	(16.4)	(20.2)
1	10.8	11.2	17.6	7.7	11.2	0.8	-11.8	-3.7
	(11.8)	(12.0)	(11.5)	(10.7)	(13.4)	(15.1)	(16.9)	(22.9)
2	22.1	23.9	30.8*	14.6	21.9	20.4	8.6	-11.0
	(18.2)	(17.7)	(17.1)	(15.9)	(19.7)	(18.7)	(21.7)	(33.7)
3	37.0	39.7*	46.7**	25.9	36.5	13.5	3.0	26.9
	(25.0)	(23.7)	(23.4)	(21.7)	(26.8)	(26.5)	(35.1)	(53.4)
4	51.3*	52.8*	60.6**	36.5	49.1	44.0	44.0	53.5
	(30.4)	(28.4)	(28.2)	(26.2)	(32.0)	(30.6)	(38.9)	(62.1)
5	87.7**	87.3**	95.4**	64.5*	84.7**	85.6*	107.1**	113.8
	(41.7)	(37.9)	(37.1)	(34.6)	(41.9)	(47.3)	(54.2)	(90.4)
Observations	40,811,971	40,506,377	40,506,377	31,641,694	24,642,705	17,781,787	16,384,023	14,368,359
Year FEs	Y	Y	Y	Y	Y	Y	Y	Y
CZ FEs	Y	Y	Y	Y	Y	Y	Y	Y
Individual FEs	N	Y	Y	Y	Y	Y	Y	Y
Medicaid expansion	N	N	Y	Y	Y	Y	Y	Y

*** p<0.01; ** p<0.05; * p<0.10

Notes: Robust standard errors clustered at the CZ level in parenthesis. The dependent variable is the total amount of outstanding medical debts in collection, measured in 2017 USD. Columns (1) to (3) use the full sample and sequentially add fixed effects and a binary control for the state-level Medicaid expansion from 2014 onward. Column (4) restricts the sample to never movers and column (5) further restricts it to never mover under the age of 65. Column (6) goes back to the full sample as the starting point and then restricts the treatment group to CZs where the merger-induced HHI increase exceeded 150; column (7) further restricts the treatment group to CZs where the post-merger HHI exceeded 1800 and column (8) increases the post-merger HHI threshold to 2500.

F. Alternative control group

Figure F.I. EVENT STUDY PLOTS FOR THE EFFECT OF HOSPITAL MERGER EXPOSURE ON MEDICAL DEBT, ALTERNATIVE CONTROL GROUP

F.I.1 Likelihood of having medical debt: baseline

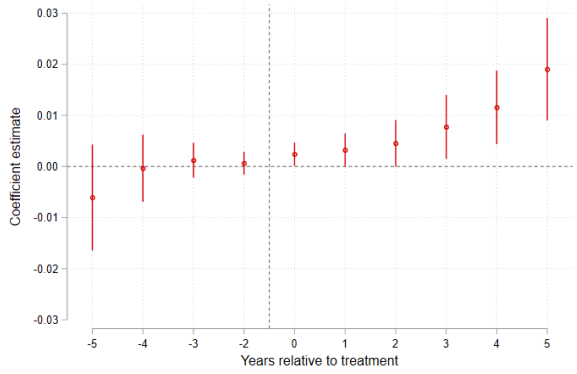

F.I.2 Likelihood of having medical debt: alternative control group

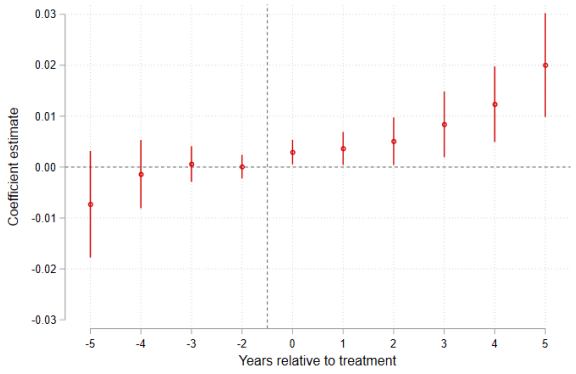

F.I.3 New medical debt (last 12 months): baseline

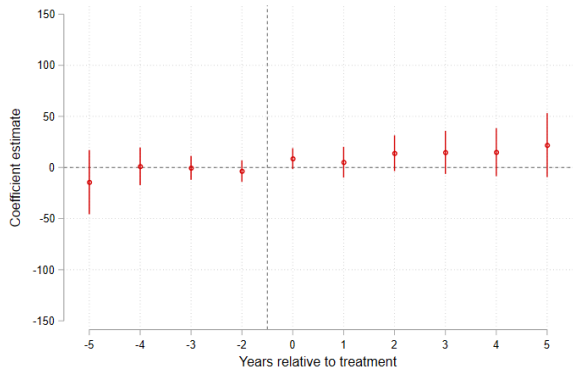

F.I.4 New medical debt (last 12 months): alternative control group

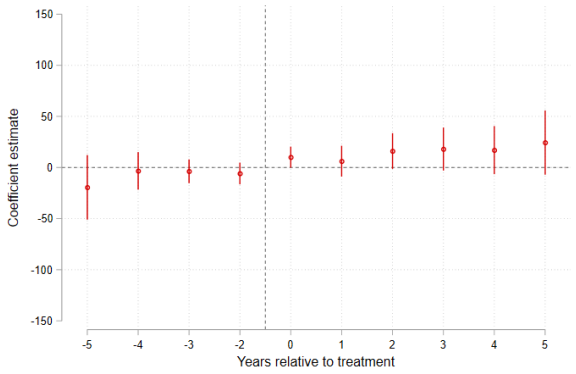

F.I.5 Total medical debt in collections: baseline

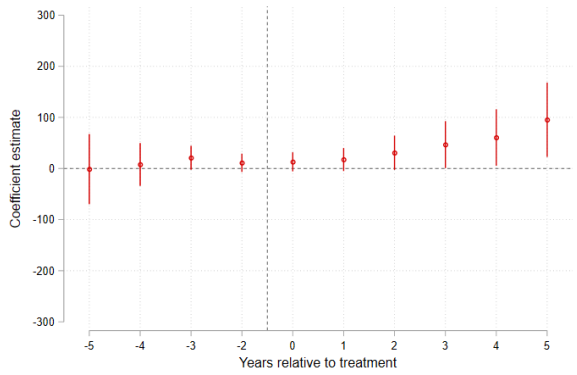

F.I.6 Total medical debt in collections: alternative control group

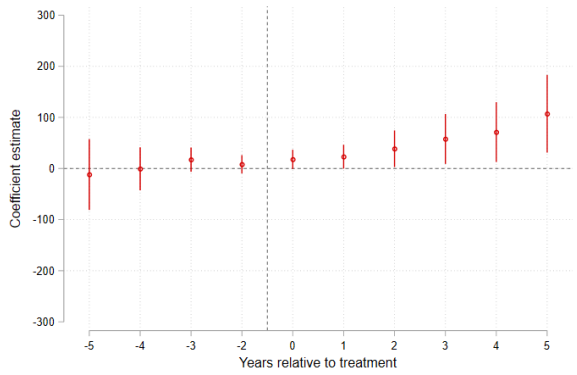

Panels F.I.1, F.I.3, and F.I.5 replicate the baseline estimates also displayed in Figure III in the main draft, for comparison purposes. Panels F.I.2, F.I.4, and F.I.6 show the estimates for each of the corresponding outcomes of interest when using the alternative control group, where control CZs include only the 507 that did not have a merger predicted to increase HHI in any year between 2011 and 2019. Each panel shows the point estimates and 95% confidence intervals for β_e, obtained from implementing Equation 2, and illustrates the effect of exposure to hospital mergers on medical debt outcomes. The panels, outcomes, specification fixed effects and controls, and standard errors are otherwise as defined in Figure III. The horizontal axis denotes years relative to the merger event; event time $t = -1$ is omitted as the reference period, with the dashed vertical line indicating the period just before the merger took place.

G. Event study regression tables for medical debt, specifications with lagged dependent variable

Table G.I. Effect of merger exposure on likelihood of having medical debt (adjusted) in collections (2011-2019), specifications with a lagged dependent variable.

	(1) Has medical debt (adj) Full sample	(2) Has medical debt (adj) Full sample	(3) Has medical debt (adj) Full sample	(4) Has medical debt (adj) Never movers	(5) Has medical debt (adj) Never movers Under 65	(6) Has medical debt (adj) HHI change >150	(7) Has medical debt (adj) HHI change >150 level>1800	(8) Has medical debt (adj) HHI change >150 level>2500
-5	-0.0021	-0.0029	-0.0027	-0.0022	-0.0041	-0.0094*	-0.0088	-0.0034
	(0.0046)	(0.0063)	(0.0063)	(0.0062)	(0.0070)	(0.0050)	(0.0054)	(0.0034)
-4	0.0016	-0.0007	-0.0007	-0.0003	-0.0017	-0.0042	-0.0037	-0.0007
	(0.0034)	(0.0041)	(0.0041)	(0.0039)	(0.0044)	(0.0036)	(0.0038)	(0.0035)
-3	0.0023	0.0018	0.0018	0.0022	0.0020	-0.0027	-0.0013	0.0013
	(0.0018)	(0.0021)	(0.0021)	(0.0021)	(0.0025)	(0.0029)	(0.0031)	(0.0041)
-2	0.0004	0.0004	0.0004	0.0005	0.0005	-0.0006	-0.0005	-0.0035
	(0.0013)	(0.0012)	(0.0012)	(0.0012)	(0.0013)	(0.0019)	(0.0023)	(0.0033)
0	0.0029**	0.0027**	0.0025**	0.0022**	0.0031**	-0.0004	-0.0024	-0.0087**
	(0.0011)	(0.0011)	(0.0011)	(0.0011)	(0.0013)	(0.0022)	(0.0026)	(0.0036)
1	0.0018	0.0026**	0.0024*	0.0020	0.0031**	0.0003	-0.0008	-0.0011
	(0.0012)	(0.0013)	(0.0013)	(0.0013)	(0.0016)	(0.0019)	(0.0021)	(0.0022)
2	0.0020	0.0034*	0.0032*	0.0022	0.0037*	0.0040*	0.0030	-0.0020
	(0.0016)	(0.0019)	(0.0019)	(0.0018)	(0.0022)	(0.0022)	(0.0028)	(0.0040)
3	0.0050***	0.0067***	0.0065***	0.0051**	0.0072***	0.0019	0.0002	0.0027
	(0.0018)	(0.0024)	(0.0024)	(0.0023)	(0.0028)	(0.0029)	(0.0038)	(0.0043)
4	0.0049***	0.0085***	0.0083***	0.0068***	0.0091***	0.0036	0.0045	0.0052
	(0.0018)	(0.0025)	(0.0026)	(0.0025)	(0.0030)	(0.0033)	(0.0045)	(0.0058)
5	0.0081***	0.0140***	0.0138***	0.0121***	0.0159***	0.0046	0.0071	-0.0008
	(0.0025)	(0.0037)	(0.0037)	(0.0036)	(0.0043)	(0.0046)	(0.0058)	(0.0053)
Observations	35,350,980	35,100,004	35,100,004	27,400,191	21,088,821	15,395,283	14,184,857	12,394,759
Year FEs	Y	Y	Y	Y	Y	Y	Y	Y
CZ FEs	Y	Y	Y	Y	Y	Y	Y	Y
Individual FEs	N	Y	Y	Y	Y	Y	Y	Y
Medicaid expansion	N	N	Y	Y	Y	Y	Y	Y

*** p<0.01; ** p<0.05; * p<0.10

Notes: Robust standard errors clustered at the CZ level in parenthesis. The dependent variable is a binary indicator for having medical debt in collections. All specifications include a lagged dependent variable as a control. Columns (1) to (3) use the full sample and sequentially add fixed effects and a binary control for the state-level Medicaid expansion from 2014 onward. Column (4) restricts the sample to never movers and column (5) further restricts it to never mover under the age of 65. Column (6) goes back to the full sample as the starting point and then restricts the treatment group to CZs where the merger-induced HHI increase exceeded 150; column (7) further restricts the treatment group to CZs where the post-merger HHI exceeded 1800 and column (8) increases the post-merger HHI threshold to 2500.

Table G.II. Effect of merger exposure on new medical debt in collections (adjusted) created over the previous 12 months (2011-2019), specifications with a lagged dependent variable.

	(1) New med debt last 12m adj) Full sample	(2) New med debt last 12m adj Full sample	(3) New med debt last 12m adj Full sample	(4) New med debt last 12m adj Never movers	(5) New med debt last 12m adj Never movers Under 65	(6) New med debt last 12m adj HHI change >150	(7) New med debt last 12m adj HHI change >150 level>1800	(8) New med debt last 12m adj HHI change >150 level>2500
-5	-26.4**	-30.8**	-33.6*	-31.8**	-40.4**	-41.1***	-36.8**	-3.3
	(12.0)	(15.7)	(17.5)	(15.4)	(18.0)	(15.8)	(17.4)	(24.5)
-4	-4.1	-9.6	-11.3	-8.6	-13.7	-27.0	-17.7	-10.6
	(13.5)	(15.8)	(15.3)	(13.2)	(14.9)	(20.0)	(14.2)	(9.3)
-3	-3.2	-4.5	-3.5	-2.4	-3.4	-9.5	-3.9	-3.0
	(6.1)	(7.3)	(7.1)	(6.2)	(7.6)	(11.5)	(10.6)	(9.6)
-2	-5.4	-5.2	-4.2	-4.4	-4.8	-5.5	-1.1	-2.4
	(5.1)	(5.3)	(5.4)	(4.9)	(5.9)	(7.2)	(8.3)	(9.6)
0	1.4	2.7	7.7	3.1	5.3	-6.5	-8.3	-2.2
	(5.3)	(5.0)	(5.1)	(4.6)	(5.7)	(6.5)	(7.5)	(13.3)
1	-3.2	-0.6	4.3	-2.6	-1.9	-6.1	-7.5	13.1
	(7.3)	(8.3)	(7.7)	(7.1)	(8.7)	(10.0)	(11.4)	(16.5)
2	4.9	7.7	12.9	5.2	8.8	2.1	3.5	19.5
	(7.2)	(9.1)	(8.8)	(8.0)	(9.9)	(10.9)	(13.6)	(20.6)
3	2.9	8.6	14.0	4.8	8.2	-8.4	0.1	26.9
	(8.4)	(11.0)	(10.7)	(9.8)	(12.2)	(12.2)	(16.2)	(23.0)
4	1.0	8.4	14.4	3.8	5.9	7.4	15.3	46.2
	(9.5)	(12.3)	(11.9)	(10.9)	(13.5)	(15.4)	(20.1)	(31.8)
5	6.4	14.7	20.9	11.1	15.3	32.0	48.0*	74.4
	(12.7)	(16.6)	(15.9)	(14.5)	(17.9)	(24.8)	(26.8)	(50.2)
Observations	35,350,980	35,100,004	35,100,004	27,400,191	21,088,821	15,395,283	14,184,857	12,394,759
Year FEs	Y	Y	Y	Y	Y	Y	Y	Y
CZ FEs	Y	Y	Y	Y	Y	Y	Y	Y
Individual FEs	N	Y	Y	Y	Y	Y	Y	Y
Medicaid expansion	N	N	Y	Y	Y	Y	Y	Y

*** p<0.01; ** p<0.05; * p<0.10

Notes: Robust standard errors clustered at the CZ level in parenthesis. The dependent variable is the flow of new medical debt over the previous 12 months, measured in 2017 USD. All specifications include a lagged dependent variable as a control. Columns (1) to (3) use the full sample and sequentially add fixed effects and a binary control for the state-level Medicaid expansion from 2014 onward. Column (4) restricts the sample to never movers and column (5) further restricts it to never mover under the age of 65. Column (6) goes back to the full sample as the starting point and then restricts the treatment group to CZs where the merger-induced HHI increase exceeded 150; column (7) further restricts the treatment group to CZs where the post-merger HHI exceeded 1800 and column (8) increases the post-merger HHI threshold to 2500.

Table G.III. Effect of merger exposure on the total amount (adjusted) of medical debt in collections (2011-2019), specifications with a lagged dependent variable.

	(1) Total (adj) balance	(2) Total (adj) balance	(3) Total (adj) balance	(4) Total (adj) balance	(5) Total (adj) balance	(6) Total (adj) balance	(7) Total (adj) balance	(8) Total (adj) balance
	Full sample	Full sample	Full sample	Never movers	Never movers Under 65	HHI change >150	HHI change >150 level>1800	HHI change >150 level>2500
-5	-30.7	-21.9	-23.9	-16.3	-22.3	-57.7	-56.3	-15.9
	(23.2)	(33.0)	(34.5)	(29.6)	(36.1)	(35.5)	(38.2)	(57.5)
-4	3.0	1.2	0.0	7.4	4.7	-37.9	-38.2	-25.0
	(29.9)	(31.4)	(31.1)	(27.0)	(32.6)	(28.8)	(27.4)	(37.4)
-3	-2.2	-1.1	-0.5	2.8	1.6	-19.6	-11.4	-1.6
	(14.0)	(15.1)	(15.0)	(13.4)	(16.3)	(19.8)	(18.6)	(31.9)
-2	2.1	4.1	4.8	5.5	6.4	-10.2	-5.7	-10.2
	(10.1)	(9.8)	(10.0)	(9.0)	(10.8)	(13.7)	(14.8)	(25.4)
0	14.9	12.4	15.9*	9.4	13.1	-2.2	-10.5	-26.8
	(9.2)	(9.1)	(9.6)	(8.7)	(10.6)	(13.4)	(17.2)	(23.9)
1	15.3*	15.4*	18.9**	10.4	14.4	-3.2	-10.0	8.7
	(8.7)	(9.1)	(9.0)	(8.2)	(10.2)	(12.2)	(14.3)	(17.4)
2	22.9*	25.4*	29.1**	16.2	24.5	15.1	10.6	-19.2
	(13.7)	(15.1)	(14.8)	(13.6)	(16.7)	(14.5)	(17.0)	(30.1)
3	36.2**	41.0**	44.9**	29.6*	41.0**	11.6	10.4	28.9
	(14.3)	(17.7)	(17.7)	(16.3)	(20.1)	(19.6)	(25.9)	(38.2)
4	27.4*	41.6**	45.8**	29.4	40.9*	23.2	25.1	27.9
	(15.1)	(19.9)	(19.9)	(18.3)	(22.4)	(19.8)	(26.5)	(41.2)
5	51.8***	70.6***	75.0***	54.9**	74.4***	54.9*	61.0*	53.9
	(19.0)	(25.9)	(25.5)	(23.5)	(28.7)	(28.2)	(36.5)	(49.5)
Observations	35,350,980	35,100,004	35,100,004	27,400,191	21,088,821	15,395,283	14,184,857	12,394,759
Year FEs	Y	Y	Y	Y	Y	Y	Y	Y
CZ FEs	Y	Y	Y	Y	Y	Y	Y	Y
Individual FEs	N	Y	Y	Y	Y	Y	Y	Y
Medicaid expansion	N	N	Y	Y	Y	Y	Y	Y

*** p<0.01; ** p<0.05; * p<0.10

Notes: Robust standard errors clustered at the CZ level in parenthesis. The dependent variable is the total amount of outstanding medical debts in collection, measured in 2017 USD. All specifications include a lagged dependent variable as a control. Columns (1) to (3) use the full sample and sequentially add fixed effects and a binary control for the state-level Medicaid expansion from 2014 onward. Column (4) restricts the sample to never movers and column (5) further restricts it to never mover under the age of 65. Column (6) goes back to the full sample as the starting point and then restricts the treatment group to CZs where the merger-induced HHI increase exceeded 150; column (7) further restricts the treatment group to CZs where the post-merger HHI exceeded 1800 and column (8) increases the post-merger HHI threshold to 2500.

H. Heterogeneity dimensions

H.1 Gender

Figure H.I. EVENT STUDY PLOTS FOR THE EFFECT OF HOSPITAL MERGER EXPOSURE ON MEDICAL DEBT, GENDER SUBSAMPLES

H.I.1 Likelihood of having medical debt: female

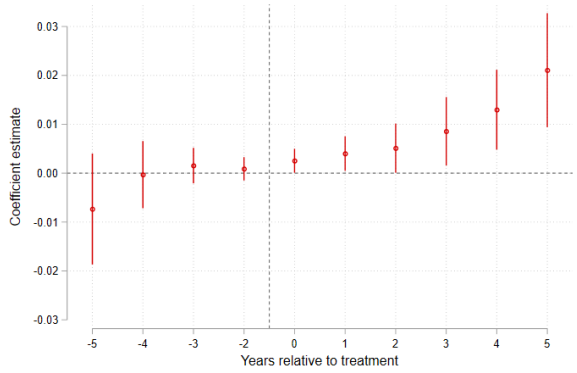

H.I.2 Likelihood of having medical debt: male

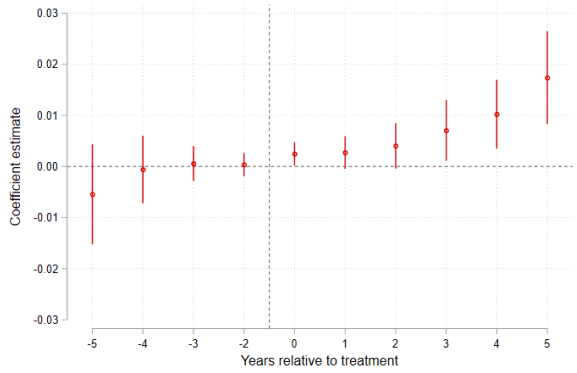

H.I.3 New medical debt (last 12 months): female

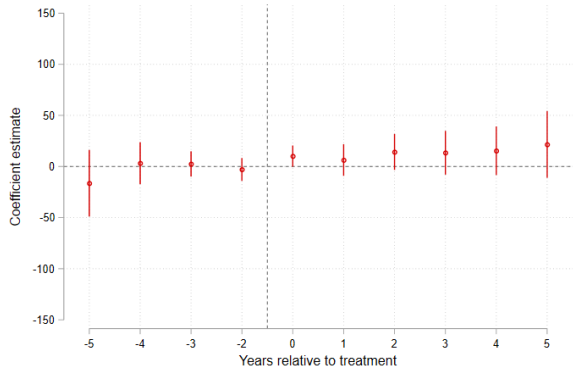

H.I.4 New medical debt (last 12 months): male

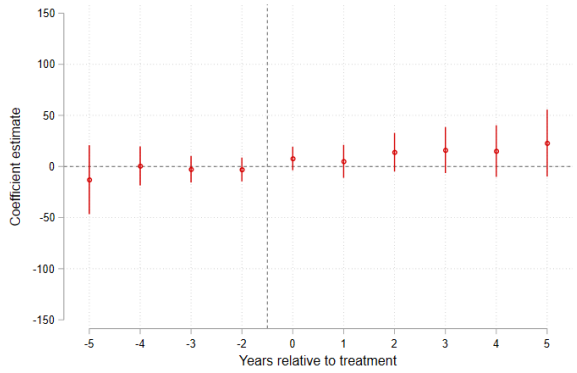

H.I.5 Total medical debt in collections: female

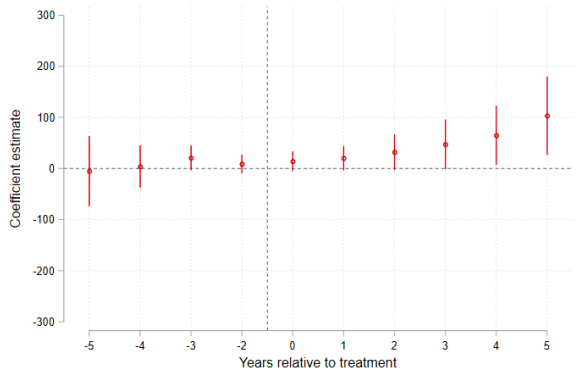

H.I.6 Total medical debt in collections: male

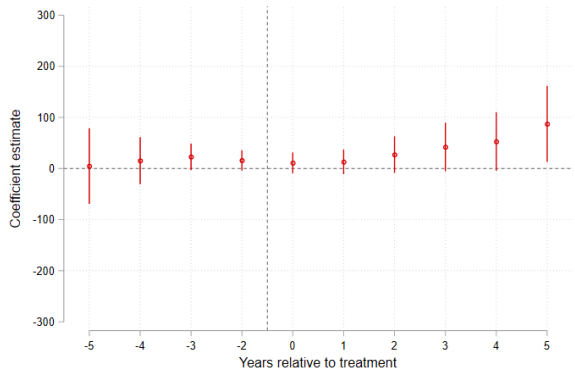

This figure replicates Figure III in the main draft, but splitting the sample based on gender. Each panel shows the point estimates and 95% confidence intervals for β_e, obtained from implementing Equation 2, and illustrates the effect of exposure to hospital mergers on medical debt outcomes. The panels, outcomes, specification fixed effects and controls, and standard errors are otherwise as defined in Figure III. The horizontal axis denotes years relative to the merger event; event time $t = -1$ is omitted as the reference period, with the dashed vertical line indicating the period just before the merger took place.

H.2 Medicaid expansion states

Figure H.II. EVENT STUDY PLOTS FOR THE EFFECT OF HOSPITAL MERGER EXPOSURE ON MEDICAL DEBT, SUBSAMPLE DEFINED BASED ON WHETHER STATES EXPANDED MEDICAID COVERAGE BETWEEN 2014-2019

H.II.1 Likelihood of having medical debt: expansion states

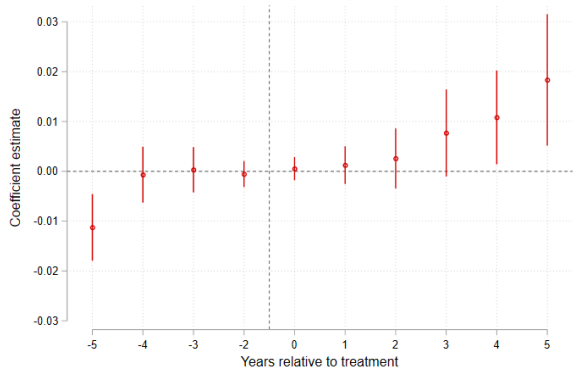

H.II.2 Likelihood of having medical debt: non-expansion states

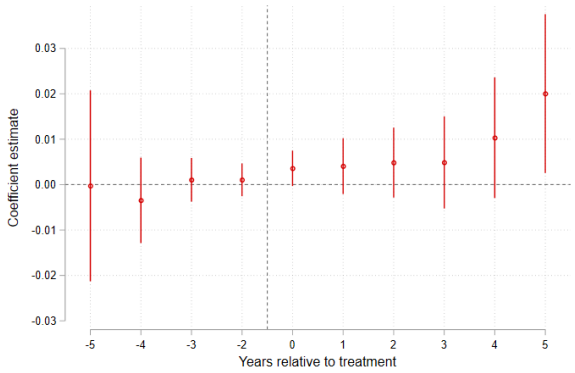

H.II.3 New medical debt (last 12 months): expansion states

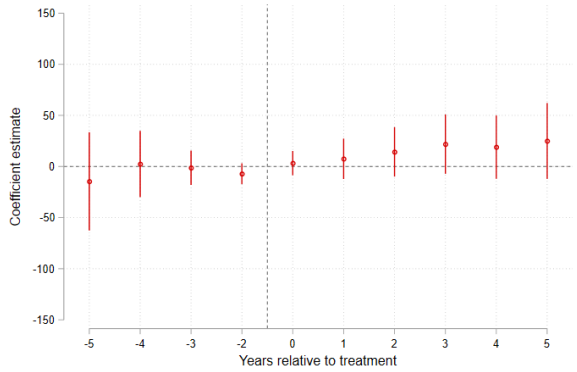

H.II.4 New medical debt (last 12 months): non-expansion states

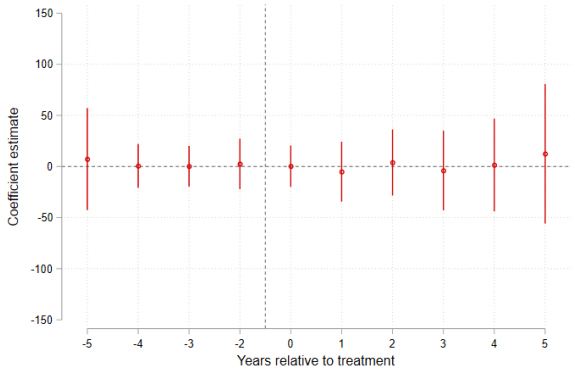

H.II.5 Total medical debt in collections: expansion states

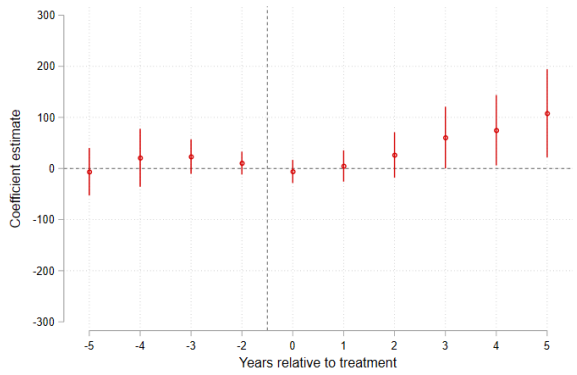

H.II.6 Total medical debt in collections: non-expansion states

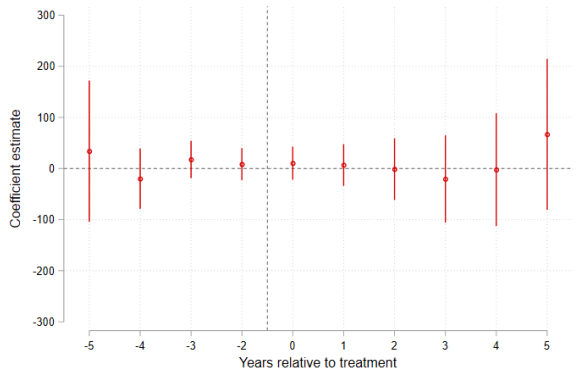

This figure replicates Figure III in the main draft, but splitting the sample based on states that expanded Medicaid coverage between 2014-2019 and those that did not. Each panel shows the point estimates and 95% confidence intervals for β_e, obtained from implementing Equation 2, and illustrates the effect of exposure to hospital mergers on medical debt outcomes. The panels, outcomes, specification fixed effects and controls, and standard errors are otherwise as defined in Figure III. The horizontal axis denotes years relative to the merger event; event time $t = -1$ is omitted as the reference period, with the dashed vertical line indicating the period just before the merger took place.

H.3 Subset with medical debt in collections in 2013

Figure H.III. EVENT STUDY PLOTS FOR THE EFFECT OF HOSPITAL MERGER EXPOSURE ON MEDICAL DEBT, SUBSAMPLE WITH MEDICAL DEBT IN COLLECTIONS IN 2013

H.III.1 New medical debt (last 12 months): baseline

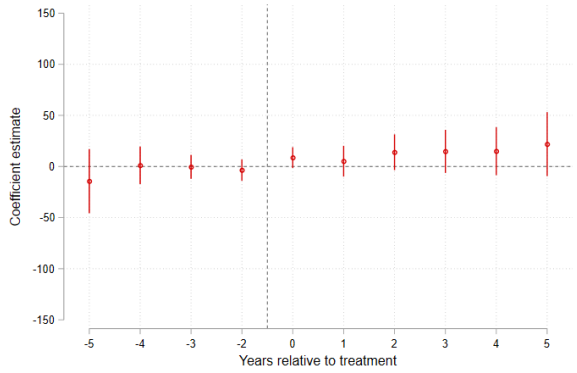

H.III.2 New medical debt (last 12 months): subsample with debt in 2013

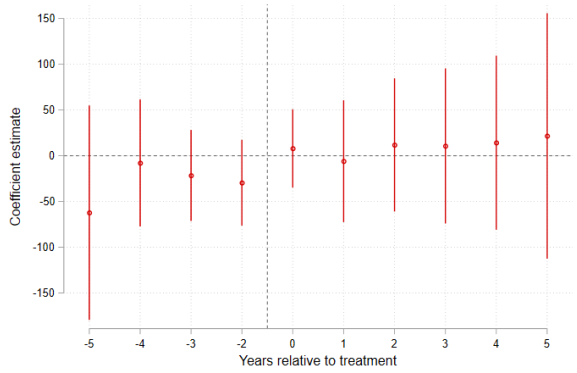

H.III.3 Total medical debt in collections: baseline

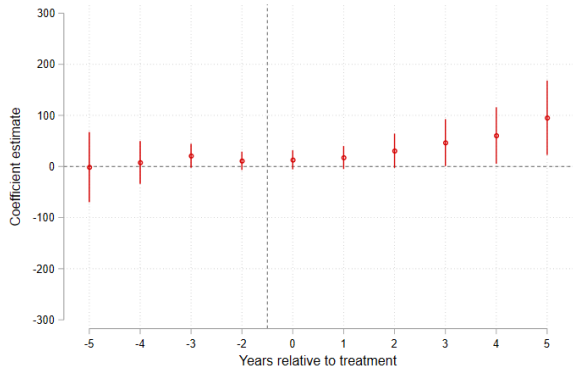

H.III.4 Total medical debt in collections: subsample with debt in 2013

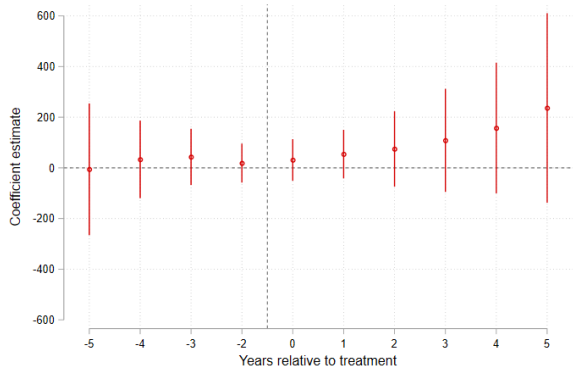

This figure replicates Figure III in the main draft, but restricting the sample to individuals that had medical debt in collections in 2013. Each panel shows the point estimates and 95% confidence intervals for β_e, obtained from implementing Equation 2, and illustrates the effect of exposure to hospital mergers on medical debt outcomes. Panels H.III.1 and H.III.3 replicate the corresponding baseline panels in Figure III. Panels H.III.2 and H.III.4 show the results for the corresponding outcomes when restricting the sample to individuals with medical debt in collections in 2013. The specification fixed effects, controls, and standard errors are as defined in Figure III. The horizontal axis denotes years relative to the merger event; event time $t = -1$ is omitted as the reference period, with the dashed vertical line indicating the period just before the merger took place.

J. Event study regression tables and supplemental results for the IHME mechanisms analysis

J.1 Event study regression tables

Tables J.I, J.II, and J.III report the event study regression estimates underlying Figures VI, VII, and VIII in the main text, respectively. Each table corresponds to one figure, and each column reports the estimates plotted in the corresponding panel of that figure. All columns report the Sun and Abraham (2021) interaction-weighted event study coefficients from Equation 1, estimated on the CZ-level IHME Disease Expenditure panel with CZ and year fixed effects, a binary control for state-level Medicaid expansion, and time-varying controls for the log of CZ population and the log of CZ population aged 65 and over. Standard errors are clustered at the CZ level, and event time $t = -1$ is the omitted reference period.

Table J.I. Event study estimates for the effect of merger exposure on private insurance spending and utilization (2010–2019).

	(1) Log private ins. spending	(2) Log private ins. encounters	(3) Log spending per private enrollee	(4) Log spending per private encounter
-5	0.0052	0.0079	0.0084	-0.0027
	(0.0098)	(0.0110)	(0.0100)	(0.0049)
-4	-0.0022	-0.0032	0.0002	0.0010
	(0.0063)	(0.0069)	(0.0062)	(0.0033)
-3	-0.0101*	-0.0064	-0.0085	-0.0037
	(0.0055)	(0.0055)	(0.0054)	(0.0027)
-2	-0.0014	-0.0024	-0.0023	0.0010
	(0.0037)	(0.0036)	(0.0035)	(0.0019)
0	0.0003	-0.0051	0.0002	0.0054***
	(0.0037)	(0.0036)	(0.0035)	(0.0019)
1	0.0042	-0.0055	0.0038	0.0097***
	(0.0055)	(0.0052)	(0.0052)	(0.0026)
2	0.0106	-0.0026	0.0106	0.0132***
	(0.0072)	(0.0066)	(0.0067)	(0.0033)
3	0.0189**	-0.0061	0.0183**	0.0251***
	(0.0086)	(0.0084)	(0.0082)	(0.0046)
4	0.0259***	-0.0020	0.0248***	0.0279***
	(0.0099)	(0.0103)	(0.0094)	(0.0060)
5	0.0281**	-0.0008	0.0258**	0.0289***
	(0.0131)	(0.0139)	(0.0126)	(0.0076)
Observations	6,854	6,854	6,854	6,854
Year FEs	Y	Y	Y	Y
CZ FEs	Y	Y	Y	Y
Medicaid expansion	Y	Y	Y	Y
Population controls	Y	Y	Y	Y

*** p<0.01; ** p<0.05; * p<0.10

Notes: Robust standard errors clustered at the CZ level in parentheses. Each column reports the Sun and Abraham (2021) interaction-weighted event study coefficients plotted in the corresponding panel of Figure VI in the main text, estimated from equation 1 using the IHME Disease Expenditure data. Column (1) uses log total private insurance spending as the outcome; column (2) log privately-insured encounters; column (3) log spending per privately-insured enrollee; column (4) log spending per privately-insured encounter. All regressions include CZ and year fixed effects, a binary control for state-level Medicaid expansion, and time-varying controls for the log of CZ population and the log of CZ population aged 65 and over. Event time t = -1 is the omitted reference period.

Table J.II. Event study estimates for the effect of merger exposure on out-of-pocket spending and utilization (2010–2019).

	(1) Log OOP spending	(2) Log OOP encounters	(3) Log OOP spending per encounter	(4) OOP spending share
-5	-0.0204	-0.0076	-0.0128	-0.0049***
	(0.0136)	(0.0138)	(0.0089)	(0.0017)
-4	-0.0018	0.0035	-0.0053	-0.0017
	(0.0085)	(0.0086)	(0.0060)	(0.0011)
-3	-0.0051	0.0031	-0.0082*	-0.0012
	(0.0068)	(0.0062)	(0.0045)	(0.0009)
-2	-0.0032	-0.0047	0.0015	-0.0005
	(0.0047)	(0.0051)	(0.0031)	(0.0006)
0	0.0142***	0.0124**	0.0018	0.0024***
	(0.0051)	(0.0062)	(0.0030)	(0.0007)
1	0.0102	0.0054	0.0048	0.0021**
	(0.0070)	(0.0079)	(0.0038)	(0.0009)
2	0.0126	0.0145	-0.0019	0.0029**
	(0.0093)	(0.0100)	(0.0048)	(0.0012)
3	0.0229**	0.0294***	-0.0065	0.0040***
	(0.0109)	(0.0108)	(0.0062)	(0.0014)
4	0.0093	0.0185*	-0.0092	0.0024
	(0.0123)	(0.0111)	(0.0084)	(0.0016)
5	0.0207	0.0364**	-0.0156	0.0040*
	(0.0185)	(0.0159)	(0.0125)	(0.0022)
Observations	6,854	6,854	6,854	6,854
Year FEs	Y	Y	Y	Y
CZ FEs	Y	Y	Y	Y
Medicaid expansion	Y	Y	Y	Y
Population controls	Y	Y	Y	Y

*** p<0.01; ** p<0.05; * p<0.10

Notes: Robust standard errors clustered at the CZ level in parentheses. Each column reports the Sun and Abraham (2021) interaction-weighted event study coefficients plotted in the corresponding panel of Figure VII in the main text, estimated from equation 1 using the IHME Disease Expenditure data. Column (1) uses log total out-of-pocket spending as the outcome; column (2) log out-of-pocket encounters; column (3) log out-of-pocket spending per encounter; column (4) the out-of-pocket share of total healthcare spending, in levels. All regressions include CZ and year fixed effects, a binary control for state-level Medicaid expansion, and time-varying controls for the log of CZ population and the log of CZ population aged 65 and over. Event time t = -1 is the omitted reference period.

Table J.III. Event study estimates for the effect of merger exposure on Medicare coverage and payer composition (2010–2019).

	(1) Log Medicare enrollment	(2) Log private ins. enrollment	(3) Medicare spending share	(4) Private ins. spending share
-5	-0.0057***	-0.0031	0.0048***	-0.0051*
	(0.0017)	(0.0031)	(0.0014)	(0.0028)
-4	-0.0029***	-0.0023	0.0040***	-0.0054***
	(0.0011)	(0.0017)	(0.0009)	(0.0018)
-3	-0.0014*	-0.0016	0.0032***	-0.0058***
	(0.0008)	(0.0014)	(0.0008)	(0.0016)
-2	-0.0001	0.0009	0.0015***	-0.0008
	(0.0005)	(0.0011)	(0.0005)	(0.0011)
0	0.0000	0.0001	-0.0010*	0.0029**
	(0.0007)	(0.0013)	(0.0006)	(0.0011)
1	-0.0015	0.0004	-0.0040***	0.0050***
	(0.0011)	(0.0017)	(0.0008)	(0.0015)
2	-0.0027*	-0.0001	-0.0059***	0.0086***
	(0.0014)	(0.0021)	(0.0010)	(0.0019)
3	-0.0040**	0.0006	-0.0095***	0.0121***
	(0.0019)	(0.0026)	(0.0012)	(0.0022)
4	-0.0058**	0.0011	-0.0116***	0.0159***
	(0.0024)	(0.0030)	(0.0015)	(0.0025)
5	-0.0064*	0.0023	-0.0120***	0.0177***
	(0.0034)	(0.0040)	(0.0021)	(0.0033)
Observations	6,854	6,854	6,854	6,854
Year FEs	Y	Y	Y	Y
CZ FEs	Y	Y	Y	Y
Medicaid expansion	Y	Y	Y	Y
Population controls	Y	Y	Y	Y

*** p<0.01; ** p<0.05; * p<0.10

Notes: Robust standard errors clustered at the CZ level in parentheses. Each column reports the Sun and Abraham (2021) interaction-weighted event study coefficients plotted in the corresponding panel of Figure VIII in the main text, estimated from equation 1 using the IHME Disease Expenditure data. Column (1) uses log Medicare enrollment as the outcome; column (2) log private insurance enrollment; column (3) the Medicare share of total healthcare spending, in levels; column (4) the private insurance share of total healthcare spending, in levels. Medicare is inclusive of Medicare Advantage. All regressions include CZ and year fixed effects, a binary control for state-level Medicaid expansion, and time-varying controls for the log of CZ population and the log of CZ population aged 65 and over. Event time t = -1 is the omitted reference period.

J.2 Hospital mergers and the uninsured population

The payer-mix results in Subsection V.3 document that hospital mergers reduce Medicare enrollment and Medicare's share of local healthcare spending while increasing the private insurance share. A natural question is whether this recomposition partly reflects an extensive-margin deterioration in coverage: if merger exposure pushed patients out of insurance altogether, the shift in payer shares could reflect coverage loss rather than a reallocation of covered care across payers. Rising uninsurance would also constitute an independent, direct channel from mergers to medical debt.

Figure J.I addresses this question by estimating Equation 1 with the log of the CZ-level uninsured population, from the Census Bureau's Small Area Health Insurance Estimates (SAHIE) program, as the outcome. The estimates provide no evidence that hospital mergers increase the number of uninsured residents in treated commuting zones. If anything, the post-treatment point estimates are negative, but they are imprecisely estimated and are preceded by mildly positive pre-treatment estimates, so we do not interpret them as evidence of a merger-induced decline in uninsurance. What the figure does rule out is an increase: there is no post-treatment horizon at which the estimated effect on the uninsured population is positive. The payer-mix channel documented in Subsection V.3 therefore operates on the intensive margin of insurance coverage—a reallocation of covered care from Medicare toward commercial payers—rather than on the extensive margin of coverage itself.

Figure J.I. EVENT STUDY PLOT FOR THE EFFECT OF HOSPITAL MERGER EXPOSURE ON THE CZ-LEVEL UNINSURED POPULATION

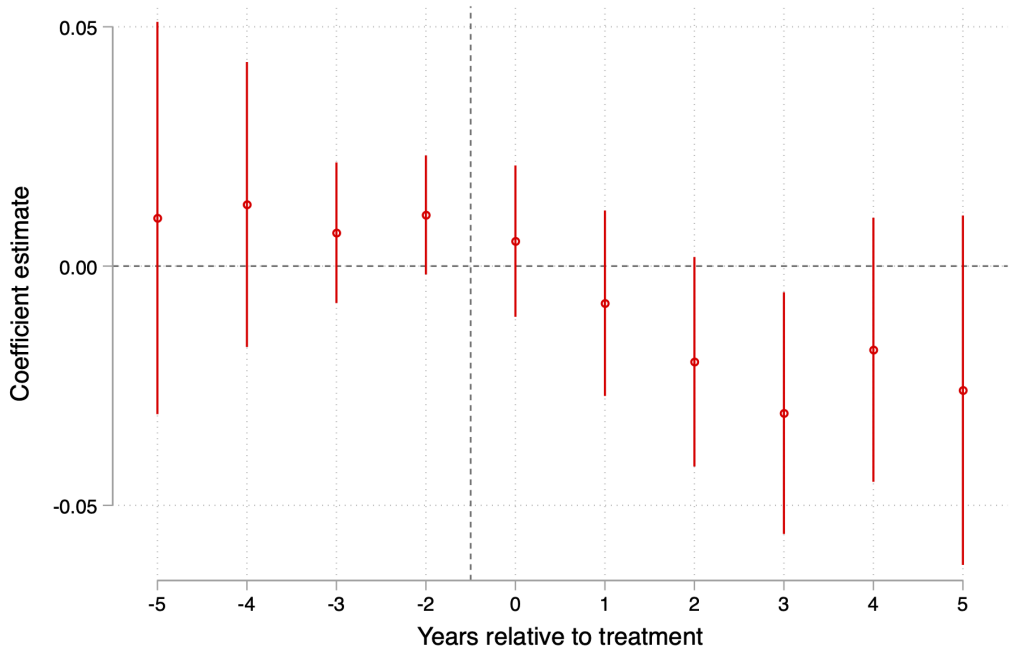

This figure shows the point estimates and 95% confidence intervals for $\hat{\beta}_e$, obtained from estimating Equation 1 with the log of the CZ-level uninsured population as the outcome of interest. The uninsured population count comes from the Census Bureau's Small Area Health Insurance Estimates (SAHIE) program. The regression includes CZ and year fixed effects, a binary control for state-level Medicaid expansion, and time-varying controls for the log of CZ population and the log of CZ population aged 65 and over; standard errors are clustered at the CZ level. The horizontal axis denotes years relative to the merger event; event time $t = -1$ is omitted as the reference period, with the dashed vertical line indicating the period just before the merger took place.

K. The Labor Market Channel

Here we consider whether the medical debt effects documented in Section IV are partly the result of lower local labor demand or employer power resulting from hospital mergers, which could reduce earnings and/or employment in local populations and thereby render them less able to pay medical bills.

Several recent papers have found hospital mergers negatively impacted healthcare workers. For example, Prager and Schmitt (2021) find that hospital mergers which result in large increases in local concentration negatively impact wage growth for workers with hospital industry-specific skills, while Allegretto and Graham-Squire (2023) find that nurses' wages are negatively impacted by hospital system consolidation in smaller markets. Andreyeva et al. (2024) find large cost savings result from systems acquiring independent hospitals, primarily through employee reductions for support and back-office staff. Such effects of hospital mergers could financially stress hospital industry workers and make them less able to service their debts.

Conversely, Brot et al. (2024) use hospital mergers to instrument for increases in local healthcare prices and find the consequent increases in local non-wage labor compensation (driven by more expensive employer-sponsored health insurance) have no effects within healthcare employers but significantly reduce payroll and employment at non-healthcare industry employers, driving significantly lower average earnings and higher unemployment at the county level. In effect, employers move up and to the left on the local aggregate labor demand function. The price channel documented in Subsection V.1—which shows an approximately 2.8% increase in private insurance spending five years after treatment—is the type of commercial insurance cost increase that Brot et al. (2024) argue passes through to employer-sponsored health insurance premiums and depresses local labor demand. If operative, this channel could contribute to our medical debt findings by reducing workers' capacity to pay for healthcare, independently of the direct price and cost-sharing effects on patients.

We test for evidence of hospital merger impacts on local earnings and employment, as well

as on local hiring and separations, for CZ-wide aggregate effects and within the Health Care and Social Assistance (NAICS 62) sector using the county-by-industry-by-quarter Quarterly Workforce Indicators (QWI) data from the U.S. Census Bureau, which we aggregate at the CZ-by-industry-by-year level. Further disaggregation of the Health Care and Social Assistance Sector results in sharply higher numbers of suppressed cells, which recommends against doing so.

To dig deeper into the impacts of hospital mergers on hospital industry-specific employment, we utilize hospital-level annual full-time equivalent (FTE) employment counts from the American Hospital Association. Earnings are not observed in this data set. We estimate Equation 1 with these as our outcomes variables.

Figure K.I shows the estimated impacts on the log of average earnings and employment by sector. There is no evidence of a significant effect on CZ-wide average earnings (Panel A). There is some evidence of a very small, negative, effect on average earnings in the Health Care and Social Assistance sector (Panel B), growing to nearly -2 log points by year five, though a modest differential trend cannot be ruled out. There is also some evidence of a small, positive impact on CZ-wide employment (Panel C), reaching just over 1 log point in year two and then stabilizing. Panel D shows a clear increase in employment in the Health Care and Social Assistance sector (statistically significant by the second year after treatment, reaching over 2.5 log points by year five), which taken with the sectoral average earnings effect in Panel B suggests that hospital mergers resulted in an increased number of lower-wage or part-time workers being hired within the sector. However, the pre-trends for this result are concerning and suggest it should be taken with some skepticism. Interestingly, Panel E suggests no evidence that mergers increased FTE employment within hospitals (though there is also no evidence of a decrease in hospital FTE employment).

To further understand the employment patterns described above, Figure K.II shows our estimated impacts of hospital mergers on hires and separations, by sector. For all sectors within treated CZs there is evidence of an increase in both hires (Panel A) and separations (Panel C), though that outcome exhibits severe pre-trends which are suggestive of a null effect. These results further cast doubt on the marginally positive estimated effects on all-sector employment in

Figure K.I. Focusing on the Health Care and Social Assistance sector, there is evidence of larger increases in both hires (Panel B, reaching nearly 7 log points by year five) and separations (Panel D, over 6 log points by year five), both with good pre-trends and statistically significant by the second year after treatment, which hint that notably higher churn (which is defined at the level of the employer (Lazear and McCue, 2018), so not directly measurable here) is driving the sectoral increase in employment and decrease in average earnings seen in Figure K.I.

In sum, despite the fact that the price channel findings in Subsection V.1 provide the upstream condition identified by Brot et al. (2024)—higher commercial insurance costs—we do not find evidence that this translates into the downstream labor market contraction they document. The absence of a CZ-wide earnings decline suggests that the labor demand channel is not a primary driver of the medical debt results documented in Section IV. It is possible that increased churn within the Health Care and Social Assistance sector may reflect worsening labor standards for workers in that sector, but the earnings and employment results suggest any such impact would have been modest. The medical debt effects are therefore more likely attributable to the direct price, cost-sharing, and coverage channels documented in Section V than to an indirect labor market pathway.

Figure K.I. EVENT STUDY PLOTS FOR THE EFFECT OF HOSPITAL MERGER EXPOSURE ON EARNINGS AND EMPLOYMENT, BY SECTOR

K.I.1 Average earnings:
All Sectors

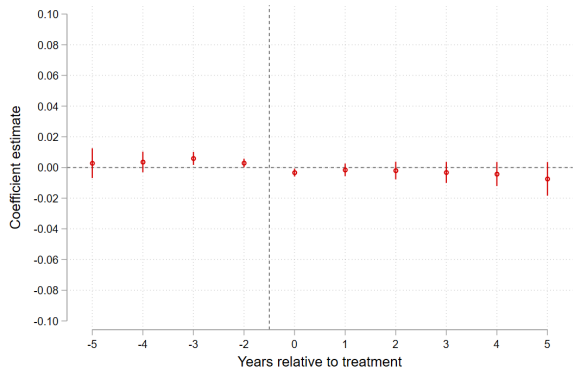

K.I.2 Average earnings:
Health Care and Social Assistance

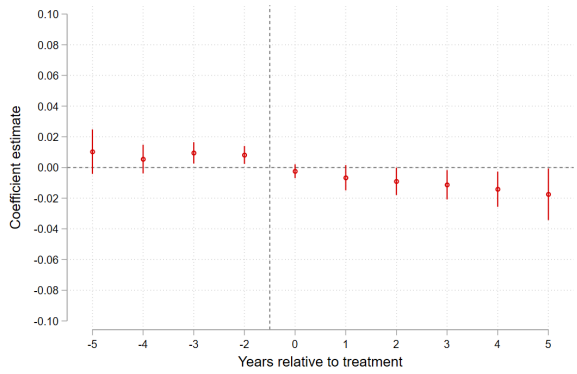

K.I.3 Employment:
All Sectors

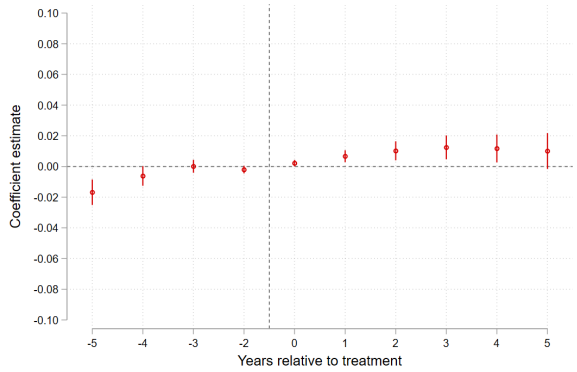

K.I.4 Employment:
Health Care and Social Assistance

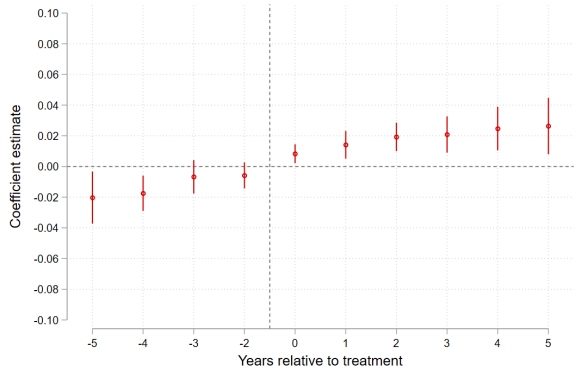

K.I.5 Employment (FTE): Hospitals

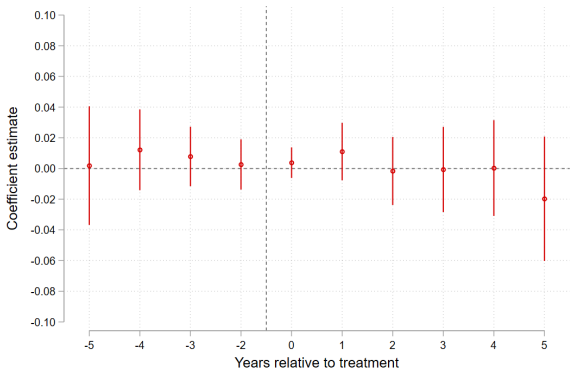

Each panel shows the point estimates and 95% confidence intervals for β_e, obtained from implementing Equation 1, and illustrates the effect of exposure to hospital mergers. The estimates included year and CZ fixed effects, as well as a binary control that accounts for the timing of state-level Medicaid expansion, and the logs of the total population and the population aged 65+, as outlined in Section III.1; standard errors are clustered at the CZ level. The outcome for Panel K.I.1 is average CZ-level earnings. The outcome for Panel K.I.2 is average CZ-level earnings in the Health Care and Social Assistance sector. The outcome for Panel K.I.3 is average CZ-level employment. The outcome for Panel K.I.4 is average CZ-level employment in the Health Care and Social Assistance sector. The outcome for Panel K.I.5 is CZ-level full-time-equivalent employment in hospitals. The horizontal axis denotes years relative to the merger event; event time $t = -1$ is omitted as the reference period, with the dashed vertical line indicating the period just before the merger took place.

Figure K.II. EVENT STUDY PLOTS FOR THE EFFECT OF HOSPITAL MERGER EXPOSURE ON NEW HIRES, SEPARATIONS, AND LABOR MARKET TIGHTNESS, BY SECTOR

K.II.1 Hires:
All Sectors

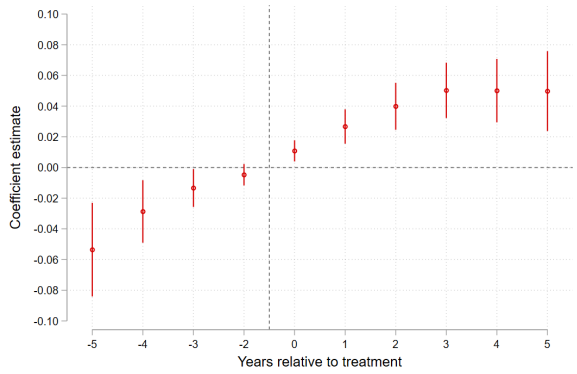

K.II.2 Hires:
Health Care and Social Assistance

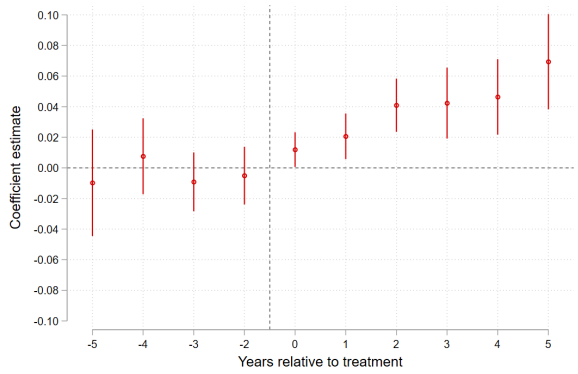

K.II.3 Separations:
All Sectors

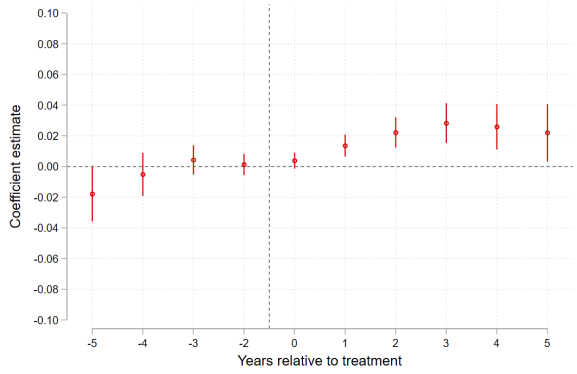

K.II.4 Separations:
Health Care and Social Assistance

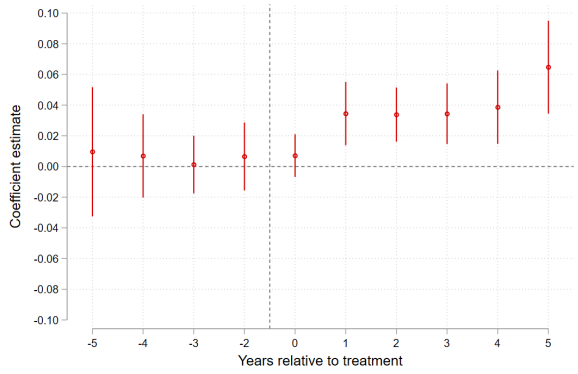

Each panel shows the point estimates and 95% confidence intervals for β_e, obtained from implementing Equation 1, and illustrates the effect of exposure to hospital mergers. The estimates included year and CZ fixed effects, as well as a binary control that accounts for the timing of state-level Medicaid expansion, and the logs of the total population and the population aged 65+, as outlined in Section III.1; standard errors are clustered at the CZ level. The outcome for Panel A is total CZ-level hires. The outcome for Panel B is CZ-level hires in the Health Care and Social Assistance sector. The outcome for Panel C is total CZ-level separations. The outcome for Panel D is total CZ-level separations in the Health Care and Social Assistance sector. The horizontal axis denotes years relative to the merger event; event time $t = -1$ is omitted as the reference period, with the dashed vertical line indicating the period just before the merger took place.

L. The Financial Distress Channel

The results in Section IV indicate that hospital mergers increase medical debt. But if financially distressed hospitals are more likely to get acquired and are also more likely to resort to debt collection and credit reporting to enforce payment, then it is possible that hospital financial distress and not mergers drive the increased medical debt we document.

To address this identification concern, we obtained the American Hospital Association (AHA) Annual Survey data underlying the Cooper et al. (2024) dataset used in our main analysis. The AHA survey contains measures of hospital financial health that are not available in the Cooper et al. (2024) data, specifically interest expense and total expenses.[28] We acquired the survey for alternating years between 2011 and 2019 (2011, 2013, 2015, 2017, and 2019) and linearly interpolated interest expense and total expense values for the intervening years. Our final sample is a hospital-level dataset that includes all hospitals except the acquiring hospital/system for our treated units. For treated units, we limit observations to the two years leading up to and including the acquisition year since post-merger financial distress of acquired hospitals can be conflated with the acquirer's financial health. We then use this data to estimate the following model:

$$target_{h,z,t} = \beta \text{distress}_{h,z,t-x} + \eta_{h,z,t} + \gamma_z + \lambda_t + \phi x_{z,t} + \varepsilon_{z,t} \tag{3}$$

where $target_{h,z,t}$ is an indicator variable that takes the value one when hospital h is acquired in commuting zone z in year t. If hospital h is acquired more than once in our analysis period, we include only the first acquisition in our analysis. $\text{distress}_{h,z,t-x}$ is a measure of hospital h financial health in commuting zone z, x years prior to the merger at time t. Distress is constructed by taking a hospital's interest expense as a share of its total expense.[29] And since financial distress *prior*

[28] 90% of AHA survey data matches to Cooper et al. (2024) data.

[29] Note that interest expenses can be negative for one of three reasons: 1. The hospital is a net-creditor such that its gains from interest is higher than what it paid, 2. capitalized interest, which means that the hospital is doing a major construction and the interest it is paying on that construction goes under cost such that its interest expenses are negative, 3. financial operations like swapping/refinancing/premium effects. 10% of our sample has negative interest expenses, which we set to zero so as not to drop them from our analysis while simultaneously (correctly) classifying them as non-distressed hospitals. The results presented in this section are robust to dropping these observations and to

to the merger rather than the year of the merger is the relevant predictor, we report the results of including a one year lag and a two year lag, but not both of them simultaneously given they are highly correlated. $\eta_{h,z,t}$ is the log of hospital h's bed count in z at time t to control for the role of hospital size in acquisition probability. γ_z and λ_t are CZ and year fixed effects, respectively. $x_{z,t}$ represents time-varying CZ-level control variables: i. a binary variable indicating whether the ACA Medicaid expansion was enacted in z at t, ii. the log of commuting zone population, and iii. the log of commuting zone population aged 65 and over. Standard errors are clustered at the CZ-level.

leaving them as negative values in our analysis.

As presented in Table L.I, we find that one standard deviation increase in a hospital's distress measure ($\approx$ 1.78 percentage points of interest expenses relative to total expenses) is associated with roughly a 0.7-1 percentage point increase in the probability of being acquired, depending on whether we look at a hospital's financial distress one or two years prior to the merger. These findings suggest that financial distress is a statistically-significant predictor of mergers, lending credence to the identification concern on which this analysis is premised.

Table L.I. FINANCIAL DISTRESS AND THE PROBABILITY OF BEING TARGETED

	Target	Target
Distress, $t-1$	0.00713***	
	(0.00191)	
Distress, $t-2$		0.00984***
		(0.00232)
Commuting zone FE	Yes	Yes
Year FE	Yes	Yes
N	22,111	19,191
R^2	0.074	0.083

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Estimates of the linear probability model presented in equation 3; coefficients scaled per one standard deviation of financial distress, which is equal to approximately 1.78 percentage point. All columns include commuting zone and year fixed effects. Both specifications have the following controls: an indicator for whether the state had expanded Medicaid by the given year, log of hospital bed count, log of total population and log population aged 65 and over. Robust standard errors clustered at the commuting zone level in parentheses.

The next question is whether the distress that increases the likelihood a hospital is an acquisition target is responsible for post-acquisition medical debt. In order to answer that question, we investigate the financial impact of mergers on distressed target hospitals. We use the Sun and Abraham (2021) event study estimator to evaluate the following model:

$$\text{distress}_{h,z,t} = \sum_{e=-5}^{E_i} \beta_e \mathbb{1}\{\ell_{h,z,c}\}\mathbb{1}\{t - c_z = e\} + \eta_{h,z,t} + \gamma_z + \lambda_t + \phi x_{z,t} + \varepsilon_{z,t} \tag{4}$$

where $distress_{h,z,t}$ is now our dependent variable and $\mathbb{1}\{\ell_{h,z,c}\}$ indicates that hospital h was acquired in commuting zone z in treatment year c_z. $\eta_{h,z,t}$ is a log of hospital h's bed count in z at time t to control for hospital size's role in magnitude of financial distress. γ_z and λ_t are CZ and

year fixed effects, respectively. $x_{z,t}$ represents time-varying CZ-level control variables. Standard errors are clustered at CZ level.

The results presented in Figure L.I below show that acquired hospitals were distressed compared to control hospitals in the three years leading up to the merger, including the merger year. But financial distress eases after the merger such that coefficients are not statistically different from zero for the four years following the merger. This tells us that while distressed hospitals are more likely to be acquired, the merger appears to remedy their financial difficulties such that targets are no different from control hospitals two years after the acquisition. We therefore conclude that it is not financial distress that is driving our medical debt findings: pre-treatment coefficients are zero and post-treatment coefficients magnify and are statistically significant five years post-merger. This inference is strengthened by the results in the following Appendix M, which show that target hospitals are not more likely to use the apparatus of credit reporting than control hospitals.

Figure L.I. MERGERS' ROLE IN FINANCIAL DISTRESS

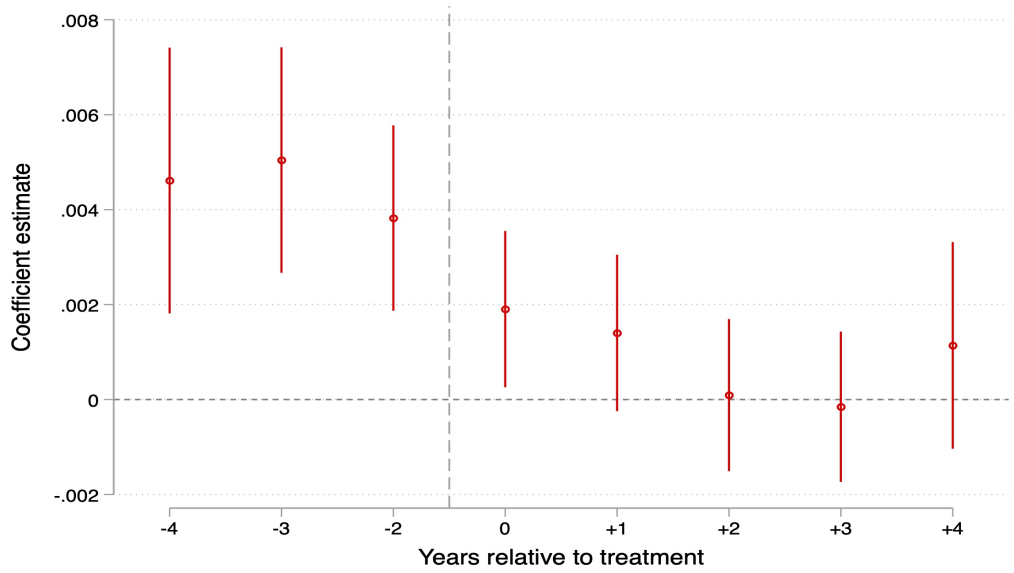

Notes: This figure shows the point estimates and 95% confidence intervals for β_e, obtained from estimating Equation 4 when the outcome of interest is hospital distress. The estimates include year and CZ fixed effects, as well as a binary control that accounts for the timing of state level Medicaid expansion, and the logs of hospital bed count, the total population and the population aged 65+; standard errors are clustered at the CZ level. The horizontal axis denotes years relative to the merger event; event time $t = 1$ is omitted as the reference period, with the dashed vertical line indicating the period just before the merger took place.

M. Hospital Debt-Collection Practices

The results in Section IV establish that hospital mergers increase the medical debt carried by consumers in affected markets. A natural follow-on question concerns the hospital conduct that generates that debt, beyond the pricing and coverage outcomes addressed in section V: do merged hospitals differ from other hospitals in how they pursue payment once bills go unpaid? This appendix links the hospital-level ownership panel underlying our treatment definition to newly-collected data on hospital billing and debt-collection practices.

Two findings organize what follows. First, merger targets are *not* more likely than other hospitals to report patient debt to credit bureaus or to sell patient debt to collectors. Those are the two practices that place unpaid bills onto credit reports and thereby generate the collections tradelines we study in the body of the paper, so the explanation for our main findings is *not* likely to be that mergers increase reporting and debt sales, as opposed to the unpaid bills that underlie them. Second, targets are less likely to use the more aggressive collection measures (suing patients, placing liens, garnishing wages), which reflects a broader cross-sectional regularity: independent hospitals use these measures far more than hospitals belonging to multi-hospital systems, and an acquisition converts an independent hospital into a system hospital.

Unfortunately, these findings cannot support a causal interpretation. Unlike every other outcome we study, the collections practices data are a single post-period cross-section observed after every merger in the ownership panel. No event study is possible, and selection into merging is uncontrolled. We report these findings because they discipline the interpretation of our main results: they help rule out one candidate mechanism (increased reporting and debt sales by targets) and characterize the margin of hospital conduct, court-based collection enforcement, that our credit-panel outcomes do not directly measure.

M.1 Data and linkage

The Lown Institute has assembled a hospital-level dataset on billing and collection practices covering 2,500 hospitals, collected from hospitals' publicly posted policy documents during 2024–2026

and confirmed, where possible, by direct contact with hospital staff (Lown Institute, 2026). For each hospital, the data record whether the hospital reports patient debt to credit bureaus, sells patient debt, denies care to patients with outstanding bills, sues patients, places liens on property, and garnishes wages, along with the hospital's financial-assistance policies for indigent patients. Each item is recorded as an affirmative yes, an explicit no, or "no information online." We code each item accordingly and treat non-disclosure as its own outcome: across hospitals, 40 percent of item-responses are undisclosed, and the share of a hospital's items that are undisclosed, its billing-policy *opacity*, is our seventh outcome. We examined the financial-assistance items as well, comparing acquiring, target, and never-treated hospitals, and found no robust differences. Those estimates are imprecise and inconsistently signed across items and specifications that could possibly explain the paper's main findings.

The Lown Institute collections practices data are keyed by CMS Certification Number (CCN), while the hospital ownership panel from Cooper et al. (2024) is keyed by AHA identifier. We construct a crosswalk from the AHA annual surveys (2011–2024), which report both identifiers, assigning each AHA hospital its most recent CCN. Hospitals whose provider number changed across survey years, typically reflecting post-merger provider-number consolidation, are matched on their most recent CCN, supplemented by matches on historical CCNs and by manually reviewed name-and-state matches for hospitals whose current AHA record lacks a provider number. The result links 2,399 of the 2,500 hospitals in the Lown Institute sample to the ownership panel. Inclusion in the Lown Institute sample is strongly predicted by hospital size but is uncorrelated with merger-target status: in a logit of Lown Institute inclusion on post-2013 target status, acquirer status, and log beds among panel hospitals operating in 2014 or later, the target coefficient is 0.03 ($p = 0.70$) while log beds enters at 0.33 ($p < 0.01$).

Because this analysis does not rely on the credit panel, we are not restricted to the 2014–2018 treatment cohorts that are the subject of the main analysis. We instead use every merger observable in the hospital-level ownership panel, 2000–2022. Among the 2,390 matched hospitals, 960 were acquired at least once ("ever-targets"), 703 belonged to a system that acquired at least one hospital

("ever-acquirers"), 340 played both roles at some point, and 1,076 were never party to a merger. But since the collections practices outcomes are observed only once, after the end of the ownership panel, there is no way to estimate an event study. We also classify each hospital by whether it belongs to a multi-hospital system in its last observed panel year. Among matched never-treated hospitals, 684 are independent and 392 are system members.

M.2 Empirical approach

We estimate cross-sectional regressions on the full matched sample with separate indicators for the two merger roles:

$$y_h = \beta_T \,\text{Target}_h + \beta_A \,\text{Acquirer}_h + \gamma \log(\text{beds}_h) + \delta_{o(h)} + \gamma_{z(h)} + \varepsilon_h, \tag{5}$$

where y_h is a collections practices outcome (or the opacity share) for hospital h, Target_h indicates that h was acquired in a merger during 2000–2022, Acquirer_h indicates that h belonged to an acquiring system during that period (the two indicators can both equal one), $\delta_{o(h)}$ are ownership-type fixed effects (nonprofit, public, for-profit), and $\gamma_{z(h)}$ are commuting-zone fixed effects, the geographic market definition used throughout the paper, so that comparisons are among hospitals serving the same local market. Never-treated hospitals are the implicit reference group. Standard errors are clustered by the hospital system to which h belonged in its last observed panel year, for reasons Table M.III makes apparent. We emphasize what this design cannot rule out: selection into merging is uncontrolled, and with no pre-period data we cannot distinguish practices that changed upon acquisition from pre-existing differences among hospitals that merged.

M.3 Results

Table M.I and Figure M.I report the estimates of Equation 5 for all seven outcomes.

Start with the two practices most relevant to the rest of the paper. *Reporting debt to credit bureaus* shows no difference for targets (0.003, SE 0.071) or acquirers (0.001, SE 0.032): merger participants are statistically indistinguishable from never-treated hospitals in their same commuting

zone in the practice that places unpaid bills onto credit reports. *Selling patient debt* is positive but statistically insignificant for targets (0.110, SE 0.100) and acquirers (0.016, SE 0.043). Whatever generates the medical debt increases we document in Section IV, it is not that merged hospitals adopt reporting or debt-sale practices that inflate collections tradelines, holding constant the level of unpaid bills.

Denying care over unpaid bills bears on our credit card results. In Section IV.3.1 we find that merger exposure raises credit card balances along the intensive margin and note that hospitals conditioning access to non-emergency care on the resolution of outstanding balances could drive patients to finance medical bills on credit cards (Kliff and Silver-Greenberg, 2023). The practice is indeed widespread: 37 percent of disclosing never-treated hospitals report it. But neither targets nor acquirers are any more likely to report it than never-treated hospitals in the same commuting zone. Our finding with respect to credit card debt is therefore not due to merged hospitals being more likely to use this mechanism. As we show below, care denial is most prevalent among independent hospitals.

The court-based enforcement items are where targets differ. Targets are 15 percentage points less likely to place liens (SE 0.057) and 15 points less likely to garnish wages (SE 0.073), with a marginally significant 12-point difference for suing patients (SE 0.072), against never-treated means near 0.70. The acquirer coefficients on all three are tight zeros: acquiring systems' hospitals collect like never-treated hospitals (conditional on covariates), while the hospitals they absorb pursue court-based collection less than either. Finally, both merger roles predict greater transparency. The undisclosed share is 6.1 points lower for targets ($p = 0.05$) and 4.1 points lower for acquirers ($p < 0.01$), against a never-treated mean of 48 percent.

Table M.II splits the target coefficient by acquisition cohort. The credit-reporting and debt-sale nulls hold in every cohort; no cohort uses either practice significantly less than never-treated hospitals. The court-enforcement differences are concentrated among hospitals acquired in 2014–18, the treatment window of the main analysis, and among the 2019–22 cohort, while targets of the 2000s merger waves differ little from never-treated hospitals. Under a causal reading, recently

Table M.I. DEBT-COLLECTION PRACTICES: EVER-TARGETS AND EVER-ACQUIRERS VS. NEVER-TREATED HOSPITALS, 2000–2022 MERGERS

	Reports debt	Sells debt	Denies care	Sues patients	Places liens	Garnishes wages	Share no info
Ever target	0.003	0.110	-0.046	-0.122*	-0.146**	-0.147**	-0.061*
	(0.071)	(0.100)	(0.060)	(0.072)	(0.057)	(0.073)	(0.031)
Ever acquirer	0.001	0.016	0.031	-0.013	-0.006	-0.023	-0.041***
	(0.032)	(0.043)	(0.040)	(0.033)	(0.034)	(0.039)	(0.014)
Never-treated mean	0.721	0.436	0.372	0.716	0.708	0.695	0.477
N	1,472	1,005	945	1,246	1,360	1,228	2,399

Notes: *** p<0.01; ** p<0.05; * p<0.10. Each column reports Equation 5 for one outcome: coefficients on ever-target and ever-acquirer status (2000–2022; the indicators can overlap), with log beds, ownership-type fixed effects, and commuting-zone fixed effects; standard errors clustered by hospital system in parentheses. Item outcomes are binary indicators coded from the Lown data with "no information online" treated as missing; the final column is the share of a hospital's items that are undisclosed. "Never-treated mean" is the mean among hospitals party to no merger. Sample: 2,399 matched hospitals; *N* varies with item-level non-disclosure.

Figure M.I. MERGER ROLES AND DEBT-COLLECTION PRACTICES

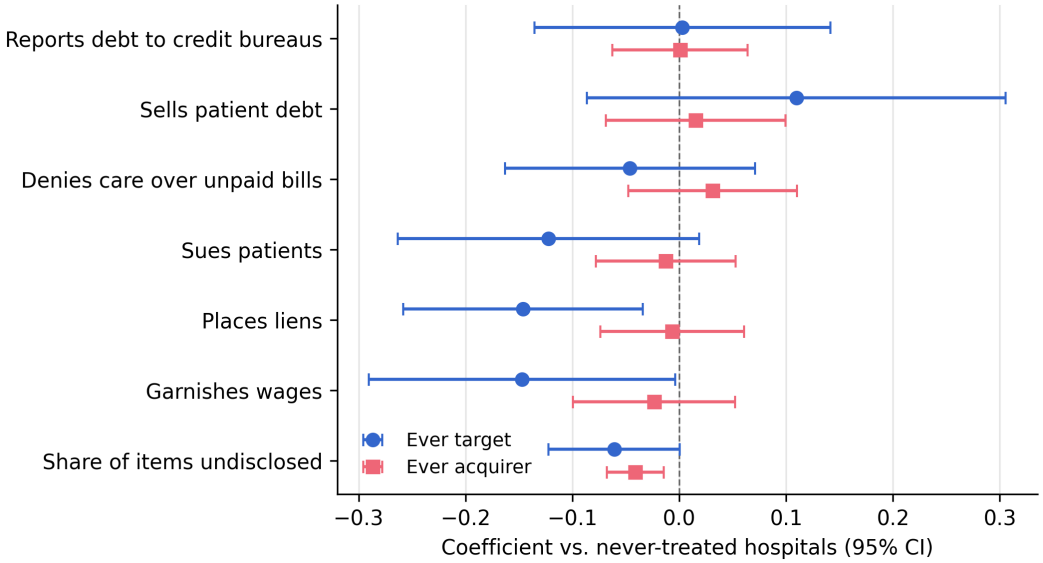

Notes: Point estimates and 95% confidence intervals for the ever-target and ever-acquirer coefficients reported in Table M.I.

acquired hospitals are brought onto their acquirers' contemporary collection regimes, which have softened amid mounting public and regulatory scrutiny of hospital collection lawsuits. Under a selection reading, the composition of targets has shifted across waves. Either way, the mergers driving our main results are exactly the ones for which target collection conduct is most distinctly *less* aggressive.

Table M.II. TARGET PRACTICE DIFFERENCES BY ACQUISITION COHORT

	Reports debt	Sells debt	Denies care	Sues patients	Places liens	Garnishes wages	Share no info
Acquired 2001–04	0.074	0.191**	0.068	-0.059	-0.080	-0.028	-0.015
	(0.078)	(0.095)	(0.093)	(0.089)	(0.085)	(0.086)	(0.046)
Acquired 2005–09	0.000	0.058	0.107	-0.040	-0.049	-0.092	-0.030
	(0.082)	(0.136)	(0.122)	(0.093)	(0.082)	(0.117)	(0.045)
Acquired 2010–13	0.037	0.059	0.007	0.001	-0.013	0.002	-0.023
	(0.078)	(0.087)	(0.065)	(0.064)	(0.062)	(0.065)	(0.041)
Acquired 2014–18	-0.127	-0.015	-0.128*	-0.216**	-0.264***	-0.266***	-0.091***
	(0.092)	(0.101)	(0.075)	(0.091)	(0.074)	(0.088)	(0.029)
Acquired 2019–22	0.100	0.291*	-0.116*	-0.213*	-0.238**	-0.224**	-0.126***
	(0.114)	(0.153)	(0.064)	(0.116)	(0.100)	(0.101)	(0.028)

Notes: *** p<0.01; ** p<0.05; * p<0.10. Coefficients on acquisition-cohort indicators (never-treated omitted) from Equation 5 with the target indicator replaced by cohort indicators (cohort = year of first acquisition); the ever-acquirer indicator and all controls as in Table M.I are retained; standard errors clustered by hospital system in parentheses.

M.4 Independent hospitals collect unpaid bills differently than system hospitals

Why would acquired hospitals pursue court-based collection less? One possibility is simply that system hospitals use these measures less than independent hospitals, and what an acquisition does is convert an independent hospital into a system hospital. Figure M.II makes the case directly. Among never-treated hospitals that disclose a given practice, independents are systematically more likely to report using every collection measure: 86 percent garnish wages, 85 percent place liens, and 85 percent sue patients, versus 55, 60, and 62 percent of never-treated system members, with gaps of similar size for credit reporting (84 vs. 61 percent), debt sales (60 vs. 32 percent), and care denial (57 vs. 27 percent). Ever-targets and ever-acquirers cluster with the never-treated system

hospitals on the court-enforcement items, so the merger-role differences in Table M.I are, to a first approximation, the system–independent difference. On reporting and debt sales, targets sit at or slightly above the system hospitals. Independent hospitals are disproportionately the ones using aggressive collections practices.

Figure M.II. COLLECTIONS PRACTICES BY ORGANIZATIONAL STATUS

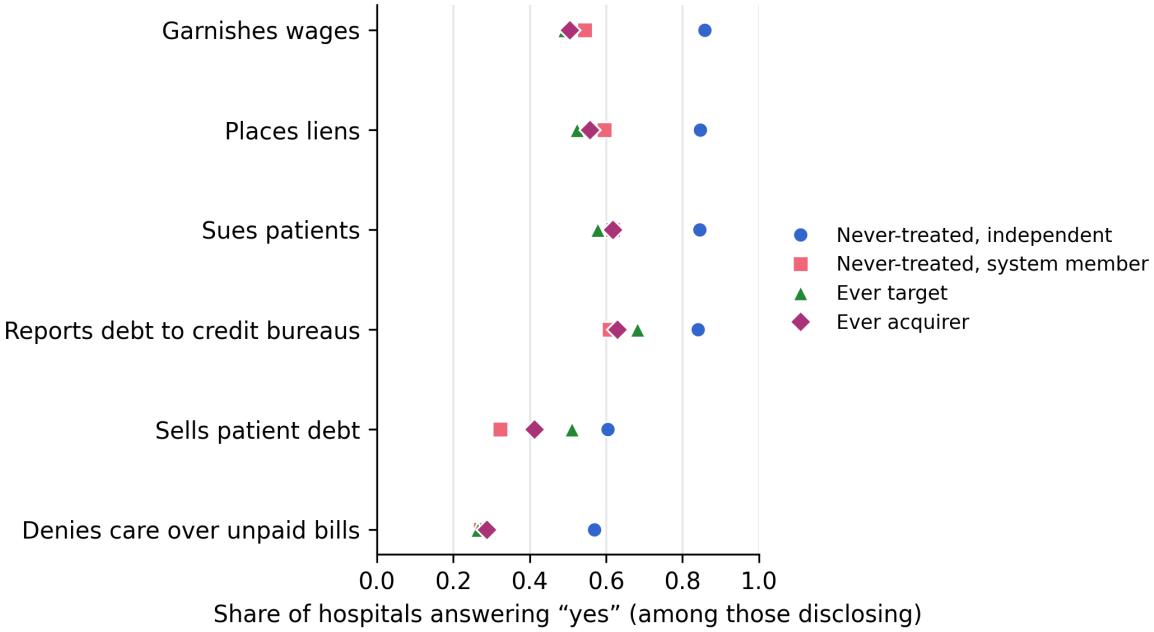

Notes: Share of hospitals answering "yes" to each collection-practice item in the Lown data, among hospitals disclosing that item, by organizational group: never-treated independent hospitals ($n = 684$), never-treated system members ($n = 392$), ever-targets ($n = 960$), and ever-acquirers ($n = 703$; the target and acquirer groups overlap for 340 hospitals). Unadjusted means.

Figure M.II may appear to sit in tension with Table M.I: ever-acquirers show much lower prevalence than never-treated independents on every practice in the figure, yet their coefficients in the table are near-zeroes. However, the apparent contradiction can be rectified. Figure M.II compares acquirers to never-treated *independent* hospitals, while Equation 5 compares them to the full never-treated pool, roughly two thirds of which is independent, and does so conditional on size, ownership, and geographic market, which amounts to a comparison of acquiring system hospitals to non-acquiring system hospitals. Against the blended never-treated benchmark, the unconditional acquirer coefficient for wage garnishment is -0.10 rather than the 35 percentage points visible in the figure; adding log beds cuts it to -0.05, and the ownership and commuting-zone

fixed effects reduce it to a statistical zero. Acquirers are large (a median of 202 beds, against 45 among never-treated independents), and size predicts less stringent collection even among never-treated independent hospitals themselves: within that group alone, the log-beds gradient for wage garnishment is -0.11 ($p < 0.01$).

The target differences in Table M.I, by contrast, are not an observables story. They are estimated conditional on the same controls. In other words, an acquiring system's hospital collects like an observably similar never-treated hospital, whereas a target collects less aggressively than observably similar never-treated hospitals. One further reason for caution in reading the figure is that independents disclose these items at roughly half the rate of system hospitals (for example, 32 versus 60 percent for wage garnishment), so the plotted shares condition on disclosure rates that differ sharply across groups.

M.5 Collection policy is set at the system level

Table M.III locates where collection policy is set. Panel B shows that these practices are overwhelmingly system-level attributes: among hospitals belonging to multi-hospital systems, between 89 and 97 percent of the variance in each item lies between rather than within systems (ICCs of 0.77–0.97), and in most systems every member hospital reports identical values. Panel A compares ever-targets to the acquirer-side hospitals of their own systems, with system fixed effects. On the court-enforcement items and care denial, the differences are tight zeros. Once acquired, targets are indistinguishable from their acquirers, which is what the variance decomposition implies: joining a system means adopting its regime.

The two exceptions are telling. Within their own systems, targets remain marginally *more* likely than their acquirers to report debt to credit bureaus ($+2.7$ percentage points, $p = 0.06$) and to sell patient debt ($+1.2$ points, $p = 0.06$ (Figure M.II). The natural reading is that targets arrive with the independent sector's high reporting and debt-sale propensities and converge toward, but not fully to, their acquirers' regimes. Implicitly, they became *less* likely to report and sell debt after acquisition, even though the single cross-section cannot document the transition directly. That

reading reinforces the conclusion from Table M.I: nothing in these data suggests acquisition raises a hospital's propensity to place unpaid bills on credit reports. The residual within-system gap runs in the direction of a legacy of aggressive collection being shed post-acquisition, which means it doesn't explain our baseline findings.

Table M.III. COLLECTION POLICY AS A SYSTEM ATTRIBUTE

Panel A: Targets vs. own-system acquirers (system fixed effects)				
	Coefficient	SE	*N*	Systems
Reports debt to credit bureaus	0.027*	(0.015)	702	94
Sells patient debt	0.012*	(0.007)	514	69
Denies care over unpaid bills	0.005	(0.010)	488	53
Sues patients	0.011	(0.017)	623	84
Places liens	0.014	(0.011)	686	89
Garnishes wages	0.011	(0.019)	573	83
Share of items undisclosed	0.007	(0.009)	998	130
Panel B: Variance decomposition across multi-hospital systems				
	R^2 between	ICC	*N*	Systems
Reports debt to credit bureaus	0.937	0.927	1,116	168
Sells patient debt	0.970	0.966	801	114
Denies care over unpaid bills	0.948	0.943	765	96
Sues patients	0.916	0.903	974	147
Places liens	0.945	0.936	1,053	152
Garnishes wages	0.909	0.894	920	142
Share of items undisclosed	0.804	0.772	1,531	225

Notes: *** p<0.01; ** p<0.05; * p<0.10. Panel A: coefficient on target role among matched ever-targets and ever-acquirer-side hospitals, restricted to systems containing at least one of each; regressions include log beds and system fixed effects, with heteroskedasticity-robust standard errors (the fixed effect absorbs the clustering dimension). Hospitals that were both target and acquirer are assigned their first role. Panel B: share of outcome variance between systems (R^2 from a regression on system indicators) and the ANOVA intraclass correlation, among matched hospitals in systems with at least two matched members.

M.6 Interpretation

These results discipline the interpretation of the paper's main findings in two ways.

First, they rule out a mechanical reporting channel. If mergers raised medical debt in collections because merged hospitals became more likely to report unpaid bills to credit bureaus or to sell them to debt buyers, we would expect positive target coefficients on those two items. The reporting coefficient is a tight zero, the debt-sale coefficient is, if anything, positive but insignificant, and

within their own systems the residual target–acquirer gap on these items points to targets shedding an independent-sector legacy rather than adopting new reporting practices. Target hospitals are not more likely to do the things that would be thought to generate the results in the rest of the paper. The medical debt increases we document must instead originate upstream, in the price and prevalence of patient obligations: mergers raise commercial prices, expand the set of encounters with out-of-pocket exposure, and shift the payer mix toward commercial insurance (Section V). The same logic applies to the credit card results. Merged hospitals are no more likely to deny care over unpaid bills, so the intensive-margin credit card borrowing we document in Section IV.3.1 is not explained by target hospitals adopting that practice upon acquisition.

Second, the margin on which targets do differ, court-based enforcement, reflects the organizational structure of the hospital industry rather than anything specific to the merger event. Independent hospitals dominate the aggressive tail of the collection distribution (Figure M.II); collection policy is set at the system level and adopted essentially completely by acquired hospitals (Table M.III); and mergers move hospitals into systems. Target hospitals are less likely to take the more aggressive collection measures because system hospitals are less likely to take them than independent hospitals. In this respect, cross-sectional variation between hospitals with respect to collections practices mirrors variation between employers with respect to labor standards violations (Bernhardt et al., 2009; Weil, 2014): smaller, independent hospitals, potentially operating on thinner margins, are more likely to cut corners, in this case utilizing aggressive collections practices toward their patients.

Two measurement caveats deserve emphasis. The Lown Institute data record *stated* policies scraped from hospital websites, largely at the system level. The tight within-system convergence in Table M.III is consistent both with genuine policy adoption and with member hospitals posting a common system policy document whose enforcement varies locally, and the sharp independent–system disclosure gap partly reflects that systems maintain professionalized, centralized policy websites that independents lack. Court records on collection lawsuits would distinguish stated policy from realized practice, and we regard that as the natural next step for this line of inquiry.

In addition, the practices are observed in 2024–2026, after our credit-panel study period ends; the collection regimes that prevailed during 2011–2019 are unobserved. For these reasons, and because selection into merging is uncontrolled, we present these findings as a descriptive complement to, rather than a test of, the causal analysis in the body of the paper.