Who Pays for Rising Health Care Prices? Evidence from Hospital Mergers*

Zarek Brot, University of Chicago and NBER
Zack Cooper, Yale University and NBER
Stuart V. Craig, University of Wisconsin-Madison
Lev Klarnet, Harvard University
Ithai Lurie, U.S. Department of Treasury
Corbin Miller, Internal Revenue Service

January 2026

Abstract

We analyze the economic consequences of rising health care prices in the United States. By increasing the cost of employer-sponsored health insurance, rising prices serve as a de facto payroll tax on labor. Using exposure to hospital mergers as an instrument, we show that health care price increases lead to greater-than-dollar-for-dollar reductions in payroll and decreases in employment at employers outside the health sector. At the county level, rising health care prices reduce aggregate labor income, increase flows into unemployment, and lower federal income tax revenue. These disemployment effects are concentrated among lower- and middle-income workers.

*We thank Joseph Altonji, Steven Berry, Zachary Bleemer, Anne Case, Angus Deaton, Amy Finkelstein, Joshua Gottlieb, Jason Hockenberry, Anders Humlum, Dmitri Koustas, Neale Mahoney, Alex Mas, Costas Meghir, Fiona Scott Morton, Chima Ndumele, Seth Zimmerman, and many seminar participants for extremely valuable feedback. We benefited enormously from excellent research assistance provided by Felix Aidala, Krista Duncan, James Han, Mirko De Maria, Kelly Qiu, Shambhavi Tiwari, and Mai-Anh Tran. This project received financial support from Arnold Ventures and the National Institute on Aging (Grant P01-AG019783). We acknowledge the assistance of the Health Care Cost Institute (HCCI) and its data contributors, Aetna, Humana, and UnitedHealthcare, in providing the claims data analyzed in this study. HCCI had a right to review this research to guarantee we adhered to reporting requirements for the data related to patient confidentiality and the ban on identifying individual providers. Neither HCCI nor the data contributors could limit publication for reasons other than the violation of confidentiality requirements around patients and providers, nor could they require edits to the manuscript as a condition of publication. The opinions expressed in this article and any errors are those of the authors alone. This research was conducted while some of the authors were employees at the U.S. Department of the Treasury and the Internal Revenue Service (IRS). The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors and do not necessarily reflect the views or the official positions of the U.S. Department of the Treasury or the IRS. Any taxpayer data used in this research was kept in a secured Treasury or IRS data repository, and all results have been reviewed to ensure that no confidential information has been disclosed.

1 Introduction

The prices of health care goods and services are a central driver of the variation in and growth of health spending in the United States (US) (Cooper et al., 2019a; Health Care Cost Institute, 2020). Since 2000, prices for health care services, medical devices, and pharmaceuticals have grown markedly faster than prices for goods and services outside the health sector (Bureau of Labor Statistics, 2022). Over this period, for example, prices in the hospital sector — a $1.3 trillion industry — rose faster than prices in virtually any other sector of the economy (Bureau of Labor Statistics, 2022). While price growth need not be a problem if it reflects improvements in quality, a wide-ranging literature has illustrated that much of the growth in prices in the US health care sector over the last two decades has arisen from unproductive rent-seeking activities.[1]

In most markets, rising prices are paid directly by consumers. In the US health care sector, however, the majority of working-age adults purchase health care via employer-sponsored health insurance (ESI) (Kaiser Family Foundation, 2019). Employers either provide health insurance via third-party insurers or self-insure and contract with managed care organizations to administer their health benefits. As the cost of ESI rises, retaining workers becomes more expensive for employers. As a result, rising health care prices function as a de facto payroll tax on employers outside the health care sector. The incidence of this "health care tax" could be borne by employers or workers themselves.

In this paper, we study the downstream economic consequences of rising health care prices—this "health care tax"—by tracing how increases in health care prices affect ESI premiums, wages, and employment. Simply correlating changes in health care prices with changes in insurance premiums and labor market outcomes would generate severely biased estimates because rising local incomes may increase demand for health care services and thereby raise local prices. To address this concern, we use local hospital mergers as a shock to health care prices and then measure the subsequent effects of the merger-driven price increases.

Our analysis draws on four key data sources. We measure health care prices and spending using HCCI insurance claims, which cover 28% of ESI enrollees nationwide. We measure premiums for fully insured employers using DOL Form 5500 filings, and rely on a hospital merger database constructed by Brot et al. (2024a) to identify price shocks. We measure employment and earnings for workers using individual tax records from the Internal Revenue Service (IRS). In the majority of our analysis, we focus on analyzing the impact of rising hospital prices on workers outside the health care sector (e.g., individuals whose employment is unaffected by changing monopsony power in the hospital sector).

[1] See, e.g., Cooper et al. (2019a), Brand et al. (2023), and Brot et al. (2024b) on hospital mergers; Capps et al. (2018) and Lin et al. (2021) on hospital-physician vertical integration; Cooper et al. (2020) on surprise billing; Dafny (2005) on hospital upcoding; and Dafny et al. (2024) on drug copayment coupons.

Our empirical approach exploits the fact that some hospital mergers generate substantial increases in market concentration–and thus prices–while others do not. To estimate merger effects, we implement a difference-in-differences design comparing prices at merging hospitals with those at non-merging hospitals before and after each merger. The post-merger price increases we estimate vary substantially across transactions, and much of the variation in post-merger price increases can be explained by the changes in market structure that specific mergers generate. Transactions predicted, ex ante, to substantially lessen competition raise prices by, on average, 5.9%, whereas mergers that do not appear to meaningfully reduce competition result in minimal price changes. We find no evidence that the variation in the merger-induced price changes is correlated with trends in local economic activity prior to the merger. As a robustness check, we also estimate the change in predicted markups generated by each transaction, following Capps et al. (2003).

It would be intuitively appealing to compare employers exposed to mergers with those unexposed. However, during our sample period, hospital mergers were so widespread that virtually every US employer was proximate to at least one merging hospital, and many were exposed to several mergers within short time spans. As a result, rather than using a difference-in-differences design, we instead exploit variation in the *intensity* of employers' exposure to merger-driven hospital price increases to estimate the downstream effects of rising health care price.

To do so, we exploit the fact that employers outside the health care sector face varying degrees of exposure to mergers based on their workers' ex ante demand for specific hospitals that eventually merged, as well as variation in the extent to which each merger raised prices at merging hospitals. An employer, for example, is heavily exposed if its employees receive a large share of their hospital care from a merging hospital that experienced a substantial merger-induced price increase during our sample period. To capture this variation, we construct an instrument that measures each employers' exposure to merger-driven changes in the average price of care over time, holding fixed other prices and quantities of health care consumption. We use this measure as an instrument for an employer's annual health care spending. This instrument is valid as long as an employer's exposure to mergers is independent of their counterfactual outcomes had no mergers occurred and that mergers only impact non-health care labor market outcomes via increasing health care costs. Going forward, we develop multiple approaches to test this exogeneity assumption and show that employers' historical trends in payroll and the count of workers they employ are uncorrelated with the occurrence of hospital mergers and the scale of the price increases that consummated mergers produce. We also show that hospital mergers do not induce changes in labor market concentration outside the health sector, which itself could impact labor market outcomes.

We begin by showing that merger-driven increases in hospital prices raise health care spending on an approximate dollar-for-dollar basis. Using data on a restricted subset of employers for whom we can measure insurance premiums, we then show that health spending is passed through almost one-for-one into employer ESI premiums. Finally, despite facing higher premiums, we show that employers do not appear to respond to health spending increases by shifting workers towards high-deductible health plans.

We then show that rising health care prices reduce employment. When exposed to a 1% increase in health care prices, employers outside the health care industry reduce their payroll by 0.37% (our 95% confidence interval for this estimate spans 0.11% to 0.63%). We find roughly equal-sized effects on employers' total count of workers, implying that rising health care prices were not passed through as simple wage reductions but instead resulted in job separations. Event study estimates show that these employer-level effects were realized almost immediately after hospital prices increased and exhibit little dynamic impact over time.

Together, our employment and payroll results imply that for every $1 in hospital revenue generated by a merger-driven price increase, payroll in non-health care employers falls by $1.69 (with a 95% confidence interval that spans $0.49 to $2.88). This reflects the fact that employers respond to rising health spending via adjustments on the extensive margin. That is, higher health spending and insurance premiums lead firms to reduce employment, lowering total labor income by more than the direct cost increase. These extensive margin responses generate deadweight loss from to the "health care tax" in much the same way that other employer taxes (e.g., payroll taxes) generate deadweight loss: by reducing employment. Our results imply an approximate 1.8% reduction in employment in response to a 1 percentage point increase in labor costs (with a 95% confidence interval of 0.6% to 3.0%), an elasticity that lies within the 0.7% to 2.4% range of estimates obtained from prior studies estimating the effect of US payroll tax increases on employment(Anderson and Meyer, 1997; Johnston, 2021; Guo, 2024).

While we estimate large disemployment effects *at specific employers*, the job losses triggered by rising health care prices may have simply resulted in reallocations of workers to other employers or gains in employment in the health care industry. Therefore, we estimate the broader regional effects of price increases—general equilibrium effects—by aggregating our employer-specific instrument to the county level. At the county level, we find that a 1% increase in *county-level* health care prices reduced *county-level* labor income per capita by 0.28% (with a 95% confidence interval of -0.02% to 0.58%) and increased flows into unemployment by 0.09 percentage points (1%, with a 95% confidence interval of 0.01 to 0.17). We find no economically or statistically significant effects on self-employment or on migration across counties. Additionally, while we estimate small (but statistically insignificant) increases in income per capita of health care workers, we still find that rising prices lead to aggregate reductions in income per capita and increases in unemployment across all workers. As with our employer-level results, these effects occurred immediately after hospital prices increased. Ultimately, these labor market changes had substantial fiscal consequences for federal and state governments: a 1% increase in health care prices reduced income tax withholdings by approximately 0.37% (95% confidence interval of -0.01% to 0.75%), while increasing unemployment insurance (UI) payments by approximately 2.42% (95% confidence interval of -0.51% to 5.35%).

Because premium increases are largely uniform across workers within a firm, theory suggests they should function as a regressive "head tax," generating larger disemployment effects for lower-wage workers (Finkelstein et al., 2023). Consistent with this prediction, we find that unemployment effects are approximately zero for workers at the top of the income distribution (i.e., those previously earning more than $100,000 annually). In addition, we also find no employment effects for workers at the bottom of the income distribution (i.e., those previously earning less than $20,000 per year), which aligns with the fact that these workers rarely receive health insurance from their employers and therefore do not become more costly to employ when health care prices rise (Lurie and Miller, 2023). Instead, the job losses we observe are concentrated among low- and middle-income workers who previously earned between $20,000 and $100,000 per year.

Our work makes three distinct contributions to the literature. First, we contribute to a broad literature that studies the incidence of employer-sponsored health insurance and work documenting how changes in the ESI market impact labor markets (Cutler and Madrian, 1998; Baicker and Chandra, 2006; Kolstad and Kowalski, 2016). Prior research has primarily examined the effects of adding new fringe benefits (Summers, 1989; Gruber and Krueger, 1991). Since new benefits are valued by workers, their costs can potentially be passed through into wages with little or no reduction in employment. By contrast, when the cost of existing fringe benefits rises, as is the case in our setting, workers are unlikely to value the additional cost borne by employers and therefore will not readily accept corresponding wage reductions. This implies—as we demonstrate empirically—that rising health care prices should lead to job losses rather than full pass-through into lower wages.[2]

Our findings echo contemporaneous studies that exploit alternative sources of variation in employers' health insurance costs—including binding insurer profit-cap regulations (Gao et al., 2025) and insurer mergers (Min, 2025)—to study downstream labor market effects. Indeed, the estimates found in Gao et al. (2025) and Min (2025), who both use Census micro data to measure employment, are quantitatively very similar to ours. Most closely related to our study is Arnold and Whaley (2024), who also use hospital mergers as a shock to local health care prices. They focus on measuring the effects of rising prices on hourly wages (using data from the American Community Survey) and find that hospital mergers lead to a large reduction in non-health care wages.[3]

[2] Indeed, while prior research on new benefits has found little disemployment response (Gruber and Krueger, 1991; Gruber, 1994; Kolstad and Kowalski, 2016), other work on the rising cost of existing benefits has typically found nontrivial disemployment responses (Cutler and Madrian, 1998; Baicker and Chandra, 2006; Gao et al., 2025; Min, 2025), though with exceptions (Arnold and Whaley, 2024).

[3] In studying the effect of hospital mergers on labor market outcomes, this paper is related to Prager and Schmitt (2021). While we focus on the channel through which merger-induced price increases are passed through to *non*-health care workers, they focus on the channel through which mergers result in the accumulation of monopsony power in

Second, we offer empirical evidence consistent with Saez and Zucman (2019), Case and Deaton (2020), and Finkelstein et al. (2023), who emphasize that, because ESI premiums are largely uniform within firms and thus constitute a larger share of compensation for low-income workers, rising health care costs exacerbate inequality and disproportionately reduce employment among lower-wage workers. We complement their conceptual analysis by demonstrating this result within a large, national sample of employers and employees, which allows us to examine heterogeneous effects even within narrower subgroups. In doing so, we provide new quasi-experimental evidence of the *uneven* distributional consequences of rising health costs. We show that the job losses that occur after increases in ESI impact workers earning between $20,000 and $100,000 a year, with no impact on higher-income workers. Our findings support the idea that increases in ESI premiums function as a regressive "head tax" on employment. Our findings complement similar contemporaneous findings by Gao et al. (2025) and Min (2025).

Finally, we provide new evidence on the macroeconomic consequences of rising domestic health care spending. Economists have previously focused on the federal tax burden of financing public health insurance (Baicker and Skinner, 2011). Our results show that the rising cost of financing private health insurance also generates important macroeconomic consequences by raising employers' labor costs. A macroeconomic literature emphasizes how sectoral shocks can propagate throughout the broader when the affected sector is central to the production in many other sectors (Baumol, 1967; Baqaee and Farhi, 2019). Because virtually all major U.S. firms offer health insurance benefits, rising health care costs can slow local economic activity. Under strong assumptions, our estimates suggest that growth in hospital prices from 2007 to 2014 reduced aggregate labor income in 2014 by approximately 5% (with a 95% confidence interval of -0.3% to 10.3%) in real terms.

Going forward, this paper is structured as follows: in Section 2, we provide background on the market for ESI in the US and an overview of the changes that have occurred in hospital markets in the last two decades. In Section 3, we describe how increases in health care prices can impact labor market outcomes. We describe the data used in this analysis in Section 4 and provide a discussion of our analytic strategy in Section 5. We present our employer-level results in Section 6 and county-level results in Section 7. We conclude in Section 8.

hospital labor markets. Like them we also find economically meaningful—though statistically insignificant—reductions in payroll among health care employers.

2 Background

2.1 Employer-Sponsored Health Insurance

Employer-sponsored health insurance (ESI) in the United States emerged during the 1940s in response to federal wage controls implemented under President Franklin Roosevelt. Unable to raise wages, firms began offering workers non-monetary compensation, including health insurance. In 1943, the IRS ruled that employer payments for health insurance were exempt from income taxes, effectively lowering the after-tax price of coverage and accelerating the spread of ESI (Starr, 1982). Today, insurers charge an annual premium that is typically experience-rated at the employer level, meaning that employers are charged premiums proportional to the expected costs of insuring their specific employees (Craig, 2022). These premiums are financed through a combination of employer and employee contributions, both of which are generally paid using pre-tax dollars.

Approximately 91% of workers are employed by firms that offer health benefits, and 59% of firms offer health benefits to at least some portion of workers (Claxon et al., 2021). As a result, 54.3% of the US population and 63% of the adult population under age 65 have ESI coverage (Keisler-Starkey and Bunch, 2022). Of those receiving ESI, 64% are covered by administrative services only (ASO) policies, under which an insurer administers the health benefits but the employer bears the financial risk. The remaining 36% are enrolled in fully insured health insurance plans offered by their employers, in which the insurer bears financial risk (Claxon et al., 2021). ASO plans are more prevalent among larger firms.

ESI coverage is notably higher among higher-income individuals. While 10% of individuals in the 0^{th} through 25^{th} percentiles of income have ESI coverage, 84% of individuals in the 95^{th} through 99^{th} percentiles have employer-based coverage (Lurie and Miller, 2023). Conditional on being insured through ESI, higher-income individuals are enrolled in substantially more expensive policies (Lurie and Miller, 2023).

Private insurers form networks of providers accessible to their enrollees and design benefits packages, including the degree of cost-sharing patients face. Insurers negotiate over prices with hospitals and physicians. In exchange for favorable prices, insurers include those providers in their coverage networks, giving enrollees access to those providers at a discount (Handel and Ho, 2021). Hospital and physician prices vary substantially across geographic regions and within regions across providers (Cooper et al., 2019a).

2.2 Consolidation in the US Hospital Sector

From 2000 to 2020, there were over 1,000 hospital mergers among the nation's approximately 5,000 hospitals (Brot et al., 2024b). During that period, there were only 13 enforcement actions taken against hospital mergers by the Federal Trade Commission (FTC), the federal agency tasked with preserving competition among hospitals in the US (Brot et al., 2024b). A growing literature demonstrates that mergers, particularly among hospitals that are close substitutes, can lead to increases in prices via a lessening of competition (Capps et al., 2003; Dafny, 2009; Tenn, 2011; Haas-Wilson and Garmon, 2011; Gaynor et al., 2015; Gowrisankaran et al., 2015; Cooper et al., 2019a; Brand et al., 2023). Furthermore, there is little evidence that mergers raise hospitals' quality (Beaulieu et al., 2020). Because of merger activity, market concentration in the hospital industry has been rising steadily since 2000 (Fulton, 2017). Given of the scale of the hospital industry and the scale of merger activity in the sector, hospital mergers have become a topic of interest and importance for policymakers and elected officials (Biden, 2021).

The joint 2010 Department of Justice (DOJ)/FTC Horizontal Merger Guidelines specify that mergers that result in increases in the Herfindahl-Hirschman Index (HHI) of at least 200 points and lead to a post-merger HHI of over 2,500 should be "presumed to be likely to enhance market power" (U.S. Department of Justice and the Federal Trade Commission, 2010). Brot et al. (2024b) show that, from 2010 to 2015, approximately 20% of hospital mergers could have been predicted *ex ante* to meaningfully lessen competition based on the screening thresholds established in the Horizontal Merger Guidelines and find that these mergers resulted in price increases of over 5%.[4]

3 Theoretical Framework

We present a simple model that illustrates the mechanism through which rising health care prices, by increasing employer-sponsored insurance (ESI) premiums, can affect local labor market outcomes.

Hospital Markets: Consider a hospital, h, providing a single service. It faces a residual (firm-specific) demand curve of $D^h(p)$ for its services, given its own price p. Consider a rent-seeking activity, such as a merger, that raises the hospital's price by Δp without shifting demand (i.e., without shifting patients' preferences for getting care at that hospital).

In Panel A of Figure 1, we graph this event. As prices rise, quantities decline. The price increase generates a deadweight loss given by the red shaded triangle. As in Harberger (1954), the deadweight loss from the merger is proportional to the reduction in quantity. In Figure 1, as is typically the case, hospital-specific demand is extremely inelastic, because patients do not tend to substitute to alternative hospitals when a hospital's relative price increases (Gowrisankaran et al., 2015; Brot-Goldberg et al., 2017; Lieber, 2017). Therefore, the deadweight loss generated in the hospital market is small. However, this deadweight loss is not the only effect. There is also a large

[4]In 2023, the DOJ and FTC revised their merger guidelines and defined problematic transactions as those that increased HHI by greater than 100 points and led to a post-merger HHI of greater than 1,800 points (U.S. Department of Justice and Federal Trade Commission, 2023). Brot et al. (2024b) also show that mergers exceeding that threshold markedly increase hospital prices.

transfer of surplus from health care payers to hospitals, which is represented by the blue rectangle.

Insurance Markets: In Kaldor-Hicks terms, the transfer to hospitals has no effect on the surplus of the hospital market. However, since patients often have health insurance, most of the transfer is paid by insurers rather than patients themselves. This raises the costs of providing insurance.

Consider a market for ESI where insurers sell contracts with actuarial value C.[5] We assume (without loss of generality) that insurance markets are perfectly competitive, and therefore the premium charged, ϕ, is determined by the cost of providing insurance coverage of generosity C: $\phi = P \times C$, where P is the general price of care. Insurance is purchased by employers on behalf of their employees, with demand $D^{\text{Ins}}(\phi)$. In Panel B of Figure 1, we plot a supply and demand model of such an insurance market.

In Panel B of Figure 1, we model how the insurance market changes as insurers pay higher prices to rent-seeking hospitals. As the price of care rises, it rotates the insurance supply curve counter-clockwise. A share of the price increase is passed through into higher premiums. Employers may respond by lowering the quality of the insurance they procure, offsetting some of the premium increase. Price increases specifically *rotate* the supply curve because the cost increase is proportional to the baseline pre-shock equilibrium coverage level. That is, consider a \$100 price increase affecting two insurers, one of which covers 30% of expenses and the other of which covers 80%. The former will increase their premiums by \$30, while the latter will increase their premiums by \$80. As a result, employers that offer greater benefits at baseline will generally face a greater premium increase.

Labor Markets: Increasing hospital prices raise costs for insurers, who pass through those costs to employers via increases in ESI premiums. To illustrate how employers respond, we set up a simple labor market model in the spirit of Gruber and Krueger (1991). We consider a single employer facing a labor market for homogeneously-skilled workers. The employer hires L workers according to its total cost, which includes both the wage w and insurance premiums, ϕ, with its labor demand thus given by $D^L(w+\phi)$. Workers only value the wage they receive, not what their employer has paid for insurance on their behalf; thus, their labor supply is given by $S^L(w)$.

In Panel C of Figure 1, we plot this market. Increases in the price of health care raise ESI premiums by $\Delta\phi$. For now, we assume that employers are perfectly inelastic with respect to their insurance purchasing, so C is fixed. The premium increase shifts labor demand down by exactly $\Delta\phi$. However, since the insurance value has not changed, labor supply remains the same. The result is that both wages w and employment L decline. The premium increase effectively serves as a tax on hiring — what we call a "health care tax" — and therefore generates deadweight loss given by a

[5] $C = 0$ represents the purchase of no insurance coverage.

Harberger triangle (highlighted in red in the figure) the same way other taxes on employers, such as payroll taxes, would.

In the Figure, we present the deadweight loss as a triangle. However, if there are other distortions that exist in the labor market, the deadweight loss from a marginal change will be proportional to the net distortion from those other pre-existing factors. For example, the fact that labor income is already taxed means that the initial employment level will be inefficiently low and the surplus of the marginal worker will be relatively high, and therefore effectively taxing labor through higher health insurance costs will sever even more valuable employer-employee relationships.

Why do both wages and employment go down, in contrast to the predictions of Summers (1989), who posits that the costs of mandated employer fringe benefits are fully passed through into wages? One key factor is that, while workers value the provision of ESI, they (likely) do *not* value increases in its price that are not associated with changes in the quality of benefits they receive. As a result, while premiums have increased, the value of the employment relationship to them has not changed, and so many workers will not accept a reduced wage. Therefore the employer's cost of retaining workers will increase, and they will lay off workers. The proportional changes in wages, employment and total labor earnings ($I = w \times L$) are given by:

$$
\begin{aligned}
\frac{\Delta w}{w} &= -\frac{\eta_D}{\eta_D + \eta_S}\frac{1}{w}\Delta\phi \\
\frac{\Delta L}{L} &= -\frac{\eta_D \eta_S}{\eta_D + \eta_S}\frac{1}{w}\Delta\phi \\
\frac{\Delta I}{I} &= -\frac{\eta^D(1+\eta^S)}{\eta^S - \eta^D}\frac{1}{w}\Delta\phi
\end{aligned}
$$

where η_D and η_S are the absolute values of the elasticities of demand and supply, respectively, meaning that each of these effects is negative.

The effect of rising health care prices on wages is similar to a simple tax: wage reductions will be greater when labor demand is more wage-elastic than labor supply. On the other hand, the changes in employment and total labor income are greater when either side is more elastic. For example, if workers are more substitutable with capital (and so labor demand is very elastic), employers will replace workers as they become costlier.

A key feature of the response is that it is proportionally larger when the pre-shock equilibrium wage level is lower. This occurs because, in contrast to income and payroll taxes, the change in premiums $\Delta\phi$ is not (directly) a function of wages. It is instead, in the terminology of Finkelstein et al. (2023), a "head tax" that is paid per worker rather than per dollar. This means that, in lower-wage labor markets, the change will be a larger proportional change in the cost of retaining a worker. All else equal, health care price increases will have much larger effects in lower-wage markets.

4 Data and Measurement

4.1 Employer Panel and Labor Outcomes

We construct a panel of employers using Employer Identification Numbers (EINs) as reported on workers' W-2 forms. We begin with all employers issuing W-2s in 2009 and restrict the sample to EINs with at least 50 employees located in the contiguous United States. We also require 95% of employers' workers to reside in a county where we have at least one HCCI beneficiary. From this group of employers, we enforce a balanced panel by keeping only those employers who appeared on at least one W-2 form in every year from 2008 to 2017. Our final sample includes 140,301 unique employers.

We identify employers in the health care industry by whether they have a North American Industry Classification System (NAICS) code starting in "62," as reported to the IRS. Approximately 7% of our employers have such a NAICS code starting with "62." While we only have HCCI data from 2008 to 2017, for our event studies, we add an additional three years of IRS data and look outcomes for our treated employers and counties from 2005 to 2017. This allows us to have a five-year pre-period. Because we require a balanced panel, our extended sample has 125,362 unique employers.

For each employer, we measure total payroll and the annual count of workers using information from W-2 forms. In each calendar year, we collect all W-2s that list a given employer and sum wages subject to Medicare tax (Box 5) to construct total payroll.[6] We use the number of unique W-2s at each EIN to measure employers count of workers employed.

Because workers may not be employed by a single firm for the full year, we adjust worker counts to reflect partial-year employment. When an individual files multiple W-2s in a year, we allocate that individual across employers in proportion to the wages earned at each employer. For example, if a worker earns $10,000 from employer A and $30,000 from employer B in the same year, we assign one-quarter of a worker to employer A and three-quarters to employer B.

4.2 Health Care Prices, Health Spending, and Hospital Quality

We use insurance claims data from 2008 to 2017 provided by HCCI to measure health care prices and utilization for individuals enrolled in employer-sponsored coverage. The HCCI data are composed of health insurance claims from Aetna, Humana, and UnitedHealth. The data capture approximately 28% of individuals in the US with employer-sponsored coverage and cover more than $125 billion in health care spending annually. Crucially, the data include the transaction prices that hospitals were actually paid by insurers, not merely their "chargemaster" (list) prices. We focus our analysis on

[6] Wages subject to Medicare tax include cash wages net of employee contributions to tax-excluded benefits, such as retirement accounts and—importantly—employee contributions toward employer-sponsored health insurance premiums.

individuals under age 65 for whom an HCCI payer is their primary insurer (e.g., the individual does not receive primary insurance coverage from their spouse's employer). This population includes both the policyholders of employer-sponsored insurance plans and their dependents (spouses and children).

These data allow us to measure prices and quantities for health care services obtained at hospitals. We construct visit-level data containing prices for both inpatient and outpatient hospital care, where an observation is a hospital admission or an outpatient visit, respectively. We also measure average spending per beneficiary annually across all medical claims (e.g., physician claims, inpatient claims, and outpatient claims). We exclude pharmacy-dispensed pharmaceutical spending from our analysis.[7] For individuals who are not enrolled in coverage for a full calendar year, we take their average monthly spending and multiply it by 12 to construct an annualized measure.

We cannot directly link employers in the HCCI data with those in our IRS employer panel. We instead assume that an employer's expected average health care spending is the weighted average of county-level HCCI spending per beneficiary, where the weights are ω_{ic0}, the share of a employer's employees that live in each county c in our base period of 2009, as measured by their listed ZIP codes on their filed W-2 forms.[8]

In practice, we measure an employer's expected health care spending as the product of where an employer's employees lived in 2009 and HCCI prices and utilization for beneficiaries in those locations across all health care providers, including doctors, hospitals, outpatient clinics, and physical therapists. Our constructed measure for employer i in year t is:

$$S_{it} = \sum_{c,k} \underbrace{\omega_{ic0}}_{\substack{\text{Employee share} \\ \text{in county } c \\ \text{at baseline}}} \times \underbrace{S_{ckt}}_{\substack{\text{Average HCCI spend} \\ \text{in county } c\text{, year } t \\ \text{at provider } k}} \tag{1}$$

Finally, we use the 100% sample of Medicare Fee For Service Claims data to measure changes in hospital quality following hospital mergers. We focus on 30-day mortality for individuals age 65 and older that were admitted to the hospital with a non-deferrable condition (e.g., patients with a condition where admissions are as likely on weekends as on weekdays). As Doyle et al. (2015) note, focusing on this cohort of severely ill patients—patients who likely have little discretion over the timing or location of their care—helps address potential selection concerns. We used data from Traditional Medicare for this analysis because the HCCI data does not capture patient mortality.

[7]Prescription drug claims are often offset by large rebates negotiated between payers and drug manufacturers. These rebates are not included in the HCCI data, rendering any claims-based spending measurement inaccurate.

[8]Based on our discussions with industry participants, this aggregation mimics the way health insurers would price premiums for an employer with beneficiaries spanning multiple geographies.

4.3 Insurance Market Outcomes

We also construct measures of employers' insurance premiums. Data on insurance premiums are scarcely available and are often subject to reporting error or provided at high levels of geographic aggregation (Dafny et al., 2011). We measure employers' insurance premiums directly using data from Form 5500, which is a regulatory filing collected by the DOL in cooperation with the IRS. The data contain measures of total premiums, covered lives, and plan characteristics for employer-sponsored insurance groups covering at least 50 employees enrolled in fully-insured insurance plans. We use the data to construct a measure of average premiums per covered life at the employer level. We provide additional details on the construction of our premium series in Appendix C. Because Form 5500 identifies employers, we can link them directly to our panel in the IRS data.[9] As a result of the limited number of employers with 5500 data on premiums, our premiums analysis is carried out on a sub-sample of our main analytic sample.

We also use the IRS data to construct a proxy measure of the share of employers' employees enrolled in a high deductible health plan (HDHP), since this information is not available in the 5500 data. To do so, we proxy for employees who have a HDHP based on whether or not they report contributions to a Health Savings Account (HSA) in Box 12 of their W-2 filing. Employees can only legally be enrolled in an HSA if they are enrolled in an HDHP, so this serves as a lower bound on HDHP enrollment. For each employer, we measure the share of employees who report any individual or employer contributions to an HSA.

4.4 County-Level Outcomes

To examine the general equilibrium effects of rising health care prices, we also conduct a county-level analysis. Our primary county-level outcomes are per capita labor income; the share of individuals who become unemployed; the share who become self-employed; tax receipts and UI payments; and the share of individuals who move to another county. We focus on individuals ages 25–64, who are most likely to receive employer-sponsored insurance directly rather than through Medicare or as dependents on commercial plans.

To measure labor income, we combine information from individuals' W-2 forms and Schedule SE filings, capturing both wage earnings and self-employment income when present. We measure unemployment using two criteria: (i) whether an individual receives unemployment insurance (UI), identified by positive income reported in Box 1 of Form 1099-G; or (ii) whether the individual earns no positive income from either W-2 or Schedule SE forms. We apply these measures only to

[9] While the 5500 data provide a granular measure of insurance premiums at the employer level, these data are limited to small- to medium-sized employers that purchase fully-insured plans. The premiums for self-funded plans are not well-documented in Form 5500, and they do not reliably reflect the total cost of insurance provision in the way that they do for the fully-insured.

individuals who were employed in the prior year, thereby capturing *flows into unemployment*, rather than the stock.[10]

We classify individuals as self-employed if they file a Schedule SE form with any positive income in a given year. We identify movers as individuals whose county of residence changes between year $t-1$ and t. We measure tax payments as total income tax withheld, as reported on W-2 forms.[11] Finally, we measure UI payments as the total amount reported on Form 1099-G. We focus on labor market outcomes for individuals age 25 through 64 (i.e., working-age adults).

We measure per capita outcomes as the sum of the outcomes described above divided by a county-level population measure for individuals age 25 to 64. We measure the county population using population files from the IRS that contain information from all tax returns that provide individual-level information with location identifiers (e.g., W-2 forms, 1099s, etc.). We identify individuals' county of residence using multiple tax forms in a hierarchical order. In practice, we use the Form 1040 filed for the prior year (typically filed between January and April of the reference year). If an individual did not file a 1040 for the prior tax year, we use information from the W-2 form and the Form 1099-G in the reference year. A virtue of this approach is that it allows us to measure population by income group.

4.5 Summary Statistics

We present descriptive statistics for our employer and county samples in Panels A and B of Table 1.[12] Our primary employer sample contains 140,301 unique employers. Across 2008 to 2017, the average employer in our sample had 298 employees and average annual wages per worker of $41,340. The average employer spent (according to our constructed measure) $4,099 per employee on health care. Among our employers, 7% (9,471) were involved the health care industry, with the remaining 93% (130,830) involvehccd in other, non-health care industries.

Our primary county sample includes 1,709 of 3,182 US continental counties. These 1,709 counties make up 160 million individuals aged 25 to 64 annually, which is roughly 96% of the total US population within that age group. There were approximately 31 million HCCI beneficiaries in those counties annually in our sample period. The average income across 2008 to 2017 in counties in our analytic sample was $42,050, the share of the population unemployed was 8.9% (3.6% with UI and 5.4% with zero income),[13] the self-employed share was 11.7%, average federal income tax

[10]We do so since UI take-up explains much of the variation in our unemployment measure, and most states have time limits on UI receipt, meaning that previous recipients may be unable to claim it in the future. Therefore, our approach is likely to be a lower bound on the stock of unemployed individuals in a given year.

[11]Withholding does not always equal final tax liability—particularly given common over-withholding—but any resulting measurement error should be orthogonal to exposure to hospital-merger-induced price changes.

[12]For a summary of the datasets used in this analysis and key outcomes see Appendix Table A.1.

[13]Note that our definition of the share of the population unemployed is different from standard unemployment rates as defined by the Bureau of Labor Statistics and other statistical agencies because we define an individual as being

payments per capita were $7,009, and average UI payments were $482.

In Appendix Table A.2, we show how the composition of our sample of employers changes as we introduce additional sample restrictions. Relative to the universe of employers with at least one employee in 2009, our overall analytic sample is comprised of employers that are on average larger and have slightly higher average income per worker. As we illustrate in Column (3), employers outside the health sector are modestly larger than health care employers. Likewise, when we limit our analysis to employers where we have insurance premiums recorded via Form 5500, the average employer is substantially larger and has higher average income, which is not surprising given that the Form 5500 is a regulatory filing required of all employers that offer a benefit plan to at least 100 employees.[14] Finally, our extended 2005 to 2017 sample, which we use for our event studies, is very similar to the overall analytic sample. In Appendix Table A.3, we show how our analytic sample of counties differs from the universe of counties in the US. The population in our analytic sample has modestly higher incomes and a lower share of the population that is self employed.

5 Empirical Strategy

Our goal is to trace out the causal impact of increases in health care prices on downstream economic and labor market outcomes. The central empirical challenge in this paper is the threat of reverse causality: that health care prices may rise in response to changes in local economic conditions, such as shifts in wages, employment, or local insurance premiums. Because health care is a normal good, income growth tends to increase the demand for care and can generate a spurious positive relationship between health care prices and local economic activity. For example, if demand for software production rises and software firms expand employment and increase salaries and benefit generosity, insurers may face higher demand for care and negotiate higher hospital prices, biasing estimates of the effect of prices on labor market outcomes.

We therefore require a source of variation in health care prices that is orthogonal to contemporaneous local economic shocks and changes in local economic conditions. An ideal shifter would induce changes in the *supply* of health care in a way that is independent of changes occurring in other industries. Our strategy is to instrument for the hospital prices faced by employees using exposure to hospital price increases caused by horizontal hospital mergers. Ultimately, employers vary markedly in the price increases from mergers faced by their employees.

As we show below, much of the merger-driven price variation arises from differential price changes across transactions. These differences in post-merger prices are shaped by pre-merger hospital market structure, which influences the extent to which local mergers can lead to meaningful changes in concentration. Prior work has shown that this market structure was heavily influenced

unemployed if they are ever without a job and receiving UI throughout the year.

[14] This comparison excludes employers in the 5500 sample that we cannot match to the IRS data.

by hospital construction and subsidization patterns in the early twentieth century (Zabinski, 2015; Chung et al., 2017), making it plausibly exogenous to contemporary local economic dynamics. We verify that this variation is not driven by recent local economic shocks that might confound our approach and we illustrate that changes in hospital market concentration do not lead to changes in labor market concentration outside the health sector.

Our instrumental variables strategy proceeds in two steps. First, we estimate the impact of hospital mergers on hospital prices, exploiting the fact that mergers vary in the extent to which they reduce competition and, correspondingly, in the extent to which they raise prices. Second, we translate these hospital-specific price effects into employer-level and county-level measures of exposure to health care price increases. These measures of exposure integrate information on where employers' employees chose to receive care prior to mergers occurring.

5.1 Hospital Mergers, Competition, and Price Changes

We focus on mergers involving hospitals located less than 50 miles from their merging counterparts between 2010 and 2015.[15] Likewise, we only include the hospitals within a merger that are within 50 miles of one another. That is, if Hospital A merges with a system that includes B and C, and B is 10 miles from A but C is 75 miles from A, we include A and B but not C in our sample. Finally, we restrict our sample to hospitals for which we have sufficient data from HCCI to measure prices. After these restrictions, we are left with 305 mergers involving 656 hospitals during our sample period. These mergers represent 63% of all hospital mergers which took place between 2010-2015 and 80% of those with parties within 50 miles of one another. We discuss the selection of these mergers and our approach to measuring their effect on prices in more detail in Appendix B. Appendix Figure A.1 presents a map of the mergers we include in our analysis.

Across our analytic sample, the average merger generated a change in the HHI of 385 points and led to a post-merger HHI of 6,038 points.[16] As the DOJ and FTC noted in their 2010 Merger Guidelines (which applied the transactions in our sample), mergers that raise the relevant market's HHI by 200 points or more and lead to a post-merger HHI of over 2,500 should be "presumed to be likely to enhance market power" (U.S. Department of Justice and the Federal Trade Commission, 2010).[17] Twenty-three percent of the mergers in our sample (69) generated a change in HHI of over 200 and led to a post-merger HHI of over 2,500.

An alternative method for characterizing the change in market power generated by mergers is the change willingness-to-pay (WTP). WTP is a standard screening tool used in hospital antitrust

[15] We do so since Cooper et al. (2019a) do not find meaningful price increases after transactions involving merging parties farther apart.

[16] We describe our approach to measuring hospital market HHI in Appendix B.3.

[17] In 2023, the DOJ and FTC introduced revised Merger Guidelines that define a problematic transaction as one that increases HHI by more than 100 and leads to a post-merger HHI of greater than 1,800.

enforcement (Capps et al., 2019). The WTP approach first measures the substitutability of merging hospitals, then uses this substitution estimate to predict the change in negotiated prices between when the hospitals are independent and when they jointly bargain as a merged unit. We describe the model and procedure we use to estimate WTP in Appendix B.4. Approximately 26% of our transactions generate a change in WTP of over 2.5%, and 14% involve a change in WTP of over 5%. As Garmon (2017) and Brot et al. (2024b) illustrate, the increase in WTP generated by a merger approximately predicts the increase in price that the merger will produce.

5.2 Measuring the Effect of Mergers on Hospital Prices and mortality

Our primary approach to estimating the effect of mergers on hospital prices is a difference-in-differences design that follows Brot et al. (2024b) and compares prices at merging hospitals before and after a merger is consummated to prices at control hospitals. We begin by constructing hospital-by-year price indices for inpatient and outpatient care. Because hospitals are multi-product firms, we construct hedonic price indices that adjust for the mix of services offered at each hospital and the age and sex of hospitals' patients. We provide a detailed description of how we measure hospital prices in Appendix A.

In practice, we compare prices at merging hospitals pre- and post-merger to prices at comparable matched non-merging hospitals over the same time period. We use propensity scores to find the 25 most similar non-merging hospitals based on a range of observable hospital and market characteristics. We describe this matching procedure in more detail in Appendix B.2. We limit our analysis to the period covering two years before and after each merger, carving out the year the merger was consummated from our estimation. We then estimate a regression for each merging hospital:

$$\log(p_{eht}) = \lambda_{eh} \times \mathbb{1}\{\text{post-merger}\}_{ht} + \eta_{eh} + \eta_{et} + \varepsilon_{eht}, \tag{2}$$

where e is a merger event and t is a year. Our target parameter is λ_{eh}, the percent increase in price at hospital h due to merger event e. After estimating λ_{eh}, we shrink each estimate towards the mean to reduce measurement error, using the standard empirical Bayes approach. Our approach is robust to alternative matching procedures, including different matching algorithms, using fewer matches, and removing the caliper restriction. See Appendix B.2 (and Appendix C of Brot et al. (2024b)) for more details on the matching approaches used in our merger analysis. In Appendix Table Table A.4, we compare merging to non-merging hospitals before and after matching. Relative to the pool of eligible control hospitals, our matched control hospitals are substantially larger, have more technologies, a larger HCCI market share, are more likely to be a teaching hospital, and are located in less concentrated markets.

We also measure the effect of hospital mergers on hospital quality by substituting a measure of patient mortality for the price measure we used in Equation (2). Our quality measure is risk-adjusted, hospital-level, 30-day mortality among Medicare Fee-For-Service beneficiaries age 65 and older with a nondeferrable condition. Nondeferrable conditions are those where patients are admitted to the hospital on weekends at the same frequency as they are on weekdays. As noted by Doyle et al. (2015), relying on nondeferrable conditions to measure hospital quality helps address selection effects (e.g., that healthier or sicker patients differentially select certain hospitals).

5.3 Hospital Mergers, Hospital Prices, and Hospital Quality

We estimate substantial variation in the scale of the post-merger price increases generated across our sample of transactions — the λ_{eh}s which we estimated in Equation (2. The the average merger in our sample raised prices by 0.9% (the 95% confidence interval spans -0.1% to 1.8%). In Appendix Figure Fig. A.2, we present an event study of our baseline estimates of the effect of hospital mergers on prices. As we illustrate in Appendix Figure Fig. A.3, this result is robust to estimating merger effects without matching. In Appendix Figure Fig. A.4, we estimate Equation (2), substitute 30-day hospital mortality for patients with a nondeferrable condition for our price measure, and show that, consistent with Beaulieu et al. (2020), hospital mergers do not lead to a statistically or economically significant change in patient mortality.[18]

Notably, 23% of transactions in our sample raised the HHI of a merging party by 200 or more and resulted in a post-merger HHI of 2,500 or more. These are transactions, which, according to the 2010 DOJ/FTC Merger guidelines that should be ”presumed to be likely to enhance market power.”

In Figure 2, we present event studies plotting the average difference in prices for merging and non-merging hospitals before and after mergers separately for transactions that were presumed likely to lessen competition and those that were not. Mergers that did not lead to substantial changes in market concentration did not lead to economically or statistically significant price increases. By contrast, the 69 mergers that would be flagged as likely to lessen competition per the 2010 Merger Guidelines led to an average price increase of 5.9% in the two years after the transactions occurred (with a 95% confidence interval of 3.6% to 8.2%). Consistent with Cooper et al. (2019a) and Brot et al. (2024b), we find that prices increased immediately after mergers were consummated.

In Appendix Figure A.5, we replicate this analysis but instead group mergers by their resulting post-merger increases in WTP. Mergers that increased WTP by 5% or more (14% of mergers) resulted in average inpatient price increases of 5.3%, while mergers that increased WTP by 2.5% or

[18]For the mortality regressions, we apply two additional sample restrictions. First, we exclude hospitals with too few observations to estimate the risk-adjusted mortality rate for nondeferrable conditions in at least one of the two years prior and one in the two years following the merger. Second, hospitals that consolidate their Medicare Provider Numbers are excluded. These restrictions reduce the number of merging hospitals from 656 in our main sample to 442 in the mortality sample.

more (26% of mergers) resulted in average inpatient price increases of 4.1%.[19]

5.4 Instrument Construction

It might appear natural to take a similar difference-in-differences strategy to compare the outcomes of employers near a merging hospital to other, non-exposed employers. However, there are simply too many mergers in this period to make such an approach tractable. Virtually every employer in our sample had a hospital merger occur nearby (e.g., in the same commuting zone), and the small number of completely unexposed employers are located in unusually rural areas. Moreover, many employers were exposed to multiple mergers, making it untenable to neatly distinguish between pre- and post-merger treatment periods.

As a result, we use a continuous measure of employers' exposure to mergers across time as an instrument for health care prices. This allows us to measure the causal effect of rising prices on employer-level and county-level labor market outcomes.[20] Specifically, we construct a single instrument that summarizes how the price of care consumed by an employer's workers would change over time if the only thing that occurred during our sample period were the mergers that we observe during our sample period. Our approach holds fixed the prices for other non-hospital health care services, as well as the mix and quantity of health care that each employer's workers consumed. We then use this measure of employers' exposure to price changes as an instrument for employers' measured prices. This approach allows us to think about mergers as local shocks that generate price changes of varying sizes, which differentially raise health care prices across employers and over time.

We define an employer i's exposure by time t to a merger event e occurring at a specific hospital h as:

$$Z_{ieht} = \underbrace{\lambda_{eh}}_{\substack{\text{\% price change} \\ \text{at hospital } h}} \times \underbrace{\sigma_{ih0}}_{\substack{\text{share of spending at} \\ \text{hospital } h \\ \text{at baseline}}} \times \underbrace{1[t \geq \tau_e]}_{\text{Timing of merger}}, \tag{3}$$

Here Z_{ieht} captures the exposure each employer has to a specific merger. Exposure is a function of 1) the scale of the price increases generated by mergers at hospitals where their employees receive care, given by λ_{eh};[21] 2) the timing of when hospital mergers occur during our sample period, given

[19] We focus on inpatient hospital prices since WTP is constructed using demand for inpatient services. Changes in WTP measured this way are less predictive of post-merger increases in outpatient prices.

[20] Our research design, in treating the employer as a single unit, assumes that benefit design and hiring decisions are made jointly at this level. However, this might not be true if, for instance, such decisions are made at sub-organizational levels within the employer. Hegland (2025) shows that multi-state employers whose employees live in a state covered by the Affordable Care Act Medicaid expansions respond by cutting back on health insurance offers to both those employees *and* employees in non-expansion states, suggesting that employers do coordinate benefits across multiple sites of employment.

[21] We assume that a merger only affects prices once it occurs, and has a constant effect on the price of care at that hospital for that year and every year following.

by $1[t \geq \tau_e]$; and 3) the extent to which an employer's employees receive care at merging hospitals, given by σ_{ih0}. We cannot match employees in the IRS data to utilization in the HCCI data. Instead, we identify the share of hospital spending that residents from each county spend at each hospital in the US, match those county-level shares to each employee in the IRS data based on the county where the employee lives, and then take an employer-level average across employees. We model the primary effect of the merger as raising prices, and assume that this price increase occurs the year the merger is consummated. We model the post-merger price increase as generating a constant and permanent increase in the hospital's price level for every year following.

Rather than capturing the pure effect on price, Z_{ieht} simulates the expected change in *spending* as a result of the price increase. We use spending rather than prices alone for two reasons: first, by incorporating quantities, spending effectively weights different price changes by their relative importance to the employer. Second, spending, unlike price levels, can be measured per-person (which we do here), and therefore normalizes our instrument by the same unit (people) as is used by our outcome measures. To make sure that our instrument is driven *only* by changes in price, rather than changes in quantities, we measure spending shares in 2008 and 2009, in notation as "time 0," and hold them fixed across time within our instrument. Holding quantities fixed in our instrument is important, since the quantity of care that individuals consume is likely correlated with labor market outcomes. Incorporating such endogenous changes to health care consumption quantities into our instrument would bias our estimates.

To construct a single employer-year-level instrument, we simply sum these Z_{ieht} across the set of all hospital merger events $\mathscr{E}$ nationwide for all hospitals $\mathscr{H}$:

$$z_{it} = \sum_{e \in \mathscr{E}, h \in \mathscr{H}} Z_{ieht}. \tag{4}$$

where z_{it} is "simulated spending," i.e., how health care spending (per-covered-life) would evolve for employer i if the only thing that changed between 2009 and t were price changes resulting from hospital mergers. Across our sample, the mean 2009 to 2015 change in employer-level simulated spending was 0.1%, with a standard deviation of 0.92 percentage points. The top 10% of employers experienced an increase in simulated spending of 0.9%.

With our instrument z_{it} constructed, we estimate our model using two-stage least-squares (2SLS), regressing outcomes y_{it} on per-person spending S_{it}, instrumenting for x_{it} with simulated spending z_{it}, controlling for employer and year fixed effects:

$$S_{it} = \delta \times z_{it} + \Theta_i + K_t + u_{it} \tag{5}$$

$$y_{it} = \beta \times S_{it} + \theta_i + \kappa_t + \varepsilon_{it}. \tag{6}$$

where S_{it} is our measure of employer health care spending per covered life, as defined in Section 4.2. We cluster our standard errors in this analysis at the employer level.

The first column of Table 2 shows our estimates of Equation (5). Our standard employer-level first-stage regression has a coefficient of 0.64, a 95% confidence interval that spans 0.60 to 0.69, and an F-statistic of 844. This coefficient illustrates the extent to which merger-driven price increases are passed through into total health spending. Our measure of simulated spending reflects a level change, since it is measured in percentage points. Therefore, to estimate pass-through, we must exponentiate the estimated coefficient. After doing so, our results imply that 90% of the price increases from hospital mergers are reflected in spending changes, implying limited quantity responses. In Appendix Figure A.7, we present binned scatterplots of the relationship between the change in simulated and true spending from 2010 to 2015.[22]

5.5 Identification

Our simulated spending instrument relies on two standard assumptions common to instrumental variables approaches: relevance and exclusion. The F-statistic from our first-stage regression presented above suggests that our instrument is relevant and strong.

With respect to exclusion, it is helpful to regard our IV approach is akin to a shift-share instrument, with the merger-driven price increases serving as a 'shift' and the employer-hospital exposure serving as a 'share.' We therefore face the same identification challenges that are common in standard shift-share designs. Borusyak et al. (2025) note that, in research designs like ours, we can proceed by assuming that *either* the shift or the share is exogenous, and need not make any assumptions about the other component. In our setting, we assume that the magnitudes of the post-merger price changes are exogenous, what Borusyak et al. (2025) refer to as the "exogenous shocks" design.

To demonstrate that these shocks are plausibly exogenous, in Appendix Figure A.6, we regress merger characteristics (e.g., whether mergers occurred, the post merger price increases transactions generated, and whether whether led to meaningful reductions in competition) on local economic trends: recent year-to-year changes in health care spending, income, and unemployment. In all cases, we find no statistically or economically significant relationships between our assumed-exogenous merger characteristics and local economic trends.

We can also demonstrate that the bulk of our identifying variation is indeed driven by the price changes that occur post merger, rather than the variation in the spending shares that define workers' exposure to merging hospitals. In Appendix D, we conduct an exercise inspired by Borusyak and

[22] In various locations in the paper, we use alternative versions of our first-stage (for example, when we run our analysis of the impact of rising health care prices on insurance premiums and have to use a restricted sample of employers). In Appendix Table A.6, we include the alternative first-stage estimates we use throughout the analysis.

Hull (2023), in which we construct alternative versions of our instrument that purge the specific sources of variation from the instrument. Measuring the strength of our instrument after sequentially removing each source of variation allows us to decompose the relative contribution of merger timing, merger intensity, and employer-to-hospital exposure to the variation in our instrument. As we illustrate in Appendix Table A.6, we find that our instrumental variation is primarily driven by variation in the scale of the post-merger price increases across mergers that employees are exposed to. This is illustrated by the fact that Column (4), which estimates a first stage excluding variation from merger-driven price increases, has a relatively weaker F-statistic. When we use *only* variation in the price increases, we estimate a first-stage coefficient similar to that of our primary instrument and retain a strong F-statistic of 509.

Second, as a more direct test, in Figure Fig. 3, we present event studies that show simulated spending for the 25% of employers with the largest change in simulated spending from 2009 to 2015 and the 25% of employers with the smallest change in simulated spending during that period. For each employer, we define the treatment year as the year in which that unit recorded its largest single-year jump in simulated spending. This anchors the event time at the moment the merger-driven price shock is most intense. [23] We then estimate a standard event study regression, centering around the timing of this enumerated event, comparing treated to control units. Our event study in Figure Fig. 3 shows that health spending increases in the year prices increase and then steadily increases over time. We see no substantial difference in trends in spending between our "treated" and "control" group in the two years before the merger-driven price increases occurred.

One concern for our exclusion restriction is that hospital mergers may have direct effects on local health by changing the quality of care. This might, put downward pressure on labor productivity, decreasing wages or employment. Alternatively, if mergers raise the quality of care, the interpretation of our results as indicating deadweight loss might change if workers now value access to merging hospitals by more. This is why, in Appendix Fig. A.4, we estimate the effect of the mergers we study on clinical quality. We replicate our research design for estimating post-merger changes in prices, but instead use the 30-day mortality rate for non-deferrable admissions among Medicare patients as our outcome. We find, echoing prior work by Beaulieu et al. (2020), that our mergers have no significant effect on clinical quality.

5.6 Alternative Strategies for Measuring the Effect of Mergers

This analysis (and our later analysis of downstream outcomes) suggests that employer exposure to mergers is not correlated with pre-existing economic trends. However, our measure of post-merger

[23] As a robustness check, in Appendix Figure Fig. A.8, we restrict the treated group to cases where one the largest one-year jump in simulated spending accounts for at least 50% of the unit's total long-difference (i.e., where the simulated spending series exhibits a single, clearly identifiable "event").

price increases is obtained by estimating changes in prices for local hospitals. As a result, our estimated price effects may also pick up other shocks to the supply or demand of health care that happen to occur simultaneously with mergers. Even if these simultaneous shocks occur randomly, they may explain some of the variation in exposure and may themselves violate exclusion restrictions. For example, there may be simultaneous demand shocks that raise prices as a result of local income growth. Including variation from these shocks may bias our results.

To purge this bias, as an additional robustness test, we replicate our main analyses using a modified version of our instrument where, rather than using the true estimated post-merger price increase λ_{eh} for each merging hospital, we instead use the *predicted* change in price given the change in competition induced by the merger, as measured by the change in WTP described in Section 5.1.[24] Note that this approach is conservative: by using *only* changes in market structure, we are purging our instrument of any sources of error in estimated price effects. However, in turn, we are also purging it of any variation that is valid but not captured by the bargaining model used to estimate WTP (Capps et al., 2003) — e.g., differences in the negotiating skill of hospital executives. In Column (3) of Table 2, we present the coefficient from the first-stage regression using this alternative instrument, which also has a strong F-statistic.[25]

6 Employer-Level Results

6.1 Health Insurance Outcomes

We begin by estimating the effect of rising health care prices on two insurance market outcomes: ESI premiums and employee HDHP enrollment. Analyzing premiums is a vital link in our causal chain, since the market for health care and market for labor are intermediated by the price of employer-sponsored health insurance. Enrollment in an HDHP serves as a proxy for whether employers shifted their workers into plans with more out-of-pocket exposure. This measure is constructed based on the share of an employer's employees who have a non-zero employee or employer contribution to an HSA, since contributions to an HSA cannot be made without being enrolled in an HDHP. Note that, as described in Section 4.3, whereas we can analyze the use of

[24]Our estimates of predicted price increases are based on the approach established by Capps et al. (2003), and we include additional detail on the estimation in Appendix B.4. WTP measures the predicted change in profit generated by a merger. Therefore, our predictions are only proportional to the predicted price change for a given merger.

[25]We note that the first stage estimate using the WTP prediction is larger than our main first stage. This occurs, in part, because our WTP estimates are based on demand estimates and spending shares for inpatient services whereas our main first stage uses estimated price changes and spending shares for both inpatient and outpatient services. We could have instead applied the our inpatient-based predicted price changes to spending shares for both inpatient and outpatient, which would have required us to assume that these price increases, predicted using inpatient demand estimates, are a good proxy for the realized price changes for outpatient services. In unreported estimates we find that our results are robust to this approach and that, when we use this approach, our first stage becomes more similarly scaled to our baseline estimates.

HSAs for all employers, our analysis of insurance premiums requires that we use a heavily restricted panel of 3,970 employers that we can link to their Form 5500 filings.

We report estimates of the effect of rising health care prices on insurance premiums in Column (1) of Table 3. Panel A presents results from uninstrumented OLS regressions estimating Equation (6). We do not observe a economically or significantly significant relationship between health care prices and insurance premiums in our OLS estimates. By contrast, in our primary 2SLS estimates of Equation (6) presented in Panel B, we find with 10% precision that a 1% increase in health care prices leads to a 0.95% increase in ESI premiums (our 95% confidence interval spans -0.11% to 2.01%). These estimates suggest that rising health care prices are almost fully reflected in higher insurance premiums.[26] Our relatively low power for this outcome largely reflects the fact that the Form-5500-linked employer panel only contains 3,970 unique employers, compared to 140,300 in our primary sample.[27] As we demonstrate in Appendix Table A.8, when we use our alternative instrument that replaces estimated price effects with predicted price effects based on changes in market structure (measured by WTP), we find effects of similar magnitude and statistical power.

Employers could respond to rising health care premiums due to rising prices by reducing the generosity of their coverage (e.g., shifting policyholders into higher deductible plans or plans with other forms of enhanced employee cost-sharing). In Column (2) of Table 3, we present OLS and IV estimates of the effect of rising health prices on employer provision of HDHPs (as measured by whether their employees contribute to an HSA). Whereas our OLS estimates suggest a positive relationship between health care prices and the use of HDHPs, when we instrument for health care prices, we do not find an economically or statistically significant relationship between health care prices and the use of HDHPs. This suggests that employer demand for coverage is very inelastic. While this may seem surprising at first, it is important to remember that employers must offer benefits uniformly to all workers, and so benefit adjustment is a coarse response to rising prices.

6.2 Labor Market Outcomes

We next turn to estimating the effect of rising prices on labor market outcomes. We focus on employers' payroll and their count of workers (in logs) as our primary outcomes. In Columns (1) and (2) of Panel A of Table 4, we present OLS estimates of Equation (6) and find no economically or statistically significant relationship between rising health care prices and employers' payroll or count of workers. By contrast, in Panel B, we present 2SLS estimates of Equation (6) and find

[26] Note that the average actuarial value of the health plans in our sample, as measured in the HCCI data, is 0.81. As a result, our estimates imply a pass-through rate slightly above 1. However, our confidence interval on our premium response estimate includes 1.

[27] We have insurance premiums for 3,970 employers. We can merge 81% of employers from our Form 5500 insurance premium sample to the panel of employers we use to measure labor market outcomes. For employers from our insurance premium sample that do not merge into our analytic sample of EINs, we use a county-level measure of simulated spending rather than an employer-level measure.

that a 1% increase in health care prices reduces both payroll and employment by 0.36% (the 95% confidence intervals are, respectively 0.11% to 0.62% and 0.10% to 0.61%).[28]

Note that our OLS-estimated coefficients are weakly positive while our IV-estimated coefficients are strongly negative. This suggests that the relationship between health care prices and non-health care labor markets includes connections both through the supply channel that we focus on (e.g., where rising health care prices make workers more expensive to employ) and also, via reverse causality, through the demand channel (e.g., where greater labor productivity increases incomes and thus health care demand, creating a positive correlation between health care prices and non-health care labor market outcomes). This highlights the necessity of our instrumental variables strategy.

Our analytic sample includes employers from both the health sector and other sectors. Non-health care employers are likely to reduce payroll and employment following an increase in local health care prices because health care price increases raise their insurance premiums. By contrast, health care employers receive higher prices following a merger and thus higher revenue, which may *increase* their payroll and employment. In Table 4, we measure the employment effects separately at non-health care (Columns (3) and (4)) and health care employers (Columns (5) and (6)).[29] After segmenting employers by industry, we see that our overall results are entirely driven by changes in payroll and employment at non-health care employers. Among non-health care employers, we find that a 1% increase in health care prices lowers payroll by 0.37% and lowers the count of workers employed by 0.40% (the 95% confidence intervals are 0.11% to 0.63% and 0.14% to 0.66%, respectively). Conversely, we do not find that an increase in health care prices generates any statistically significant changes in payroll or employment at health care employers.

We present event study estimates of the effect of rising health care prices on payroll and the count of workers (in logs) at non-health employers in Panels A and B of Figure 4.[30] The event studies show that there are flat trends in labor market outcomes in the two years before treatment, compared to significant effects after the instrument increases in value. This helps rule out differential trends in pre-merger labor outcomes for employers who are more versus less exposed to mergers. The coefficients for event time 0 are sizable, suggesting that employers respond immediately to increases

[28]That we find equal payroll and employment effects suggests that there is minimal wage reductions. We caution, however, against such a strict interpretation of this result for two reasons. First, our estimates average over many worker types. This result is consistent with higher-wage workers taking a wage cut while lower-wage workers see employment cuts, resulting in no effect on average income per retained worker despite wage pass-through. Second, our payroll measure only includes wage compensation. It is also important to note that our outcome measure of Medicare wages in the tax data is already adjusted to exclude employee contributions to health insurance premiums. As a result, our pass-through measure *does* effectively incorporate potential employer strategies to pass through rising insurance premiums in terms of higher employee contributions. Although we cannot distinguish between nominal wage adjustments and increases in the share of premiums paid, these estimates capture the combined effect of both.

[29]We identify employers in the health care industry by whether they have a reported NAICS code starting in "62" as reported to the IRS.

[30]In Appendix Figure A.10, we reproduce these event studies and limit to treated observations where the largest one-year jump in simulated spending accounts for at least 50% of the unit's total long difference.

in health care prices.[31] While this may seem surprising, it likely arises from the fact there is often a substantial lag between when a hospital merger get announced and the time the transaction is legally consummated and the parties can begin to negotiate jointly. This lag provides an anticipatory window for insurers and employers to forecast treatment effects and respond accordingly. Because insurers set personalized premium rates for employers, they can raise rates as soon as a merger occurs, anticipating higher hospital prices in the same year.

In Appendix Figure A.12, we show our baseline employer-level estimates are similar in both magnitude and precision across a range of alternative specifications: 1) using only the identifying variation in our instrument coming from cross-hospital variation in post-merger price increases; 2) replacing our estimates of post-merger price increases with predicted price increases given post-merger changes in market structure (measured via WTP); 3) forcing our estimator to only compare outcomes for employers within comparable subsets of firms by interacting our year fixed effects with quartiles of employer size (measured by employee counts in 2009), industry (measured as NAICS code in 2009), quartiles of payroll growth (measured as the percent change between 2006 and 2009), or quartiles of size growth (measured as the percent change in employee count between 2006 and 2009); and 4) excluding employers in the bottom or top quartiles of payroll or size growth between 2006 and 2009.[32] Finally, in Appendix Table A.5, we show robustness to looser restrictions on our sample of employers. Our main specification limits to a balanced panel of employers, so we do not have to account for taking log transformations of zero outcomes if employers exit or reincorporate under a different tax identification number. We also chose to focus on employers with above 50 employees at baseline because they are less likely to exit or reincorporation during our sample period. However, as we illustrate, we find very similar results when we 1) expand our sample to employers with over 10 employees at baseline; 2) focus on employers with between 10 and 50 employees; and 3) include employers who exit during the our sample period.

6.3 Scaling Employers' Response Rising Health Care Prices

Scaling our primary point estimates on a dollar-for-dollar basis, a 1% increase in health care spending translates into an approximately $40 increase in spending per insurance plan member. The average employee in our sample, according to our HCCI data, has an insurance plan that includes

[31] The estimated coefficients approximately double in magnitude between event time 0 and event time 2, which would initially suggest that effects are increasing over time. We would caution against such a conclusion, however, since these are reduced-form estimates, and, as Figure 3 shows, the first-stage coefficient also approximately doubles over this event time window. If we appropriately rescaled these estimates to account for the size of the first-stage, we would conclude that price increases primarily have immediate level effects on payroll and employment.

[32] The lone estimates where we lose precision are when we exclude the 25% of employers with the greatest wage growth between 2006 and 2009 and the top 25% of employers with the biggest increase in employment between 2006 and 2009. When we exclude those cohorts of employers, our point estimates drop modestly and become marginally less precise.

one dependent (two total members), generating a roughly \$81 spending increase per employee. The median employer in our sample has 133 employees, so a 1% increase in in spending at an employer, summed across its employees and employees' dependents, is \$10,709. Our point estimate in Column (1) of Table 4 implies that a 1% increase in health spending leads to a 0.36% reduction in payroll, which is equivalent to an approximately \$18,046 decrease at the median employer.

This implies that, for each \$1 of hospital revenue generated by mergers, employer payroll declines by \$1.69 (95% confidence interval spans \$0.49 to \$2.88). The fact that the total payroll reduction is greater than the spending increase does not necessarily reflect greater than one-for-one pass-through of per-employee health insurance costs into wage levels. Instead, it reflects the extensive margin responses in the number of workers employed, which can reduce total worker pay by far more than the direct cost increase. These extensive margin responses generate deadweight loss due to the "hospital tax" imposed by rising prices much in the same way that other forms of employer taxation generate deadweight loss by reducing employment.

To contextualize the size of the employment effect, we benchmark our estimates against studies of employer-level responses to payroll tax changes, which enter into employer decisions in a way most similar to health insurance costs. In our setting, a 1 percentage point increase in the payroll tax would cost the median employer \$365 per worker. The median employer spends \$4,039 per covered life (and thus \$8,078 per worker), meaning that a 1 percentage point increase in the payroll tax is roughly equivalent to a 4.5% increase in health care spending. Since our estimated coefficient (from Column (4) of Table 4) says that a 1% increase in health care spending reduces non-health care employment by 0.40%, a 4.5% change should reduce employment by 1.8%. In Appendix Table A.9, we compile a set of prior estimates of responses to payroll taxation. Estimates from US studies include 0.7 to 0.9% (Anderson and Meyer, 1997), 1.5% (Johnston, 2021), and 1.1 to 2.4% (Guo, 2024). We are reassured that our employment effect estimate sits squarely in the range of what is found in the existing literature, suggesting that the market power exercised through hospital mergers results in employment shifts that are similar in scale to the distortions induced by payroll taxes.

6.4 Cross-Employer Heterogeneity

Larger employers may be able to resist pass-through of cost increases in insurer-employer bargaining to a greater extent than smaller employers. Indeed, Lin and Zhou (2024) show that insurer mergers generate larger post-merger premium increases for smaller employers. To test this in our setting, we segment employers by whether they have an above- or below-median count of workers in our base year, and estimate our regression model separately for each groups. We present these results in Appendix Table A.10. Consistent with Lin and Zhou (2024), we find very large effects for below-median-size employers and more modest, less statistically precise results for above-median employers. We can statistically reject equal effects on worker counts across the our estimates at above- and below-median size employers.

7 County-Level Results

Our results from Section 6 show that employers respond to rising health care prices by reducing the number of workers they employ. However, the large effects we observe on employment may be driven, at least partially, by a re-sorting of workers across existing and new employers, as well as by a shift of workers away from wage employment to self-employment. As a result, in this section, we explore whether rising health care prices lead to aggregate effects on income per capita, self-employment, and unemployment. To do so, we aggregate our employer-level instrument up to the county level and focus on a range of new county-level outcomes.

Because the effects of rising health care prices are always intermediated by employers, we define county exposure to health care prices and mergers as the weighted average of employer-level exposure, with employer-by-county weights equal to the extent to which the employer typically hires workers from that county. Specifically, in 2009, we take every worker at an employer in our employer sample and assign them a county according to their address of residence, as reported on their W-2.[33] We then construct employer-by-county weights, ω_{ic0}, as the share of workers in a given county c who worked for employer i in 2009. We can then construct our endogenous regressor S_{ct} (per-worker spending) and instrument z_{ct} as weighted averages:

$$S_{ct} = \sum_i \omega_{ic0} S_{it} \tag{7}$$

$$z_{ct} = \sum_i \omega_{ic0} z_{it} \tag{8}$$

Effectively, we are stipulating that county-level exposure to health care price increases is a function of how employers who hire in that county are exposed to price increases.[34] At the county level, our mean change in simulated spending is 0.1% and the standard deviation is 0.5 percentage points. The top 25% of counties have an increase in simulated spending of 0.2% and the top 5% have an increase in simulated spending of 0.9%. In Figure A.9, we plot a map of county-specific changes to simulated spending from 2009 to 2015.

[33]In cases where a given worker reports multiple residences on multiple W-2s, we choose according to the W-2 with the highest wages.

[34]We are using the quantities of *employers* rather than *specific individuals* within the county. Therefore, a county's effective health care price exposure may depend on individuals from other counties, through co-working relationships and the ESI channel.

We then estimate the county-level analog of our main 2SLS regressions:

$$S_{ct} = \delta \times z_{ct} + \Theta_c + K_t + u_{ct} \tag{9}$$

$$y_{ct} = \beta \times S_{ct} + \theta_c + \kappa_t + \varepsilon_{ct} \tag{10}$$

i.e., we regress outcomes at the county c and year t level on county-year health spending, instrumenting with county-year simulated spending and including county and year fixed effects. We cluster all regressions at the county level.

In Column (1) of Appendix Table A.7, we present the first-stage regression results, estimated using Equation (9). Since we have collapsed our data down to the county level, we have fewer observations and thus less power. The F-statistic on our first-stage from this approach is 41, which is above standard thresholds for weak instrument tests.[35]

7.1 Labor Market Outcomes and Federal Tax Revenue Collected

In Table 5, we present estimates of Equation (10) and show the impact of rising health care prices on county-level overall income per capita, the share of the population unemployed, and federal income tax revenues per capita. In this analysis, we focus on all workers, including those employed in the health care industry. This allows us to assess the impact of rising health care prices on aggregate income and unemployment. As we illustrate in Column (1) of Panel B, a 1% increase in health care prices leads to a 0.28% decrease in overall county-level income per capita with a 95% confidence interval that spans -0.02% to 0.58%. Note that our measure of income per capita captures both self-employment income and W-2 income for workers inside and outside the health sector. The fact that we find non-zero effects on overall income per capita suggests that the losses in labor income among non-health-care workers is not made up for by health care workers, nor by workers shifting from other sectors to the health care industry.

However, as we illustrate in Appendix Table A.11, this overall income effect is driven by changes in the income of non-health care workers.[36] We note that the magnitudes we estimate on the employer sample and this county level sample are similar when focusing on non-health care workers or employers, suggesting that confounding factors like individuals moving across space do not meaningfully contribute to our results. Indeed, we show in Appendix Table A.12 that local increases in health care prices do not induce significant increases in worker relocation.

[35] In this table, we also include alternative estimates of the first-stage regression, as we do in Appendix Table A.6. As in those estimates, the primary source of variation comes from post-merger pricing differences. Additionally, replacing price changes with predicted price changes given market structure changes results in an equally strong instrument.

[36] We classify a worker as being in the health care sector if, in the prior year, they filed a W-2 reporting work for an employer with a NAICS code indicating the health care industry. We categorize all other individuals, including those who filed no W-2s in the prior year, as workers outside the health care sector.

Notably, as we show in Panel A of Table 5, our OLS estimates produce a weakly positive relationship between health care prices and income per capita. As in our employer-level results, this highlights how, in the absence of our instrumental variables strategy, estimates pick up correlations between local income and health care demand.

In Column (2) of Panel B of Table 5, we show that a 1% increase in health care prices leads to a 0.09 percentage point (1%) increase in overall unemployment per capita (with a 95% confidence interval that ranges from 0.01 to 0.17), which we measure as the share of the population receiving unemployment insurance or receiving zero income.[37] Per our results in Appendix Table A.11, this overall increase in unemployment is being driven by increases in labor market separations among non-health care workers who gain UI when they lose their job.

In Panel A of Figure 5, we present event study estimates of the effects of rising health care prices on income per capita and unemployment. Our results are consistent with our employer-level results: we find no differences in pre-trends before the increase in spending and a change in labor market outcomes immediately after the spending increases occur. As we illustrate in Panel A, labor income falls immediately to a lower level in the year that the price shock occurs and then continues to decline in the two years after the price shock. In Panel B of Figure 5, we present event study estimates of the effect of rising health care prices on unemployment. The event study illustrates that unemployment increases in the year of the shock and then peaks the year *after* prices increased. Note that this is a flow measure of unemployment rather than a stock measure.[38]

Our results also suggest that the degradation of the labor market has negative consequences for federal and state budgets. As workers see their salaries reduced and jobs cut, they will have less taxable labor income. Indeed, as we illustrate in Column (3) of Panel B of Table 5, we estimate that a 1% increase in health care prices brings federal income tax receipts down by 0.37%, with a 95% confidence interval that spans -0.01% to 0.75%. In dollar terms, this implies that a $1,000 increase in health spending per beneficiary caused by price growth lowers tax revenue by about $619 per worker.[39] Note that this number reflects the fact that the average worker has one dependent. Our results imply that the average tax rate on the marginal dollar is approximately 22%. For reference,

[37]Our measure of unemployment is conditional on an individual being employed in the prior year. As a result, our unemployment measure captures *flows* into unemployment rather than the *stock* of the unemployed. Since unemployment insurance is often time-limited, flows into it are more reliable than the stock of current recipients in terms of understanding unemployment patterns. Ideally, we would be able to measure time spent out of employment, but no relevant tax forms report hours worked in a given job or any other quantity measures.

[38]In Appendix Figure A.11, we reproduce these event studies on our more limited sample of counties where the largest change in simulated spending accounted for at least 50% of the change in the long difference in simulated spending during our sample period.

[39]This number is calculated by comparing the implied effects of a $1,000 increase in spending per beneficiary on income and taxes as implied by the estimated elasticities. Specifically, the tax change is calculated as tax elasticity $*$ taxes/capita $* \frac{1000}{\text{spending/beneficiary}} = -0.37 * 7,009 * \frac{1000}{4,211} = -619$. Estimating the change in income using the same method gives a $2,815. Taking the average change in taxes divided by the average change in income results in the 22% implied average marginal tax rate.

the average tax rate for individual filers earning $35,350 per year was 25%.

Likewise, since the employment effects we observe result in individuals receiving unemployment insurance, as we show in Column (5) of Appendix Table A.12, we find that a 1% increase in health care prices leads to a 2.42% increase in transfers to former workers from the UI system (the 95% confidence interval spans -0.51% to 5.35%). Collectively, these estimates highlight that the incidence of rising health care prices also falls on state and federal governments and not just on employers and workers.

In Appendix Figure A.13, we show how our county-level income per capita and unemployment results shift when we shift our source of variation and alter our sample of counties. First, we show our results are robust to relying exclusively on the variation in post-merger price increases across transactions to drive our instrument. Second, we show that our income result is robust to replacing post-merger price effects with predicted post-merger markup increases (as measured by WTP).[40] Third, we show our results are robust to interacting our time fixed effects with income quartile fixed effects so that all comparisons are made within quartiles. Fourth, our results are robust to excluding the 25% of counties with the lowest income per capita in 2009. Excluding counties in the bottom 25% of change in income per capita from 2006 to 2009 lowers our point estimates and they become imprecise. Fifth, we show that our results are robust to excluding the top 25% of counties with the highest unemployment in 2009 and change in unemployment from 2006 to 2009 largely does not shift our results. Sixth, we show that our results are robust to excluding the 25% most rural counties does not change our results. Seventh, we show that our results are robust to excluding the 25% of counties which, based on Autor et al. (2013), were the most exposed to import competition from China in the 2000s does not shift our main point estimates (though it does modestly increase our standard errors). Finally, in Appendix Table A.13, we present estimates of Equation (10) for measures of labor concentration for potential hospital input suppliers. We measure county-level labor market HHI for the food services (NAICS code 722), legal services (NAICS code 5411), and accounting services (NAICS code 5412). We find that merger-driven increases in prices do not induce economically or statistically significant changes in labor market concentration in these industries. Moreover, we also estimate a specification where we take county-level, 2-digit NAICS industry labor market concentration measures for all non-health-care industries (32 industries in total), and stack them to estimate a pooled regression that includes industry and industry by year fixed effects. The resulting coefficient can be interpreted as the effect of rising prices on the concentration of the average industry. Once again, we estimate no economically or statistically significant changes.

In Section 6.2, we estimated that each $1 increase in employer health care costs caused reductions in employer payroll of $1.69. We can undertake a similar exercise for our county-level estimates.

[40]Note that our employment results are not robust to measuring the effect of mergers using WTP.

Our primary county-level estimate finds that a 1% increase in health care spending reduces labor income by 0.28% (our 95% confidence interval for this estimate spans -0.02% to 0.58%). Average county-level spending per capita is $4,211. As before, we know that each employee's plan, on average, covers one dependent in addition to themselves, on average, so a 1% increase in average spending per *worker* is $84. Average county-level income per person in our sample is $42,050, so a 0.28% reduction therefore lowers average earnings by $119. Comparing these, we get that a $1 increase in hospital revenue reduces local labor income by approximately $1.41 (95% confidence interval for this estimate spans $-0.09 to $2.91); to shift $1 of revenue to hospitals through the "hospital tax," workers must give up more than $1 of earnings. As before, we can interpret the fact that this ratio is greater than one-for-one as evidence of deadweight loss from the "hospital tax:" this can only arise so long as there are distortionary reductions in employment as a result of rising health care costs. Our county-level estimate is likely smaller than our employer-level estimate ($1.69 per dollar) because it accounts for potential reallocation of workers across employers and, in particular, between the health care and non-health care sectors, whereas our employer-level estimate only accounts for the loss at the specific affected employer.[41] This estimate, however, is estimated with less statistical power than our employer-level estimate and, as a result, the two are not statistically distinguishable.

In Appendix E, we take our estimates and scale the effect of hospital mergers on income, employment, and federal income tax revenue. To do so, for the mergers of interest (e.g., mergers that the FTC/DOJ guidelines suggest would meaningfully lessen competition), as we describe in Appendix E, we compute the change in our instrument induced by those mergers for every county in the year the mergers occurred. Focusing on the 69 anticompetitive mergers in our analytic sample, for example — those that generated a change in HHI over 200 points and led to a post-merger increase in HHI of over 2,500 points — we find that the average anticompetitive merger led to a $18.2 million reduction in income (95% confidence interval $0.2M to $36.1M), approximately 112 job losses (95% confidence interval 17 to 207), and about $3.8 million reduction in income tax payments (95% confidence interval $0.3M to $7.4M).

7.2 Distributional Effects

As we discussed in Section 3, rising premiums are costlier, proportionally, for lower-wage workers. Therefore, there should be heterogeneous effects on workers across the income distribution. To test

[41] A second difference is that this estimate accounts for all workers, whereas our employer-level estimate only accounts for non-health care workers. If we use the reduction in non-health care income estimates from Table A.11, we estimate a $1.76 (95% confidence interval for this estimate spans $0.14 to $3.38) reduction in non-health-care labor income per $1 of hospital revenue received, and a $0.90 (95% confidence interval for this estimate spans $-1.22 to $3.03) increase in health care labor income. The potential for reallocation between sectors means that the comparison between this number and our employer-level estimate should be treated with caution, because the county where *patients* (or their coworkers) are located is not the same as the county where *health care workers* are located.

this, for each year, we measure workers' W-2 income in the prior year and segment workers into \$10,000 bins up to \$100,000 in income, after which we use \$50,000 bins, up to a bin for workers who earned \$200,000 and over. We exclude any individuals who received UI in the prior year because their prior income may be an incomplete measure of their relative position in the labor force.[42] We then estimate the effect of price increases on unemployment for each of these groups separately.

We present the results from this exercise in Figure 6. As the figure illustrates, we find close to zero unemployment effects at the top of the income distribution (i.e., those previously earning above \$100,000—approximately the 85^{th} percentile of the individual income distribution). This is consistent with the notion that increases in premiums are small relative to overall compensation for this group of workers.

Likewise, we find close to zero effects for those at the bottom of the income distribution (i.e., those previously earning below \$20,000). These results are consistent with the fact that lower-wage workers tend not to receive employer-sponsored health insurance benefits (Lurie and Miller, 2023). Because they do not receive ESI, rising health insurance premiums will not make these workers more expensive to retain.

By contrast, we find relatively uniform and negative employment effects among workers who previously earned between \$20,000 and \$100,000. We attribute this finding to two forces. First, the "head tax" pushes unemployment effects to be relatively regressive. Second, however, there is a positive correlation between wages and health insurance generosity: employers that offer high average wages also, on average, offer more comprehensive health insurance coverage (Lurie and Miller, 2023). The sum of these two forces implies that the groups hardest-hit by the unemployment effects of rising health care prices are lower- and middle-income workers.

8 Discussion and Concluding Thoughts

Over half of Americans are covered by an ESI plan. In this paper, we have shown that ESI creates a pathway through which rent-seeking and inefficiency in the health care industry can cause harm to local labor markets. We find that the "hospital tax" that rising health care prices impose on employers lowers employment (both at individual employers and overall in local economies), reduces workers' earnings, lowers tax revenue, and decreases government budgets. Our estimates imply that, at an employer-level, a 1% increase in health care prices reduces payroll at employers outside the health care industry reduce by 0.37% (our 95% confidence interval spans 0.11% to 0.63%). At a

[42] E.g., consider a worker who should expect to make an annual salary of \$60,000. If they worked in the prior year until the end of April, then were laid off, collected UI, and did not work for the rest of the year, we will classify them as having made \$20,000. While they did indeed make this much that year, it does not reflect the labor market they participate in.

county-level, 1% increase in health care prices raises aggregate flows into unemployment by 0.09 percentage points (1%) (with a 95% confidence interval of 0.01 to 0.17). Moreover, these negative consequences are disproportionally borne by lower- and middle-income individuals.

During our period of study, prices for inpatient and outpatient hospital care for the privately insured grew by 42.3% and 25.1%, respectively (Cooper et al., 2019b).[43] Given the share of total health services that hospital care represents, this price increased caused a 18% increase in non-drug health care spending. As a result, our estimates imply that the price growth between 2007 and 2014 reduced workers' incomes by approximately 5% (with a 95% confidence interval of -0.3% to 10.3%), increased aggregate unemployment by approximately 1.5 percentage points (a 15% increase or about 1 million jobs lost) (95% confidence interval from 0.1 to 3.0 percentage points), and lowered federal income tax revenues by approximately 7% (95% confidence interval of -0.1% to 13.3%). To be sure, some of this increase in hospital prices likely reflects quality increases that improved social welfare. However, those quality improvements would need to be substantial to offset harms on the scale we estimate.

In this paper, we relied on hospital mergers to generate shocks in local hospital prices. While the focus of this paper is on analyzing the labor market consequences of rising health care prices, our identification strategy —relying on the price increases generated by hospital mergers— allows us to analyze how hospital mergers impact local communities. From 2002 to 2020, there were over 1,000 hospital mergers in the US. Our results demonstrate that the average merger in our sample —a merger that raised prices by 0.9%—led to a $6.3 million reduction in income (95% confidence interval $0.1M to $12.6M), approximately 39 job losses (95% confidence interval 6 to 72), and about $1.3 million reduction in income tax payments (95% confidence interval $0.1M to $2.6M). Likewise, the 23% of consummated mergers in our sample that could have been flagged using standard screening tools — mergers that generated a change in HHI over 200 points and led to a post-merger increase in HHI of over 2,500 points — led to, on average, a $18.2 million reduction in income (95% confidence interval $0.2M to $36.1M), approximately 112 job losses (95% confidence interval 17 to 207), and about $3.8 million reduction in income tax payments (95% confidence interval $0.3M to $7.4M). Note that these job losses reflect changes in overall employment, so are net of any gains in employment that the mergers generated in the health sector. Collectively, the harm from the 305 mergers in our sample that occurred between 2010 and 2015 was substantial: we estimate these mergers led to approximately $1.9 billion in forgone wages (95% confidence interval $0.0B to $3.8B) and approximately 12,000 job losses (95% confidence interval 1,809 to 21,998).

Ultimately, this work documents the central link between health care price growth and broader macroeconomic outcomes in the US. Going forward, we hope this work motivates future research on the absolute and distributional effects of rising health care prices and health spending. For

[43] Note that the Cooper et al. (2019b) estimates are net of inflation.

example, while we have focused on outcomes for workers outside the health care sector, future research should examine who receives the gains from health care price increases, and how those gains are distributed across the health care workforce. We also hope this work encourages future study of how rising health spending growth shapes regional economic growth and productivity changes across the nation.

References

American Hospital Association (AHA), "Annual Survey of Hospitals," 2019. Accessed September 5, 2019.

Anderson, Patricia M. and Bruce D. Meyer, "The Effects of Firm Specific Taxes and Government Mandates with an Application to the US Unemployment Insurance Program," *Journal of Public Economics*, 1997, *65* (2), 119–145.

Arnold, Daniel and Christopher Whaley, "Who Pays for Health Care Costs? The Effects of Health Care Prices on Wages," 2024.

Autor, David H., David Dorn, and Gordon H. Hanson, "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 2013, *6*, 2121–68.

Baicker, Katherine and Amitabh Chandra, "The Labor Market Effects of Rising Health Insurance Premiums," *Journal of Labor Economics*, 2006, *24* (3), 609–634.

Baicker, Katherine and Jonathan Skinner, "Health Care Spending Growth and the Future of U.S. Tax Rates," *Tax Policy and the Economy*, 2011, *25* (1), 39–68.

Baqaee, David Rezza and Emmanuel Farhi, "The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem," *Econometrica*, 2019, *87* (4), 1155–1203.

Baumol, William J., "Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis," *American Economic Review*, 1967, *57* (3), 415–426.

Beaulieu, Nancy D., Leemore S. Dafny, Bruce E. Landon, Jesse B. Dalton, Ifedayo Kuye, and J. Michael McWilliams, "Changes in Quality of Care after Hospital Mergers and Acquisitions," *New England Journal of Medicine*, 2020, *382* (1), 51–59.

Benzarti, Youssef and Jarkko Harju, "Using Payroll Tax Variation to Unpack the Black Box of Firm-Level Production," *Journal of the European Economic Association*, 2021, *19* (5), 2737–2764.

Biden, Joseph R., "Executive Order on Promoting Competition in the American Economy," 2021.

Bíró, Anikó, Réka Branyiczki, Attila Lindner, Lili Márk, and Dániel Prinz, "Firm Heterogeneity and the Impact of Payroll Taxes," 2022. World Bank Policy Research Working Paper 10265.

Borusyak, Kirill and Peter Hull, "Non-Random Exposure to Exogenous Shocks," *Econometrica*, 2023, *91*, 2155–2185.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "A Practical Guide to Shift-Share Instruments," *Journal of Economic Perspectives*, 2025, *39* (1), 181–204.

Brand, Keith, Christopher Garmon, and Ted Rosenbaum, "In the Shadow of Antitrust Enforcement: Price Effects of Hospital Mergers from 2009-2016," *Journal of Law and Economics*, 2023, *66*, 639–669.

Brot-Goldberg, Zarek C., Amitabh Chandra, Benjamin R. Handel, and Jonathan T. Kolstad, "What Does a Deductible Do? The Impact of Cost-Sharing on Health Care Prices, Quantities, and Spending Dynamics," *Quarterly Journal of Economics*, 2017, *132*, 1261–1318.

Brot, Zarek, Zack Cooper, Stuart Craig, and Lev Klarnet, "Amalgamated Merger Database 2001-2020," 2024.

Brot, Zarek, Zack Cooper, Stuart V. Craig, and Lev Klarnet, "Is There Too Little Antitrust Enforcement in the US Hospital Sector?," *American Economic Review: Insights*, 2024, *6* (4), 526–42.

Bureau of Labor Statistics, "Consumer Price Index Database," 2022.

Capps, Cory, David Dranove, and Christopher Ody, "The Effects of Hospital Acquisitions of Physician Practices on Prices and Spending," *Journal of Health Economics*, 2018, *59*, 139–52.

Capps, Cory, David Dranove, and Mark Satterthwaite, "Competition and Market Power in Option Demand Markets," *RAND Journal of Economics*, 2003, pp. 737–763.

Capps, Cory, Laura Kmitch, Zenon Zabinski, and Slava Zayats, "The Continuing Saga of Hospital Merger Enforcement," *Antitrust Law Journal*, 2019, *82*, 441–496.

Case, Anne and Angus Deaton, *Deaths of Despair and the Future of Capitalism*, Princeton University Press, 2020.

Chung, Andrea P., Martin Gaynor, and Seth Richards-Shubik, "Subsidies and Structure: The Lasting Impact of the Hill-Burton Program on the Hospital Industry," *Review of Economics and Statistics*, 2017, *99*, 926–943.

Claxon, Gary, Matthew Rae, Gregory Young, Nisha Kurani, Heidi Witmore, Jason Kerns, Jackie Cifuentes, Greg Shmavonian, and Anthony Damico, "2022 Employer Health Benefits Survey," 2021.

Cooper, Zack, Fiona Scott Morton, and Nathan Shekita, "Surprise! Out-of-Network Billing for Emergency Care in the United States," *Journal of Political Economy*, 2020, *128*, 3226–77.

Cooper, Zack, Stuart Craig, Martin Gaynor, and John Van Reenen, "The Price Ain't Right? Hospital Prices and Health Spending on the Privately Insured," *Quarterly Journal of Economics*, 2019, *134* (1), 51–107.

Cooper, Zack, Stuart Craig, Martin Gaynor, Nir J. Harish, Harlan M. Krumholz, and John Van Reenen, "Hospital Prices Grew Substantially Faster than Physician Prices for Hospital-Based Care in 2007-14," *Health Affairs*, 2019, *38*, 184–89.

Craig, Stuart, "Competition in Employer-Sponsored Health Insurance: Implications for a Public Option," 2022.

Cutler, David M. and Brigitte C. Madrian, "Labor Market Responses to Rising Health Insurance Costs," *RAND Journal of Economics*, 1998, *29*, 509–530.

Dafny, Leemore, "Estimation and Identification of Merger Effects: An Application to Hospital Mergers," *Journal of Law and Economics*, 2009, *52* (3), 523–550.

Dafny, Leemore, David Dranove, Frank Limbrock, and Fiona Scott Morton, "Data impediments to empirical work on health insurance markets," *The BE Journal of Economic Analysis & Policy*, 2011, *11* (2), 1–24.

Dafny, Leemore, Kate Ho, and Edward Kong, "How Do Copayment Coupons Affect Branded Drug Prices and Quantities Purchased?," *American Economic Journal: Economic Policy*, 2024, *16* (3), 314–46.

Dafny, Leemore S., "How Do Hospitals Respond to Price Changes?," *American Economic Review*, 2005, *95* (5), 1525–1547.

Doyle, Joseph J., John A. Graves, Jonathan Gruber, and Samuel A. Kleiner, "Measuring Returns to Hospital Care: Evidence From Ambulance Referral Patterns," *Journal of Political Economy*, 2015, *123* (1), 170–214.

FactSet Research Systems, "FactSet Research Systems database," 2020. Accessed November 2021.

Finkelstein, Amy, Casey McQuillan, Owen Zidar, and Eric Zwick, "The Health Wedge and Labor Market Inequality," *Brookings Papers on Economic Activity*, 2023, *Spring*, 425–475.

Fisher, Linda T. and Mary B. Andersen, *5500 Preparer's Manual for 2018 Plan Years*, Wolters Kluwer, 2019.

Fulton, Brett D., "Health Care Market Concentration Trends in the United States: Evidence and Policy Responses," *Health Affairs*, 2017, *36* (9), 1530–1538.

Gao, Janet, Shan Ge, Lawrence D.W. Schmidt, and Cristina Tello-Trillo, "How Do Health Insurance Costs Affect Low- and High-Income Workers?," 2025.

Garmon, Christopher, "The Accuracy of Hospital Merger Screening Methods," *RAND Journal of Economics*, 2017, *48* (4), 1068–1102.

Gaynor, Martin, Kate Ho, and Robert J. Town, "The Industrial Organization of Health Care Markets," *Journal of Economic Literature*, 2015, *53* (2), 235–84.

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

Gruber, Jonathan, "The Incidence of Mandated Maternity Benefits," *American Economic Review*, 1994, pp. 622–641.

Gruber, Jonathan, "The Incidence of Payroll Taxation: Evidence from Chile," *Journal of Labor Economics*, 1997, *15* (3), S72–S101.

Gruber, Jonathan and Alan B. Krueger, "The Incidence of Mandated Employer-Provided Insurance: Lessons from Workers' Compensation Insurance," *Tax Policy and the Economy*, 1991, *5*, 111–143.

Guo, Audrey, "Payroll Tax Incidence: Evidence from Unemployment Insurance," *Journal of Public Economics*, 2024, *239*, 105209.

Haas-Wilson, Deborah and Christopher Garmon, "Hospital Mergers and Competitive Effects: Two Retrospective Analyses," *International Journal of the Economics of Business*, 2011, *18* (1), 17–32.

Handel, Ben and Kate Ho, "The Industrial Organization of Health Care Markets," in Kate Ho, Ali Hortaçsu, and Alessandro Lizzeri, eds., *Handbook of Industrial Organization*, Vol. 5, Elsevier, 2021, pp. 521–614.

Harberger, Arnold C., "Monopoly and Resource Allocation," *American Economic Review*, 1954, *44* (2), 77–87.

Health Care Cost Institute, “2019 Health Care Cost and Utilization Report,” 2020.

Hegland, Thomas A., “How do Multistate Firms Design their Compensation Packages When Local Conditions Vary? Evidence from Interstate Spillover Effects of the Affordable Care Act Medicaid Expansions,” 2025.

Ho, Kate and Robin S. Lee, “Insurer Competition in Health Care Markets,” *Econometrica*, 2017, *85* (2), 379–417.

Johnston, Andrew C., “Unemployment Insurance Taxes and Labor Demand: Quasi-Experimental Evidence from Administrative Data,” *American Economic Journal: Economic Policy*, 2021, *13* (1), 266–293.

Kaiser Family Foundation, “Health Insurance Coverage of the Total Population - State Health Facts,” 2019.

Keisler-Starkey, Katherine and Lisa N. Bunch, “Health Insurance Coverage in the United States: 2021 - Current Population Reports,” 2022.

Kolstad, Jonathan T. and Amanda E. Kowalski, “Mandate-Based Health Reform and the Labor Market: Evidence from the Massachusetts Reform,” *Journal of Health Economics*, 2016, *47*, 81–106.

Lieber, Ethan M.J., “Does it pay to know prices in health care?,” *American Economic Journal: Economic Policy*, 2017, *9* (1), 154–79.

Lin, Haizhen and Yaying Zhou, “Distributional effect of insurer concentration on premiums: Evidence from employer-sponsored health insurance,” 2024.

Lin, Haizhen, Ian M. McCarthy, and Michael Richards, “Hospital Pricing Following Integration with Physician Practices,” *Journal of Health Economics*, 2021, *77.*

Lobel, Felipe, “Who Benefits from Payroll Tax Cuts? Market Power, Tax Incidence and Efficiency,” 2024.

Lurie, Ithai Z. and Corbin L. Miller, “Employer-Sponsored Health Insurance Premiums and Income in US Tax Data,” *Journal of Public Economics*, 2023, *224*, 104942.

Min, Jessica, “Causes and Consequences of Rising Employer-Sponsored Health Insurance Costs: Evidence from Insurer Mergers,” 2025.

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

Raval, Devesh, Ted Rosenbaum, and Steven A. Tenn, "A Semiparametric Discrete Choice Model: An Application to Hospital Mergers," *Economic Inquiry*, 2017, *55* (4), 1919–1944.

Saez, Emmanuel and Gabriel Zucman, *The Triumph of Injustice: How the Rich Dodge Taxes and How to Make Them Pay*, WW Norton & Company, 2019.

Saez, Emmanuel, Benjamin Schoefer, and David Seim, "Payroll Taxes, Firm Behavior, and Rent Sharing: Evidence from a Young Workers' Tax Cut in Sweden," *American Economic Review*, 2019, *109* (5), 1717–1763.

Securities Data Company Platinum, "Securities Data Company Platinum database," https://data.bls.gov/cew/doc/layouts/csv_quarterly_layout.htm 2020. Accessed December 2021.

Starr, Paul, *The Social Transformation of American Medicine: The Rise of a Sovereign Profession and the Making of a Vast Industry*, Basic Books, 1982.

Summers, Lawrence, "Some Simple Economics of Mandated Benefits," *American Economic Review*, 1989, *79* (2), 177–83.

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

United States Census Bureau (Census), "Census Bureau's Small Area Health Insurance Estimates," https://www.census.gov/data/datasets/time-series/demo/sahie/estimates-acs.html 2017. Accessed November 2021.

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

U.S. Department of Justice and the Federal Trade Commission, "Horizontal Merger Guidelines," 2010.

Zabinski, Zenon, "Hospital Market Structure and Access to Medical Care: The Long-Term Impact of the Hill-Burton Act," 2015.

Table 1: Employer-Level and County-Level Summary Statistics

Panel A: Employer Characteristics

	Mean (1)	SD (2)	P25 (3)	P50 (4)	P75 (5)	N (6)
Employer Total Payroll ($/year)*	12,721,000	25,129,000	2,391,000	4,977,000	11,418,000	140,301
Employer Count of Workers	298	509	76	133	282	140,301
Employer Average Wages per Worker ($/year)	41,340	25,431	23,408	36,526	52,810	140,301
Share of Employees with Premiums	0.51	0.31	0.21	0.60	0.77	39,341
Share of Employees with a Health Savings Account	0.04	0.12	0.00	0.00	0.00	140,301
Health Spending per Beneficiary ($/year)	4,099	705	3,649	4,039	4,478	140,301
Premiums from 5500 Data ($/year)	5,037	1,574	3,943	4,930	6,001	3,970

Panel B: County Characteristics

	Mean (1)	SD (2)	P25 (3)	P50 (4)	P75 (5)	N (6)
Income Per Capita ($/year)	42,050	9,218	36,021	40,055	45,507	1,709
Share with Unemployment Insurance	0.04	0.02	0.02	0.03	0.04	1,709
Share with Zero Income	0.05	0.01	0.04	0.05	0.06	1,709
Share Unemployed	0.09	0.02	0.07	0.09	0.10	1,709
Unemployment Insurance Payments per Capita ($/year)	482	395	186	363	665	1,709
Share Self-Employed	0.12	0.03	0.10	0.11	0.13	1,709
Share Moving Annually	0.07	0.02	0.05	0.06	0.08	1,709
Income Tax Withholdings per Capita ($/year)	7,009	2,088	5,634	6,584	7,836	1,709
Health Spending Per Beneficiary ($/year)	4,211	445	3,910	4,182	4,480	1,709

Notes: This table presents employer-level and county-level descriptive statistics for our main analytic samples, from 2008 to 2017. In Panel A, employer payroll, employer counts of workers, employer wages, the share of employees with premiums, and the share of employees with a health savings account come from the Internal Revenue Service (IRS). Data on health spending per beneficiary come from the Health Care Cost Institute (HCCI). Data on insurance premiums comes from the Department of Labor's 5500 forms. In Panel B, income per capita, the share with unemployment insurance, unemployment insurance payments per capita, the share of the population self-employed, the share of the population moving out of the county annually, and income tax withholdings per capita come from IRS returns. Income per capita is measured as the sum of wage (W-2) income and self-employment (Schedule SE) income. We define the share unemployed as the share of individuals with either positive unemployment insurance receipts and/or with zero income in the year.
* Rounded to $1,000 to preserve privacy.

Table 2: First Stage: Regressing Annual Health Care Spending on Simulated Health Care Spending

	Log(Health Spending per Beneficiary)		
	(1)	(2)	(3)
Simulated Spending	0.64***	0.51***	2.90***
	(0.02)	(0.02)	(0.05)
Instrument	Baseline	Using Only Price Change Variation	Using Predicted Price Changes
Mean Dependent Variable	4,099	4,099	4,099
Observations	1,403,010	1,403,010	1,403,010
Number of Unique Employers	140,301	140,301	140,301
F-Statistic on First Stage	844	509	2,932

Notes: This table presents coefficient estimates from a regression of employer-level annual health spending per beneficiary on employer-level simulated spending per beneficiary, as given in Equation (5). Each estimate includes employer and year fixed effects. Each column presents estimates from a regression using a different instrument. Column (1) presents estimates using our baseline instrument. Column (2) presents estimates using a modified version of our baseline instrument that purges any variation other than that coming from differences in post-merger price changes across mergers. Column (3) presents estimates using a modified version of our baseline instrument that replaces the estimated post-merger price increases with the estimated post-merger change in WTP. Data on health spending and simulated spending come from the Health Care Cost Institute. Means are reported in levels rather than in logs. Standard errors are reported in parentheses and are clustered at the employer level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 3: The Impact of Rising Health Care Prices on Health Insurance Market Outcomes

Panel A: OLS Estimates

	Log(Insurance Premiums) (1)	Share of Employees with a Health Savings Account (2)
Log(Spending per Beneficiary)	0.03	0.01***
	(0.03)	(0.00)

Panel B: IV Estimates

	Log(Insurance Premiums) (1)	Share of Employees with a Health Savings Account (2)
Log(Spending per Beneficiary)	0.95*	–0.01
(Instrumented using Merger-Driven Price Increases)	(0.54)	(0.04)
Mean Dependent Variable	5,036	0.04
Observations	39,700	1,403,010
Number of Unique Employers	3,970	140,301
F-Statistic on First Stage	44	844

Notes: This table presents ordinary least squares (Panel A) and instrumental variables (Panel B) estimates from regressions of annual employer-level log health insurance premiums per enrollee (Column (1)) and the share of employees with contributions to a health savings account (Column (2)) on employer-level annual spending per beneficiary, instrumenting for annual spending per beneficiary with employer-level simulated spending per beneficiary. Each estimate includes employer and year fixed effects. Data on insurance premiums come from the Department of Labor Form 5500 filings. Data on an employer's share of enrollees with a health savings account come from the Internal Revenue Service. Means are reported in levels rather than logs. Standard errors are in parentheses and are clustered at the employer level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 4: The Impact of Rising Health Care Prices on Employer Payroll and Employment

Panel A: OLS Estimates

	All Employers		Non-Health Care Employers		Health Care Employers	
	Log(Payroll) (1)	Log(Workers) (2)	Log(Payroll) (3)	Log(Workers) (4)	Log(Payroll) (5)	Log(Workers) (6)
Log(Spending per Beneficiary)	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01	< 0.01
	(< 0.01)	(< 0.01)	(< 0.01)	(< 0.01)	(0.02)	(0.02)

Panel B: IV Estimates

	All Employers		Non-Health Care Employers		Health Care Employers	
	Log(Payroll) (1)	Log(Workers) (2)	Log(Payroll) (3)	Log(Workers) (4)	Log(Payroll) (5)	Log(Workers) (6)
Log(Spending per Beneficiary) *(Instrumented using Merger-Driven Price Increases)*	–0.36***	–0.36***	–0.37***	–0.40***	–0.30	0.23
	(0.13)	(0.13)	(0.14)	(0.13)	(0.56)	(0.55)
Mean Dependent Variable*	12,721,000	298	13,045,000	304	8,242,000	204
Observations	1,403,010	1,403,010	1,308,300	1,308,300	94,710	94,710
Number of Unique Employers	140,301	140,301	130,830	130,830	9,471	9,471
F-Statistic on First Stage	844	844	793	793	52	52

Notes: This table presents ordinary least squares (Panel A) and instrumental variables (Panel B) estimates from regressions of annual employer-level log payroll (Columns (1), (3), (5)) and log worker counts (Columns (2), (4), (6)) on employer-level annual spending per beneficiary, instrumenting for annual spending per beneficiary with employer-level simulated spending per beneficiary. In Columns (1) and (2), we include all employers. In Columns (3) through (6), we include only those employers categorized as not being in the health care industry (Columns (3) and (4)) or being in the health care industry (Columns (5) and (6)), as determined by their reported NAICS code. Each estimate includes employer and year fixed effects. Our labor market data come from the Internal Revenue Service. Means are reported in levels rather than logs. Standard errors are in parentheses and are clustered at the employer level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.
* Rounded to $1,000 for preserve privacy.

Table 5: The Impact of Rising Health Care Prices on County-Level Labor Income and Employment

Panel A: OLS Estimates

	Log(Income per Capita) (1)	Share Unemployed (2)	Log(Income Tax per Capita) (3)
Log(Spending per Beneficiary)	0.02*	0.01***	0.02
	(0.01)	(0.00)	(0.01)

Panel B: IV Estimates

	Log(Income per Capita) (1)	Share Unemployed (2)	Log(Income Tax per Capita) (3)
Log(Spending per Beneficiary)	–0.28*	0.09**	–0.37*
(Instrumented using Merger-Driven Price Increases)	(0.15)	(0.04)	(0.19)
Mean Dependent Variable	42,050	0.09	7,009
Observations	17,090	17,090	17,090
Number of Unique Counties	1,709	1,709	1,709
F-Statistic on First Stage	41	41	41

Notes: This table presents ordinary least squares (Panel A) and instrumental variables (Panel B) coefficient estimates from regressions of annual county-level log income per capita (Column (1)), share of the population collecting unemployment insurance or earning zero labor income (Column (2)), and log federal income tax receipts per capita (Column (3)) on county-level annual spending per beneficiary, instrumenting for annual spending per beneficiary with county-level simulated spending per beneficiary. Each estimate includes county and year fixed effects. Our labor market and tax revenue data come from the Internal Revenue Service. Means are reported in levels rather than logs. Standard errors are in parentheses and are clustered at the county level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure 1: The Impact of Hospital Mergers in Hospital Markets, Insurance Markets, and Labor Markets

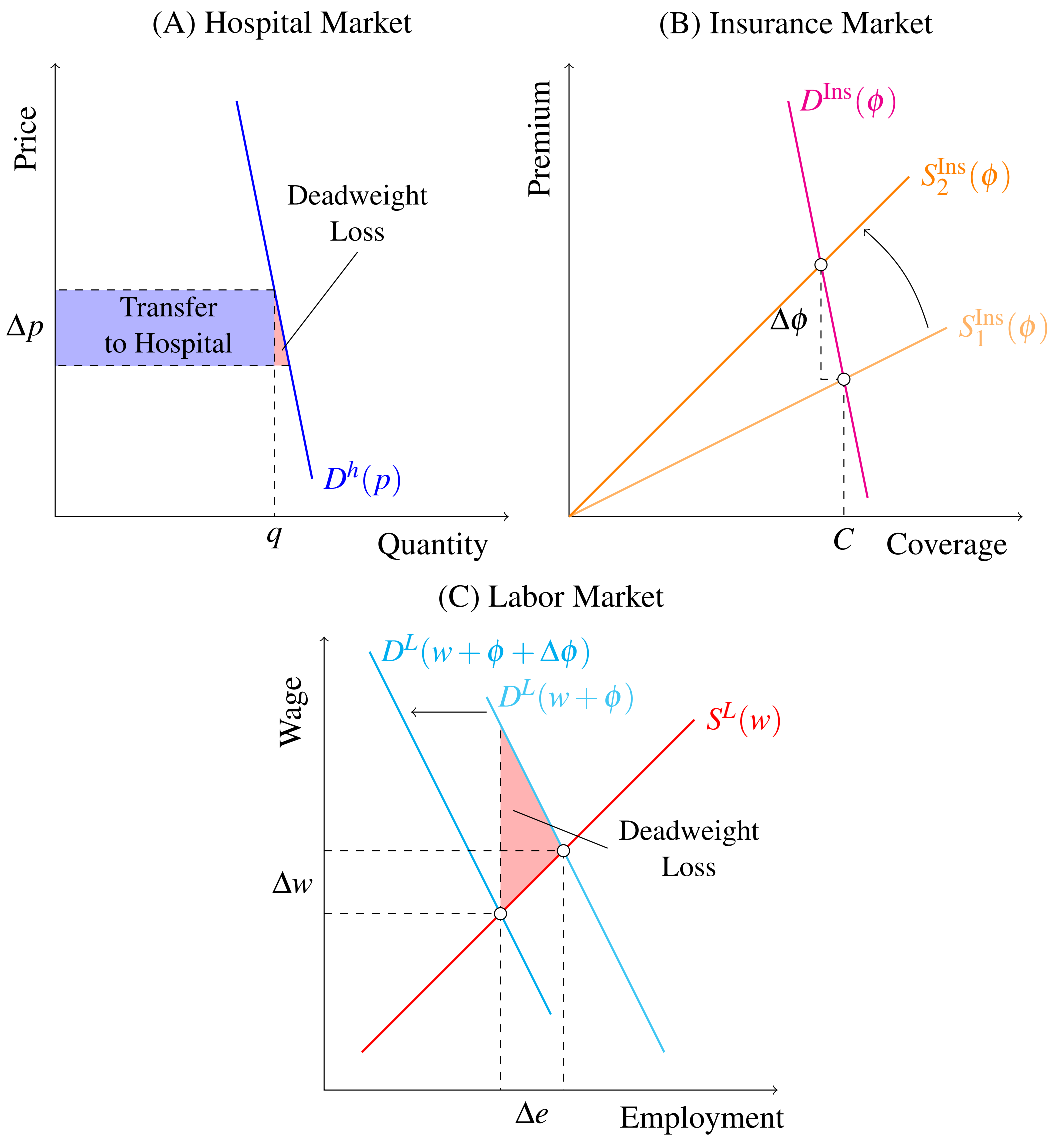

Notes: This illustrates our theory of how rising health care prices impacts insurance premiums and downstream changes in labor market outcomes. In Panel A, we highlight the fact that, after a merger (or other source of rent-seeking), prices rise. This generates deadweight loss (the red triangle), and a transfer from payers to the merging hospital (the blue rectangle). In Panel B, we highlight the effect on the market for health insurance coverage. The insurance supply curve rotates around the origin, from $S_1^{Ins}(\phi)$ to $S_2^{Ins}(\phi)$. In Panel C, we highlight the effects on the market for labor. The insurance premium increase shifts labor demand down by $\Delta\phi$, the change in premiums. This results in a fall in equilibrium wages and employment, Δw and ΔL. It also results in deadweight loss, given by the red triangle.

Figure 2: The Impact of Mergers on Hospital Prices by FTC Reporting Guidelines

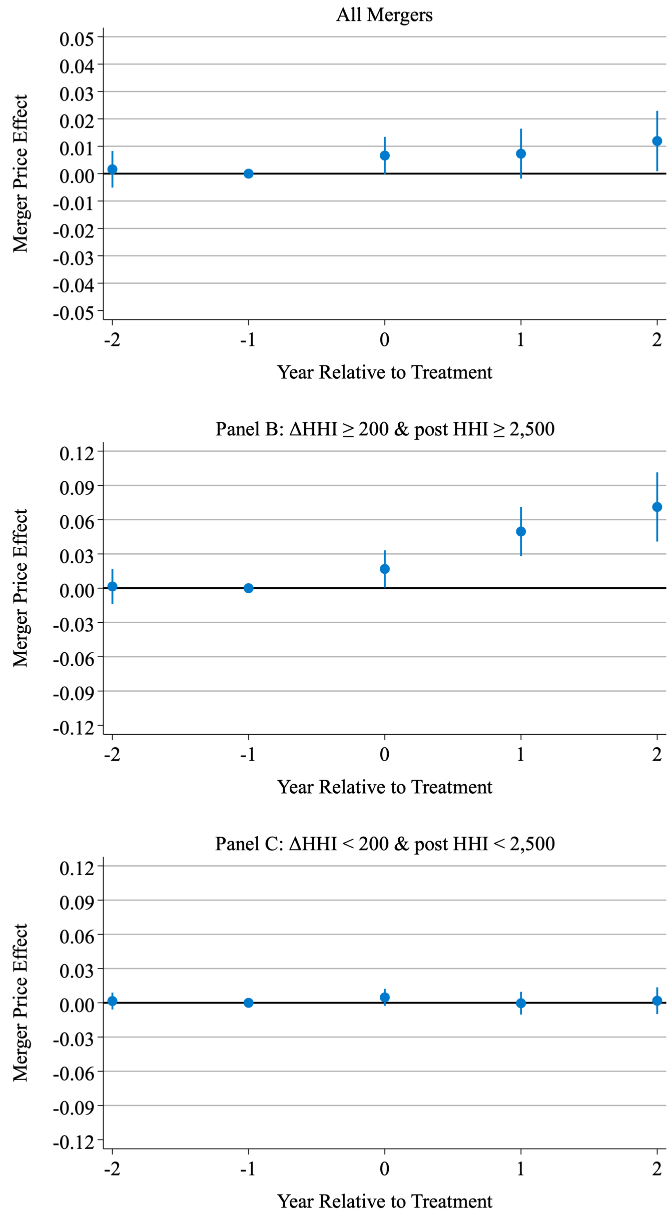

Notes: This figure presents difference-in-difference estimates of the effect of 2010 to 2015 hospital mergers on hospital prices. We estimate a pooled version of Equation (2) that estimates the average price increase within a group of mergers. We do so for the following groups: all mergers (Panel A), mergers that generated a $\Delta HHI \geq 200$ and post-merger $HHI \geq 2,500$ (Panel B) and mergers that generated a $\Delta HHI < 200$ or post-merger $HHI < 2,500$ (Panel C). Hospital prices are a weighted average of inpatient and outpatient prices. The weights are constructed as the average share of inpatient and outpatient revenue for the treated hospital in 2008 and 2009, the two years prior to the first merger in our sample. Each dot represents a point estimate, and the vertical line displays the corresponding 95% confidence interval. Standard errors are clustered at the hospital level. Hospital pricing data come from the Health Care Cost Institute. The average difference-in-differences estimates comparing the two years prior to the merger with two years post for mergers in Panel A is 0.009 (0.005), in Panel B is 0.059 (0.012), and in Panel C is -0.000 (0.005).

Figure 3: Event Study Estimates of First Stage: Regressing Employer Spending on Simulated Employer Spending

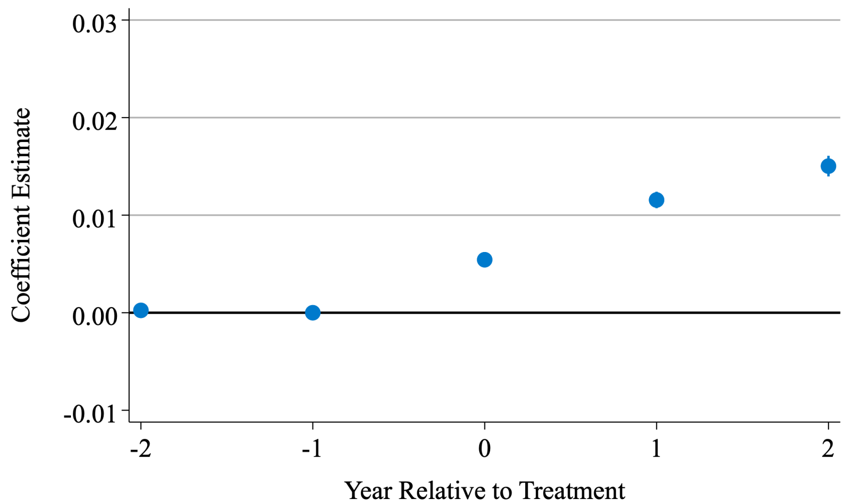

Notes: This figure presents event-study estimates from our first-stage regression using data from 2008 to 2017. “Treated” employers are in the top quartile of the cumulative change in simulated spending between the first and last years of the sample, and “control” employers are those in the bottom quartile. For each treated employer, the treatment year is the year in which it experienced its largest single-year increase in simulated spending. Dots denote point estimates; vertical lines denote 95% confidence intervals.

Figure 4: Event Study Estimates of the Impact of Rising Health Care Prices on Employer Payroll and Worker Count at Non-Health Care Employers

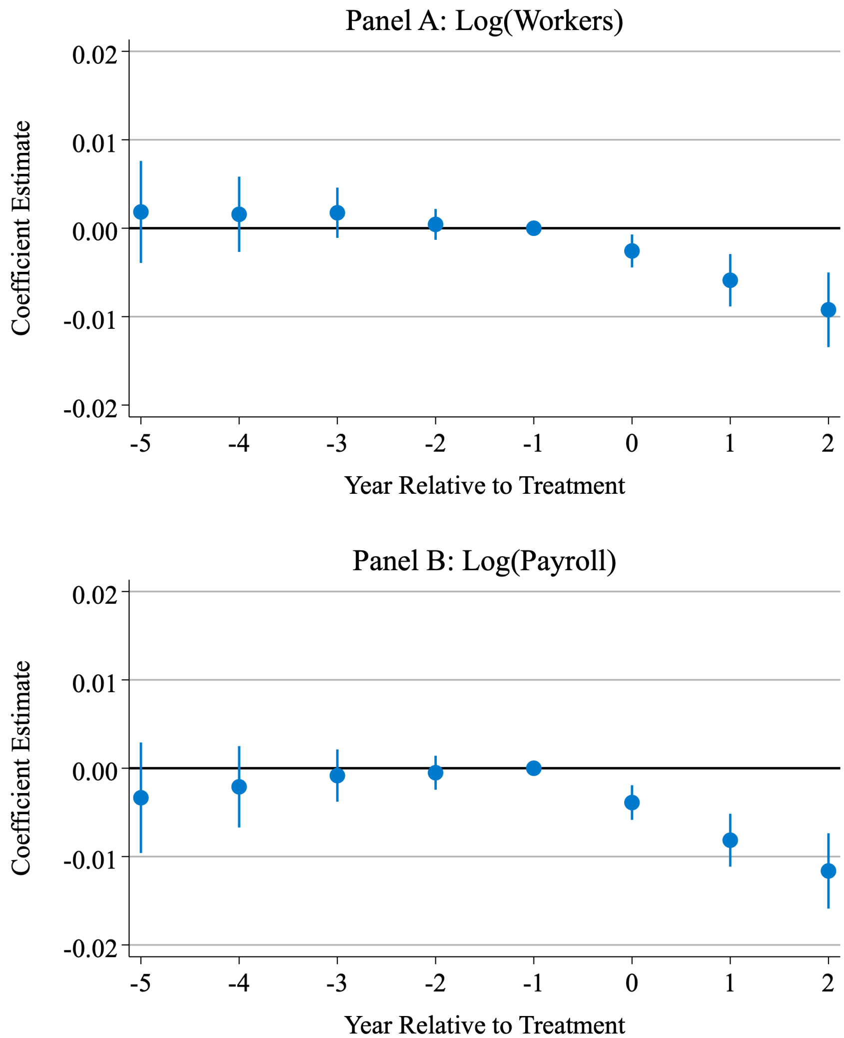

Notes: This figure presents event-study estimates of the impact of rising health-care prices on employer labor-market outcomes using data from 2005 to 2017. "Treated" employers are in the top quartile of the cumulative change in simulated spending between the first and last years of the sample, and "control" employers are those in the bottom quartile. For each treated employers, the treatment year is the year in which it experienced its largest single-year increase in simulated spending. Dots denote point estimates; vertical lines denote 95% confidence intervals.

Figure 5: Event Study Estimates of the Impact of Rising Health Care Prices on County-Level Income per Capita and Employment

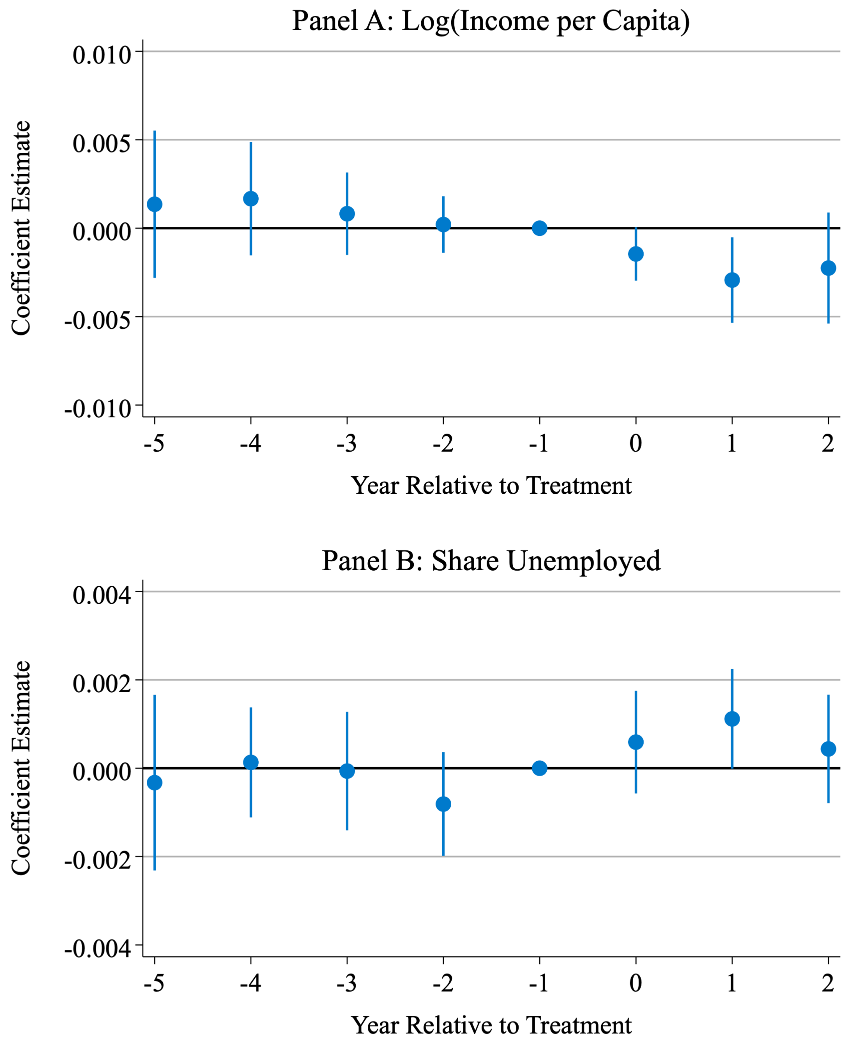

Notes: This figure presents event-study estimates of the impact of rising health-care prices on county labor-market outcomes from 2005 to 2017. “Treated” counties are in the top quartile of the cumulative change in simulated spending between the first and last years of the sample, and “control” counties are those in the bottom quartile. For each treated county, the treatment year is the year in which it experienced its largest single-year increase in simulated spending. Dots denote point estimates; vertical lines denote 95% confidence intervals.

Figure 6: The Impact of Rising Health Care Prices on Changes in Unemployment across the Income Distribution

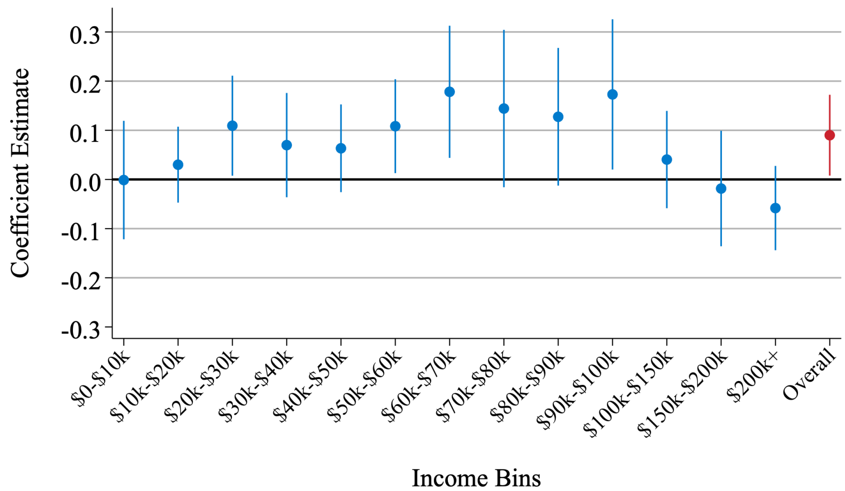

Notes: This figure shows estimates of Equation (10) on the share of the population collecting unemployment insurance. Effects are estimated separately for individuals based on bins of their income measured in the prior year. The dots represent point estimates and vertical lines are 95% confidence intervals.

ONLINE APPENDICES

A Measuring Prices in the HCCI Data

Following Cooper et al. (2019a), we take the approach of constructing a sample of patient visits, or "cases." For each case, we observe the negotiated transaction price. We then use clinical codes indicating the procedure performed during a case and the severity of a patient's illness, along with demographic characteristics of the patient to adjust for the mix of services provided by each hospital. Specifically, we estimate a regression of the form:

$$ln(p_{idht}) = \alpha_{ht} + \beta X_i + \pi_{dt} + \varepsilon_{idht}. \tag{11}$$

X_i includes controls for each patient's age and gender. π_{dt} is a year-specific fixed-effect, which flexibly adjusts for the complexity of a patients' condition and treatment by year.

We then use the estimates from Equation (11) to generate predicted values for each hospital-year, holding fixed the coefficients accounting for patient characteristics and clinical severity at the average levels in the data:

$$p_{ht}^{INDEX} = \hat{\alpha}_{ht} + \hat{\beta}\overline{X} + \hat{\pi}_{dt}\overline{dt}. \tag{12}$$

We generate price indices for both inpatient and outpatient samples separately. For the inpatient sample, we adhere closely to the methodology described in Cooper et al. (2019a), limiting the sample to individuals who are between 18 and 65 years old. We also limit our sample to individuals with a valid Diagnostic-Related Group (DRG) code, which we use to define the π_{dt} fixed effect.

For the outpatient sample, we apply similar restrictions, limiting our sample to patients between 18 and 65 years old. To ensure the prices we measure are complete — not payments negotiated as a bundle of services — we limit our analysis to patient days in which there is only one outpatient visit. We also limit our analysis to patient days in which there is only one CPT procedure code that maps to a valid Medicare APC payment rate. Although this restriction limits the data to approximately 30% of patient days, we view this sample as one that provides a clean distinction between price and quantity. We use APC codes to define the π_{dt} fixed effects for the outpatient price index.

For both samples, we restrict our analysis to the subset of hospitals that match to an acute general surgical hospital in the roster of hospitals we derive from the American Hospital Association's Annual Survey Data.

B Measuring Merger Activity and the Price Effects of Mergers

This appendix describes how we construct the panel of mergers we use, as well as how we measure potential competitive effects for mergers. Much of this material is reproduced from Brot et al. (2024b).

B.1 Merger Panel

The primary data we use to measure merger activity come from the American Hospital Association's (AHA's) Annual Survey of Hospitals. These data contain biographical information on the near universe of general acute care hospitals in the US, including a measure of system ownership. Our final roster contains 4,846 hospitals in the continental US. We track mergers in our hospital panel using changes to the system identifier provided by the AHA for 2002 to 2020. We leverage several additional data sources — the FactSet Research Systems database, the Irving Levin Associates' Health Care Services Acquisition Reports, and Securities Data Company Platinum — to verify the existence and timing of mergers.[44] We observe 1,164 mergers of general acute care hospitals from 2002-2020. Of these, we use 305 in our analysis, that (1) are consummated between 2010 and 2015; (2) have at least two parties within 50 miles of each other and (3) have at least one case in the HCCI data during both the two years before the merger and the two years after. Appendix Figure A.1 includes a map of the mergers we focus on in this analysis.

B.2 Matching Merging and Non-Merging Hospitals

In order to identify the treatment effect of mergers on hospital prices, we need a control group of non-merging hospitals to estimate counterfactual trends in prices for merging hospitals. Our control group is composed of hospitals that did not experience a merger between 2008 and $t+2$, where t is the year that the focal merging hospital merged.

To ensure that our control hospitals represent plausible counterfactual trends, we use propensity scores to match comparison hospitals to merging hospitals on pre-merger observable characteristics. We estimate a probit regression of the form:

$$\mathbb{P}\{h \text{ merged}\} = X'\beta + \varepsilon_h \tag{13}$$

where X contains a vector of hospital characteristics — drawn from the AHA data and measured in the year before our first merger (2009) — that may meaningfully determine price trends at hospitals: total number of hospital beds; total inpatient admissions; full time equivalents; number of unique technologies; share of Medicare patients; share of Medicaid patients; whether the hospital is a teaching hospital; a non-profit hospital; or a government hospital; the distance to the hospital's nearest competitor; the distance to the hospital's nearest hospital in its system or not; and whether the hospital is independent or part of a system. X also includes measures of local area characteristics around the hospital. First, we include the local HHI, measured as described in Appendix B.3. Second, we include the share of the hospital's county covered by private insurance, which we construct using data from the Census's Small Area Health Insurance Estimates (SAHIE).

[44] For more information on how we use these datasets to track hospital ownership, see Appendix D of Cooper et al. (2019a).

Finally, we include the share of the county insured by HCCI payors specifically, using data from HCCI to form the numerator and data from SAHIE to form the denominator.

We then use this model to construct propensity scores for each hospital. For each merging hospital, we find its 25 nearest neighbors (in terms of propensity scores) from the set of potential control hospitals. We also impose a caliper restriction so that the propensity scores of matched controls must be within 20% of a standard deviation from the treated merging hospital, even if this requires that the control group contain fewer than 25 hospitals.

B.3 Measuring the Herfindahl-Hirschman Index

To measure the Herfindahl-Hirschman Index (HHI), assume that a market M includes many hospital systems $S \in \mathscr{S}(M)$, where $\mathscr{S}(M)$ is the set of systems in M. Each system is defined as a set of one or more hospitals h, which have a collective owner. Formally,

$$HHI_M = 10,000 \times \sum_{S \in \mathscr{S}(M)} \left(\sum_{h \in S} s_{hM} \right)^2$$

where s_{hM} is h's market share within M. A monopoly market has an HHI of 10,000; if instead there are many small independent hospitals, the HHI will be closer to 0.

Measuring HHIs requires us to define relevant geographic markets and measure hospitals' market shares. We assume that a hospital's relevant market includes every hospital within a 30-minute drive time from their facility. We measure a hospital's market share as its share of inpatient hospital beds. We use hospital beds rather than activity to define concentration because, unlike hospital activity, changes in bed volume in the short run are unlikely to be highly correlated with changes in hospital quality or prices. We measure the change in HHI for a hospital h due to merger e as the difference between the HHI in its market in the year before the merger and a computed counterfactual where we change system membership to reflect the merger, holding bed counts and the system membership of non-participating hospitals fixed.

B.4 Measuring Willingness-to-Pay

Theory predicts that the extent to which mergers raise prices depends on the extent to which merging hospitals are good substitutes for one another, and whether their patients are unwilling to go to another hospital. As pointed out by Capps et al. (2003) and Gowrisankaran et al. (2015), since demand for hospitals is very inelastic, a standard model of Nash-Bertrand pricing would predict extremely high prices following mergers and suggest mergers could raise prices by implausibly large amounts. Instead, these prior studies have developed a theory of price-setting in which prices are bilaterally negotiated between hospital systems and insurers, which bargain on behalf of their enrollees. In these models, prices are not determined by patients' price elasticities but are instead driven by what is effectively the insurer's elasticity — in terms of how much insurers can subsequently raise premiums if the hospital system is included in the insurer's preferred network of hospitals. In this way, hospital prices are determined by patients' *ex ante* willingness to pay for the option to go to the hospital when buying an insurance plan (Ho and Lee, 2017).

B.4.1 Hospital–Insurer Bargaining and ΔWTP

Capps et al. (2003) model the *ex post* utility of patient i at hospital h as $U_{ih} = U(X_{ih}) + \varepsilon_{ih}$, where $U(\cdot)$ denotes expected utility at the hospital and ε_{ih} represents idiosyncratic patient preferences at specific hospitals with ε distributed i.i.d. standard Gumbel. X_{ih} contains patient and hospital characteristics that determine preferences for a given hospital, including the patient's specific health care needs as well as the distance between them and the hospital.

If a patient faces a hospital network $\mathcal{N}$ that limits what hospitals she has access to, then the patient's *ex ante* expected utility of access to a network $\mathcal{N}$ is:

$$\begin{aligned} EU_i(\mathcal{N}) &= E[\max_{j\in\mathcal{N}} U_{ij}] \\ &= ln\left(\sum_{j\in\mathcal{N}} \exp(U_{ij})\right) \end{aligned}$$

Moreover, say that a hospital h is dropped from the network. Capps et al. (2003) show that the change in expected utility as a result of this network change is:

$$\begin{aligned} \Delta EU_{ih} &= EU_i(\mathcal{N}) - EU_i(\mathcal{N}\setminus h) \\ &= ln\left(\sum_{j\in\mathcal{N}} \exp(U_{ij})\right) - ln\left(\sum_{j\in\mathcal{N}\setminus h} \exp(U_{ij})\right) \\ &= ln\left(\frac{1}{1-s_{ih}}\right) \end{aligned}$$

where s_{ih} is that hospital's expected market share from patient i under network $\mathcal{N}$. If consumers are always indifferent between receiving a 1-point increase in EU and a $\$\gamma_i$ payment, then we can describe patients' *ex ante* "willingness-to-pay' for hospital h as $W_{ih} = \gamma_i \Delta EU_{ih}$. We integrate over the distribution of consumers F_i to calculate market-level willingness to pay as $W_h = \int_i \gamma_i ln\left(\frac{1}{1-s_{ih}}\right) dF_i$. Where W_h represents the amount that the average consumer is willing to pay for access to hospital h. Both Capps et al. (2003) and Gowrisankaran et al. (2015) show that, in standard models of bargaining (either pure Nash or Nash-in-Nash), the price for h's services that will be negotiated jointly by the hospital and insurer is proportional to W_h.

The above notation assumes that all hospitals are independent. If, instead, hospitals are part of some system S, the hospitals will bargain jointly. That is, prices will be determined by the willingness to pay for the *entire system*, $W_S = \int_i \gamma_i ln\left(\frac{1}{1-s_{iS}}\right) dF_i$, with $s_{iS} = \sum_{j\in S} s_{ij}$. Systems are able to exert greater leverage than individual hospitals because they can threaten to hold out the entire system from the insurer's network if a deal on prices fails to be realized.[45]

We model the case of a merger (m) between two hospitals h and h'.[46] The impact of the merger

[45] In practice, we consider the relevant bargaining entity to be the only the hospitals in a system within a given HRR, to avoid diffusing local changes in bargaining leverage over large acquiring systems.

[46] This is without a loss of generality and can be replaced with systems.

on the bargaining leverage of the two hospitals is the difference between the willingness to pay of the merged system and the sum of the willingness to pay for h and h' individually. Due to a lack of data on individual insurance take-up, we follow Capps et al. (2003) and assume that $\gamma_i = \gamma$ for all patients. The percent change in willingness to pay due to the merger is:

$$\Delta WTP_m = \frac{\int_i \left[ln\left(\frac{1}{1-(s_{ih}+s_{ih'})}\right) - \left(ln\left(\frac{1}{1-s_{ih}}\right) + ln\left(\frac{1}{1-s_{ih'}}\right)\right)\right] dF_i}{\int_i \left[ln\left(\frac{1}{1-s_{ih}}\right) + ln\left(\frac{1}{1-s_{ih'}}\right)\right] dF_i}$$

where γ_i drops out of the equation under the assumption of homogeneity. Importantly, we focus on the joint complementarities created by a merger, excluding the effect of pure scale increases for each participating hospital. In unpublished results, we found that, particularly for mergers in which a single independent hospital was acquired by a large chain, allowing the scale effects to enter into the change in WTP predicted implausibly (and incorrectly) large post-merger price increases.

Under these assumptions, the potential price changes due to a merger should be proportional to ΔWTP_m.

B.4.2 Estimating Demand for Hospitals

Measuring ΔWTP_m requires us to estimate substitution patterns in the relevant market. Capps et al. (2003) underscore the importance of patient heterogeneity in this calculation — heart attack patients may care much more about hospital closeness than patients undergoing elective surgeries.

We therefore take the semiparametric approach to demand estimation developed by Raval et al. (2017). That is, we estimate $U(X_{ih})$ by assuming we can partition patients into groups $g \in G$ based on their characteristics, such that:

$$U_{ih} = U_{g(i)h} = \delta_{g(i)h} + \varepsilon_{ih}$$

Patients within the same groups are assumed to have the same *ex ante* expected utility for any particular hospital, but patients across groups may have different preferences in an unrestricted way. It is then true that, for patients within the same group, expected market shares at each hospital are equal within groups, such that:

$$s_{ih} = s_{g(i)h} = \frac{\exp(\delta_{g(i)h})}{\sum_{j \in \mathcal{N}} \exp(\delta_{g(i)j})}$$

Using this procedure, a valid partition of patients allows us to use the observed group-level market shares as an equivalent measure to individual-specific choice probabilities, and therefore patient utility for each hospital-by-group pair.

We calculate group-specific market shares for each hospital using every inpatient hospitalization for Medicare patients (in the group) during our relevant time period.[47] We exclude any hospitalization in which a patient attended a hospital more than 100 miles from their home. The Raval et al. (2017) approach provides an algorithm that partitions patients into increasingly small groups until the resulting groups are no smaller than S_{min}. This minimum group size parameter is

[47] That is, we assume that there is no relevant extensive margin substitution to no hospitalization as a result of changes in market structure.

set to balance a bias-variance trade-off: allowing for smaller groups reduces bias by allowing us to capture consumers' heterogeneous preferences for hospitals. However, smaller bins also increase variance by estimating preferences over smaller samples, where market shares may be estimated with error.

The algorithm proceeds as follows:

Step 1: The econometrician first establishes a set of discrete patient characteristics, ordered by "importance." Specifically, we group according to the following characteristics:

1. Patient home county
2. Patient home 5-digit zip code
3. Major Diagnostic Category of the patient's illness
4. Binary indicator for whether the patient's illness was such that the hospitalization was likely to be discretionary (rather than an emergency)
5. Binary indicator for whether the patient's illness was likely to require a surgical treatment (rather than a purely medical treatment)
6. Quartiles of the weight placed on the Diagnosis-Related Group for the patient's illness[48]
7. The Diagnosis-Related Group for a patient's illness (as measured by their primary diagnosis code)
8. Patient age, in 10-year buckets
9. Patient sex

Step 2: We partition patients into groups based on their unique values for every characteristic (e.g., if the characteristics are gender, race, and county, there will be one group for black female patients in New York County, another group for white male patients in Cook County, etc.).

Step 3: We assign groups based on any partitions from Step 2, as long as the partition has a size above S_{min}. Any patients in partitions with size below S_{min} are left ungrouped.

Step 4: We then disregard the lowest-priority characteristic.

Step 5: We repeat Steps 2-4 until we reach a single characteristic (the patient's home county).

The partitions this algorithm produces vary in granularity to allow for more heterogeneity among patients characteristics when larger sample sizes are available. For example, denser counties will have more groups, subdivided by illness and patient demographics. By contrast all patients will be grouped together in smaller counties where data are sparser.

We run the algorithm separately for each year of mergers in our data. To ensure that we capture finer partition of groups — and therefore flexible substitution patterns — we pool data from the two years prior for each year of mergers. We then compute patient choice probabilities for each hospital ($\hat{s}_{gh}$) for each group. To compute expected proportional changes in price, we compute the percent change in willingness-to-pay,

[48] This DRG weight is used to determine hospital payments under Medicare's reimbursement system.

$$\Delta WTP_m = \frac{\sum_g P_g \left[ln\left(\frac{1}{1-(\hat{s}_{gh}+\hat{s}_{gh'})}\right) - \left(ln\left(\frac{1}{1-\hat{s}_{gh}}\right) + ln\left(\frac{1}{1-\hat{s}_{gh'}}\right)\right)\right]}{\sum_g P_g \left[ln\left(\frac{1}{1-\hat{s}_{gh}}\right) + ln\left(\frac{1}{1-\hat{s}_{gh'}}\right)\right]} \tag{14}$$

where P_g is the share of patients within group g. Our primary specification uses a minimum group size of 50, resulting in 27,525 groups sized between 50 and 1,449.

C Measuring Premiums in the Form 5500

C.1 Sample Construction

Form 5500 is a regulatory filing required of all employers that offer a benefit plan to at least 100 employees. Although the data provide a rich source of employer-level data on premiums, there are many idiosyncrasies in the filing process that obfuscate true levels of premiums. Following Craig (2022), we implement a series of data restrictions and cleaning steps described below.

For fully insured benefits, employers are required to file a Schedule A form, which reports enrollment, premiums, plan type, and insurer for the plan. We use these data to construct an employer-year measure of average premiums per covered life. Self-insured employers are required to submit separate forms related to the administration of their plans. These forms pertain primarily to the financial details of the trust that is established to maintain plan funding. Self-insured employers are inconsistent in the degree to which their plans are funded through such trusts or the employers' general assets, making premium measurement for these employers intractable. However, levels and trends for fully and self-insured premiums are broadly similar (Craig, 2022).

We construct a panel of employers based on a combination of 5500 base forms and Schedule A forms. We exclude groups reporting on behalf of multi-employer plans, employers that operate plans in multiple states, and voluntary filers – i.e., employers that file despite the fact that their plans fall below filing thresholds (100 employees) in all years of the data.

C.2 Premium Measurement

We measure each plan's total premiums directly from the Schedule A form. However, employers can change in both absolute size and health plan enrollment from year to year, and a "per-person" measure of premiums lends itself more closely to comparisons with the scale of our health spending measure from the first-stage. We, therefore, standardize premiums using the total number of plan participants to compute average premiums per covered life per year. This ensures that the premium fluctuations we observe are related to changes in price, rather than fluctuations in employer size or plan participation.

The best measure of plan coverage comes from Line 1e of the Schedule A, which requires employers to report the "approximate number of persons covered at end of policy or contract year." However, it is clear from the data and documentation that employers are inconsistent in whether they include dependents in this field. From a preparer's manual for the 5500, Fisher and Andersen (2019) note that:

> *"The DOL says dependents should be included in the count reported on Line 1e, although whether dependents are include or excluded in the data provided by an insurance company varies depending on the carrier's own internal procedures. Generally, preparers simply use the information provided by the insurance company without further analysis. Dependents are not counted for any other purposes on the Form 5500 or its schedules."*

Line 6a of the 5500 base form asks employers to report the number of active plan participants (employees) at the beginning and end of the year. Although Line 6a consistently excludes dependents, it does not pertain to a single insurance policy, whereas the Schedule A Line 1e measure does. Instead, Line 6a typically represents the super set of enrolled employees across a number of benefit plans (e.g., life insurance, dental insurance, accidental death and dismemberment). Following Craig (2022), we standardize reporting across employer-years by identifying observations in which Line 1e reflects employee participation, rather than total plan coverage. We then adjust the coverage measure from Line 1e to match the average ratio of health plan coverage to overall benefit participation in adjacent years.

The approach from Craig (2022) requires manually reviewing rosters using a range of information in the other 5500 filings an employer submits including Schedule A filings for non-health plan contracts to adjudicate whether Line 1e reflects dependents or not. In order to scale this process, we manually classify Line 1e observations for four states: Massachusetts, Montana, North Carolina, and Texas. We then use these states to train a random forest algorithm to perform the remaining assignments.

The random forest algorithm classifies observations as to whether or not the coverage figure from Line 1e of the Schedule A includes dependents. In order to replicate the information set used to perform the manual assignments, we include the following measures as potential predictors:

- r, which is defined as the ratio of Schedule A coverage (Line 1e) to base form plan participation (Line 6a)
- Whether r is large enough to suggest clear dependent reporting ($r > 1.1$)
- The change in r from the previous year of an employers' reporting
- The standard deviation of r within employer
- The value and first difference in Line 6a reporting
- The largest Line 1e value reported among non-health plans in the year
- The measure and first difference of total premiums per person implied by naive use of Line 1e

The random forest prediction summarizes the average prediction over a large number of decision trees. At each node within a given decision tree, the sample is split to optimally categorize each side as high or low probability of including dependent coverage, resulting two "leaves." This decision is made by evaluating possible splits among a randomly selected subset of the potential predictors. At the next node, a new variable split is defined and this process repeats until the groups of observations within each "leaf" reach a minimum threshold. We use 10-fold cross validation to choose the hyper-parameters that minimize our classification error: we choose among eight variables at each node, allow trees to deepen until a minimum leaf size of 40, and average over 200 individual decision trees.

We then use these predictions to adjust observations classified as reporting covered employees to reflect dependents as described above. Finally, we deal with remaining classification error by implementing trims at the 5th and 95th percentiles of the premium distribution.

D Decomposing Instrumental Variation

In Section 5.5, we are interested in measuring the extent to which the three components of our primary instrument—merger timing, merger price effects, and employer-hospital exposure—drive the identifying variation that we use to estimate the effects of rising health care prices, so that we can focus our tests of identification assumptions on the right components. Specifically, we are interested in the following question: If we prevented one of these components from providing identifying variation, what would we be left with? This is especially important since some of the components of our instrument might face individual exogeneity issues. For instance, an employer may be more or less exposed to merging hospitals based on where they are; in turn, locations with differing exposure may be on different economic trends.

To diagnose this, we borrow a method from Borusyak and Hull (2023). They show that one can remove the identifying variation from one or more potentially endogenous components (while preserving the ability of those components to appropriately scale the instrument) by explicitly specifying the data-generating process of the plausibly exogenous components. An econometrician can then take the expected value of the instrument given the potentially endogenous components (averaged over values of the exogenous components). By subtracting this "expected instrument" from the primary instrument, the remainder is purged of any identifying variation from the potentially endogenous components while retaining that of the plausibly exogenous components. This method also allows us to check our estimates for robustness to using these alternative instruments if one is worried about exogeneity violations.

We model two components: The post-merger price increases, and the merger timing. It is much more difficult to think about an explicit model of the process that generates variation in employer exposure, which involves employer location choice. Moreover, employer exposure variation likely has the greatest threat to exogeneity, as described above. We therefore always allow this variation to be purged in our recentered instruments. When we allow merger timing to be exogenous, we assume that any merger we observe could have potentially occurred between the two years before and two years after the observed merger year with equal probability. When we allow post-merger price effects to be exogenous, we assume that they are drawn randomly from the empirical distribution.

In Columns 2 through 4 of Appendix Table A.6 and Appendix Table A.7, we present versions of our first stage with these recentered instruments. In Column 2, we allow price effects and timing to be modeled as exogenous. In Column 3, we allow *only* price effects to be modeled as exogenous. In Column 4, we allow *only* timing to be modeled as exogenous (i.e., the specifications in 4 and 3 are non-nested). We see that excluding employer exposure (going from column 1 to 2 in Appendix Table A.6) reduces the instrumental strength somewhat (moving the F-statistic from 844 to 554) and reduces the magnitude of the first-stage coefficient slightly. Removing the variation from timing as well (moving from column 2 to 3) does very little to either, implying that timing provides little instrumental variation. Indeed, when we *only* use variation from timing (column 4), our instrument is significantly weaker. These descriptions are true for both the employer-level and county-level specifications.

We conclude from this exercise that the cross-merger variation is the primary source of variation driving our instrument. The timing provides very little variation; employer exposure provides a large amount of variation, but is not essential given the price effect variation, and is more troublesome if we are worried about violations of the exclusion restriction.

E Scaling the Effect of Hospital Mergers

In this paper, we have used hospital mergers as an instrument for rising health care prices. In this section, we use our estimates to quantify the average effect of individual mergers on aggregate income, employment, tax revenue, and mortality. To do so, for the mergers of interest, we compute the change in our instrument induced by those mergers for every county in the year the mergers occurred. This change is different across counties, since each county is differentially exposed to a given hospital by virtue of the frequency that its residents tended to go to that hospital before the merger. We multiply this quantity by our first-stage estimate and then by our IV estimate for the relevant outcome. To convert our estimates where the measured outcome is in logs or shares to levels, we multiply the estimate by the baseline county average (for logs) or by the baseline county population (for shares). This process produces estimated effects of mergers on levels of outcomes for each individual county. We then sum over counties to estimate the total effect. This effectively measures the consequences of mergers for one year after they occur.

Across all merging hospitals in our analytic sample, we find that the average post-merger price increase is 0.9% (our confidence interval on this estimate spans from -0.1% to 1.8%). Our estimates from Section 7 imply that, on average, one of these mergers would have led to a $6.3 million reduction in income (95% confidence interval $0.1M to $12.6M), approximately 39 job losses (95% confidence interval 6 to 72), and about $1.3 million reduction in income tax payments (95% confidence interval $0.1M to $2.6M).[49] Because we obtain these figures by integrating over the set of observed price changes, population totals, and hospital spending, they reflect each hospital's observed post-merger price change and the degree to which those changes translate into dollars of additional health care spending.

Focusing on the 69 anticompetitive mergers in our analytic sample — those that generated a change in HHI over 200 points and led to a post-merger increase in HHI of over 2,500 points — we find that the average anticompetitive merger led to a $18.2 million reduction in income (95% confidence interval $0.2M to $36.1M), approximately 112 job losses (95% confidence interval 17 to 207), and about $3.8 million reduction in income tax payments (95% confidence interval $0.3M to $7.4M).[50]

Collectively, based on our estimates, the 69 anticompetitive mergers in our sample from 2010 to 2015 generated income losses of approximately $1.3 billion (95% confidence interval $0.0B to $2.5B), led to 7,742 job losses (95% confidence interval 1,177 to 14,308), and lowered federal income tax revenue by $265 million (95% confidence interval $21M to $510M). Across all 305 mergers in our sample, the aggregate harm in the year following a merger added up to $1.9 billion in forgone wages (95% confidence interval $0.0B to $3.8B), led to 11,903 job losses (95% confidence interval 1,809 to 21,998), and lowered federal income tax revenue by $409 million (95% confidence interval $32M to $786M).

[49] Note that, unlike in our employer-level analysis, our county-level measure of aggregate job losses is derived from a flag for UI receipt. Since not all of those who become unemployed take up UI, this serves as a lower bound on the total effects of mergers on changes in employment.

[50] The particular anticompetitive mergers in our sample are among smaller hospitals and labor markets than the average merger. If we instead assume a 5% price increase across all mergers in our sample, we estimate a $33.5 million reduction in income (95% confidence interval $0.4M to $66.6M), a $7.0 million reduction in federal tax revenue (95% confidence interval $0.5M to $13.5M), and 206 job losses (95% confidence interval 31 to 381).

F Additional Tables and Figures

Table A.1: Datasets, outcomes, sample sizes, and key limitations

Dataset	**Key outcomes & measures**	**Sample size & restrictions**	**Key limitations**
Health Care Cost Institute (HCCI) claims	Hospital prices; average health spending per beneficiary	28% of individuals with employer-sponsored health insurance in the United States, age < 65, 2008–2017	Data contributors are Aetna, Humana, and UnitedHealthcare (2008–2017).
U.S. Department of Labor Form 5500 Schedule A	Average health insurance premiums per beneficiary	3,970 fully insured employers (subset of main employer panel with balanced panel), 2008–2017	Restricted to fully insured plans; self-insured premiums unobserved; filings required only for employers with $\geq$ 100 employees; reports aggregated employer and employee contributions only.
IRS individual tax returns (W-2, 1099-G, Schedule SE) *(employer-level outcomes)*	Total payroll; worker counts; share of W-2s with HSA contributions	140,301 employers with $\geq$ 50 W-2s in 2009; balanced panel 2005–2017	Multi-employer workers are prorated across firms proportional to income; income excludes fringe benefits; HSA contributions provide a lower-bound measure of HDHP enrollment.
IRS individual tax returns (W-2, 1099-G, Schedule SE) *(county-level outcomes)*	Income (W-2 + Schedule SE) per capita; share unemployed; share with any self-employment earnings; tax withholdings; migration flows	1,709 counties with $\geq$ 500 HCCI beneficiaries, 2005–2017	UI-based unemployment captures flows, not duration; non-filers and those exiting the labor force may be missed; migration measure captures year-to-year moves only.
Roster of hospital ownership transitions from Brot et al. (2024b)	Used to derive hospital mergers and hospital characteristics	305 mergers involving 656 hospitals, 2010–2015	—

Note: All outcomes are measured annually. See the data section for additional details.

Table A.2: Employer Characteristics by Sample

Descriptive Statistics

	All Employers	Overall Analytic Sample	Non-Healthcare Analytic Sample	Healthcare Analytic Sample	Event Study Analytic Sample	5500 Analytic Sample
	Mean (1)	Mean (2)	Mean (3)	Mean (4)	Mean (5)	Mean (6)
Health Spending per Beneficiary ($/year)		4,099	4,101	4,079	4,098	4,024
Share of Employees with a Health Savings Account	0.02	0.04	0.04	0.03	0.04	0.06
Employer Total Payroll ($/year)*	2,779,000	12,721,000	13,045,000	8,242,000	13,250,000	29,789,000
Employer Count of Workers	62	298	304	204	308	491
Employer Average Income per Worker ($/year)	34,966	41,340	41,464	39,624	41,831	59,521
Premiums per Beneficiary ($/year)						5,037
Observations	3,024,169	140,301	130,830	9,471	125,362	3,970

Notes: This table presents descriptive statistics for our sample of employers with the various additional sample restrictions we add to cohorts of firms for analysis, based on data from 2008 to 2017. Data on employer payroll, employer counts of workers, employer wages, and the share of employees with a health savings account come from the Internal Revenue Service (IRS). Data on health spending per beneficiary come from the Health Care Cost Institute (HCCI). Data on insurance premiums come from the Department of Labor's 5500 forms.
* Rounded to $1,000 for privacy protection.

Table A.3: Comparison of Analytic Sample of Counties to Universe of Counties – 2008–2017

Descriptive Statistics

	All Counties	Analytic Sample
	Mean (1)	Mean (2)
Income Per Capita ($/year)	38,532	42,050
Share with Unemployment Insurance	0.03	0.04
Share with Zero Income	0.06	0.05
Share Unemployed	0.09	0.09
Unemployment Insurance Payments per Capita ($/year)	397	482
Share Self-Employed	0.13	0.12
Share Moving Annually	0.07	0.07
Income Tax Withholdings per Capita ($/year)	6,055	7,009
Observations	3,182	1,709

Notes: This table compares our analytic sample of counties with the universe of counties in the United States using data from 2008 to 2017. Our analytic sample captures approximately 96% of the U.S. population. Data on income per capita, the share with unemployment insurance, unemployment insurance payments per capita, the share of the population self-employed, the share of the population moving out of the county annually, and income tax withholdings per capita come from Internal Revenue Service returns. Income per capita is measured as the sum of wage (W-2) income and self-employment (Schedule SE) income. We define the share unemployed as the share of individuals with either positive unemployment insurance receipts and/or with zero income in the year.

Table A.4: Balance Table for Hospital-Level Analysis

	All Hospitals	Merging Hospitals	Eligible Control Hospitals		Matched Control Hospitals	
	Mean	Mean	Mean	Difference (2) - (3)	Mean	Difference (2) - (5)
	(1)	(2)	(3)	(4)	(5)	(6)
Beds (Total)	168	261	156	117***	270	–9
Admissions (Total)	7,560	12,799	6,824	6,678***	13,170	–371
Full Time Equivalents	943	1,565	903	751***	1,633	–67
Number of Unique Technologies	50	65	48	18***	67	–2
Share Medicaid Patients	0.17	0.18	0.17	0.01*	0.19	–0.01
Share Medicare Patients	0.49	0.46	0.50	–0.04***	0.43	0.02***
County Share with HCCI Insurance	0.14	0.17	0.13	0.05***	0.18	–0.01*
Teaching Hospital	0.23	0.38	0.21	0.19***	0.40	–0.01
Non-Profit Hospital	0.61	0.75	0.60	0.18***	0.79	–0.04*
Government Hospital	0.24	0.08	0.34	–0.28***	0.18	–0.10***
Independent Hospital	0.44	0.31	0.64	–0.36***	0.56	–0.25***
HHI (25-mile radius)	0.41	0.28	0.43	–0.17***	0.29	–0.01
HHI (50-mile radius)	0.20	0.13	0.20	–0.07***	0.15	–0.01*
Price Index	1,491	1,473	1,459	18	1,475	–2
Unique Hospitals	4,524	656	2,810		712	

Notes: This table presents hospital characteristics for the various samples of hospitals used in this analysis. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.5: Impact of Rising Health-Care Prices on Employer Payroll and Employment: Robustness to Alternative Sample Restrictions

	Firm Size in 2009: 10+		Firm Size in 2009: 10–50		Unbalanced Panel	
	Log(Payroll) (1)	Log(Workers) (2)	Log(Payroll) (3)	Log(Workers) (4)	Log(Payroll) (5)	Log(Workers) (6)
Panel A: OLS Estimates						
Log(Spending per Beneficiary)	0.00	–0.01**	–0.01**	–0.01***	0.02***	0.02***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Panel B: IV Estimates						
Log(Spending per Beneficiary) *(Instrumented using Merger-Driven Price Increases)*	–0.47***	–0.45***	–0.52***	–0.47***	–0.32**	–0.32**
	(0.10)	(0.10)	(0.15)	(0.13)	(0.15)	(0.15)
Mean Dependent Variable*	4,171,000	102	961,000	24	11,663,000	278
Observations	3,553,850	3,553,850	2,271,450	2,271,450	1,635,571	1,635,571
Number of Unique Employers	355,385	355,385	227,145	227,145	184,443	184,443
F-Statistic on First Stage	1,636	1,636	887	887	885	885

Notes: Each column reports coefficient estimates from regressions of annual employer-level log payroll or log worker counts on log(Spending per Beneficiary). The spending variable is instrumented with employer-level simulated spending per beneficiary constructed from merger-driven hospital price increases. Columns (1)–(2) restrict the sample to employers that filed at least 10 W-2s in 2009 (our baseline balanced-panel sample restricts to employers that filed at least 50 W-2s in 2009). Columns (3)–(4) further narrow the sample to employers with 10–50 W-2s in 2009. Columns (5)–(6) relax the balanced-panel requirement, retaining all employers that appear in at least one sample year ("unbalanced panel"). All specifications include employer and year fixed effects. Standard errors (in parentheses) are clustered at the employer level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.
* Rounded to $1,000 to protect privacy.

Table A.6: Alternative Employer-Level First-Stage Estimates

	Baseline (1)	Only Price and Timing Variation (2)	Only Price Variation (3)	Only Timing Variation (4)	Using Predicted Price Changes (5)	Health Care Employers (6)	Non-Health Care Employers (7)
Simulated Spending	0.64***	0.53***	0.51***	0.30***	2.90***	0.61***	0.65***
	(0.02)	(0.02)	(0.02)	(0.03)	(0.05)	(0.08)	(0.02)
Observations	1,403,010	1,403,010	1,403,010	1,403,010	1,403,010	94,710	1,308,300
Number of Unique Employers	140,301	140,301	140,301	140,301	140,301	9,471	130,830
F-Statistic on First Stage	844	554	509	84	2,932	52	793

Notes: This table presents coefficient estimates from a regression of employer-level annual health spending per beneficiary on employer-level simulated spending per beneficiary, as given in Equation (5). Each estimate includes employer and year fixed effects. Each column presents estimates from a different regression. Column (1) presents estimates using our baseline instrument. Columns (2)–(4) present estimates using a modified version of our baseline instrument that purges any variation other than that coming from differences in post-merger price changes across mergers and timing across mergers (2), only differences in post-merger price changes (3), and only differences in merger timing (4). Column (5) presents estimates using a modified version of our baseline instrument that replaces the estimated post-merger price increases with the estimated post-merger change in WTP. Column (6) presents estimates using our baseline instrument but restricting to only employers identified as being in the health care sector. Column (7) presents estimates using our baseline instrument but restricting to only employers identified as not being in the health care sector. Data on health spending and simulated spending come from the Health Care Cost Institute. Means are reported in levels rather than in logs. Standard errors are reported in parentheses and are clustered at the employer level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.7: County-Level First-Stage Estimates

	Baseline (1)	Only Price and Timing Variation (2)	Only Price Variation (3)	Only Timing Variation (4)	Using Predicted Price Changes (5)
Simulated Spending	1.03***	0.90***	0.88***	0.60***	5.85***
	(0.16)	(0.16)	(0.16)	(0.21)	(0.95)
Observations	17,090	17,090	17,090	17,090	17,090
Number of Unique Counties	1,709	1,709	1,709	1,709	1,709
F-Statistic on First Stage	41	32	30	8	38

Notes: This table presents coefficient estimates from a regression of county-level annual health spending per beneficiary on county-level simulated spending per beneficiary, as given in Equation (9). Each estimate includes county and year fixed effects. Each column presents estimates from a different regression. Column (1) presents estimates using our baseline instrument. Columns (2)–(4) present estimates using a modified version of our baseline instrument that purges any variation other than that coming from differences in post-merger price changes across mergers and timing across mergers (2), only differences in post-merger price changes (3), and only differences in merger timing (4). Column (5) presents estimates using a modified version of our baseline instrument that replaces the estimated post-merger price increases with the estimated post-merger change in WTP. Data on health spending and simulated spending come from the Health Care Cost Institute. Means are reported in levels rather than in logs. Standard errors are reported in parentheses and are clustered at the county level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.8: Alternative Insurance Premiums and Health Savings Account Results using Predicted Price Changes (i.e., Changes in Willingness-to-Pay Estimation)

Panel A: OLS Estimates

	Log(Insurance Premiums) (1)	Share of Employees with a Health Savings Account (2)
Log(Spending per Beneficiary)	0.03	0.01***
	(0.03)	(0.00)

Panel B: IV Estimates

	Log(Insurance Premiums) (1)	Share of Employees with a Health Savings Account (2)
Log(Spending per Beneficiary)	0.87*	–0.03
(Instrumented using Merger-Driven Δ WTP)	(0.47)	(0.03)
Mean Dependent Variable	5,036	0.04
Observations	39,700	1,403,010
Number of Unique Employers	3,970	140,301
F-Statistic on First Stage	90	2,932

Notes: This table presents ordinary least squares (Panel A) and instrumental variables (Panel B) coefficient estimates from regressions of annual employer-level log health insurance premiums per enrollee (Column (1)) and the share of employees with contributions to a health savings account (Column (2)) on employer-level annual spending per beneficiary, instrumenting for annual spending per beneficiary with employer-level simulated spending per beneficiary. However, when constructing our measure of simulated health spending, in lieu of estimating our merger-driven price increases using difference-in-differences regression, we use predicted price changes measured by changes in willingness-to-pay estimation to estimate the gains in market power for each merger. Each estimate includes employer and year fixed effects. Data on insurance premiums come from the Department of Labor Form 5500 filings. Data on an employer's share of enrollees with a health savings account come from the Internal Revenue Service. Means are reported in levels rather than logs. Standard errors are in parentheses and are clustered at the employer level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.9: Comparison of Our Employment Estimate to Prior Payroll Tax Elasticity Estimates

	Setting	Disemployment Response to 1pp Payroll Tax Increase
Our estimate†	US, 2008-2017	1.8%
US Studies		
Anderson and Meyer (1997)	US, 1978-1984	0.7-0.9%
Johnston (2021)	US, 2003-2012	1.5%
Guo (2024)	US, 2008-2013	1.1-2.4%
Other Studies		
Gruber (1997)	Chile, 1979-1985	0.0-0.3%
Saez et al. (2019)	Sweden, 2003-2013	1.0%
Benzarti and Harju (2021)	Finland, 1996-2015	3.4%
Bíró et al. (2022)	Hungary, 2010-2015	0.3%
Lobel (2024)	Brazil, 2008-2017	0.5%

Notes: We present estimates of the implied average employer-level reduction in the number of workers employed as a result of a 1 percentage point increase in payroll taxation from prior studies, as well as the equivalent as implied from our estimate of the effect of rising health costs on wages.
† We compute this estimate in two steps. First, we note that HCCI spending per covered life at the median employer is $4,039 and the average worker has 1 dependent, meaning employers spend $8,078 per worker. Second, we note that payroll per worker at the median employer is $36,526. For this employer, a 1pp payroll tax increase is $365 per worker, equivalent in dollar terms to a 4.5% increase in health care spending. Our estimate of the employment elasticity of health care prices for employers outside the health care industry is 0.40 (from Column (4) of Table 4). Multiplying these quantities together, our results imply that a 1pp payroll tax would reduce employment by 1.8%.

Table A.10: Impact of Rising Health Care Prices on Employer Payroll and Employment, by Employer Size

Panel A: OLS Estimates

	All Employers		Employers Below Median Size		Employers Above Median Size	
	Log(Payroll) (1)	Log(Workers) (2)	Log(Payroll) (3)	Log(Workers) (4)	Log(Payroll) (5)	Log(Workers) (6)
Log(Spending per Beneficiary)	0.00	0.00	0.00	0.00	0.01	0.01*
	(0.00)	(0.00)	(0.01)	(0.01)	(0.01)	(0.01)

Panel B: IV Estimates

	All Employers		Employers Below Median Size		Employers Above Median Size	
	Log(Payroll) (1)	Log(Workers) (2)	Log(Payroll) (3)	Log(Workers) (4)	Log(Payroll) (5)	Log(Workers) (6)
Log(Spending per Beneficiary)	–0.37***	–0.40***	–0.62***	–0.78***	–0.18	–0.11
(Instrumented using Merger-Driven Price Increases)	(0.14)	(0.13)	(0.23)	(0.22)	(0.16)	(0.16)
Mean Dependent Variable*	13,045,000	304	3,478,000	84	22,526,000	523
Observations	1,308,300	1,308,300	651,170	651,170	657,130	657,130
Number of Unique Employers	130,830	130,830	65,117	65,117	65,713	65,713
F-Statistic on First Stage	793	793	291	291	519	519

Notes: This table reports ordinary least squares (Panel A) and instrumental variables (Panel B) coefficient estimates from regressions of annual employer-level log payroll and log worker counts on log(Spending per Beneficiary). The endogenous spending variable is instrumented with employer-level simulated spending per beneficiary, constructed from predicted merger-driven price increases. Employer size is classified by the number of W-2s filed in 2009: Columns (3)–(4) restrict the sample to small employers (below-median W-2 count) and Columns (5)–(6) to large employers (above-median); Columns (1)–(2) use the full sample. All specifications include employer and year fixed effects. Payroll and worker counts come from IRS W-2 data. Standard errors (in parentheses) are clustered at the employer level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.
* Rounded to $1,000 to protect privacy.

Table A.11: Impact of Rising Health-Care Prices on County Income and Unemployment: All, Non-Healthcare, and Healthcare Workers

Panel A: OLS Estimates

	All Workers		Non-Health Care Workers		Health Care Workers	
	Log(Income per Capita) (1)	Share Unemployed (2)	Log(Income per Capita) (3)	Share Unemployed (4)	Log(Income per Capita) (5)	Share Unemployed (6)
Log(Spending per Beneficiary)	0.01	0.01***	0.02	0.01***	0.00	0.01**
	(0.01)	(0.00)	(0.01)	(0.00)	(0.02)	(0.01)

Panel B: IV Estimates

	All Workers		Non-Health Care Workers		Health Care Workers	
	Log(Income per Capita) (1)	Share Unemployed (2)	Log(Income per Capita) (3)	Share Unemployed (4)	Log(Income per Capita) (5)	Share Unemployed (6)
Log(Spending per Beneficiary)	–0.28*	0.09**	–0.36**	0.09**	0.19	0.03
	(0.15)	(0.04)	(0.17)	(0.04)	(0.23)	(0.03)
Mean Dependent Variable	39,535	0.09	40,253	0.09	46,629	0.06
Observations	17,090	17,090	17,090	17,090	17,090	17,090
Number of Unique Counties	1,709	1,709	1,709	1,709	1,709	1,709
F-Statistic on First Stage	41	41	41	41	41	41

Notes: Each column reports an instrumental-variables coefficient from a county-level regression of either log(Income per Capita) or the Share Unemployed on log(Spending per Beneficiary). The spending variable is instrumented with county-level simulated spending per beneficiary constructed from merger-driven hospital price increases. Columns (1)–(2) use employment and earnings for all workers; Columns (3)–(4) restrict the outcome sample to workers outside the health-care sector (NAICS 62 excluded); Columns (5)–(6) consider only health-care workers (NAICS 62). All specifications include county and year fixed effects. Standard errors (in parentheses) are clustered at the county level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.12: Additional County-Level Labor Market Outcomes

	Share Moving to Another County	Share Self-employed	Share with Zero Income	Share Receiving Unemployment Insurance	Log (Unemployment Insurance Payments per Capita)
	(1)	(2)	(3)	(4)	(5)
Log(Spending per Beneficiary)	0.00	0.00	0.00	0.09**	2.42
(Instrumented using Merger-Driven Price Increases)	(0.01)	(0.03)	(0.01)	(0.04)	(1.50)
Mean Dependent Variable	0.07	0.12	0.05	0.04	482
Observations	17,090	17,090	17,090	17,090	17,090
Number of Unique Counties	1,709	1,709	1,709	1,709	1,709
F-Statistic on First Stage	41	41	41	41	41

Notes: This table presents instrumental variables coefficient estimates from regressions of annual county-level share of the population who moved to another county (Column (1)), the share of the population receiving self-employment income (Column (2)), the share of the population who received zero income (Column (3)), the share of the population receiving unemployment insurance (Column (4)), and the log unemployment insurance payments per capita (Column (5)). Each estimate includes county and year fixed effects. Our labor market and tax revenue data come from the Internal Revenue Service. Means are reported in levels rather than logs. Standard errors are in parentheses and are clustered at the county level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.13: Effect of Rising Health-Care Prices on Non-Health-Care Labor-Market Concentration

	All Industries HHI	Non-Health Care HHI	Food Service NAICS 722 HHI	Legal Services NAICS 5411 HHI	Accounting NAICS 5412 HHI
	(1)	(2)	(3)	(4)	(5)
Log(Spending per Beneficiary)	–340	–384	–860	–4,545	1,287
(Instrumented Using Merger-Driven Price Increases)	(676)	(680)	(585)	(3,392)	(2,958)
Dependent Variable Mean	1,750	1,799	343	1,263	1,272
Observations	485,106	468,016	17,090	17,062	17,090
Number of Unique Counties	1,709	1,709	1,709	1,709	1,709
F-Statistic on First Stage	41	41	41	41	41

Notes: Each column reports an instrumental-variables coefficient from a county-level regression of the indicated Herfindahl–Hirschman Index (HHI) on log(Spending per Beneficiary). The spending variable is instrumented with county-level simulated spending per beneficiary constructed from merger-driven hospital price increases. Column (1) stacks all 2-digit NAICS codes and estimates the model at the 2-digit-NAICS-by-county-by-year level, including 2-digit-NAICS-by-county and 2-digit-NAICS-by-year fixed effects. Column (2) is identical to Column (1) but removes health-care NAICS codes (62). Columns (3)–(5) focus on service sectors potentially affected by hospital mergers: Food Services (NAICS 722), Legal Services (NAICS 5411), and Accounting (NAICS 5412). Standard errors (in parentheses) are clustered at the county level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.14: IV–DiD Estimates of the Impact of Rising Health Care Prices on Labor-Market Outcomes

Panel A: Firm-Level Outcomes

	Baseline Top- vs. Bottom-Quartile Simulated Spending		Robustness ≥ 50% Single-Year Simulated Spending Jump	
	Log(Payroll) (1)	Log(Workers) (2)	Log(Payroll) (3)	Log(Workers) (4)
Log(Spending per Beneficiary)	–0.79***	–0.63***	–0.81***	–0.65***
(Instrumented Using Merger-Driven Price Increases)	(0.15)	(0.15)	(0.16)	(0.16)
Observations	1,175,316	1,175,316	1,163,384	1,163,384
F-Statistic on First Stage	899	899	700	700

Panel B: County-Level Outcomes

	Baseline Top- vs. Bottom-Quartile Simulated Spending		Robustness ≥ 50% Single-Year Simulated Spending Jump	
	Log(Income per Capita) (1)	Share Unemployed (2)	Log(Income per Capita) (3)	Share Unemployed (4)
Log(Spending per Beneficiary)	–0.74	0.34	–0.97	0.40
(Instrumented Using Merger-Driven Price Increases)	(0.57)	(0.23)	(0.82)	(0.31)
Observations	17,120	17,120	16,992	16,992
F-Statistic on First Stage	3	3	2	2

Notes: Each coefficient is an IV–difference-in-differences estimate, summarizing the event-study analyses in Figures 4, 5, A.10, and A.11. The instrument is an indicator that equals 1 for units in the top quartile of the cumulative change in simulated spending between the first and last sample years and 0 for units in the bottom quartile. Columns labeled "Baseline" replicate the treatment definition used in Figures 4 and 5 (top- vs. bottom-quartile growth). Columns labeled "Robustness" follow Figure A.10 and A.11, further restricting treated units to those that experience at least one single-year increase in simulated spending that is 50% or more of their final simulated-spending value. Standard errors (in parentheses) are clustered at the unit (employer or county) level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure A.1: Map of Hospital Mergers

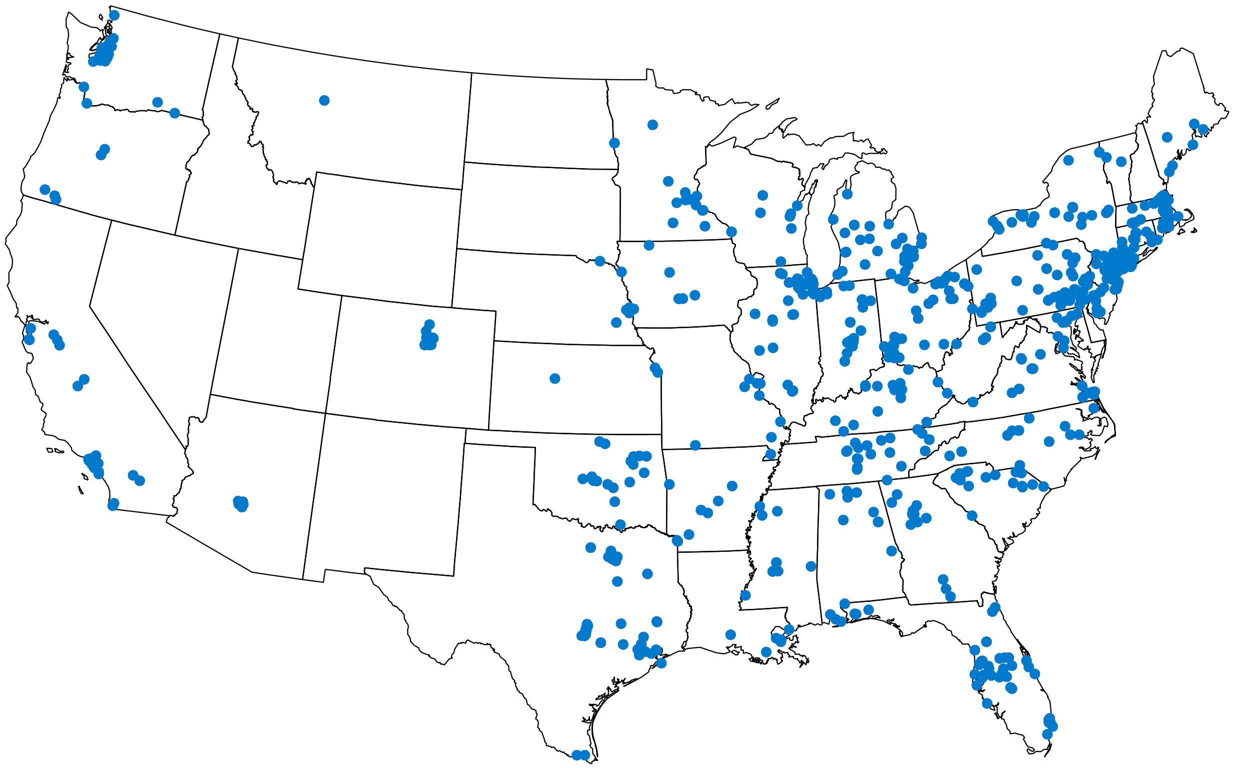

Notes: This presents the hospitals involved in the 304 hospital mergers from 2010 to 2015 we used in our analysis.

Figure A.2: The Impact of Mergers on Hospital Prices

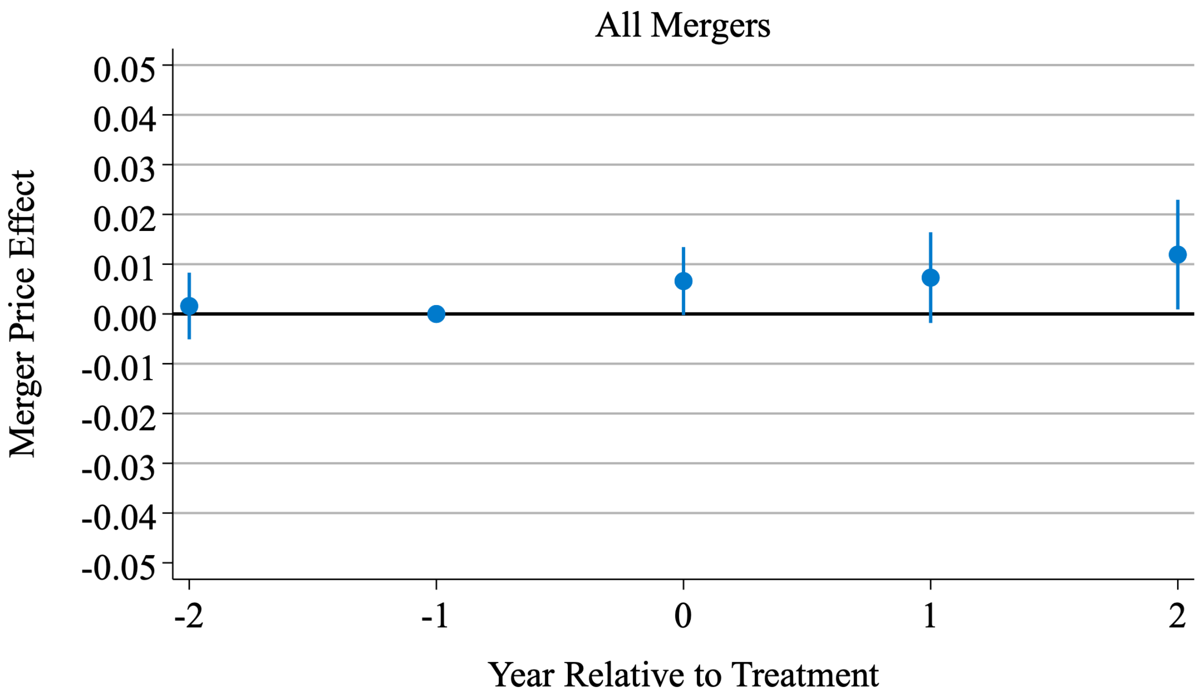

Notes: This figure presents difference-in-differences estimates of the effect of 2010-2015 hospital mergers on hospital prices. Hospital prices are a weighted average of inpatient and outpatient prices. The weights are constructed as the average share of inpatient and outpatient revenue for the treated hospital in 2008 and 2009, the two years prior to the first merger in our sample. We estimate Equation (2) to recover the average price change by merger group for transactions occurring between 2010 and 2015. Each dot is a yearly point estimate, and vertical bars denote 95% confidence intervals. Standard errors are clustered at the hospital level.

Figure A.3: The Impact of Mergers on Hospital Prices, No Matching

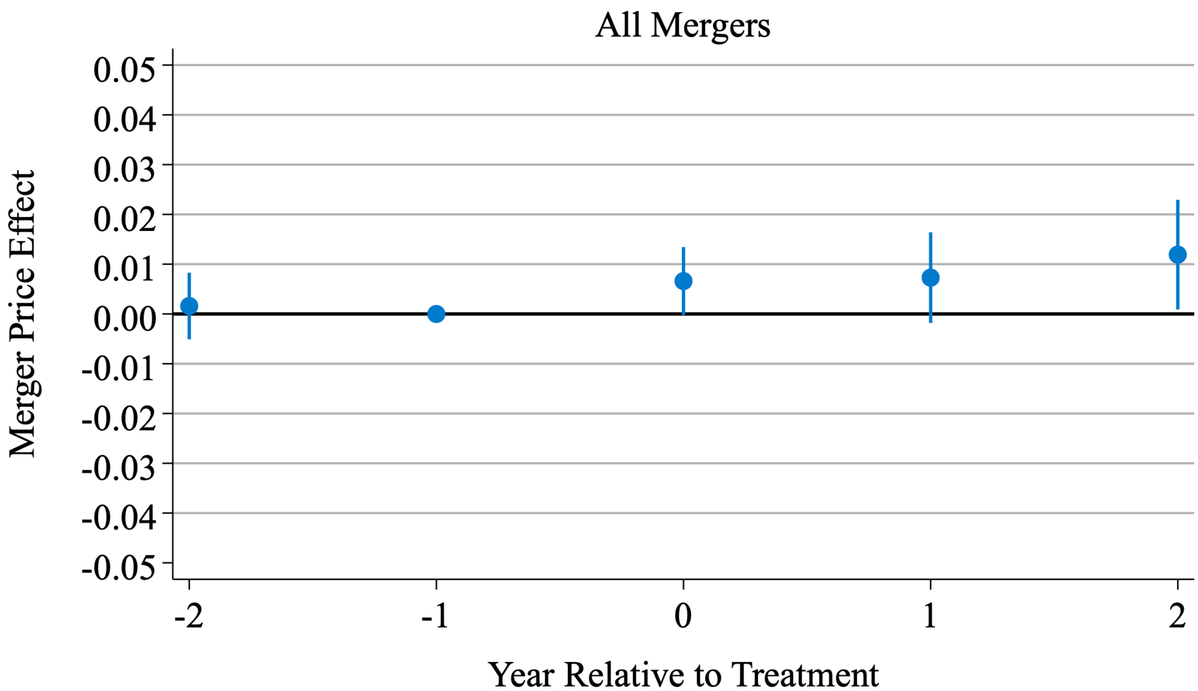

Notes: This figure replicates the hospital–level difference-in-differences analysis from Figure A.2 but without the propensity-score matching step used there. Hospital prices are a weighted average of inpatient and outpatient prices. The weights are constructed as the average share of inpatient and outpatient revenue for the treated hospital in 2008 and 2009, the two years prior to the first merger in our sample. We estimate Equation (2) to recover the average price change by merger group for transactions occurring between 2010 and 2015. Each dot is a yearly point estimate, and vertical bars denote 95% confidence intervals. Standard errors are clustered at the hospital level.

Figure A.4: Impact of Hospital Mergers on Risk-Adjusted 30-Day Mortality for Non-Deferrable Admissions

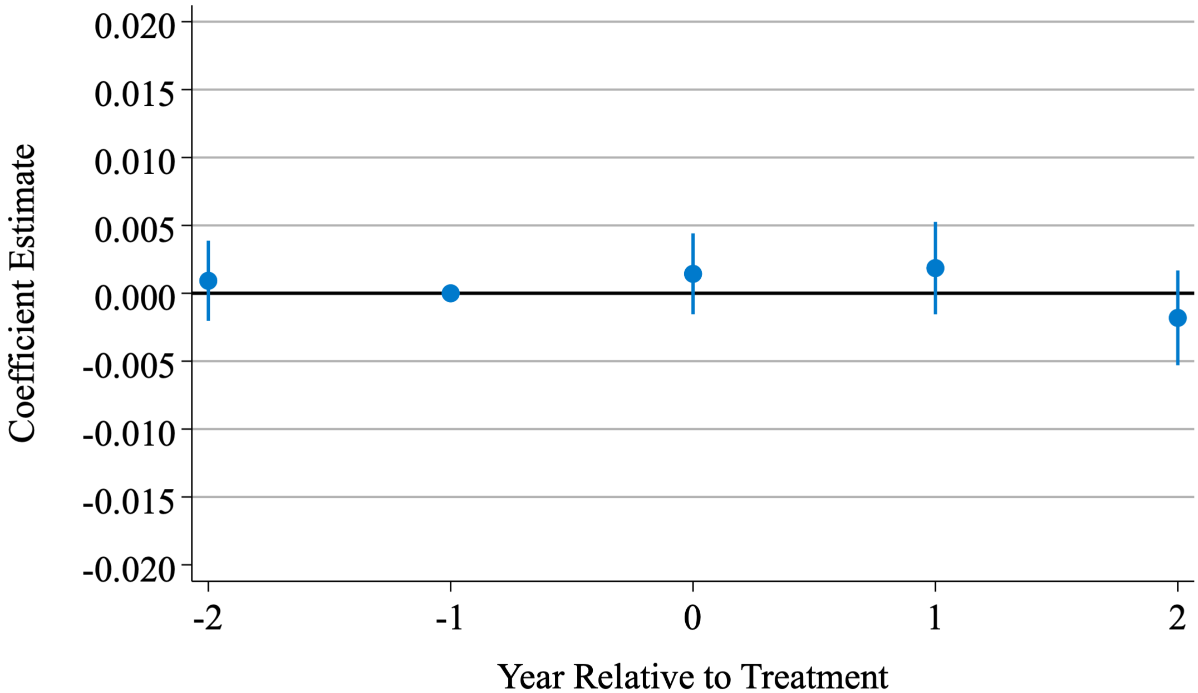

Notes: This figure presents difference-in-difference estimates of the effect of 2010–2015 hospital mergers on 30-day all-cause mortality for non-deferrable admissions in the Medicare Fee-for-Service population. We estimate Equation (2) that delivers the average mortality change within each merger group, following the approach used for hospital prices in Figure 2. Each dot is a yearly point estimate, and vertical bars denote 95% confidence intervals. Standard errors are clustered at the hospital level.

Figure A.5: The Impact of Mergers on Inpatient Hospital Prices by Willingness-to-Pay

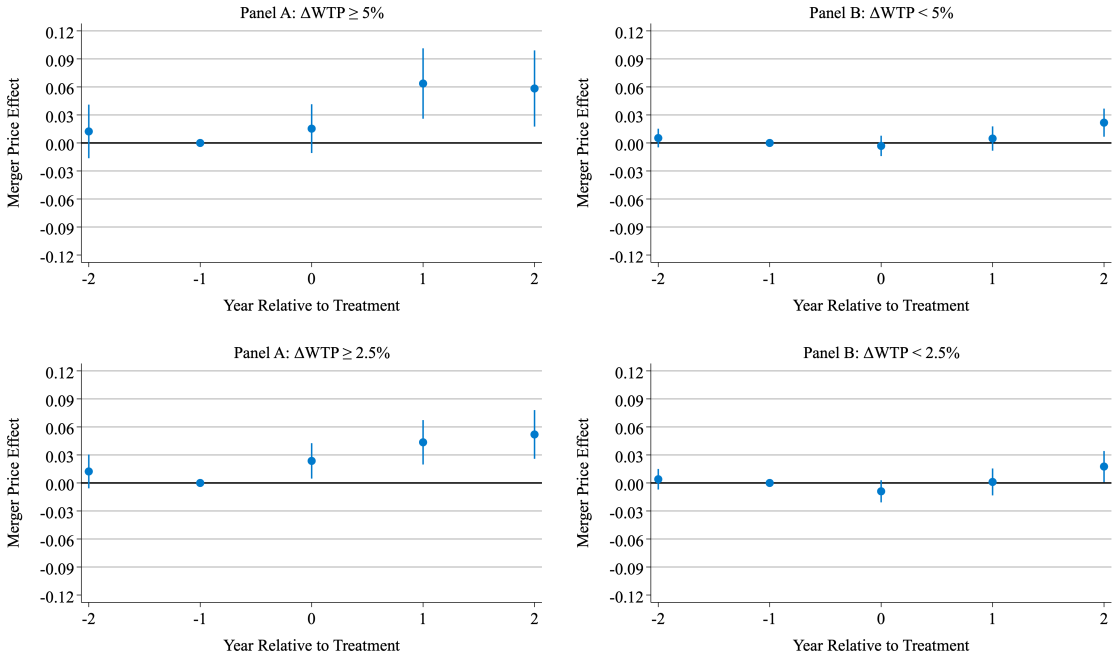

Notes: This figure presents difference-in-difference estimates of the effect of 2010 to 2015 hospital mergers on inpatient hospital prices. We estimate a pooled version of Equation (2) that estimates the average price increase within a group of mergers. We do so for the following groups: mergers that generated a $\Delta WTP \geq 5\%$ (Panel A), mergers that generated a $\Delta WTP < 5\%$ (Panel B), mergers that generated a $\Delta WTP \geq 2.5\%$ (Panel C), and mergers that generated a $\Delta WTP < 2.5\%$ (Panel D). Each dot represents a point estimate, and the vertical line displays the corresponding 95% confidence interval. Standard errors are clustered at the hospital level. Hospital pricing data come from the Health Care Cost Institute. The average difference-in-differences estimates comparing the two years prior to the merger with two years after the mergers in Panel A is 0.053 (0.026), in Panel B is 0.011 (0.009), in Panel C is 0.041 (0.016), and in Panel D is 0.008 (0.010).

Figure A.6: Relationship Between Trends in Economic Activity Prior to Mergers and Merger Location and Price Effects

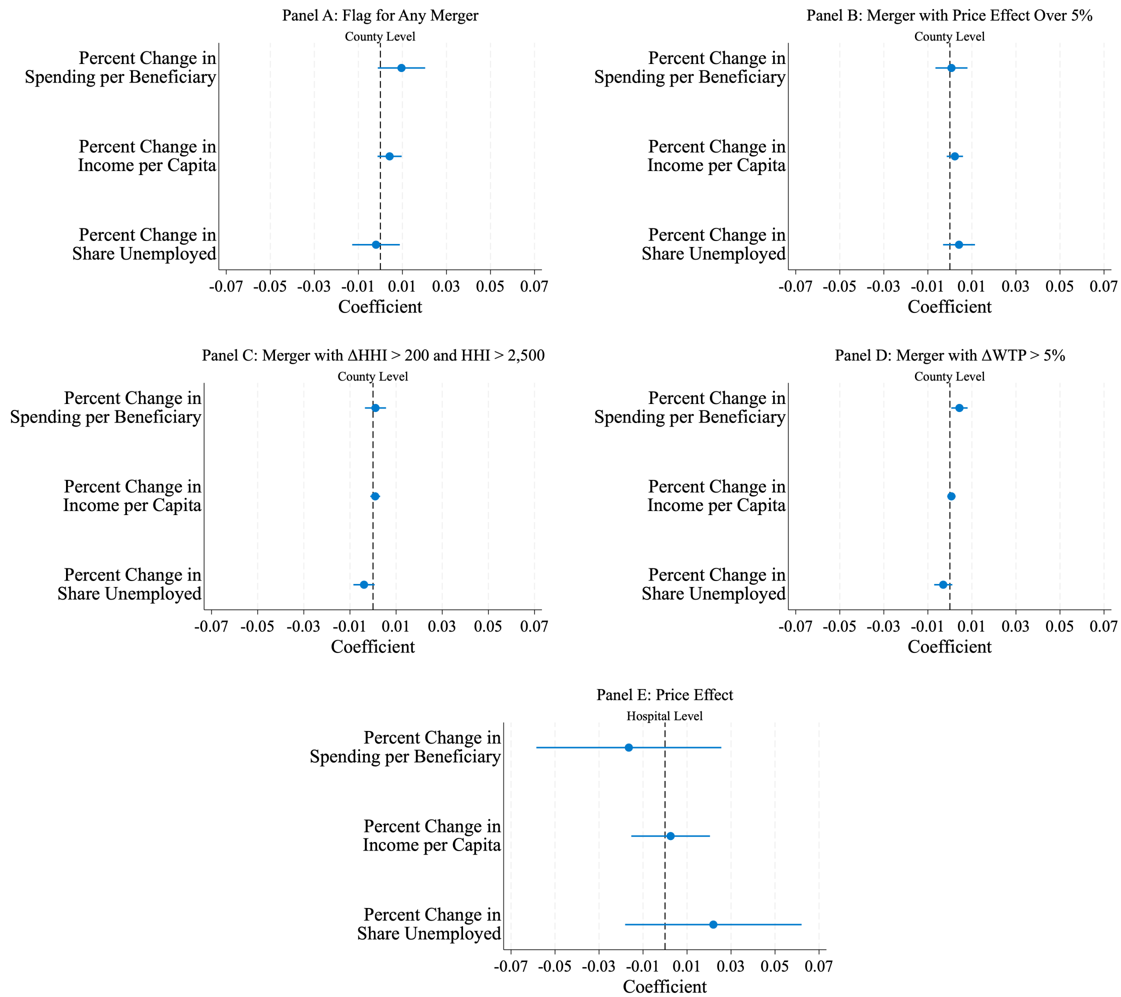

Notes: This figure presents the relationship between changes health spending per beneficiary and economic trends in the year prior to the merger and merger location and price effects. The independent variables are measured as the lagged percent change and Z scored. The regressions also control for the lagged levels of spending per beneficiary, income per capita, and share unemployed. Panel A shows a probit regression at the county-year level, where the dependent variable is an indicator that equals one if the county experienced a merger that year within its borders, while Panel B presents a similar regression with the dependent variable indicating whether the merger led to a price increase of 5% or more. Panels C and D present similar regressions where the dependent variable indicates whether a merger could be flagged using FTC guidelines and led to an increase in WTP of 5% or more, respectively. Panel E displays a hospital-level regression for the 656 merging hospitals in our sample, with the dependent variable being the merger's price effect, λ_{eh}, as estimated in Equation 2. The data for Panels A, B, C, and D cover all counties in our primary analysis sample from 2010 to 2015.

Figure A.7: Binned Scatterplot of Employer-Level Health Spending Per Beneficiary and Simulated Health Spending

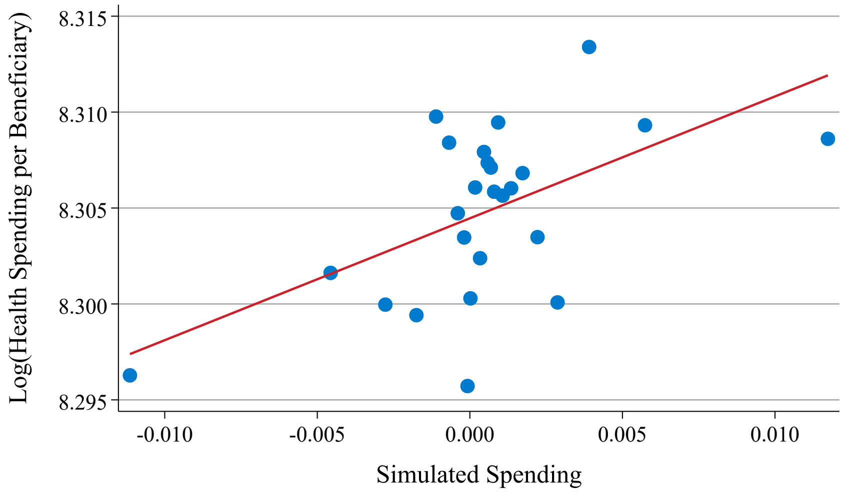

Notes: This figure presents a binned scatterplot of employer-level annual health spending per beneficiary against employer-level simulated spending per beneficiary, taking into account employer and year fixed effects.

Figure A.8: Event Study Estimates of First Stage: Regressing Employer Spending on Simulated Employer Spending, Restricted to Single Year Change at Least 50% of Final Simulated Spending

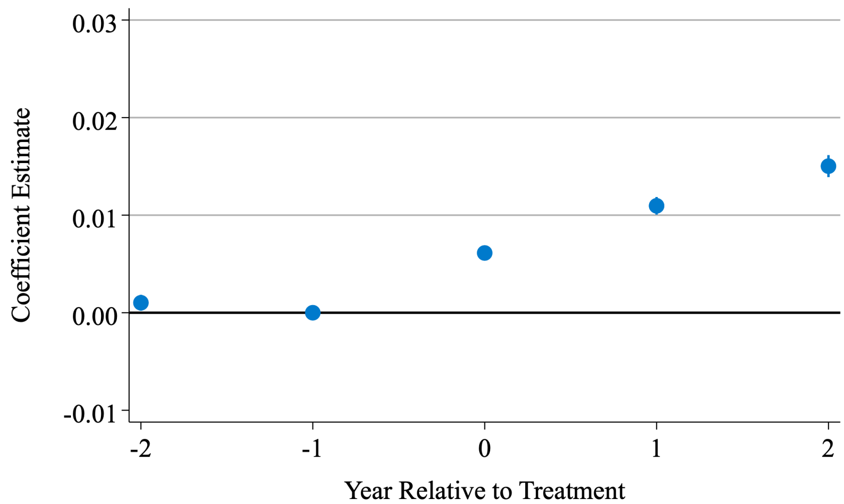

Notes: This figure replicates the event-study specification in Figure 3 except further restricts treated employers to have experienced at least one single-year jump in simulated spending that is $\geq$ 50% of their final simulated-spending value. "Treated" employers are in the top quartile of the cumulative change in simulated spending between the first and last sample years, and "control" employers are those in the bottom quartile. Dots denote point estimates; vertical lines denote 95% confidence intervals.

Figure A.9: Change in Simulated Health Spending by County, 2009 to 2015

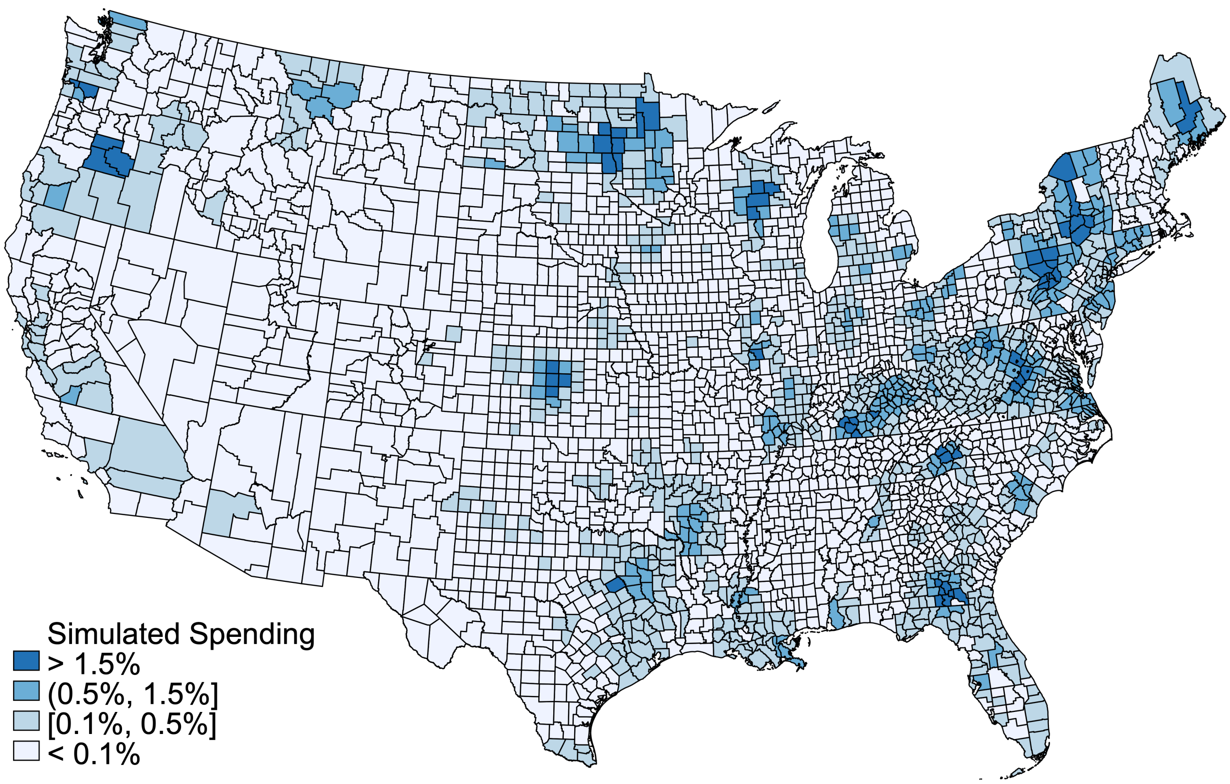

Notes: This figure presents the county-level change in simulated spending between 2009 and 2015. Darker areas are counties more exposed to the price increases generated by hospital mergers.

Figure A.10: Event Study Estimates of the Impact of Rising Health Care Prices on Employer Payroll and Worker Count at Non-Health Care Employers, Restricted to Single Year Change at Least 50% of Final Simulated Spending

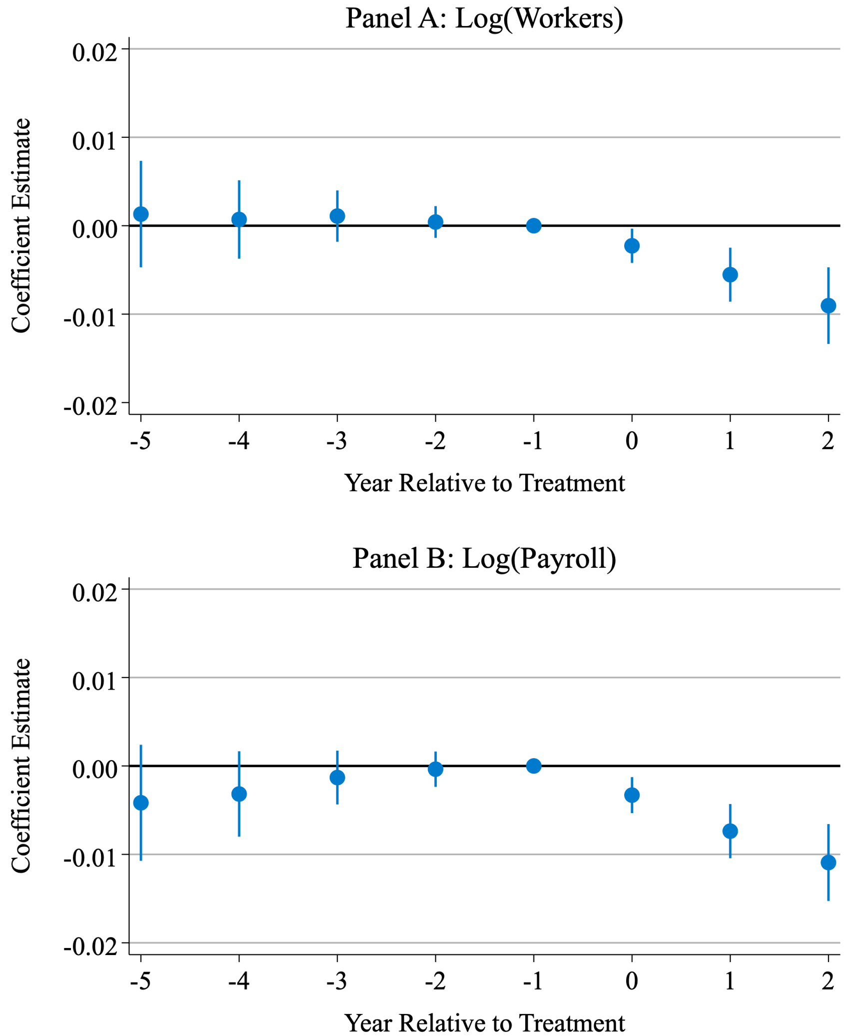

Notes: This figure replicates the event-study specification in Figure 4 except further restricts treated employers to have experienced at least one single-year jump in simulated spending that is $\geq$ 50% of their final simulated-spending value. "Treated" employers are in the top quartile of the cumulative change in simulated spending between the first and last years of the sample, and "control" employers are those in the bottom quartile. For each treated employers, the treatment year is the year in which it experienced its largest single-year increase in simulated spending. Dots denote point estimates; vertical lines denote 95% confidence intervals.

Figure A.11: Event Study Estimates of the Impact of Rising Health Care Prices on County-Level Income per Capita and Employment, Restrict to Single Year Change at Least 50% of Final Simulated Spending

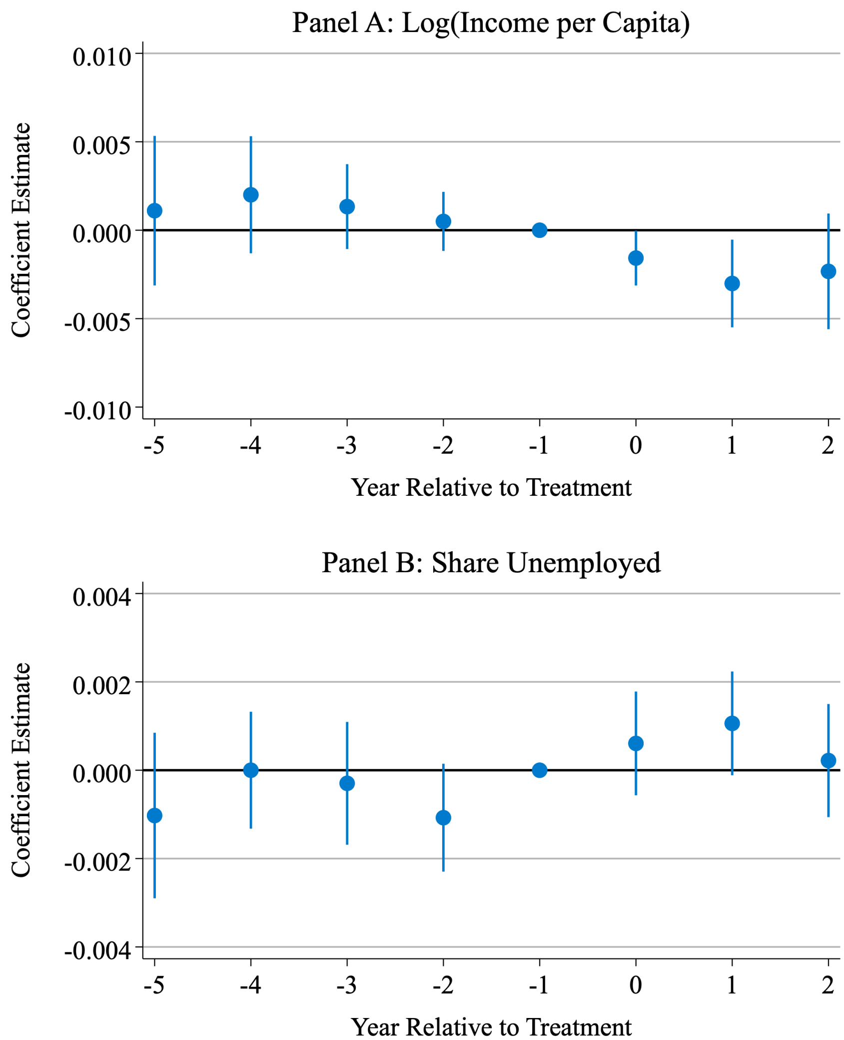

Notes: This figure replicates the event-study specification in Figure 5 except further restricts treated counties to have experienced at least one single-year jump in simulated spending that is $\geq$ 50% of their final simulated-spending value. "Treated" counties are in the top quartile of the cumulative change in simulated spending between the first and last years of the sample, and "control" counties are those in the bottom quartile. For each treated county, the treatment year is the year in which it experienced its largest single-year increase in simulated spending. Dots denote point estimates; vertical lines denote 95% confidence intervals.

Figure A.12: Robustness Tests of Employer-Level Payroll and Employment Effects

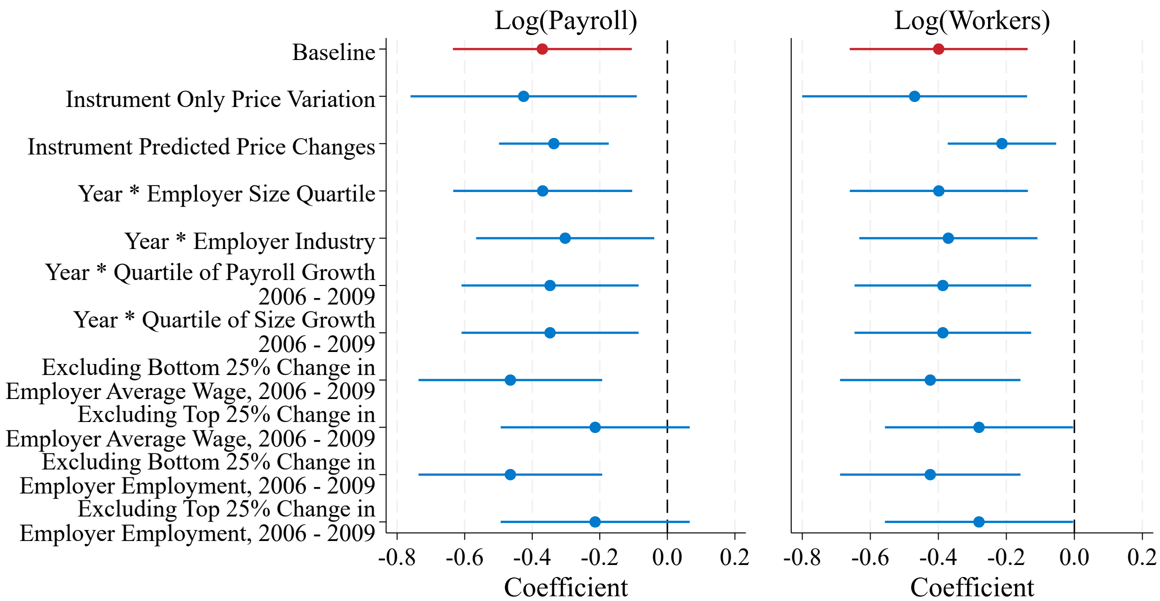

Notes: This figure presents two-stage least-squares estimates of Equation (6) of the effect of price increases on logged employer payroll and count of workers. Each estimate includes employer and year fixed effects. Observations are unique at the employer-year level. Our labor market data comes from the Internal Revenue Service. The dots represent point estimates and bars represent 95% confidence intervals.

Figure A.13: Robustness Tests of County-Level Income and Employment Effects

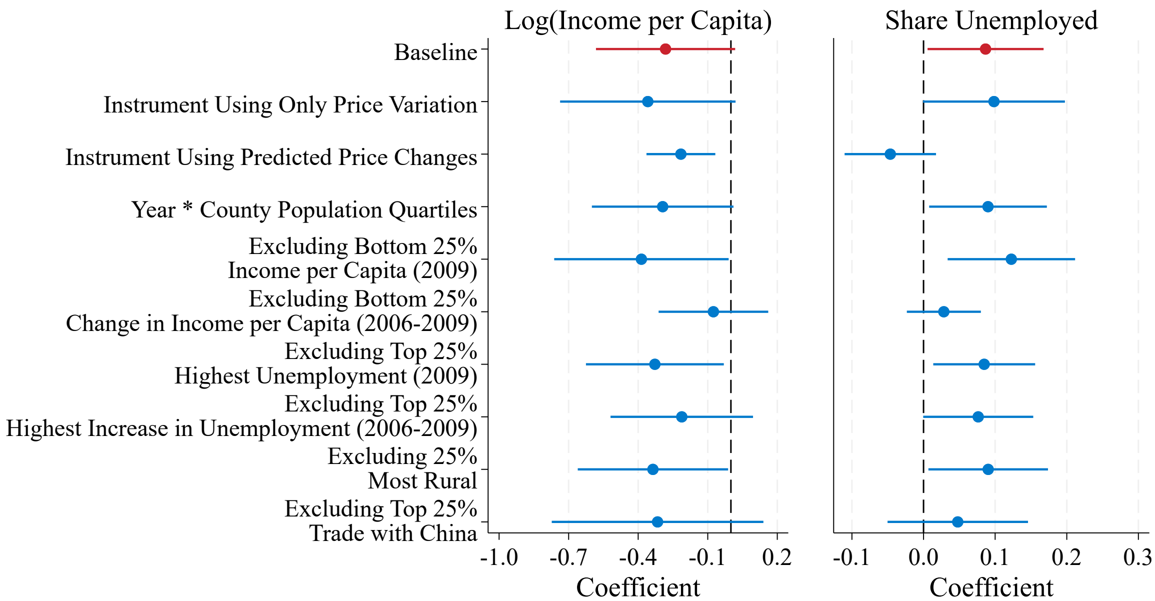

Notes: This figure presents two-stage least-squares estimates of Equation (10) of the effect of price increases on logged income per capita and share unemployed. Each estimate includes county and year fixed effects. Observations are unique at the county-year level. Our labor market data comes from the Internal Revenue Service. The dots represent point estimates and bars represent 95% confidence intervals.