

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

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

April 2024

Abstract

We analyze the economic consequences of rising health care prices in the US. Using exposure to price increases caused by horizontal hospital mergers as an instrument, we show that rising prices raise the cost of labor by increasing employer-sponsored health insurance premiums. A 1% increase in health care prices lowers both payroll and employment at firms outside the health sector by approximately 0.4%. At the county level, a 1% increase in health care prices reduces per capita labor income by 0.27%, increases flows into unemployment by approximately 0.1 percentage points (1%), lowers federal income tax receipts by 0.4%, and increases unemployment insurance payments by 2.5%. The increases in unemployment we observe are concentrated among workers earning between \$20,000 and \$100,000 annually. Finally, we estimate that a 1% increase in health care prices leads to a 1 per 100,000 population (2.7%) increase in deaths from suicides and overdoses. This implies that approximately 1 in 140 of the individuals who become fully separated from the labor market after health care prices increase die from a suicide or drug overdose.

*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. 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. 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 for health care goods and services play a central role the driving variation and growth in 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 pharmaceutical products have grown markedly faster than prices for goods and services outside the health sector (Bureau of Labor Statistics, 2022). During that period, for example, prices in the hospital sector — a \$1.3 trillion industry — increased faster than prices in 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 has arisen from unproductive rent-seeking activities.¹

In most markets, rising prices are paid directly by consumers. However, in the US health care sector, care for the majority of working-age adults is insured via employer-sponsored health insurance (ESI) (Kaiser Family Foundation, 2019). Prior research suggests that the cost of ESI, like other fringe benefits, is financed out of reductions in workers’ wages (Summers, 1989; Gruber and Krueger, 1991; Baicker and Chandra, 2006). As a result, ESI creates a tight link between changes in the health care industry and labor market outcomes outside the health care sector.

In this paper, we estimate the effect of rising prices in the health care sector on health insurers, workers, employers, and the federal government. To do so, we empirically trace out the entire causal chain that links changes in the market for health care services that raise prices through to downstream changes in labor market outcomes and federal income tax revenue. We show that rising health care prices increase health spending and are passed on to employers via higher ESI premiums. In turn, these higher premiums decrease labor demand, lowering employment.

We begin by developing a model that links health care prices to labor market outcomes through ESI. We highlight that when workers do not fully value increases in the price of health care services, they will not accept complete pass-through of ESI premium increases into wages. Because employers cannot fully reduce wages, they respond to price increases by laying off workers. To the extent that employers pay the same insurance premiums for different types of workers, a uniform increase in those premiums will make low-wage workers relatively more expensive to retain than high-wage workers. As a result, rising health spending caused by price growth can serve as a regressive “head tax” that has larger disemployment effects for lower-wage workers (Saez and Zucman, 2019; Case and Deaton, 2020; Finkelstein et al., 2023).

To quantify the effect of rising health care prices on labor market outcomes, we bring together a range of data. We leverage insurance claims data from the Health Care Cost Institute (HCCI),

¹See, e.g., Cooper et al. (2019a), Brand et al. (forthcoming), and Brot-Goldberg et al. (forthcoming) 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. (forthcoming) on drug copayment coupons.

which contain data on 28% of US with ESI coverage, to measure prices and utilization of health care services. We measure ESI premiums for fully-insured employers using data from Department of Labor (DOL) Form 5500 filings. Finally, we measure labor market outcomes using individual tax filings from the Internal Revenue Service (IRS).

Simply correlating changes in health care prices and labor market outcomes in this analysis is likely to be confounded by the fact that local health care prices may rise as workers become wealthier. As a result, we utilize price increases generated by horizontal mergers of hospitals to serve as a shock to local health care prices and, in turn, health care spending. From 2000 to 2020, there were over 1,000 hospital mergers in the US. The literature suggests hospital mergers lead to hospital price increases without any apparent improvement in quality or reductions in the quantity of care delivered.² Therefore, in addition to being a useful instrument, given the scale of consolidation that occurred in the hospital industry, studying the downstream consequences of these mergers is important in its own right.

To measure the effect of hospital mergers on prices, we use a panel of 304 mergers that were consummated between 2010 and 2015. Leveraging the analysis in [Brot-Goldberg et al. \(forthcoming\)](#), we take a difference-in-difference approach and compare prices at merging hospitals with prices at non-merging hospitals before and after mergers occurred. The average hospital merger led to price increases of 1.2% within two years after it occurred. However, as we illustrate, this average masks substantial variation in the post-merger price increases across transactions.

While it would be intuitively appealing to use a similar difference-in-difference approach to compare the outcomes of employers exposed to merging hospitals to the outcomes of those who were not exposed, this approach is not tractable in our setting. During our period of study, there were so many hospital mergers that virtually every US employer was nearby at least one merging hospital. Moreover, most employers were “treated” by multiple mergers, making it hard to neatly define “before” and “after” exposure periods.

Therefore, we measure the effect of price increases by using variation in the *intensity* of employers’ exposure to merger-driven hospital price increases. An employer is exposed to hospital price increases over time as hospital mergers occur nearby. That exposure is greater when an employer’s workers receive a larger share of their care at merging hospitals and when those hospitals raise their prices by a greater extent following their merger. To implement this approach, we construct an employer-by-year-level instrument that measures how the average price of care for an employer would have evolved over time if the *only* thing that had occurred were hospital mergers that raised prices, with other prices and quantities held fixed. We use this as an instrument for annual health spending by employers.

²See, for example, [Gowrisankaran et al. \(2015\)](#), [Garmon \(2017\)](#), [Cooper et al. \(2019a\)](#), [Beaulieu et al. \(2020\)](#), [Brot-Goldberg et al. \(forthcoming\)](#), and [Brand et al. \(forthcoming\)](#).

Our instrument is valid as long as an employer's exposure to mergers is independent of their counterfactual outcomes if no mergers had occurred nearby. We rely on several methods to assess the validity of this assumption. First, we develop an event-study-style approach to assess whether changes in merger exposure are related to *pre-treatment* changes in labor market outcomes. Second, we use tools from the literature on "formula instruments" to remove the potentially endogenous components from our instrument (Borusyak and Hull, forthcoming). In addition, we demonstrate that our results are robust to excluding employers from our analytic sample that have outlier trends in labor market outcomes before the merger events. All of our results are robust to these approaches.

We use this strategy to estimate the effect of rising health care prices on multiple employer-level outcomes. We find that merger-driven increases in hospital prices are passed through close to dollar-for-dollar into health spending and, in turn, into employees' ESI premiums. Despite facing higher premiums, employers did not appear to respond to price increases by shifting workers towards high deductible health plans.

When exposed to a 1% increase in health care prices, employers outside the health care industry reduced their payroll by 0.37%. Likewise, we find roughly equal-sized effects on employers' total count of workers employed. This implies that rising prices are largely resulting in job separations and are not merely being financed by reductions in workers' wages. As a result, our estimates suggest that the deadweight loss from hospital mergers (where quantities of care do not fall post-merger) are occurring in the adjacent labor market. Event study estimates suggest that these employer-level effects were realized almost immediately after prices increased, with little dynamic impact over time.

While we estimate large disemployment effects *at specific employers*, these may have simply resulted in reallocations of workers to other employers. We therefore estimate the broader regional effects of price increases by aggregating our employer-specific instrument to the county level. We find that a 1% increase in *county-level* health care prices reduced *county-level* labor income per capita by 0.27% and increased flows into unemployment by 0.1 percentage points (1%). We do not find that health care price increases induced any economically or statistically significant changes in self-employment or migration across counties. As with our employer-level results, these effects occurred immediately after hospital prices increased. Ultimately, these labor market changes had substantial fiscal consequences for the federal and state governments: a 1% increase in health care prices reduced income tax withholdings by 0.4%, while increasing unemployment insurance payments by approximately 2.5%. Collectively, these results underscore that the incidence of rising health care prices also falls on the state and federal governments and not just on workers.

We also explore the heterogeneity of the effect of rising health care prices on unemployment among workers across the income distribution. We find that unemployment effects are approximately zero for workers at the very top of the income distribution (i.e., those previously earning more than

\$100,000 annually). This is consistent with the views of [Finkelstein et al. \(2023\)](#), [Saez and Zucman \(2019\)](#), and [Case and Deaton \(2020\)](#) who assert that premiums serve as a regressive “head tax” on workers. However, we also find close to zero effects for workers at the very bottom of the income distribution (i.e., those previously earning less than \$20,000 per year). This aligns with the fact that these low-wage workers rarely receive health insurance from their employers and therefore do not become more expensive to employ when health care prices increase ([Lurie and Miller, 2023](#)). In general, the fact that availability and generosity of health insurance is positively correlated with wages pushes the unemployment effects of rising health care prices to be relatively more progressive than they would be in a world where health insurance premiums were uniform across workers.

A growing literature has highlighted how individuals who lose their job experience markedly higher short-run mortality that stems from increases in overdoses and suicides ([Eliason and Storrie, 2009](#); [Sullivan and von Wachter, 2009](#); [Venkataramani et al., 2020](#)). Echoing these findings, we observe that a 1% increase in county-level health care prices led to approximately one additional county-level death from suicides and overdoses per 100,000 people. Taken together with our estimates of the effect of rising health care prices on county-level employment, our results imply that approximately 1 in 140 of the individuals who became separated from the labor market due to rising hospital prices died. This estimate is consistent with the magnitudes found in prior work estimating the impact of job losses on mortality ([Sullivan and von Wachter, 2009](#); [Eliason and Storrie, 2009](#); [Venkataramani et al., 2020](#)).

Based on our estimates, our results suggest that, on average, a single hospital merger in our sample would have led to 39 job losses, approximately \$6 million in forgone wages, and a \$1.3 million reduction in federal income tax revenue. Likewise, we estimate that the average merger that raised prices by 5% or more would have led to 203 job losses, about \$32 million in forgone wages, a \$6.8 million reduction in federal income tax revenue, and between 1 and 2 additional deaths from suicide and overdose. This implies that the aggregate economic harm from an individual merger that raises hospital prices by 5% or more is approximately \$42 million.

Our work makes three distinct contributions to the literature. First, we contribute to a broad literature assessing the relationship between health care spending, health insurance, and the labor market ([Currie and Madrian, 1999](#); [Gruber, 2000](#)). Consistent with this literature, we show that the cost of health care is financed out of the labor market. However, the theoretical literature has largely focused on the labor market effects of adding new fringe benefits, like mandating maternity coverage ([Summers, 1989](#); [Gruber and Krueger, 1991](#)). New fringe benefits are likely to be valued by workers, which implies that workers will accept compensating wage reductions for the new benefit without a distortionary disemployment response. By contrast, the rising price of existing fringe benefits, absent changes in quality, are unlikely to be valued by workers and, as we illustrate in our setting,

can therefore cause large disemployment effects.³ We also provide new quasi-experimental evidence of the uneven distributional consequences of the rising price of health care services. Our results support the idea that ESI premiums can serve as a regressive “head tax” on employment among workers who receive ESI (Case and Deaton, 2020; Finkelstein et al., 2023; Gao et al., 2023).

Second, we contribute to the industrial organization literature exploring the consequences of rising market power (Berry et al., 2019; De Loecker et al., 2020). Since Harberger (1954), economists have traditionally focused on measuring the harms from market power in partial equilibrium terms, with damages coming from quantity reductions in the market where the market power is exercised. Because hospitals face very price-inelastic residual demand curves, the Harberger approach would suggest that the deadweight loss from hospital mergers is low (Gowrisankaran et al., 2015; Brod-Goldberg et al., 2017). By contrast, our results suggest that, in very inelastic markets, a first-order deadweight loss may instead arise in the market for close complements: in our case, job losses and, ultimately, deaths of workers outside the health care sector.

Third, we provide new evidence on the macroeconomic consequences of rising domestic health care spending. Prior work linking rising health care costs to the broader economy has primarily focused on the tax burden of financing public health insurance (Baicker and Skinner, 2011). Our results suggest that rising health care prices generate macroeconomic consequences by raising the cost of labor to employers and lowering federal tax revenue. A literature in macroeconomics has emphasized how sectoral shocks can have outsized effects on the broader economy when the goods and services provided in the sector, like energy, are hard to substitute away from and are ubiquitous (Baumol, 1967; Baqaee and Farhi, 2019). Indeed, both of these features characterize the US health care sector.

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 quantify the cumulative harm of hospital mergers in Section 8 and conclude in Section 9.

³Indeed, empirical research on new mandates has found little employment response (Gruber and Krueger, 1991; Gruber, 1994; Kolstad and Kowalski, 2016), while work on the rising cost of existing mandates has typically found disemployment responses (Cutler and Madrian, 1998; Baicker and Chandra, 2006; Arnold and Whaley, 2023; Gao et al., 2023). Other work on existing mandates has found wage effects but has not studied employment effects (Anand, 2016).

2 Background

2.1 Employer-Sponsored Health Insurance

Employer-sponsored health insurance in the US has its genesis in 1940s wage controls introduced by President Franklin Roosevelt. In response to wage caps, firms began offering non-monetary compensation to workers in the form of ESI. In 1943, the IRS allowed health insurance to be purchased with pre-tax income, lowering the cost of obtaining health insurance coverage (Starr, 1982). Today, insurers charge an annual premium, which is generally adjusted based on each employer's own expected costs (Craig, 2022). Ultimately, insurance premiums are funded by a combination of employer contributions and employee contributions, with the employer contributions excluded from taxation.

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 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 through 25th percentiles of income have ESI coverage, 84% of individuals in the 95th through 99th percentiles have ESI coverage (Lurie and Miller, 2023). Likewise, conditional on being insured through ESI, higher-income individuals select markedly 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. During that period, there were only 13 enforcement actions taken against hospital mergers by the Federal Trade Commission (FTC),

the agency tasked with preserving competition among hospitals in the US (Brot-Goldberg et al., forthcoming). The joint Department of Justice (DOJ)/FTC Horizontal Merger Guidelines specify that mergers that result in increases in 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-Goldberg et al. (forthcoming) show that, from 2010 to 2015, 20% of hospital mergers could have been predicted *ex ante* to meaningfully lessen competition using the screening thresholds established in the Horizontal Merger Guidelines and find that these mergers resulted in price increases of over 5%.

A growing literature demonstrates that mergers, particularly among hospitals that are close substitutes, can lead to increases in prices (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). Conversely, 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 the year 2000 (Fulton, 2017). Likewise, 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).

3 Theoretical Framework

We present a simple model to demonstrate the intuition for how, via increases in ESI premiums, increases in the price of health care can generate changes in 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 single hospital’s price goes up (Gowrisankaran et al., 2015; Brot-Goldberg et al., 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 transfer from health care payers to hospitals, 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 .⁴ 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 such an insurance market.

In Panel B of Figure 1, we model how the insurance market changes as insurers pay higher prices (transfers) 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. We model two employers, A and B , who differ only in their demand for insurance such that A buys higher coverage than B . Since the cost shock to insurance markets is proportional to the initial coverage level, A will generally face a greater premium increase than B .

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 their labor demand thus given by $D^L(w + \phi)$. Workers only value the wage they receive, not what their employer 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 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, and therefore generates deadweight loss given by a Harberger triangle (highlighted in red on the Figure) the same way a tax would.

Why do both wages and employment go down, in contrast to the predictions of Summers (1989), where the cost of mandated employer fringe benefits is 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 changes in wages and employment and the percent

⁴ $C = 0$ represents the purchase of no insurance coverage.

change in total labor income ($I = w \times L$) are given by:

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

where η_D and η_S are the absolute values of the elasticities of demand and supply, respectively.

The effect of rising health care costs 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 relatively 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 employment (and total income) response is that it is 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 start by building a panel of employers. We identify employers by their Employer Identification Numbers (EINs) as recorded on individuals’ W-2 forms. We select all employers with W-2 forms in 2009 and limit our sample to those that appeared on at least 50 W-2 forms (i.e., EINs that had at least 50 employees that year) and were located in the contiguous US. 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 140,300 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 a NAICS code starting with ‘62.’

For each employer, we measure their total payroll and count of workers using W-2 forms. In

each year, we take all W-2 forms that listed the employer in question. We sum the total of wages subject to Medicare tax (Box 5) to measure total payroll and use the count of W-2s to reflect the total workers. We construct an adjustment to account for the fact that workers may not be with an employer the entire year. For individuals who file more than one W-2 in a year, when we allocate them to employers, we do so proportionally to their income for that employer. So, for example, if an individual works for employer A that year making \$10,000 and for employer B making \$30,000, we count them as $\frac{1}{4}$ towards employer A's count of workers that year and $\frac{3}{4}$ towards employer B's count of workers.

4.2 Health Care Prices and Spending

We use insurance claims data from 2008 to 2017 provided by the Health Care Cost Institute (HCCI) to measure health care prices (discussed in Section 5.1) 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 includes the transaction prices that hospitals were actually paid by insurers, not merely their “chargemaster,” or list prices. We focus our analysis on 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.⁵ 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 an employer's employees that live in each county c in our base period of 2009, as measured from their listed ZIP codes on their filed W-2 forms.⁶

⁵Prescription 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.

⁶Based our discussions with industry participants, this aggregation mimics the way health insurers would price premiums for an employer with beneficiaries spanning multiple geographies.

In practice, an employer’s expected health care spending is 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:

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

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 the Form 5500, which is a regulatory filing collected by the Department of Labor (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 100 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 B. Because the 5500 identifies employers, we can link them directly to our panel in the IRS data.⁷ 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. We proxy for whether employees have a HDHP based on whether they report contributions to a Health Savings Account (HSA) in Box 12 of their W-2 filing. 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 explore the general equilibrium effects of our employer-level estimates, we also undertake a county-level analysis. Our primary county-level outcomes of interest are per capita labor income, the share of the population who are unemployed, the share of the population who are self-employed,

⁷While 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 the Form 5500, and they do not reliably reflect the total cost of insurance provision in the way that they do for the fully-insured.

tax receipts and unemployment insurance (UI) payments, and the share of the population who moved to another county. We focus on the population of individuals aged between 25 and 64 in the year of interest, who represent the individuals likely to receive ESI directly rather than receive coverage under Medicare or as a dependent on a commercial plan.

To measure labor income, we combine information from each individual's W-2 forms and Schedule SE forms to capture earned wages and self-employment income, when present. We measure unemployment based on whether an individual is receiving UI based on the presence of any income in Box 1 in filings of Form 1099-G (which denotes the amount of UI payments received), or the individual earned no positive income from W-2 or Schedule SE forms. We allow these measures to only apply to those who were not unemployed (by this measure) in the year prior. As a result, our approach captures *flows into unemployment* rather than the stock of unemployed.⁸ We count individuals as self-employed if they filed a Schedule SE form with any positive income during the year. We identify movers as individuals whose county of residence changed between $t - 1$ and t . We measure tax payments as the total amount reported as withheld on W-2 forms.⁹ Finally, we measure UI payments as the total amount paid as reported on the 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. However, a drawback of this approach is that individuals who lose their jobs and exit the labor market can potentially go missing from IRS data. As a result, we also show our results are robust to measuring the county population among individuals age 25 to 64 using the US Population Estimates from the Census Bureau.

We additionally measure county-level death by cause using microdata from the Restricted Vital Statistics database provided by the Centers for Disease Control and Prevention's (CDC) National Center for Health Statistics. These data capture information from all US death certificates, including location and age of decedent, and cause of death. We match all recorded deaths between 2008 and

⁸We 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.

⁹Although withholding sometimes does not capture the full tax bill, especially given the over-withholding that is typical, we expect any error in this measure to be independent from exposure to hospital merger-related price increases.

2017 to their counties. Using these data, we construct a measure of deaths per 100,000 people in a county and year.

4.5 Summary Statistics

We present descriptive statistics for our employer and county samples in Panels A and B of Table 1. Our primary employer sample contains 140,300 unique employers. Across 2008-2017, the average employer in our sample had 297 employees and average wages per worker of \$41,339. The average employer spent (according to our constructed measure) \$4,099 per employee on health care. 6.8% of our employers (9,471) were in the health care industry, with the other 93.2% (130,829) in other 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. In these counties, the average income across 2008-2017 was \$41,908, the share of the population unemployed was 8.9% (3.6% with UI and 5.4% with zero income),¹⁰ the share self-employed was 11%, average federal income tax payments per capita were \$7,009, and average UI payments were \$482. The mean county experienced 237 deaths per 100,000 population age 25 to 64 annually, with 23 deaths per 100,000 from overdoses and suicides.

In Appendix Table A.1, we show how the composition of our sample of employers changes as we introduce additional sample restrictions imposed by data availability. For example, our primary employer sample is composed of 140,300 employers, whereas our sample with data on insurance offerings from Box DD on W-2 returns is composed of 39,341 employers because only firms with over 250 employees are required to complete Box DD. Employers where we have data on insurance provision are markedly larger than employers where we do not. Employers outside the health sector are modestly larger than health care employers. In Appendix Table A.2, we show how our analytic sample of counties differs from the universe of counties in the US. The population in our analytic sample have 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 outcomes. Our core empirical challenge is the potential for reverse causality: that increases in prices may

¹⁰Note 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 unemployed if they are ever without a job and receiving UI throughout the year.

be *caused* by changes in an employer’s economic conditions that affect the wages it pays and the workers it employs. For example, if demand for software spikes, the employers who produce software may increase their employees’ salaries and hire additional employees. Since health care is a normal good, both of these changes in labor market outcomes will raise demand for local health care services. This increase in demand, in turn, could raise prices and health spending, inducing a positive correlation between prices and labor market outcomes.

We, therefore, need to find a source of changes in health care prices that is not driven by local economic conditions. Our general approach is to use exposure to price increases generated by new horizontal hospital mergers as a source of variation in the price of hospital care consumed by employees. This requires two steps: first, we must estimate the effect of hospital mergers on hospital prices; second, we must find a way to map these hospital-specific price increases to increases in employers’ price of health care and health spending.

5.1 Measuring the Effect of Mergers on Hospital Prices

We follow [Brot-Goldberg et al. \(forthcoming\)](#) to estimate the effect of mergers on hospital prices. We sketch the relevant details of the methods below, and interested readers should refer to that paper for more details on our estimation strategy

We begin by constructing a hospital-level price index using the HCCI data. Hospitals are multi-product firms that offer numerous services. Each hospital differs in the mix of services they offer and the demographic profile of the patients they treat. We address this heterogeneity by creating hedonic price indices for inpatient and outpatient prices. Our price indices adjust for the mix of services hospitals offer (measured using Diagnosis Related Group (DRG) and Current Procedural Terminology (CPT) codes for inpatient and outpatient care, respectively) and the age and sex of their patients. Our final measure is a weighted average of inpatient and outpatient price indices, weighted by their share of spending at the hospital-year. We provide a detailed description of how we measure hospital prices in [Appendix A](#).

We then estimate the effect of a hospital merger on prices for the participating hospitals. We examine 304 hospital mergers occurring in the US between 2010 and 2015, where at least two of the merging parties were within 50 miles of one another. We focus on mergers of hospitals located within 50 miles of each other, since [Cooper et al. \(2019a\)](#) do not find meaningful price effects involving merging parties located more than 50-miles apart. [Appendix Figure A.1](#) includes a map of the mergers we focus on in this analysis.

We measure the post-merger changes in hospital prices using a difference-in-difference regression that compares the change in prices at merging hospitals, before and after a merger, to concurrent

changes in prices at matched non-merging hospitals over the same period.¹¹ We limit to the period covering two years before and after the merger, carving out the year the merger was consummated from our estimation. We then estimate a regression for each merging hospital:

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

where e is a merger event and t is a year. Our target parameter is λ_{eh} , the merger-induced price increase (in percent terms) at merging hospital h due to merger event e . We deviate from [Brot-Goldberg et al. \(forthcoming\)](#) and use empirical Bayes methods to shrink our λ_{eh} estimates in order to reduce noise driven by measurement error.

We observe that merging hospitals raise inpatient and outpatient prices by 2.6% and 1.0%, respectively, within two years after being involved in a transaction.¹² This implies an average price increase post-merger across inpatient and outpatient prices of 1.2%. However, as we note in [Brot-Goldberg et al. \(forthcoming\)](#), this average price increase masks wide variation in the size of the price increases generated across transactions: across the mergers we analyze, we observe a standard deviation in the price increases generated of 14.6%. Likewise, we observe that, out of 304 mergers, 125 raised prices across at least one of the merging facilities by more than 5%.¹³ In [Brot-Goldberg et al. \(forthcoming\)](#), we show that the mergers that generated the largest price increases were the transactions that involved a more substantial lessening of competition, indicated by greater increases in willingness-to-pay or the merging parties' Herfindahl-Hirschman index. [Brot-Goldberg et al. \(forthcoming\)](#) also show that the 20% mergers that generated the largest prices could have been predicted, *ex ante*, to lessen competition using the standard screening tools available to antitrust enforcement agencies.

5.2 Instrument Construction

Absent context, one appealing identification strategy might be to use a difference-in-difference design to compare outcomes of employers nearby a hospital merger against other, non-exposed employers, before and after the merger occurred. However, such an approach is not tractable in our setting because the volume of mergers that occurred in this period means there are few untreated regions and few regions with a well-delineated pre-period. Of, for example, the 709 commuting zones (CZs) in the US that have a hospital within their borders, fewer than 1% were not exposed to a hospital merger occurring in either their CZ or a neighboring CZ between 2010 and 2015. The

¹¹Our sample of mergers in this analysis differs slightly from [Brot-Goldberg et al. \(forthcoming\)](#) because, in this analysis, we limit to individuals age 18 and over.

¹²See Appendix Figure A.2 for an event study of our average post-merger price increases.

¹³In [Brot-Goldberg et al. \(forthcoming\)](#) show that these mergers do not, on average, have any effect on the volume of care delivered.

few areas without exposure to mergers tend to be much poorer and more rural than “exposed” CZs, making employers in those areas unlikely to serve as good controls. Moreover, of the CZs where a merger occurred, all but one experienced multiple mergers during this period, making it untenable to distinguish between pre- and post-merger treatment periods.

Therefore, rather than using a binary indicator of an employer’s exposure to any merger, we instead construct a continuous measure of employers’ exposure to merger-driven price increases across time and illustrate how those price increases impact health spending. We exploit the fact that an employer’s exposure to higher prices caused by hospital mergers at any given time is a function of 1) the extent to which an employer’s employees receive care at merger hospitals; 2) the scale of the price increases generated by mergers at hospitals where their employees receive care; and 3) the timing of when hospital mergers occur. This approach allows us to think about mergers as local “shocks” that generate varying sized price changes, which differentially raise health spending across employers and over time.

To implement our continuous exposure approach, we construct a single instrument summarizing the degree to which merger-driven price increases should pass through in health care spending, holding fixed the prices for other health care services, as well as the mix and quantity of health care consumption. We then use this “simulated spending” measure as an instrument for employers’ measured prices.

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

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

where exposure Z_{ieht} captures the percent change in health care spending for the employers’ employees, holding quantities fixed.

This instrument has three components. First, the term λ_{eh} represents the extent to which the merger event e raises average prices at the merging hospital h . Mergers that raise prices by more will, naturally, have a larger effect on spending. The λ_{eh} values here are estimated by our post-merger price effect regressions in Section 5.1.¹⁴ Second, the term $1[t \geq \tau_e]$ reflects the timing of the merger. 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.

The third component, σ_{ih0} , captures the share of employee spending at merging hospital h . As in Section 4.2, we construct these shares as a weighted average based on each employer’s geographic

¹⁴ λ_{eh} is zero for hospitals not involved in merger e and for providers that are not hospitals.

distribution of employees:

$$\sigma_{ih0} = \frac{\sum_c \omega_{ic0} \times S_{c0} \times \sigma_{ch0}}{\sum_c \omega_{ic0} \times S_{c0}}, \quad (4)$$

where ω_{ic0} is the share of employer i 's employees who report living in county c , σ_{ch0} is the share of spending in county c at hospital h , and S_{c0} is the average spending in county c . We measure all quantities in 2008 and 2009, in notation as “0,” which is before any mergers in our sample were consummated, and hold them fixed across time. As a result, our measure effectively holds quantities fixed across time, with only hospital prices changing over time because of mergers. In order to respect important differences in the degree to which patients tend to travel for more and less complicated services, we distinguish between prices and spending shares for inpatient and outpatient services.

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

$$z_{it} = \sum_{e \in \mathcal{E}, h \in \mathcal{H}} Z_{ieht}. \quad (5)$$

z_{it} evolves over time as mergers switch on, and increases to a greater extent if the employer had greater dependence on hospitals which merged, especially those that generated larger increases in prices.

Across our sample, the mean 2009 to 2015 change in employer-level simulated spending is 0.001 with a standard deviation of 0.0092. The top 25% of employers in our sample experience an increase in simulated spending of 0.0031. The top 5% of employers experience an increase in simulated spending of 0.015.

With our instrument z_{it} constructed, our primary empirical approach is two-stage least-squares (2SLS) estimation regressing outcomes y_{it} on spending x_{it} , instrumenting for x_{it} with simulated spending z_{it} :

$$x_{it} = \delta \times z_{it} + \Theta_i + K_t + u_{it} \quad (6)$$

$$y_{it} = \beta \times x_{it} + \theta_i + \kappa_t + \varepsilon_{it}. \quad (7)$$

We focus on spending in our first-stage because it allows us to align the scale of hospital merger-driven price increases at the level of expected spending per individual. To the extent that quantity responses are negligible, we expect spending to increase one-for-one as our price instrument increases. As a result, our second-stage estimates reveal the effect of a global increase in health care prices, assuming quantities stay fixed in response to price increases. We cluster our standard errors in this analysis at the employer-level.

In Column (1) of Table 2, we show our first-stage estimates of Equation (6). Our standard first-stage employer-level regression has a coefficient of 0.65 with an F-statistic of 864.78. In our setting, this first-stage coefficient illustrates the extent to which merger-driven price increases are passed through into total spending. Because our first-stage is a log-level regression (simulated spending is measured in percentage points), after exponentiating our first-stage coefficient, our results imply that 91% (computed as $e^{0.649} - 1$) of the price increases from hospital mergers are passed through into total spending increases. In Appendix Figure A.3, we present scatter plots of the relationship between the change in simulated and true spending from 2010 to 2015.¹⁵

5.3 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. Satisfying instrumental exclusion requires a core assumption: that idiosyncratic changes in merger-driven price increases are unrelated to idiosyncratic unit-specific trends in outcomes, except through the channel by which they make health insurance more expensive to purchase. This exclusion restriction is a continuous analogue to the sort of parallel counterfactual trends assumptions required in standard difference-in-differences research designs. One potential violation of this assumption, for instance, would occur if mergers are triggered by local economic decline, which also causes employers to lay off workers.

In our analysis, we take four approaches to dealing with identification concerns. First, we develop an approach to adapt standard event study methods for diagnosing parallel trends assumptions to our setting, where units face multiple treatments of differential intensity that occur with staggered timing. Second, we isolate specific, arguably clean, sources of variation in our instrument and test robustness to using only those sources of variation following [Borusyak and Hull \(forthcoming\)](#). Third, we exclude “failing” hospitals experiencing low bed occupancy before a merger from our panel of mergers, since their low occupancy might be a function of declines in local economies. Fourth, we show our results are robust to excluding units with outlier trends in outcomes before our treatment period begins and alternatively including year \times employer characteristics in lieu of year fixed effects.

¹⁵In 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.3, we include the alternative first-stage estimates we use throughout the analysis.

5.3.1 Quasi Event Study Approach

In standard difference-in-difference designs, researchers use event study approaches to show that a change in a binary, absorbing treatment status in year t does not appear to “affect” the difference in outcomes between treated and untreated units in $t - 1$, $t - 2$, and other years before the treatment was actually applied (i.e, the pre-period). Showing that the treatment does not impact pre-treatment outcomes provides suggestive evidence that untreated units would be likely to follow parallel trends to treated units in the absence of the treatment occurring and could thus serve as an appropriate control group. This approach is complicated in our setting. Because almost all employers are exposed to mergers and each merger differs in how intensely it effects a given employer, there is no strict delineation between treated and untreated units. Moreover, many employers are exposed to multiple merger “treatments,” meaning there is no strict delineation between pre-treatment and post-treatment periods for a unit.

As a result, we instead consider a continuous analogue to the event study. Our target parameter estimates the effect of the current year’s treatment intensity (i.e., the extent of merger-driven price increases) in a given year on outcomes in the present, past and future, net of any other effects on outcomes in that time period. We take the approach of [Schmidheiny and Siegloch \(2023\)](#) and estimate the following distributed-lag specification:

$$y_{it} = \sum_{k \in \{-7, \dots, 0, 2, \dots, 7\}} \gamma_{-k} z_{i(t+k)} + \theta_i + \kappa_t + \varepsilon_{it}. \quad (8)$$

We regress outcomes y in period t on the contemporaneous values of the instrument z , as well as lags and leads of z .¹⁶ The distributed-lag model is an inversion of the standard event study regression specification. Rather than estimate the effect of a treatment in t on an outcome in $t + k$, we estimate the effect of a treatment in $t - k$ on an outcome in t , net of any other effects.¹⁷ [Schmidheiny and Siegloch \(2023\)](#) show that this is equivalent to an event study approach in simple settings. This method benefits from not requiring us to define specific event occurrences in our setting, where there are many “events” of differing intensity.

To recover the equivalent of event study coefficients, we must cumulate the effects. That is, for $k \geq 0$, our event study estimate $\beta_k = \sum_{\ell=0}^k \gamma_{\ell}$, and for $k < 0$, $\beta_k = \sum_{\ell=-1}^k \gamma_{\ell}$. Our data only allow us to consistently estimate β_k for $k \in \{-2, 0, 1, 2\}$.¹⁸ We normalize β_{-1} to zero as in typical event

¹⁶Note that we estimate the reduced-form regression rather than the scaled two-stage least-squares regression. We do so due to the fact that, consistent with our identification assumptions, the effect of merger-driven price increases on *earlier* spending levels is close to zero. If we constructed the associated Wald estimator, it would produce an extremely large and biased estimate due to standard weak instruments bias, leading us to incorrectly reject parallel pre-merger trends.

¹⁷Note that we specifically use our notation to denote the coefficient as $-k$. The coefficient on z_{t+k} maps to the coefficient in event time $-k$ in a standard event study regression.

¹⁸This is the only range for which all mergers have a full range of relevant leads and lags. For example, we cannot

studies. With these β coefficients estimated, we can plot them in a way that is comparable to typical event study figures. The extent to which β_{-2} is different from β_{-1} serves as an analogue to a pre-trend test of parallel counterfactual trends: it measures the extent to which exposure to greater merger-driven price increases is correlated with changes in outcomes in earlier years.

In Figure 2, we present event study estimates of our first-stage as described above, regressing employer-level spending on simulated spending. We see a minimal pre-trend, with spending declining slightly in the years before a merger occurs. Once the merger is consummated (event time 0), spending increases precipitously and continues to rise in the following years.¹⁹

5.3.2 Isolating Variation

Variation in our simulated spending measure comes from three sources: the timing of when mergers occur, the scale of the price increases generated by each merger, and the extent to which an employer's workers receive care from specific hospitals. A natural concern is that some components of this variation are not plausibly independent of idiosyncratic employer-specific outcomes ε_{it} . For example, employers in more urban areas have more hospitals nearby and thus are likely to be exposed to more mergers than employers in more rural areas. If urban labor markets are on different trends than more rural labor markets, this could bias our estimates by introducing a confounding relationship between the instrument and the outcome. On the other hand, some components of the changes we observe in simulated spending are likely to be much more plausibly exogenous. Variation in merger timing, at least within a small time frame, is likely a function of idiosyncrasies in the speed of legal agreements and/or the extent of regulator involvement. Likewise, variation in the price effects generated by specific mergers is likely determined by pre-merger hospital market structure, as well as the extent to which hospital managers are skilled at exploiting market power to raise prices.²⁰ These mechanisms are much more likely to be independent of labor markets for non-hospital employment.

To the extent that some components of our instrumental variable (IV) may be contaminated by potential endogeneity concerns, we would ideally like to purge the unusable variation and rely exclusively on the plausibly exogenous variation to drive our instrument. To do so, we use a method developed by [Borusyak and Hull \(forthcoming\)](#). In their framework, bias comes from

estimate effects more than 2 years prior for mergers that occurred in 2010. While we could theoretically estimate such coefficients, they would put greater weight on units exposed to mergers in the middle of our sample window, and we do not want to mistake differences in treatment effects across event time for differences in unit-specific weights.

¹⁹The upward trend reflects institutional arrangements. Some insurer-hospital contracts are renegotiated infrequently, so even after a merger, it may take a year or two for the merging hospital's market power to be reflected in prices, as is also illustrated in Figure A.2.

²⁰We view initial hospital market structure as being determined by, primarily, the geographic distance between hospitals, much of which was determined by mid-20th century investment in hospital capacity as a result of the Hill-Burton Act in 1946 ([Chung et al., 2017](#)).

the fact that the expected value of the instrument—averaging over potential realizations of the exogenous components of the instrument—may be correlated with ε_{it} due to the presence of the potentially-endogenous components. Their solution, which we employ, is to construct that expected value, and difference it out of the primary instrument. Using this difference as an instrument will purge estimates of endogeneity bias, while still allowing for potentially-endogenous components to appropriately scale changes in other components.

In practice, this method requires that we explicitly decide which components of our instrument are plausibly exogenous, and make assumptions about their joint distributions. In our setting, the largest threat to exogeneity comes from our employer-by-hospital exposure measures, σ_{ih0} , which are a function of where employers are located and their proximity to hospital mergers. The concern with this component of our IV arises because an employer’s location determines both where its patients receive care and their exposure to other local economic trends. As a result, following [Borusyak and Hull \(forthcoming\)](#), we allow this variation in σ_{ih0} to be excluded from our IV. Following this approach means that our instrument is being driven by variation in the scale of the price increases generated by mergers and/or the timing of mergers. To implement the [Borusyak and Hull \(forthcoming\)](#) correction, we assume that the specific timing of the merger within the window we decide is random in that any merger we study could have potentially occurred in any year between 2010 and 2015 with uniform probability $\frac{1}{6}$. We additionally assume that the merger price effects, λ_{eh} , are drawn randomly from the empirical distribution.

We can then construct the expected instrument, $\tilde{z}_{it} = E[z_{it}]$, as the expected value of the instrument over the distribution of assumed-exogenous components. We then estimate the effect of health care prices on outcomes using the expression given in Equations 6 and 7, replacing our primary instrument z_{it} with a version of it purged of the expected instrument, $z_{it} - \tilde{z}_{it}$.

In Column (3) of Appendix Table A.3, we present the version of the first-stage where we assume that the only exogenous component of our instrument is the post-merger price increases λ_{eh} . We find a similar point estimate to our primary estimate given in Column (1), albeit with a somewhat lower F-statistic. This suggests that variation in price increases post-merger is an important source of variation in our instrument. We present the estimate from our first-stage regression when we assume the price effects *and* timing of mergers are exogenous (Column (2)) and when we assume only timing is exogenous (Column (4)). Using both price effects and timing to drive our instrument results in a similar coefficient and F-statistic relative to only relying on variation from price effects. By contrast, only using timing variation (Column (4)) severely reduces the F-statistic from 864.78 to 24.33.

For further robustness, instead of measuring post-merger price increases directly, we supplement our standard simulated spending measure with an alternative version where variation in the instrument is being driven exclusively by variation in merger-driven changes in hospitals’ market

power. Specifically, for each merger, we construct a structural measure of the expected post-merger increase in markups, following the “willingness to pay” (WTP) approach of [Capps et al. \(2003\)](#) and [Gowrisankaran et al. \(2015\)](#). Our process for constructing these WTP measures is described in detail in Appendix D of [Brot-Goldberg et al. \(forthcoming\)](#). We then reconstruct our instrument as described in 3, replacing λ_{eh} with ΔWTP_{eh} , the predicted change in markups. We present the results from this alternative first-stage in Column (5) of Appendix Table [A.3](#). We estimate a first-stage coefficient of 2.87 and an F-statistic of 2,924.00. The difference in coefficient magnitudes reflects the fact that markup changes are denominated in different units than prices.

Going forward, as we present results from our primary specification, we show that our results are robust to using these alternative instrument specifications.

6 Employer-Level Results

6.1 Health Insurance Outcomes

We begin by estimating the effect of rising health care prices on ESI premiums and the share of a firms’ employees enrolled in a HDHP. 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. We proxy for the share of an employer’s employees enrolled in a HDHP by measuring the share of their 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 HSAs for all employers, our analysis of insurance premiums requires that we use a 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 simple OLS results from estimating Equation (7) and shows no economically or statistically significant relationship between health care prices and insurance premiums. By contrast, we present our primary 2SLS estimates of Equation (7) in Panel B and find that a 1% increase in health care prices leads to a 1.0% increase in ESI premiums. These estimates suggest that rising health care prices are fully passed through into higher insurance premiums.²¹ Our relatively low power for this outcome largely reflects the fact that the 5500-linked employer panel only contains 3,970 unique employers compared to 140,300 in our primary sample.²² As we illustrate in Appendix Table [A.5](#), we also observe increases in insurance premiums

²¹Note that the average actuarial value of health plans in our sample, as measured in the HCCI data, is 0.81. As a result, our estimates imply a pass-through rate of above 1. However, our confidence interval on our premiums estimate is wide and includes 1.

²²We have insurance premiums for 3,970 employers. We can merge 81% of employers from our Form 5500

when we construct our instrument using WTP estimates of the effect of mergers on prices, rather than constructing our IV using realized post-merger price increases measured via a difference-in-difference regression.

Employers and insurers could respond to rising health care prices by reducing the actuarial value of their coverage (e.g., shifting policyholders onto higher deductible plans or plans with other forms of enhanced employee cost-sharing), rather than, or in addition to, raising insurance premiums. In Column (2) of Table 3, we present OLS and IV estimates of the effect of rising health prices on employer provision of HDHPs. 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.

6.2 Labor Market Outcomes

We next turn to estimating the effect of rising prices on labor market outcomes. We focus on employers' logged payroll and their count of workers as our primary outcomes. In Columns (1) and (2) of Panel A of Table 4, we present OLS estimates of Equation (7) and find no economically or statistically significant relationship between rising health care prices and employer payroll or their count of workers. By contrast, in Panel B, we present 2SLS estimates of Equation (7) and find that a 1% increase in health care prices reduces both payroll and employment by 0.36%.²³

Scaling these estimates dollar-for-dollar, a 1% increase in health care spending translates into a \$40.09 increase in spending per insurance plan member. The average employee in our sample, according to our HCCI data, has an insurance plan that includes 1 dependent, generating a \$80.18 spending increase per employee. The median employer in our sample has 132 employees, so 1% increase in in spending at an employer summed across their employees and employees' dependents is \$10,583.76. Our point estimate in Column (1) of Table 4 imply that a 1% increase in health spending leads to a 0.36% reduction in payroll, which is equivalent to an approximately \$17,900.00 decrease at the median employer. 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 health insurance costs

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 onto our analytic sample of EINs, we use a county-level measure of simulated spending rather than an employer-level measure.

²³That we find equal payroll and employment effects suggests minimal wage pass-through, implying that the incidence primarily falls on employers. 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. If employers respond to premium increases by raising the employee ESI contribution, this will lower non-wage compensation but not appear in our payroll measure.

into wage levels. Instead, it likely reflects the fact that employers respond to rising health spending by reducing the count of workers they employ and not necessarily just lowering their workers' wages. As a result, we interpret the gap between the amount of money received by hospitals and the reduction in payroll as the deadweight loss of the merger-induced transfer from other sectors to the health care sector to fund the increase in health care prices.

Our analytic sample includes employers from both the health sector and non-health sector. Non-health care employers are likely to experience reductions in 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 may experience *increases* in payroll and employment because higher health care prices raise their revenue. Likewise, health care employers in our panel may themselves be directly involved in the mergers that contribute to our instrument. As a result, in Table 4, we measure the employment effects at non-health care (Columns (3) and (4)) and health care employers (Columns (5) and (6)) separately.²⁴ 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 0.4%. 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. Going forward, in our employer-level analysis, we focus *exclusively* on employers outside the health care industry so that we can focus the effect of rising health care prices on labor market outcomes and ignore any potential direct effects of hospital mergers on monopsony power or rent-sharing, which could be present at health care employers.

We present event study estimates of the effect of rising health care prices on logged payroll and the count of workers at non-health employers in Panels A and B of Figure 3. 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 switches on. 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 in health care prices. The estimated coefficients approximately double 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, as Figure 2 shows, the first-stage coefficient also approximately doubles over this event time window. If we appropriately scale for the size of the first-stage, we would conclude that price increases primarily have immediate level effects on payroll and employment.

How large are our effects? In Appendix Table A.6, we benchmark our employment effects,

²⁴We identify employers in the health care industry by whether they have a NAICS code starting in '62' as reported to the IRS.

presented in Column (4) of Table 4, against the closest analogues in the economics literature: estimates of employers' responses to increases in payroll taxes. Like rising health care prices, payroll taxes raise the cost of retaining workers. Prior studies estimate that a 1 percentage point increase in payroll taxes reduces employer-level worker counts by between 0 and 3.4% depending on the study. Our estimates imply that an equivalent increase in labor costs due to rising health care prices would generate a 1.7% reduction in employment. Our estimate is closest to the most comparable study to our own, Johnston (2021), who estimates a 1.5% reduction in employment.

In Appendix Figure A.5, we show our baseline employer-level estimates are similar in both magnitude and precision across a range of alternative specifications: 1) using only instrumental variation in post-merger price increases across mergers; 2) using only instrumental variation in predicted markups; 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 2009.²⁵

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 partially be driven 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-2.²⁶ We then construct employer-by-county weights, ω_{ic0} , as the share of workers in a given

²⁵The 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 imprecise.

²⁶In cases where a given worker reports multiple residences on multiple W-2s, we choose according to the W-2 with the highest wages.

county c who worked for employer i in 2009. We can then construct our endogenous regressor x_{ct} and instrument z_{ct} as weighted averages:

$$x_{ct} = \sum_i \omega_{ic} x_{it}$$

$$z_{ct} = \sum_i \omega_{ic} z_{it}$$

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.²⁷ At the county-level, our mean change in simulated spending is 0.001 and the standard deviation is 0.0056. The top 25% of counties have an increase in simulated spending of 0.002 and the top 5% have an increase in simulated spending of 0.009.

With those constructed, our primary empirical approach is to run, as we do at the employer level, a series of 2SLS regressions:

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

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

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.4, 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 approximately 42, which is above standard thresholds for weak instrument tests.

7.1 Labor Market Outcomes

In Table 5, we present estimates of Equation (10) and show the impact of rising health care prices on county-level income per capita, the share of the population unemployed, and federal income tax revenues per capita. As we illustrate in Column (1) of Panel B, a 1% increase in health care prices leads to a 0.27% decrease in county-level income per capita. 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. However, as we illustrate in Appendix Table A.7, this overall income effect is driven

²⁷We 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.

by changes in the income of non-health care workers.²⁸ Notably, as we illustrate in Panel A, when we run our OLS, we observe a positive relationship between health spending and income per capita. Likewise, as our estimates in Appendix Table A.8 illustrate, we do not observe that increases in health care prices lead to an economically or statistically significant change in whether an individual receives self-employment income.

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, which we measure as the share of the population receiving unemployment insurance or receiving zero income.²⁹ Per our results in Appendix Table A.7, 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 4, 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, then it rises in later years as the local economy recovers. One concern with using IRS data to construct a population denominator is that individuals who become wholly separated from the labor market and do not receive UI or disability insurance might drop out of our population measure. As a result, in Appendix Figure A.6, we show our income event study using census data to construct a denominator. When we use this denominator in lieu of our IRS denominator, we observe a slightly larger overall effect and somewhat less of a return to baseline income two years after the spending shock. In Panel B of Figure 4, we present event study estimates of the effect of rising health care prices on unemployment. The event study illustrates that unemployment spikes immediately in the year of the shock then falls back down to zero. This reflects the fact that our measure of unemployment is a flow measure rather than a stock measure.

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 observe that a 1% increase in health care prices brings federal income tax receipts down by 0.4%. Likewise, since the employment effects we observe result in individuals receiving unemployment insurance,

²⁸We 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.

²⁹Our 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.

as we show in Columns (5) of Appendix Table A.8, we find that a 1% increase in health care prices leads to a 2.5% increase in UI spending. Collectively, these estimates highlight that the incidence of rising health care prices also falls on the state and federal governments, and not just employers and workers.

In Appendix Figure A.7, 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).³⁰ Third, we show our results are robust to interacting our year fixed effects with county population quartile fixed effects, so that all comparisons are done 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, 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, excluding the 25% most rural counties does not change our results. Finally, 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, but it does increase our standard errors.

7.2 Distributional Effects

As we discussed in Section 3, the effect of rising health spending should have a heterogeneous effect on individuals across the income distribution. To test 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.³¹ 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 5. We find close to zero unemployment effects at the top of the income distribution (i.e., those previously earning above \$100,000 - approximately the 85th 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. The

³⁰Note that our employment results are not robust to measuring the effect of mergers using WTP.

³¹E.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.

theory of health insurance serving as a “head tax” on workers would predict that these unemployment effects would be regressive, in that they would be higher for relatively lower-wage workers.

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

We also find relatively uniform effects in the middle of this range 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. However, as [Lurie and Miller \(2023\)](#) also find, there is a positive correlation between wages and health insurance generosity. As we highlight in [Section 3](#), rising health care prices generate larger increases in premiums for workers who are, *ex ante*, receiving more generous health insurance benefits. This pushes premium increases to be larger for higher-wage workers.

7.3 Mortality

Rising health care prices lead employers to lay off workers. A recent literature has documented that individuals who lose their job can face striking social consequences ([Eliason and Storrie, 2009](#)). In particular, a growing literature has found that job losses can induce increases in individuals’ risk of premature mortality, particularly via deaths from self-harm, such as suicide, overdose, and liver disease ([Eliason and Storrie, 2009](#); [Sullivan and von Wachter, 2009](#); [Pierce and Schott, 2020](#); [Venkataramani et al., 2020](#)). As a result, we analyze whether the job losses induced by rising health care prices also led to increases in mortality.

To do so, we analyze the county-level count of suicides, accidental poisonings, or poisonings of undetermined intent per 100,000 people via the CDC’s restricted mortality data and estimate Equation (10). We limit our analysis to working-age individuals aged between 25 and 64 at time of death (i.e., those who are most likely to receive ESI).³² As a placebo check, we construct three alternative measures of mortality that should be largely unaffected by job losses induced by rising health care prices: 1) suicides, accidental poisonings, or poisonings of undetermined intent among

³²Accidental poisonings and poisonings of undetermined intent typically reflect drug overdoses. We exclude liver disease as it is likely to take significant time to accumulate, whereas poisonings and suicides are acute. We define our death measures, following [Case and Deaton \(2015\)](#) and [Pierce and Schott \(2020\)](#), using International Statistical Classification of Diseases and Related Health Problems (ICD) 10 codes. We define suicides as those with codes for intentional self-poisoning (X60 - X69), other intentional self-harm (X70 - X84), and sequelae of intentional self-harm (Y87.0). We also include accidental poisonings (X40 - X45) and poisonings with undetermined intent (Y10 - Y19). We also include prescription drug complications (Y45, Y47, and Y40) and other harms with other undetermined intent (Y20 - Y25).

individuals age 65 and older, who are likely to be outside the labor market; 2) all-cause mortality excluding suicides, accidental poisonings, or poisonings of undetermined intent among individuals age 25 to 64; and 3) cancer mortality among individuals age 25 to 64.

As we illustrate in Column (1) in Panel B of Table 6, a 1% increase in health care prices increases county-level suicides and overdoses by 0.62 death per 100,000 population among working-age adults (a 2.7% increase). OLS estimates are positive as well, but much smaller in magnitude. By contrast, as we illustrate in Column (2), we do not observe a statistically significant change in suicides and overdoses for individuals aged 65 and older, who we think are outside the labor market and therefore unlikely to be impacted by the increase in health care prices. Likewise, as we illustrate in Columns (3) and (4), we do not observe that health care price increases lead to an increase in all deaths excluding suicides and overdoses or an increase in deaths from cancer.

In Figure 6, we present our event study estimates for deaths from suicides and overdoses among working-age adults. Here, we see flat trends in mortality in the years prior to the spending shock occurring. This suggests that merger shocks are uncorrelated with pre-existing mortality trends. Then, we show that the mortality effects from suicides and overdoses occur at $k = 1$, one year *after* the spending shock, with de minimis effects in the year of the shock. The implied timing of deaths from suicides and overdoses — a flow measure — tells a clear story when combined with our labor market result: when health care prices rise, there is an immediate disemployment effect. Then, a portion of those who lost their job succumb to suicide or overdose in the year after.

As we illustrate in Appendix Table A.9, the past literature has found that there is approximately 1 death per 300 to 600 job losses (Sullivan and von Wachter, 2009; Eliason and Storrie, 2009; Pierce and Schott, 2020; Venkataramani et al., 2020). The estimated death-per-job-loss rate has increased over time concurrently with the increasing intensity of the opioid epidemic in the US. We use our estimates of the effect of rising health spending on employment in Column (2) in Table 5 to scale our effect in terms of deaths per job loss induced by rising health spending. Combining our estimates from Table 5 and Table 6, we observe approximately 1 death per 140 job losses in our sample.³³ Our estimates are likely slightly higher than other studies in the literature because, whereas other studies look at the effect of job loss on mortality, our measure (receipt of unemployment insurance) captures a true separation from the labor market and not simply the loss of a job (which keeps open the possibility that a individual can maintain employment at another establishment or on their own).

In Appendix Figure A.8, we show how our county-level mortality results shift when we shift our source of variation and alter our sample of counties. First, we show that our results remain similarly scaled, but lose precision when we rely exclusively on the variation in post-merger price increases across transactions to drive our instrument. Second, we show that our results are robust to

³³We calculate this as $100,000 * (\text{IV coefficient on unemployment} / \text{IV coefficient on deaths})$.

measuring the effect of mergers via WTP. Third, we show our results are robust to interacting our year fixed effects with county population quartile fixed effects, so that all comparisons are done within county population quartiles. Fourth, we show our results are similarly scaled, but lose some precision when we exclude the 25% of counties with the lowest income per capita in 2009. By contrast, our results remain robust when we exclude the 25% of counties with the lowest income growth from 2006 to 2009. Fifth, we show our results are robust to excluding the 25% of counties with the highest unemployment in 2009 and the counties with the biggest increase in unemployment from 2006 to 2009. Sixth, we show our results are robust to excluding the 25% most rural counties. Finally, we show 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.

8 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 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 merging hospitals in our analytic sample, we find that the average post-merger price increase is 1.2%. Our estimates from Section 7 imply that, on average, these mergers led to a \$6 million reduction in income, 39 job losses, and a \$1.3 million reduction in income tax payments.³⁴ 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. If instead, we simulate a uniform, hypothetical, anticompetitive merger across all hospitals in our analytic sample, we find that a 5% price increase at the average merging hospital would have led to a \$32 million reduction in income, a \$6.8 million reduction in federal tax revenue, 203 job losses, and

³⁴Note 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.

1 to 2 deaths of despair. The US Department of Transportation (DOT) estimates the value of a statistical life in 2015 at \$9.6 million ([Department of Transportation, 2022](#)). This implies that, on average, a hospital merger that raised prices by 5% or more led to approximately \$41.6 million in economic harm (\$32 million in forgone wages and \$9.6 million from deaths). In our sample, 125 of 304 mergers led to price increases of 5% or more.

9 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 immense harm to local economies. Rising health care prices raise ESI premiums, lower employment (both at individual employers and overall in local economies), reduce workers' earnings, lower tax revenue, squeeze government budgets, and increase suicides and overdoses. Moreover, the majority of these negative consequences are borne by lower- and middle-income individuals and not by individuals earning more than \$100,000 annually.

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](#)).³⁵ Given share of total health services that hospital care represents, this price increase caused a 10% increase in overall health care spending. As a result, our estimates imply that the price growth between 2007 and 2014 reduced workers' incomes by 2.7%, increased unemployment by approximately 0.86 percentage points (a 10% increase or 1.44 million jobs lost), lowered federal income tax revenues by 3.4%, and led to an increase in suicides and overdoses of 6.2 per 100,000 population (approximately 10,000 additional deaths across working-age adults). Based on the DOT's 2015 value of a statistical life (\$9.6 million), the economic value of this loss of life would be approximately \$96 billion. 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 of the scale we estimate.

As discussed in Section 8, the average hospital merger led to reductions in local income of \$6 million and mergers that generated price increases of over 5% reduced welfare by an average of \$42 million. From 2002 to 2020, there were over 1,000 hospital mergers in the US. During this period, the FTC took enforcement actions against 13 transactions. Our results are concerning because, as [Brot-Goldberg et al. \(forthcoming\)](#) note, approximately 20% of consummated mergers in our sample could be predicted, *ex ante*, to raise prices using standard screening tools available to the FTC and, in practice, did lead to *ex post* price increases that were often far higher than 5%. This suggests that a great deal of highly-damaging hospital mergers were observed by but were not

³⁵Note that their estimates are net of inflation.

stopped by regulators, and that these mergers have had substantial effects on labor market outcomes and mortality outside the health sector.

Ultimately, this work highlights that health care price growth is generating severe macroeconomic and social consequences in the US. In the absence of concrete steps to address health care price growth, rising health spending will raise labor costs and reduce business dynamism outside the health sector, put pressure on the federal budget, and exacerbate income inequality. Rising health care spending will also precipitate suicides and overdoses. As a result, we hope this research motivates future analysis of strategies to address health care price growth in the US and ways to screen for and challenge hospital mergers that lessen competition and lead to higher prices. On the academic front, we hope this work motivates future analysis of the absolute and distributional effects of rising health care prices and health spending in the US. For instance, while we have focused on the harms to non-health care workers, it will be important for future research to assess who receives the rents accrued from health care price increases. Finally, we hope this work motivates further study of how rising health spending impacts regional growth and productivity across the US.

References

- Anand, Priyanka**, “Health Insurance Costs and Employee Compensation: Evidence from the National Compensation Survey,” *Health Economics*, 2016, 26, 1601–1616.
- 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,” 2023.
- 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.
- Baqae, 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.
- Berry, Steven, Martin Gaynor, and Fiona Scott Morton**, “Do Increasing Markups Matter? Lessons from Empirical Industrial Organization,” *Journal of Economic Perspectives*, August 2019, 33 (3), 44–68.
- 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*, forthcoming.
- 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*, forthcoming.
- 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-Goldberg, Zarek C., Zack Cooper, Stuart V. Craig, and Lev Klarnet**, “Is There Too Little Antitrust Enforcement in the US Hospital Sector?,” *American Economic Review: Insights*, forthcoming.
- 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.
- Case, Anne and Angus Deaton**, “Rising Morbidity and Mortality in Midlife Among White Non-Hispanic Americans in the 21st Century,” *Proceedings of the National Academy of Science*, 2015, *112*, 15,078–15,083.
- 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.
- Currie, Janet and Brigitte C. Madrian**, “Health, Health Insurance and the Labor Market,” in O. Ashenfelter and D. Card, eds., *Handbook of Labor Economics*, Vol. 3, Elsevier, 1999, chapter 50, pp. 3309–3416.
- 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).
- Dafny, Leemore, Kate Ho, and Edward Kong**, “How Do Copayment Coupons Affect Branded Drug Prices and Quantities Purchased?,” *American Economic Journal: Economic Policy*, forthcoming.
- Dafny, Leemore S.**, “How Do Hospitals Respond to Price Changes?,” *American Economic Review*, 2005, 95 (5), 1525–1547.
- De Loecker, Jan, Jan Eeckhout, and Gabriel Unger**, “The Rise of Market Power and Macroeconomic Implications,” *Quarterly Journal of Economics*, 2020, 135, 561–644.
- Department of Transportation**, “Departmental Guidance on Valuation of a Statistical Life in Economic Analysis,” 2022.

- Eliason, Marcus and Donald Storrie**, “Does Job Loss Shorten Life?,” *Journal of Human Resources*, 2009, 44 (2).
- Finkelstein, Amy, Casey McQuillan, Owen Zidar, and Eric Zwick**, “The Health Wedge and Labor Market Inequality,” 2023. NBER Working Paper No. 31091.
- 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 Firm Labor Composition and Technology Investment?,” 2023.
- 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**, “Health Insurance and the Labor Market,” in A. J. Culyer and J. P. Newhouse, eds., *Handbook of Health Economics*, 1 ed., Vol. 1, Elsevier, 2000, chapter 12, pp. 645–706.
- 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,” 2023.
- 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.
- 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.
- Lin, Haizhen, Ian M. McCarthy, and Michael Richards**, “Hospital Pricing Following Integration with Physician Practices,” *Journal of Health Economics*, 2021, 77.
- Lobel, Felipe**, “The Visible Hand of Firms: Consequences for Efficiency and Incidence of Payroll Taxation,” 2023.
- 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.
- Pierce, Justin R. and Peter K. Schott**, “Trade Liberalization and Mortality: Evidence from US Counties,” *American Economic Review: Insights*, March 2020, 2 (1), 47–64.
- 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.

- Schmidheiny, Kurt and Sebastian Siegloch**, “On Event Studies and Distributed-Lags in Two-Way Fixed Effects Models: Identification, Equivalence, and Generalization,” *Journal of Applied Econometrics*, 2023, 38 (5), 695–713.
- Starr, Paul**, *The Social Transformation of American Medicine: The Rise of a Sovereign Profession and the Making of a Vast Industry*, Basic Books, 1982.
- Sullivan, Daniel and Till von Wachter**, “Job Displacement and Mortality: An Analysis Using Administrative Data,” *Quarterly Journal of Economics*, 2009, 124 (3), 1265–1306.
- 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.
- U.S. Department of Justice and the Federal Trade Commission**, “Horizontal Merger Guidelines,” 2010.
- Venkataramani, Atheendar S., Elizabeth F. Bair, Rourke L. O’Brien, and Alexander C. Tsai**, “Association Between Automotive Assembly Plant Closures and Opioid Overdose Mortality in the United States: A Difference-in-Differences Analysis,” *JAMA Internal Medicine*, 2020, 180 (2), 254–262.

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*	12,721,000	25,129,000	2,391,000	4,977,000	11,418,000	140,300
Employer Count of Workers	297	509	75	132	282	140,300
Employer Average Wages per Worker	41,339	25,431	23,408	36,525	52,809	140,300
Share of Employees with Premiums	0.511	0.314	0.215	0.597	0.770	39,341
Share of Employees with a Health Savings Account	0.038	0.121	0.000	0.000	0.000	140,300
Health Spending per Beneficiary	4,099	704	3,649	4,039	4,478	140,300
Premiums from 5500 Data	5,036	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	41,908	9,205	35,872	39,911	45,369	1,709
Share with Unemployment Insurance	0.036	0.017	0.025	0.032	0.042	1,709
Share with Zero Income	0.054	0.014	0.043	0.053	0.062	1,709
Share Unemployed	0.089	0.024	0.073	0.085	0.101	1,709
Unemployment Insurance Payments per Capita	482	395	185	363	664	1,709
Share Self-Employed	0.110	0.023	0.095	0.107	0.123	1,709
Share Moving Annually	0.066	0.021	0.052	0.063	0.076	1,709
Income Tax Withholdings per Capita	7,009	2,088	5,634	6,583	7,836	1,709
Health Spending Per Beneficiary	4,210	444	3,910	4,182	4,479	1,709
All Deaths per 100k People	237	72	184	229	282	1,709
Deaths from Suicides and Overdose per 100k People	23	12	15	21	29	1,709
Deaths from Cancer per 100k People	65	19	51	63	76	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 share unemployed as share of individuals with either positive unemployment insurance receipts and/or with zero income in the year. The deaths per 100,000 measures are from the Center for Disease Control and Prevention's Restricted Mortality Database.

* Rounded to \$1,000.

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

	Log(Health Spending per Beneficiary)		
	(1)	(2)	(3)
Simulated Spending	0.649*** (0.022)	0.521*** (0.023)	2.458*** (0.051)
Instrument	Baseline	Only Price Effects	WTP
Mean Dependent Variable	4,042	4,042	4,042
Observations	1,403,000	1,403,000	1,403,000
Number of Unique Employers	140,300	140,300	140,300
F-Statistic on First Stage	864.776	528.533	2,350.093

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 (6). Each estimate includes employer and year fixed effects. 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)
Health Spending	0.025 (0.032)	0.012*** (0.001)
Panel B: IV Estimates		
	Log(Insurance Premiums) (1)	Share of Employees with a Health Savings Account (2)
Health Spending <i>(Instrumented using Merger-Driven Price Increases)</i>	0.947* (0.535)	0.0004 (0.042)
Mean Dependent Variable	5,036	0.038
Observations	39,700	1,403,000
Number of Unique Employers	3,970	140,300
F-Statistic on First Stage	43.391	864.776

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. 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 comes 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)
Health Spending	0.005 (0.004)	0.003 (0.004)	0.005 (0.004)	0.004 (0.004)	0.003 (0.019)	0.001 (0.018)

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)
Health Spending <i>(Instrumented using Merger-Driven Price Increases)</i>	-0.362*** (0.130)	-0.356*** (0.129)	-0.373*** (0.134)	-0.402*** (0.133)	-0.256 (0.558)	0.283 (0.551)
Mean Dependent Variable*	12,721,000	297	13,045,000	304	8,242,000	203
Observations	1,403,000	1,403,000	1,308,290	1,308,290	94,710	94,710
Number of Unique Employers	140,300	140,300	130,829	130,829	9,471	9,471
F-Statistic on First Stage	864.776	864.776	813.863	813.863	51.582	51.582

Notes: This table presents ordinary least squares (Panel A) and instrumental variables (Panel B) coefficient 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 comes 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.

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)
Health Spending	0.018* (0.010)	0.011*** (0.003)	0.019 (0.012)
Panel B: IV Estimates			
	Log(Income per Capita) (1)	Share Unemployed (2)	Log(Income Tax per Capita) (3)
Health Spending <i>(Instrumented using Merger-Driven Price Increases)</i>	-0.268* (0.149)	0.086** (0.041)	-0.358* (0.188)
Mean Dependent Variable	41,908	0.089	7,009
Observations	17,090	17,090	17,090
Number of Unique Counties	1,709	1,709	1,709
F-Statistic on First Stage	41.960	41.960	41.960

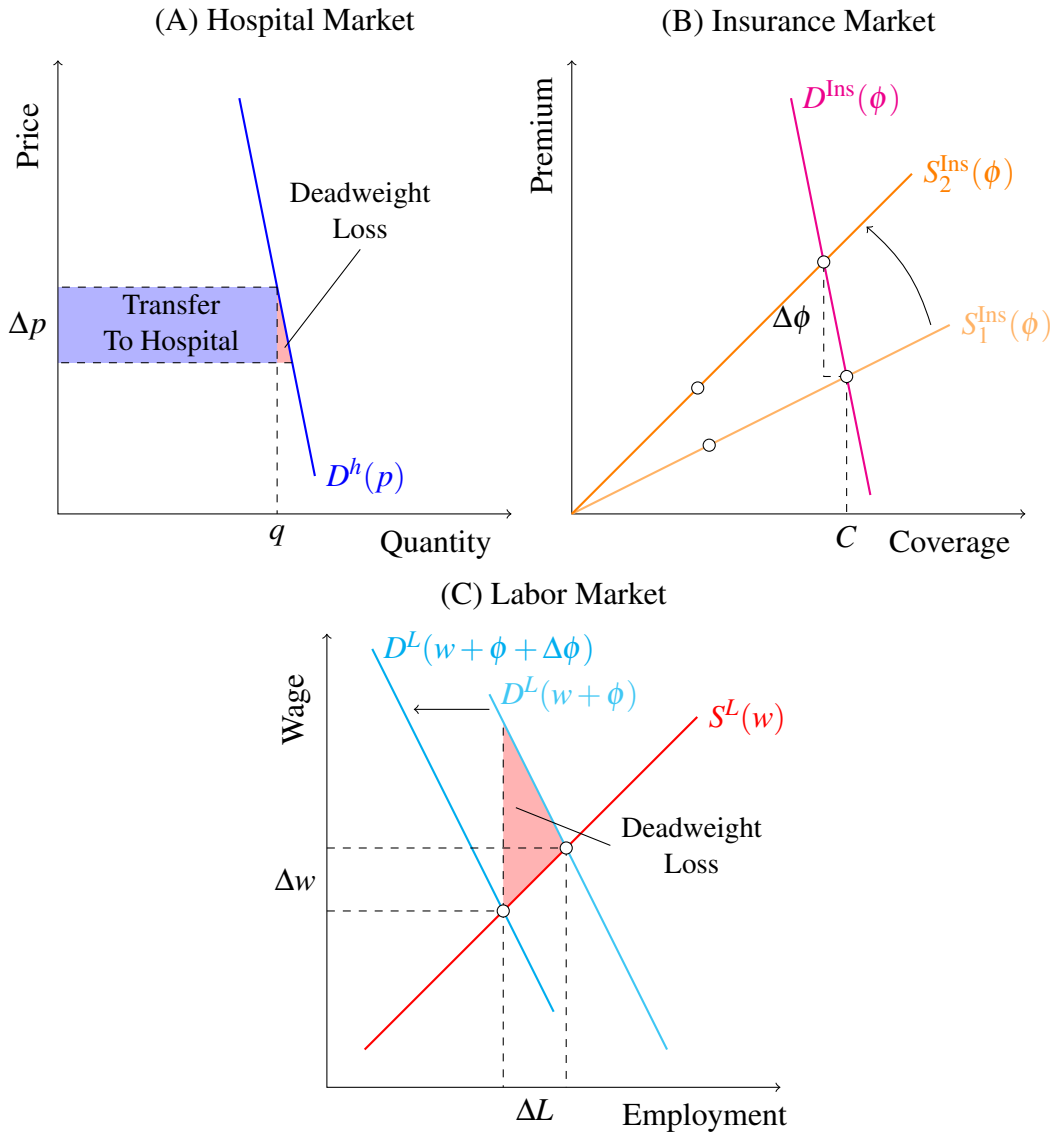
Notes: This table presents ordinary least squares (Panel A) and instrumental variables (Panel B) coefficient estimates from regressions of annual county-level log income per capita (Column (1)), share of the population collecting unemployment insurance or earning zero labor income (Column (2)), and log 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 comes 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 6: The Impact of Rising Health Care Prices on County-Level Mortality

Panel A: OLS Estimates				
	Deaths from Suicide and Overdoses Age 25-64 (1)	Deaths from Suicide and Overdoses Age 65 and Up (2)	All Deaths Excluding Suicides and Overdoses Age 25-64 (3)	Deaths from Cancer Age 25-64 (4)
Health Spending	15.131*** (2.362)	0.185 (0.737)	-4.417 (6.324)	9.059*** (3.162)
Panel B: IV Estimates				
	Deaths from Suicide and Overdoses Age 25-64 (1)	Deaths from Suicide and Overdoses Age 65 and Up (2)	All Deaths Excluding Suicides and Overdoses Age 25-64 (3)	Deaths from Cancer Age 25-64 (4)
Health Spending <i>(Instrumented using Merger- Driven Price Increases)</i>	61.873** (29.744)	-10.428 (9.589)	-39.985 (73.574)	6.362 (34.464)
Mean Dependent Variable	23.467	3.402	213.858	65.207
Observations	17,090	17,090	17,090	17,090
F-Statistic on First Stage	41.960	41.960	41.960	41.960


Notes: This table presents ordinary least squares (Panel A) and instrumental variables (Panel B) coefficient estimates from regressions of county-level deaths per 100,000 population on county-level annual spending per beneficiary, instrumenting for annual spending per beneficiary with county-level simulated spending per beneficiary. Column (1) presents our primary estimates where the outcome is deaths from suicide and overdose for those aged 25-64. Columns (2) through (4) present three placebo tests: deaths from suicide and overdose for those aged 65 and up (2), deaths from any cause *other than* suicide and overdose for those aged 25-64 (3), and deaths from cancer for those aged 25-64 (4). Each estimate includes county and year fixed effects. Mortality data come from the Centers for Disease Control and Prevention's Restricted Mortality Database. 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



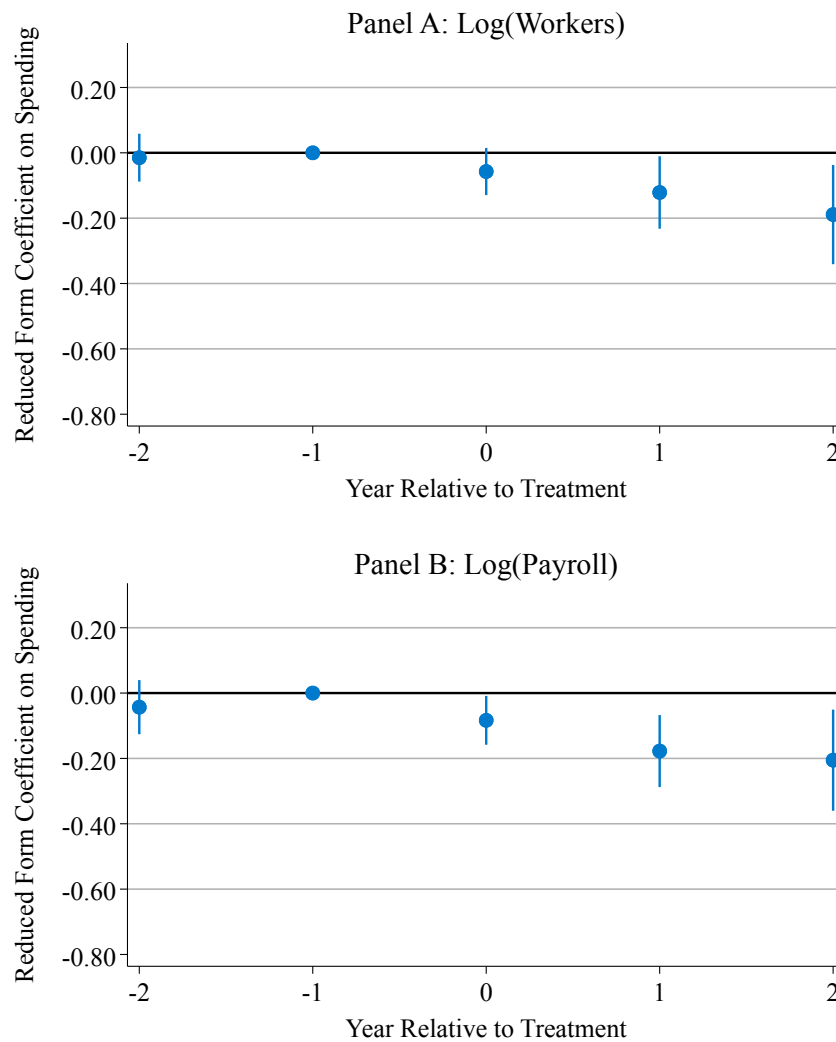
Notes: This illustrates our theory of how rising health care prices results rising insurance premiums and downstream changes in labor market outcomes. In Panel A, we highlight the fact that, after a merger (or other source of rent-seeking), prices rise. This generates deadweight loss (the red triangle), 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: Event Study Estimates of first-stage: Regressing Employer Spending on Simulated Employer Spending



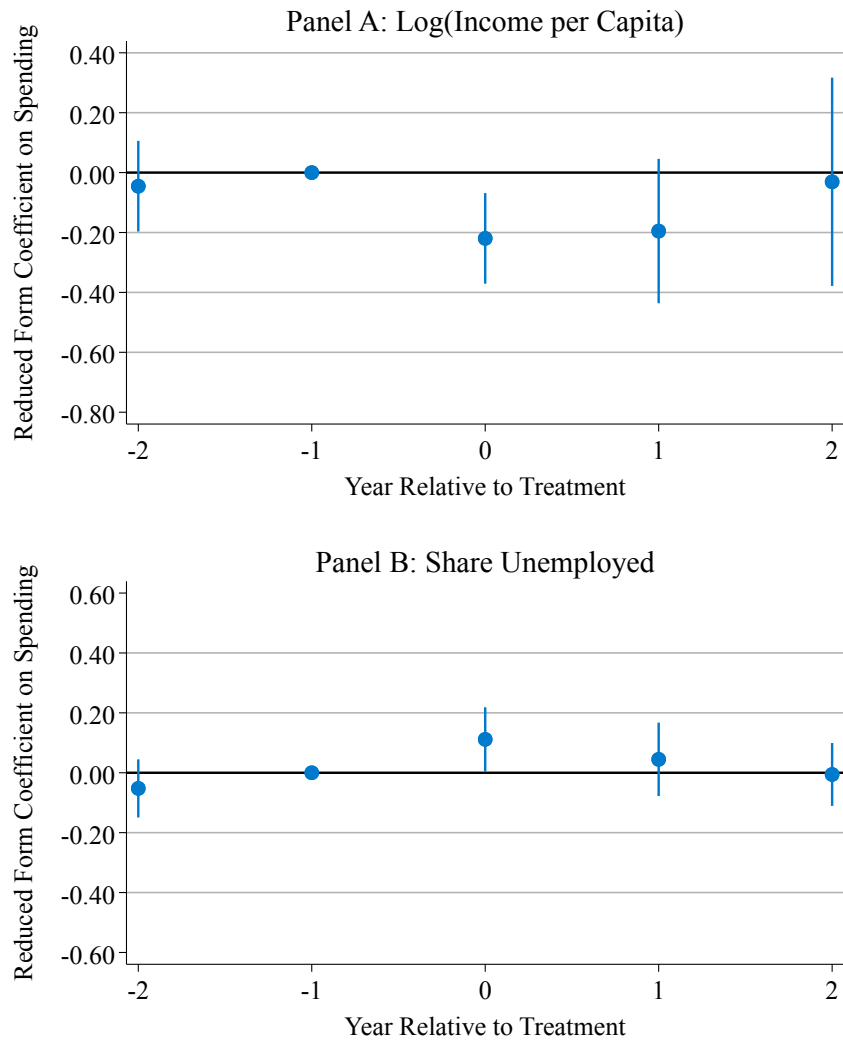
Notes: This figure presents event study-style coefficient estimates from our first-stage, as described in Section 5.3.1, where we regress employer-level annual spending per beneficiary on leads and lags of employer-level simulated spending per beneficiary. The dots are point estimates and the vertical lines are 95% confidence intervals.

Figure 3: 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-style estimates of the effect on rising health care prices on employer-level labor market outcomes, as described in Section 5.3.1, where we regress employer-level labor market outcomes on leads and lags of employer-level simulated spending per beneficiary. The dots are point estimates and the vertical lines are 95% confidence intervals.

Figure 4: 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-style estimates of the effect on rising health care prices on county-level labor market outcomes, as described in Section 5.3.1, where we regress county-level labor market outcomes on leads and lags of county-level simulated spending per beneficiary. The dots are point estimates and the vertical lines are 95% confidence intervals.

Figure 5: 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 in the prior year. The dots represent point estimates and vertical lines are 95% confidence intervals.

Figure 6: The Impact of Rising Health Care Prices on County-Level Deaths from Suicide and Overdose per 100,000 People



Notes: This figure presents event study-style estimates of the effect on rising health care prices on county-level mortality outcomes, as described in Section 5.3.1, where we regress county-level mortality outcomes on leads and lags of county-level simulated spending per beneficiary. The dots are point estimates and the 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

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

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:

$$p_{ht}^{INDEX} = \hat{\alpha}_{ht} + \hat{\beta}\bar{X} + \hat{\pi}_{dt}\bar{dt}. \quad (12)$$

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 under 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 under age 65. 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 Premiums in the Form 5500

B.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.

B.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 4 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 8 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 percentile of the premium distribution.

The Form 5500 is the only publicly available data capturing employer-level health insurance premiums. However, the Medical Expenditure Panel Survey (MEPS) reports annual state-level estimates of average premiums by firm size. In Figure [A.9](#), we plot average state-level premiums in the Form 5500 against those in the MEPS. We find that the two measures are similar, indicating that the data we construct from the Form 5500 capture the same broad levels and trends in premiums.

C Additional Tables and Figures

Table A.1: Comparison of Analytic Sample of Employers to Universe of Employers - 2008-2017

	Overall Analytic Sample Mean (1)	Non-HC Analytic Sample Mean (2)	HC Analytic Sample Mean (3)	Box DD Analytic Sample Mean (4)	5500 Analytic Sample Mean (5)
Health Spending per Beneficiary	4,099	4,100	4,079	4,261	
Share of Employees with a Health Savings Account	0.038	0.039	0.032	0.054	
Employer Total Payroll*	12,721,000	13,045,000	8,242,000	35,631,000	
Employer Count of Workers	297	304	203	784	
Employer Average Wages per Worker	41,339	41,463	39,624	45,184	
Share of Employees with Premiums from 5500 Data				0.511	5,036
Observations	140,300	130,829	9,471	39,341	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, 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.

* Rounded to \$1,000.

Table A.2: Comparison of Analytic Sample of Counties to Universe of Counties - 2008-2017

	All Counties	Analytic Sample
	Mean (1)	Mean (2)
Income Per Capita	38,378	41,908
Share with Unemployment Insurance	0.033	0.036
Share with Zero Income	0.059	0.054
Share Unemployed	0.092	0.089
Unemployment Insurance Payments per Capita	396	482
Share Self-Employed	0.126	0.110
Share Moving Annually	0.067	0.066
Income Tax Withholdings per Capita	6,054	7,009
All Deaths per 100k	241	237
Deaths from Suicides and Overdose per 100k	21	23
Deaths from Cancer per 100k	66	65
Observations	3,182	1,709

Notes: This table compares our analytic sample of counties with the universe of counties in the US using data from 2008 to 2017. Our analytic sample captures approximately 96% of the US 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 share unemployed as share of individuals with either positive unemployment insurance receipts and/or with zero income in the year. The deaths per 100,000 people measures are from the Center for Disease Control and Prevention's Restricted Mortality Database.

Table A.3: Alternative Employer-Level first-stage Estimates

	Baseline (1)	B&H Price and Timing (2)	B&H Price (3)	B&H Timing (4)	WTP (5)	Health Care Employers (6)	Non-Health Care Employers (7)
Simulated Spending	0.649*** (0.022)	0.536*** (0.022)	0.521*** (0.023)	-0.184*** (0.037)	2.458*** (0.053)	0.604*** (0.084)	0.652*** (0.023)
Observations	1,403,000	1,403,000	1,403,000	1,403,000	1,403,000	94,710	1,308,290
Number of Unique Employers	140,300	140,300	140,300	140,300	140,300	9,471	130,829
F-Statistic on First Stage	864.776	567.689	528.533	24.330	2,350.093	51.582	813.863

Notes: This table presents our employer-level first-stage estimates of Equation (6) on various sub-samples used throughout our analysis. It shows estimates of a regression of employer-level annual health spending per beneficiary (estimated based on each employee's county of residence) on employer-level simulated spending per beneficiary (also estimated based on each employee's county of residence). Each estimate includes employer and year fixed effects. 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.4: Alternative County-Level first-stage Estimates

	Baseline (1)	B&H Price and Timing (2)	B&H Price (3)	B&H Timing (4)	WTP (5)
Simulated Spending	1.034*** (0.160)	0.899*** (0.157)	0.887*** (0.158)	-0.032 (0.288)	5.799*** (0.947)
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	41.960	32.810	31.580	0.010	37.480

Notes: This table presents our county-level first-stage estimates of Equation (9) on various sub-samples used throughout our analysis. We regress county-level annual health spending per beneficiary on county-level simulated spending per beneficiary. Standard errors are presented in parentheses and clustered at the county-level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.5: Alternative Insurance Premiums and HSA Results using WTP

Panel A: OLS Estimates		
	Log(Insurance Premiums) (1)	Share of Employees with a Health Savings Account (2)
Health Spending	0.025 (0.032)	0.012*** (0.001)
Panel B: IV Estimates		
	Log(Insurance Premiums) (1)	Share of Employees with a Health Savings Account (2)
Health Spending <i>(Instrumented using Merger-Driven Δ WTP)</i>	0.868* (0.471)	-0.037 (0.025)
Mean Dependent Variable	5,036	0.038
Observations	39,700	1,403,000
Number of Unique Employers	3,970	140,300
F-Statistic on First Stage	89.600	2,923.997

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 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 comes 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.6: Comparison of Our Employment Estimate to Prior Payroll Tax Elasticity Estimates

	Setting	Disemployment Response to 1pp Payroll Tax Increase
Anderson and Meyer (1997)	US, 1978-1984	0.7-0.9%
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%
Johnston (2021)	US, 2003-2012	1.5%
Bíró et al. (2022)	Hungary, 2010-2015	0.3%
Guo (2023)	US, 2008-2013	1.1-2.4%
Lobel (2023)	Brazil, 2008-2017	0.5%
Our estimate [†]	US, 2008-2017	1.7%

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 since the average HCCI spending per worker is \$10,000 and the average payroll per worker in our sample is \$42,851, a 1pp payroll tax increase is \$428 per worker, roughly equivalent to a 4.3% increase in health care prices. Our estimate of the employment elasticity of health care prices for employers outside the health care industry is -0.4 (from Table Table 4). Multiplying these quantities together, our results imply that a 1pp payroll tax would reduce employment by 1.7%.

Table A.7: The Impact of Rising Health Care Prices on County-Level Income and Employment - Non-Health Care

Panel A: OLS Estimates			
	Log(Income per Capita) (1)	Share Unemployed (2)	Log(Income Tax per Capita) (3)
Health Spending	0.023** (0.011)	0.010*** (0.003)	0.025* (0.013)
Panel B: IV Estimates			
	Log(Income per Capita) (1)	Share Unemployed (2)	Log(Income Tax per Capita) (3)
Health Spending <i>(Instrumented using Merger-Driven Price Increases)</i>	-0.336** (0.162)	0.092** (0.043)	-0.434** (0.208)
Mean Dependent Variable	42,767	0.092	7,128
Observations	17,090	17,090	17,090
Number of Unique Counties	1,709	1,709	1,709
F-Statistic on First Stage	41.960	41.960	41.960

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. We restrict our analysis in this sample to employers outside the health sector. We identify employers in the health care industry by whether they have a North American Industry Classification System code starting in '62', as reported to the Internal Revenue Service (IRS). Each estimate includes county and year fixed effects. Our labor market and tax revenue data comes from the IRS. 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.8: Additional County-Level Labor Market Outcomes

	Share Moving	Share Self-employed	Share with Zero Income	Share Receiving Unemployment Insurance	Log (Unemployment Insurance Payments per Capita)
	(1)	(2)	(3)	(4)	(5)
Health Spending <i>(Instrumented using Merger- Driven Price Increases)</i>	-0.003 (0.014)	0.001 (0.026)	0.001 (0.014)	0.085** (0.037)	2.511* (1.479)
Mean Dependent Variable	0.066	0.110	0.054	0.036	482.128
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	41.960	41.960	41.960	41.960	41.960

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 comes 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

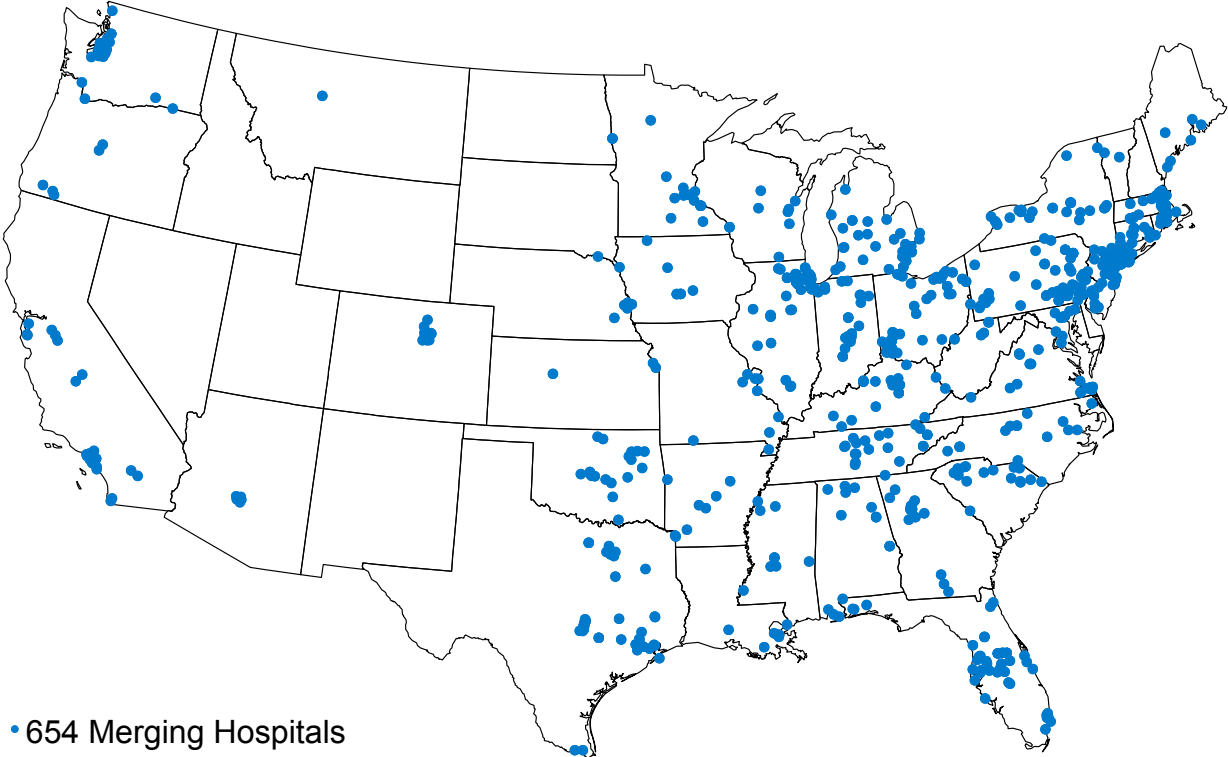
Table A.9: Comparison of Our Death Per Labor Market Separation Estimate to the Prior Deaths per Job Loss Estimates

	Setting	Time Period	Death per job losses
Sullivan and Von Wachter (2009)	USA (PA)	1970s, 1980s	1 in 546
Eliason and Storrie (2009)	Sweden	1980s	1 in 587
Pierce and Schott (2020)	USA	2000s	1 in 400
Vekataramani et al. (2020)	USA	2000s	1 in 300
Our estimate*:			1 in 140

Notes: We measure labor market separations, not job losses, which may explain higher impact.


*Calculated as: $100,000 * (\text{IV coefficient on unemployment} / \text{IV coefficient on deaths})$.

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



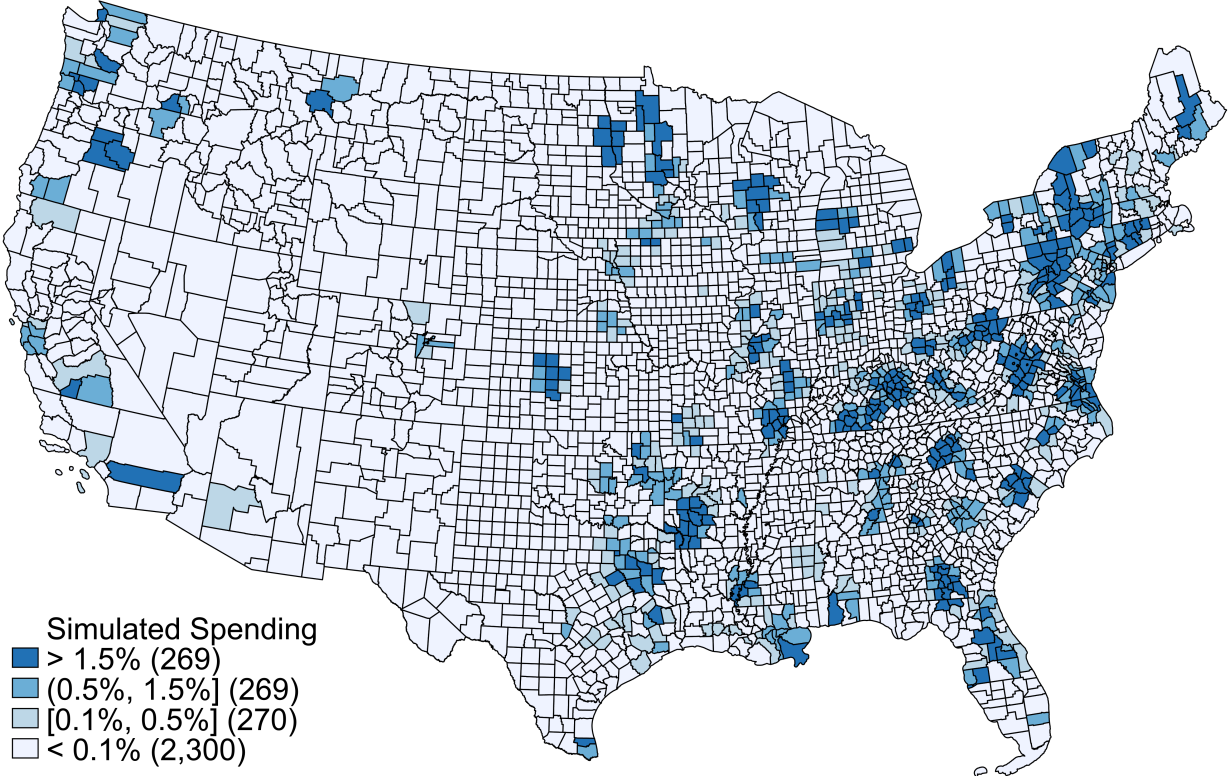
Notes: This figure presents difference-in-difference estimates of the effect of 2010 to 2015 hospital mergers on hospital prices (inpatient, outpatient, and both combined), given by Equation (2), on our sample of 304 hospital mergers that occurred between 2010 and 2015. We use empirical Bayes to shrink our λ_{eh} estimates in order to reduce noise driven by measurement error. Each dot represents a point estimate, and the vertical line displays the corresponding 95% confidence interval. Hospital pricing data come from the Health Care Cost Institute.

Figure A.3: Bin Scatter Plot of Employer-Level Health Spending Per Beneficiary and Simulated Health Spending



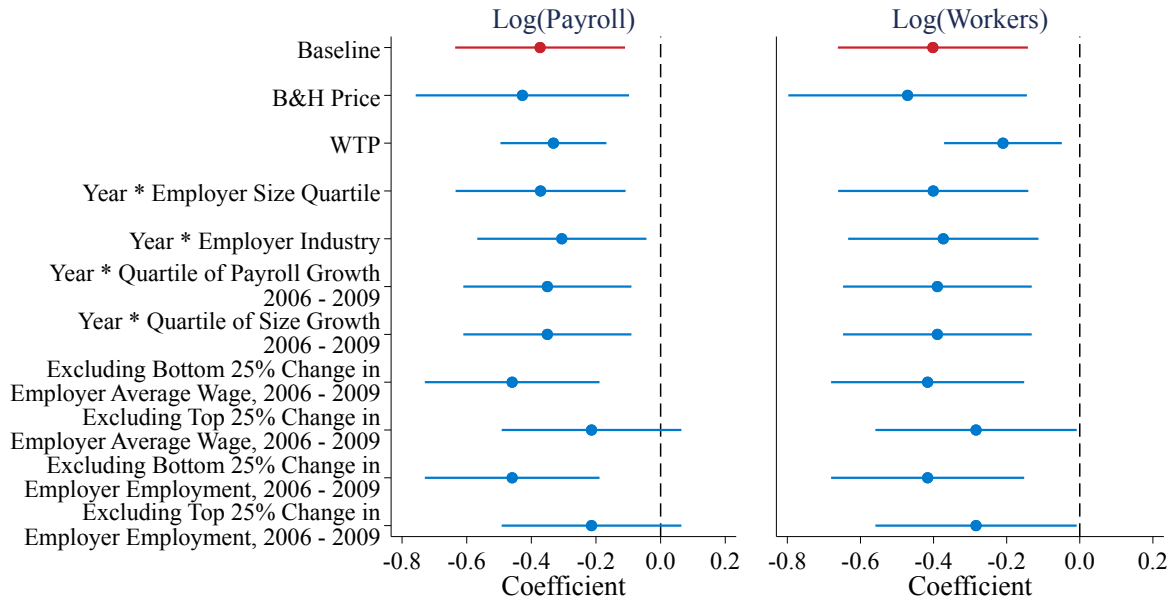
Notes: This figure presents a binned scatter plot of employer-level annual health spending per beneficiary against employer-level simulated spending per beneficiary, taking into account employer and year fixed effects.

Figure A.4: 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.5: Robustness Tests of Employer-Level Employment Effects



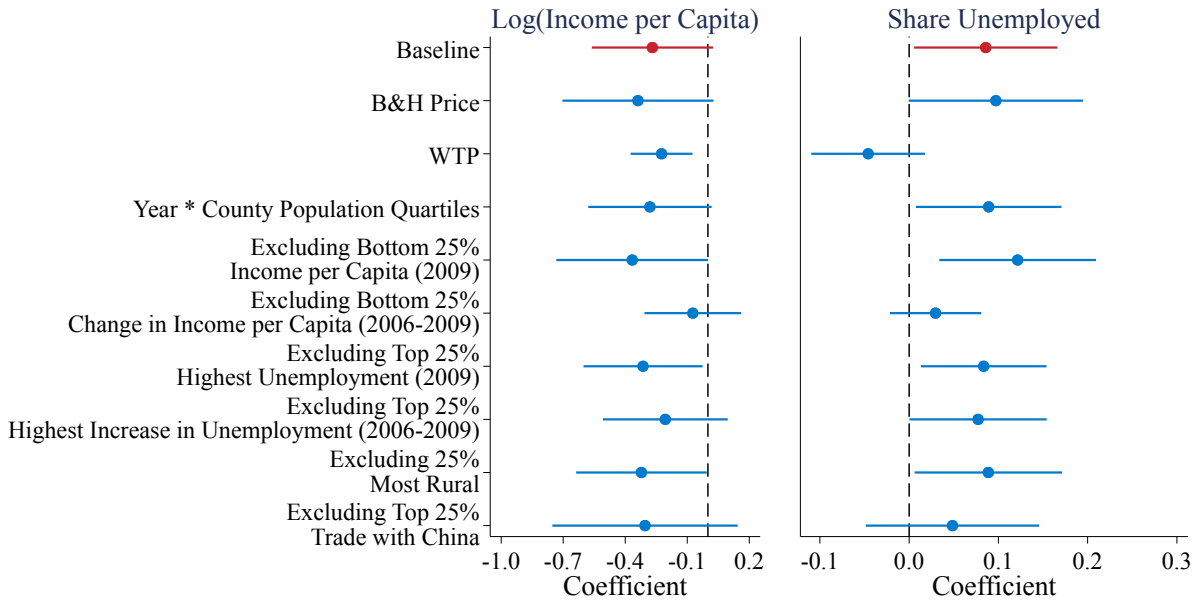
Notes: This figure presents two-stage least-squares estimates of Equation (7) of the effect of price increases on logged employer payroll and count of workers. Each estimate includes employer and year fixed effects. Observations are unique at the employer-year level. Our labor market data comes from the Internal Revenue Service. The dots represent point estimates and bars represent 95% confidence intervals.

Figure A.6: Event Studies of the Impact of Rising Health Care Prices on Counties' Income Per Capita - Overall - Census Denominator



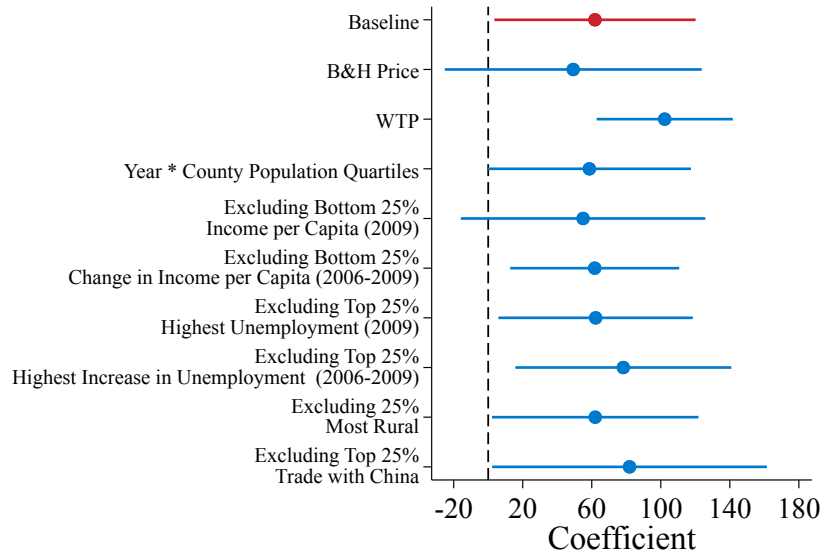
Notes: This figure presents alternative estimates of Figure 4 where we use Census-based population estimates to construct a county denominator, rather than IRS-based population estimates.

Figure A.7: Robustness Tests of County-Level Employment Effects - Overall - IRS Denominator




Notes: This figure presents two-stage least-squares estimates of Equation (10) of the effect of price increases on logged employer payroll and count of workers per employer. 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.8: Robustness Tests of Mortality Effects - Deaths from Suicides and Overdoses per 100,000 Population



Notes: This figure presents two-stage least-squares estimates of Equation (10) 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 mortality data comes from the Centers for Disease Control and Prevention's Restricted Mortality Database. The dots represent point estimates and bars represent 95% confidence intervals.

Figure A.9: Benchmarking Premiums in the Form 5500



Notes: This figure presents a binned scatter plot of state-by-year-level premiums in the 5500s and the Medical Expenditure Panel Survey (MEPS). Average premiums are measured as described in Appendix B. The MEPS premiums are state-level estimates of average premiums for single plans among employers with 50 or more employees. Each dot represents one of 50 quantiles of the premium distribution in the MEPS.