Depressing Payment: Hospital Mergers and the Wage-Benefit Tradeoff

Jason A. Gates

May 1, 2025

Abstract

This paper examines how hospital mergers affect local labor markets through their impact on the cost of employer sponsored health insurance. I estimate the effect of a hospital merger occurring within a commuting zone on the wages, employment, hours worked, part-time work, and full-time work of non-health care employees. I find evidence that mergers reduce wages by approximately 1-2% in local commuting zones where they occur. I also find evidence that they reduce average hours worked and that this can be explained in part by substitution from full-time workers to part-time workers, possibly to avoid then costlier benefit provision.

Introduction

In recent decades the United States has experienced a significant rise in the concentration of health care provider markets (Gaynor and Town 2011). In particular, there has been rapid consolidation within the hospital industry as a result of a wave of hospital mergers. At the same time, health insurance premiums paid by employers and employees have been rising quickly, with the cost of a family premium for employer-based insurance reaching $\$19,616$ in 2018, a 55% increase over the previous 10 years (Kaiser Family Foundation 2019). Seeking to identify a relationship between

these trends, a substantial literature has developed measuring the impact of hospital concentration on health care costs, with nearly all finding evidence that hospital concentration increases the prices hospitals charge (e.g. Moriya et al. 2010; Gowrisankaran et al. 2015; Cooper et al. 2019). In light of this, and given that 60 percent of non-elderly adults receive health insurance from an employer as part of their compensation (Kaiser Family Foundation 2018), these mergers may also be affecting the labor market through their impact on health care costs. For instance, while the median worker cost an employer $18.73 per hour in wages in 2019, they also cost $3.06 in health insurance premiums, making them a substantial portion of employers' compensation expenses (Bureau of Labor Statistics 2019). Any change to the cost of providing insurance will therefore directly affect how employers determine their compensation offerings. Thus, in order to fully account for the effects of these mergers, we need to understand how they affect the local labor market.

Yet to the author's knowledge no study has examined the full impact of hospital concentration on local labor markets. One previous study, Prager and Schmitt (2019), examined the impact of hospital mergers on health care workers, a market where hospitals are a major employer. This work is part of a larger recent literature studying the effect of firm concentration on local labor markets through monopsony power (Azar et al. 2017; Benmelech et al. 2018). This string of research has led to discussions of whether regulators should consider the monopsonistic effects of mergers on labor markets (Federal Trade Commission 2018). In this paper I extend this argument a step further for hospital mergers. Due to the unique way health care is financed in the United States, concentration in the health care industry reverberates not just to workers in this industry, but to workers in all industries via increased health care costs.

In this paper, I estimate the impact of hospital mergers on non-health care workers' labor outcomes including annual wages, employment status, and number of hours worked. Using both a difference-in-differences and event-study design, I compare the outcomes of workers in local labor

markets exposed to a hospital merger with those in ones that were not using both the Quarterly Census of Employment and Wages (QCEW) and the American Community Survey (ACS). I find strong evidence that hospital mergers resulted in reductions in annual wages and that these effects grew over time. The point estimates for average wage suggests that hospital mergers in the period studied decreased local annual average wages by 0.8 to 1.5 percent. I also find evidence that these mergers reduced the number of hours worked and that this may be driven by shifts away from full-time employment towards part-time employment.

Altogether the results here suggest that hospital mergers do have significant impacts on local labor markets including suppressing the annual wages of workers outside of the health care industry. Beyond its relevance to merger analyses, this study also has implications for discussions surrounding wage stagnation and inequality that have occurred in recent decades. If hospital mergers are detectably suppressing wages through their ability to increase insurance premiums, they may be playing a role in the recent stagnation of real wages. I note that the work presented here is limited in that it only estimates relatively short run effects and does not incorporate any general equilibrium effects. For example, locally reduced demand as a result of lower disposable income may have macroeconomic effects as well. Data limitations also constrain me in determining how much of the wage effects are due to changes in employment-level and hours worked rather than changes in hourly wages. Further work using data better suited to determining these relationships would allow for a more accurate accounting of the dynamics caused by these mergers at a micro-level. The remainder of the paper proceeds as follows. I first review the theoretical framework of how mergers may affect premiums and how employers may respond to these changes. I then review previous findings of how employers respond to premium cost changes. The empirical model and data are then overviewed and the last section concludes.

Conceptual Framework

Insurers

In order for hospital mergers to have an effect on workers' wages and employment through health care costs, it would need to be established that they actually have a meaningful impact on such costs. In their study of 366 mergers occurring between 2007 and 2011, Cooper et al. (2019) find that the prices merging hospitals charge to insurance companies increases by approximately 6 percent when the merging hospitals are within 5 miles of each other. This effect dissipates to just over 2 percent for those within 25 miles of each other, with no statistically detectable effect beyond that radius. They also find some evidence that these effects increase over time, with mergers up to 30 miles away increasing prices by 6 percent after 2 years.

However these are the costs the hospitals charge to insurance companies, raising the issue of how much the incidence of these costs ultimately fall on employers and how quickly. Here it is important to note that 61 percent of workers who receive insurance through their employer are enrolled in a self-insured health plan (Kaiser Family Foundation 2019). Under self-insured plans, the employers pay medical costs directly and only contract insurance companies to design and administer these plans. Therefore any increases in prices that hospitals can effect on payers would be directly borne by self-insured employers.

For the 39 percent of employees in fully-insured plans, whose health care costs are paid by insurers while employers pay a premium to them for this coverage, the impact is less easily established. Potential insight into insurer behavior in this arrangement can be gleaned from recent work in the hospital merger literature. In their structural model of the impact of a potential hospital merger on an HMO in northern Virginia, Gowrisankaran et al. (2015) estimate the merger would increase the prices that the insurer was charged by 7.2 percent and that the insurer would then

raise premiums charged to employers by 3.4 percent. While this result may not generalize to all mergers and fully-insured plans, it suggests that at least in some instances price increases faced by insurers are passed onto employers in fully-insured plans as well.

The Model

There remains the question of how employers will handle increases in health care costs that result from these mergers. The standard economic model of employer-provided benefits, and the one most frequently employed in this literature, is the compensating differential framework first posited by Summers (1989) and formalized by Gruber and Krueger (1991). Suppose labor demand is given by:

$$L_d = f_d(W + C) \tag{1}$$

Where W is wages and C is insurance costs. Furthermore suppose supply is given by

$$L_s = f_s(W + \alpha C) \tag{2}$$

Where αC is the monetary value that employees place on health insurance. Then

$$\frac{dW^*}{dC} = -\frac{\eta^d - \alpha\eta^s}{\eta^d - \eta^s} \tag{3}$$

where η^d and η^s are the elasticities of demand and supply for labor. From equation (3) we can see that if $\alpha = 1$ and employees fully value the benefit, wages fall by the full cost of the benefit. If workers do not value the benefit at all and $\alpha = 0$, then the result is identical to the incidence of a payroll tax. The impact on employment will be

$$\frac{dL}{L} = \frac{W_0 - W_2 - dC}{W_0} \cdot \eta^d \tag{4}$$

Where W_0 and W_2 are the initial and final wage levels.

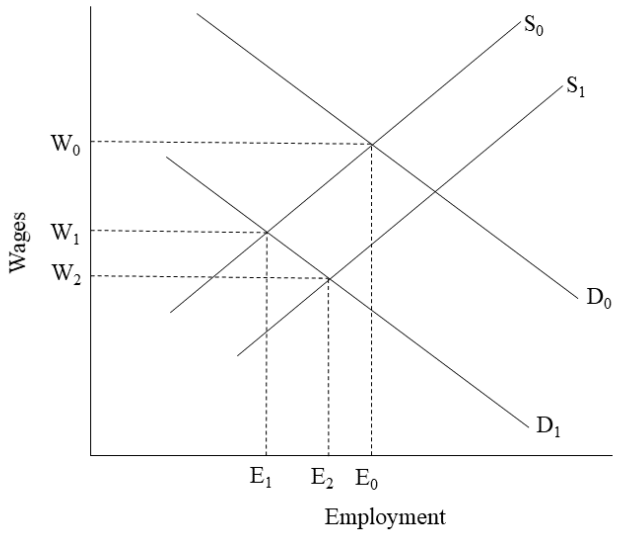

Figure (1) Labor Market Effects of Employer Sponsored Insurance

In this framework the effect of rising health insurance premiums on wages is predicted to be negative, while its effect on employment is a function of α. If workers value the insurance they receive, that is $\alpha > 0$, the effect on employment is ambiguous. The basic workings of the model are depicted in Figure 1. The increased cost to employers shifts labor demand from D_0 to D_1, moving wages downward from W_0 to W_1 and employment from E_0 to E_1. To the extent that these benefits are positively valued by workers, supply also shifts out from S_0 to S_1, further reducing wages to W_2 and moving employment to E_2.

According to the model, the effects on employment are highly dependent on worker valuation of the benefits, α. If $\alpha < 1$, and workers value change in insurance at less than its cost, employment will fall. If $\alpha = 1$ and employees value it at cost, there will be no impact on employment. If $\alpha > 1$, that is workers value insurance at greater than its monetary cost, employment will rise. One

might think in this empirical setting we would expect α to be zero, since the change in the cost of insurance I am discussing is coming from an increase in market concentration by hospitals and not from value-adding innovations. However, even in this circumstance it is likely that employees value this marginal change in insurance. If insurance costs are going up for all private payers and not just employers, the cost of purchasing insurance individually or of being uninsured and paying out of pocket has increased by more than the employer premium cost. This is because, as noted by Currie and Madrian (1999), employers are capable of purchasing insurance much more cheaply than individuals and, unlike wages, benefits are not taxed. So a worker receiving the compensating differential in the form of wages and seeking individually purchased insurance would be forced to pay a higher premium than their employer is charged (or bear all the risk of being without insurance) and this money would be taxed first. This leaves us with theoretically ambiguous impacts on employment.

Caveats

While a useful analytical tool, there are several ways in which employers and workers may behave contra the predictions of this model. For one, employment broadly defined incorporates both having a job and number of hours worked. As Cutler and Madrian (1998) note, employers may respond to increases in premiums with an increase the number of hours worked among workers with health insurance, since this is a fixed cost of employment. This may lead to total incomes rising for these workers. Furthermore, they may decide to replace full-time workers who receive employer sponsored insurance (ESI) with part-time workers they can exclude, potentially reducing average hours worked.

It is also possible that employers might respond by altering health insurance generosity and structure. For one, they could simply stop offering ESI or, for new firms, choose not to offer it

from the start. Alternatively, they could switch to lower quality insurance with higher deductibles and copays, which could keep the total premium at the same level in the face of higher health care costs. They could also increase employee contributions to premiums. This might not make sense at first glance if the incidence of insurance is ultimately borne by the worker anyway; the only difference between an employer and employee premium is that the latter are often taxed before they are paid. However, as noted by Gruber and McKnight (2003), increased employee premiums may reduce take-up rates of insurance among employees, reducing the number of employees whose health care costs the firm is responsible for. That some individuals would not take up insurance while others would indicate heterogeneity in α across employees within a firm. From the point of view of this study, I am unable to see if firms absorbed the costs through increases in employee premiums since I cannot observe employee contributions.

Interestingly, in theory, increased employee contributions may increase per-enrollee premiums even if they reduced overall firm costs. Sicker workers more likely to have higher health care costs are more willing to pay the employee premium than healthier workers - the adverse selection problem. Therefore reliance on passing on increases in health care costs through increases in employee premiums may lead to worse adverse selection, resulting in a sicker and more expensive remaining pool, increasing health insurance premiums further and so on, potentially resulting in a death spiral. This would suggest firms would be reluctant to overuse this mechanism. While employee contributions have grown rapidly, they have remained a relatively constant share of overall premiums in the past 20 years (Kaiser Family Foundation 2019), potentially lending credence to the limits of this mechanism.

With this theoretical ambiguity and potential limitations of the basic compensating-differentials model, it remains an empirical question of how labor markets respond to health care cost pressures arising from hospital mergers.

Literature Review

Wages

The wage literature has been mixed on its findings of whether employers pass through increases in the cost of health insurance to workers in the form of reduced wages as the standard incidence model predicts. A major difficulty in studying the issue is that most surveys commonly used to study labor markets such as the Current Population Survey (CPS) or Survey of Income and Program Participation (SIPP) do not report the generosity or cost of health insurance premiums workers receive. One that does is the National Compensation Survey (NCS), which collects job-level wage and benefit information. Two key studies of how wages are altered by increasing premiums are Buchmueller and Lettau (1997) and Anand (2017), both of which use the NCS. Using a first-differences specification and regressing job-level changes in wages on job-level changes in employer premiums, Buchmueller and Lettau (1997) find no relationship between changes in the cost of insurance and wages. However, the study does not account for changes in the health insurance plan design or generosity between years, making the changes in the cost of insurance for that job potentially endogenous to the changes in wages they observe.

In an update and extension of Buchmueller and Lettau (1997), Anand (2017) uses the NCS from 2003 to 2010 to again check for a wage-benefit tradeoff. Like the earlier study, she finds no evidence of changes in wages resulting from increases in employer premiums. Unlike the earlier analysis, she restricts the data only to plans that were offered multiple years in a row in order to ensure that the price change is the result of premium changes and not changes in plan characteristics. An unfortunate drawback of this method is that 31 percent of the occupation-establishments in the sample were dropped. It is possible that these plans that were offered in one year and not offered in the next were dropped because their costs grew faster than those that were not dropped. It is

also possible that plan-dropping is correlated with other employer characteristics and dynamics. Whether the third of the sample that was dropped differs in ways observable in the data is not pursued.

Gruber and Hanratty (1995) study the issue from the opposite direction, studying province- and industry-level variation in the implementation of Canada's national health insurance (NHI), which replaced employer-sponsored plans. They find that while provincial-level implementation was associated with increases in wages, workers in industries that had the highest levels of private insurance coverage prior to NHI saw wage increases lower than those with less coverage. This runs contrary to the prediction of the compensating differential model that wages would rise by the amount employer-premiums decrease (or are eliminated in this case).

Albeit not directly examining health insurance premiums, Gruber and Krueger (1991) uses the workers' compensation program, another type of employer provided insurance, to gain insight into whether health insurance premium changes would affect wages. They find that between 56 and 86 percent of increases in the cost of workers' compensation are borne by the employees in the form of lower wages. Clemens and Cutler (2014) study the impact of increasing employer premiums on public school districts. The authors find that for every $1 increase in the cost of benefits, wages fall by $0.15 and this point estimate was not statistically different from 0. Lubotsky and Olson (2015) also examined teachers, finding no evidence of reductions in wages as a result of premium increases.

While the above studies seem to contradict the incidence model predictions, there have also been a raft of studies finding evidence of a tradeoff. Bhattacharya and Bundorf (2009) find that, controlling for observables, insured obese workers have lower wages relative to insured non-obese workers, while no such wage gap exists for workers without health insurance, suggesting workers pay the incidence through wages. The magnitude of their estimates suggests employees bear nearly

all of the incidence of the cost of their coverage. Gruber (1994) finds evidence of a full shifting of the cost of mandated maternity benefits to the wages of those who benefit from it. Qin and Chernew (2014) look at public sector workers' wages and find a wage offset of 15 to roughly 50 percent depending on their specification.

Baicker and Chandra (2006) use state-level changes in medical malpractice payments as an instrument for health insurance premiums. They find a dollar-for-dollar offset of increases in premiums coming from the wages of workers with ESI, with no decline in wages for those without. Kolstad and Kowalski (2016) studied the impact of Massachusetts' health care reform on the labor market. Part of Massachusetts' reform was a mandate for certain employers to provide ESI, allowing the authors to check for a compensating differential. They find a compensating differential close to the average cost of ESI to the employer.

The conflicting findings within the literature on whether employees pay for changes in benefit costs via reductions in wages led Sommers (2005) to hypothesize that nominal wage constraints may prevent some employers from passing on the full incidence of premium changes to employers, at least in the short term. The idea is that nominal wage cuts are costly to firms due to perceptions of fairness on the part of workers. Therefore only through erosion by inflation can wages be cut, essentially putting a lower limit on how much real wages can change. He tests his hypothesis using the CPS by exploiting variation in inflation rates by region and finds that employers are constrained in passing premium costs onto workers by the inflation rate. However, the author still does find evidence of a partial wage offset even if constrained.

Employment and Hours Worked

Many of the studies that examined wages also checked for impacts on employment, particularly hours worked and binary employment measures. Similar to the findings on wages, the results

on employment also vary between studies. In their study of workers' compensation premiums, Gruber and Krueger (1991) found no statistically significant impact on employment from increases in premiums. Qin and Chernew (2014) found no impact of increased premiums on hours worked among public employees. Lubotsky and Olson (2015) also found no evidence that school districts reduced the number of teachers in response to higher health insurance costs.

However, other studies have found impacts on employment and hours worked. Gruber (1994) finds evidence of both a rise in the number of hours worked and a fall of employment among the treated group of married women as a result of mandated maternity benefits. Using the SIPP and CPS, Cutler and Madrian (1998) find that rising insurance costs increased the number hours worked of insured workers by 3 percent, suggesting firms may substitute hours per worker for the number of workers employed in response to this increase in fixed cost per worker. Montgomery and Cosgrove (1993) and Buchmueller (1999) both found an association between higher health premiums for full-time workers and the use of part-time workers within a firm.

Baicker and Chandra (2006) find a decrease in hours worked as a result of rising premiums, part of which is attributed to increases in the probability of unemployment along with increases in the probability of part-time work. In their study of the Massachusetts health reform, Kolstad and Kowalski (2016) found a small but statistically significant decrease in the number of hours worked as a result of the employer mandate for certain firms to provide coverage to their employees. Sommers (2005) found the hazard rate of unemployment was higher and statistically significant for those with health insurance when employers faced high premium growth.

Benefit Generosity

As previously mentioned, while wage and employment data are relatively easy to access, data on plan design, generosity and employee contributions is quite limited. Several studies however

have attempted to estimate an impact of premium growth on these. Perhaps the simplest way for employers to change benefits in response to higher premium costs is to drop the benefit altogether or change those eligible for them. Baicker and Chandra (2006) found increases in premiums decreased the likelihood of having ESI but were unable to see how much was attributable to declines in full-time employment as opposed to decreases in the offering of ESI.

Anand (2017) found that a $1 increase in health insurance costs for a worker resulted in just over a $0.50 increase in employee premium contributions. Lubotsky and Olson (2015) find that a $1 increase in premiums for individual health insurance increase the employee contribution to this premium by $0.17, while the number for family coverage was $0.46. This suggests that rather than reducing wages, employers may increase the employee contribution, thereby reducing the effective pay of those who use the insurance, while not affecting the pay of those who do not take it up. This would seem to contradict the idea presented above that employers are constrained in their ability to raise employee contributions while avoiding serious adverse selection problems.

Vistnes and Selden (2011) offer perhaps the most informative study on the non-wage effects of premium growth. Using the Medical Expenditure Panel Survey Insurance Component (MEPS-IC), a survey of employers, they examine the impact of metropolitan area variation in insurance costs on employer offerings of insurance, employee eligibility for this insurance, employee premium contributions, and deductibles. They find that higher insurance costs reduce the likelihood that employers offer coverage to any of their employees and, among those who continue to offer coverage, to reduce the number of workers eligible for it. Like Anand (2017) and Lubotsky and Olson (2015), they find that employers increase employee premiums in response to health care price increases. They also find that firms increase deductibles in response to cost pressure, suggesting cost changes may cause firms to change plan generosity and design. If this is indeed how employers respond, there may be no observable effect on wages or employment levels in the data used for this study.

Data

As previously mentioned, a serious issue with investigating the relationship between health care costs and wages is data availability. Ideally it would be possible to observe the price and type of service each hospital charged for each individual treated, the plan that individual was enrolled in, the cost of the premium of that plan, their employer, and a full breakdown of their compensation terms of wages, benefits and hours worked and the same breakdown for every worker in their local labor market. It would then be possible to estimate the impact the merger had on hospital prices, the extent to which these price changes increased employer premiums, how employers dealt with these premium increases, and whether there were any spillovers into the labor market of these workers (e.g. if depressed demand for high-insurance workers at affected firms have impacts on workers at unaffected firms due to depressed demand in the market).

Even short of this it would be useful to have employer-sponsored premiums available on the surveys usually used to analyze the labor market such as the CPS or SIPP. Although even if these surveys asked respondents, employees are likely not aware how much it costs their employer to insure them. Without being able to observe these costs directly, estimating the impact that mergers have on premiums along with dW^*/dC or dL/L from equations (3) and (4) respectively is impossible. In light of this shortcoming, I therefore adopt a method from the "treatment effect" literature and estimate the impact of mergers by comparing the dynamics of local labor markets that were "treated" with a hospital merger to those that were not using a difference-in-differences, two-way fixed effect methodology. The next section discusses the empirical methodology in depth.

In order to estimate this effect of hospital mergers, I rely on two sources of local labor market data, each with relative advantages. I use both the Quarterly Census of Employment and Wages (QCEW) from the Bureau of Labor Statistics (BLS) and the American Community Survey (ACS) from the Census Bureau. For both datasets, I define local labor markets using Commuting Zones

(CZs). CZs were developed by the Economic Research Service of the U.S. Department of Agriculture to more accurately delineate local labor markets than earlier methods.

The main advantage of the QCEW is that it is a near-census of worker earnings and employment at the county and industry level. It is based on ES-202 filings that every establishment is required to submit in order to calculate payroll taxes for unemployment insurance, and covers 98% of all workers. The industry level information allows for workers within the health care and insurance industries to be excluded from the sample. This is important for this analysis since I am interested in how these mergers affect labor markets through their role in the health care system rather than their role as employers in the health care labor market. For the part of the analysis using the QCEW, industries are broken out by their two-digit NAICS codes. Since it is aggregated to the county level and CZs are defined as small groups of counties, the QCEW also allows for easy aggregation to the CZ level. One drawback of the QCEW is that is only reports total payroll and total employment. Therefore, while it gives an accurate measure of average earnings of workers in a county and industry, there is no way to see the hours worked among these employees. So it will not be possible to tell whether findings of lower average annual wages are accompanied by (and potentially driven by) decreases in the number of hours worked or not.

In comparison, the ACS is an individual-level dataset with a rich set of demographic information as well as several employment outcomes including whether an individual was employed, their annual wages, and their usual hours worked per week. Furthermore, by the standard of surveys, the ACS is quite large, with over 3,000,000 observations in 2017. Responses are also legally obligated, reducing non-response issues relative to other surveys. However, to avoid issues of individual identifiability, information on low population geographies is limited. Therefore not every county is separately identifiable in the ACS. As an alternative, the ACS releases geographic information in the form of Public Use Microdata Areas (PUMAs), which are geographies that contain at least 100,000 people,

are constructed from Census tracts, and are nested within states. Therefore, while not entirely commensurate with CZs, it is possible in most cases to map PUMAs into CZs, with the exceptions being particularly rural geographies. To address this limitation of the data I assign individuals in a certain PUMA to the CZ only if at least 80% of that PUMAs population is within that CZ according the Census.

Unfortunately, PUMAs are redrawn with every decennial Census, making the geographic identifiers from 2011 and earlier incompatible with identifiers from 2012 and later. Therefore, for the ACS portion of the analysis I restrict the years from 2005 (the first year the ACS was fielded) to 2011. While restrictive, this set of years roughly corresponds to the years examined by Cooper et al. (2019) in their study of mergers effects on prices. The outcomes I examine using the ACS are logged annual wages, logged hours worked per week, whether an individual was employed or not, and whether they worked over or under 35 hours per week (in order to see if workers are shifted from full to part-time work). Unfortunately, while I have information on usual hours worked per week in the past year and income over that time, the question asking the number of weeks worked in the past year was discontinued in 2007, leaving me unable to calculate average hourly wages. I restrict the sample to adults aged 22 to 64 since these are the workers most likely to receive insurance through their employer rather than as a dependent (those under 22) or from Medicare (those 65 and older).

Data on hospital mergers are provided by Cooper et al. (2019) and are a census of all ownership changes of hospitals registered with the American Hospital Association (AHA) from 2001 to 2014. The data includes the latitude and longitude of each hospital, the hospital that was targeted in the transaction and the hospitals that were members of the acquiring system. For more information on their data cleaning process see the Online Appendix of Cooper et al. (2019).

For this analysis I consider a CZ to have experienced a merger if a hospital with a latitude

and longitude within that CZ is acquired by a hospital or system with a hospital whose latitude and longitude also place it in that CZ. It is important to note here that due to the frequency of hospital mergers in recent decades, many CZs experience multiple mergers over our time period examined here. In such cases, there is not a clearly defined the pre- and post- period to conduct a difference-in-differences analysis. Therefore, for the main analysis I restrict the treatment group to CZs that experienced one merger over the period analyzed in a manner similar to both Prager and Schmitt (2019) and Cooper et al. (2019). As a sensitivity test of this restriction, as in Cooper et al. (2019), I also estimate a model with a cumulative merger measure without dropping any CZs and find the main results remain.

Table 1. Summary Statistics

	Treated	Control
Quarterly Census of Employment and Wages 2001-2014		
Average Wages	30,981	28,412
Employment	96,011	33,666
Observations	1,946	5,502
	Treated	Control
American Community Survey 2005-2011		
Annual Wages (Dollars)	29,386	28,409
Usual hours worked per week	31.67	31.63
Employed	70.6%	70.6%
Part-Time	14.2%	14.2%
Full-Time	64.4%	64.4%
Parent	44.8%	45.2%
Age	42.24	42.42
Female	45.6%	45.5%
Less than HS	13.4%	12.1%
College	27.0%	25.4%
Foreign Born	14.2%	10.5%
Noncitizen	8.6%	5.7%
Black	8.8%	11.4%
Hispanic	13.0%	8.4%
Observations	1,177,867	1,144,141

Notes: Data for first panel is from the 2001-2014 Quarterly Census of Employment and Wages (QCEW). Sample includes all Commuting Zones (CZs) that experienced zero or one merger between 2001 and 2014. Data for second panel is from the 2005-2011 American Community Survey (ACS). Sample includes all Commuting Zones that experienced zero or one merger between 2005 and 2011.

Table 1 presents the summary statistics for the treated and untreated CZs in both datasets, with the top panel including the two measures available on the QCEW and the bottom panel including

the outcomes and covariates used in the ACS analysis. The samples are broadly similar across a variety of characteristics including employment rates, prevalence of part-time work, demographic characteristics, and education levels. One key difference between the two samples is the average size of the labor market in each. In the QCEW sample, CZs that experienced a merger had 96,011 workers on average while those that did not had 33,666 on average. This issue is mitigated by the facts that (1) this analysis uses relative changes in employment as an outcome rather than absolute changes and (2) it uses the difference-in-difference study design will difference out any fixed characteristics of the labor markets. However this disparity in size could be an issue to the extent that labor market outcomes in smaller markets have different trends relative to larger markets over the period studied. To address this concern, along with the main difference-in-differences analysis, I test for differential trends between treated and untreated CZs leading up to merger events.

Empirical Strategy

As mentioned above, for this study I adopt a difference-in-differences research design, considering individuals in labor markets that experienced a merger as "treated" and those in labor markets that were not as "controls". I will be estimating the following two-way fixed effects regression equation for individual i, CZ j, and year t using ordinary least squares

$$\text{Outcome}_{ijt} = \alpha + \beta\ \text{MERGE}_{jt} + \mathbf{\Gamma X_{it}} + \eta_j + \delta_t + \varepsilon_{ijt}. \tag{5}$$

Here Outcome is the labor market outcome of interest, either at the individual or aggregated level (in which case the subscript i is dropped). MERGE_{jt} is a binary indicator of whether the CZ experienced a hospital merger up to that point in time. For the individual-level regression, $\mathbf{X}$

is a vector of individual level covariates. Here, η_j is a CZ fixed effect, controlling for any fixed differences between commuting zones' outcomes, and δ_t is a year fixed effect, which controls for any nationwide yearly shocks experienced across all CZs. As in all difference-in-differences designs, β is identified by the assumption that absent a merger, changes in labor market outcomes in treated CZs would have been the same as changes in labor markets in control CZs had they not experienced a merger. While this assumption cannot be tested directly (as a CZ is not observable in two different states at the same time), I can test whether the outcomes being studied were evolving along the same trend prior to the merger event in treated and control CZs, which I do below.

The parameter of interest, β, is the average "treatment effect" of a merger on a local labor market. Due to the limitations listed above, this unfortunately cannot be interpreted as a parameter defined in the above incidence model. Given the data available, it is impossible to measure the impact that a given merger has on premium costs in that labor market or the impact that the subsequent cost change has on wages, dW/dC, or the other outcomes examined. Furthermore, I cannot differentiate between mergers involving larger or smaller shares of the hospital market, due to a lack of data on hospital size. Unfortunately this limits the economic interpretability of this point estimate and its applicability to particular future mergers under review. However, under the identifying assumption stated above, it can inform us retrospectively of the average overall effect that hospital mergers of any size had on each labor outcome in the study period.

Along with the basic difference-in-differences analysis, I also estimate an event-study model in order to test for both pre-trends and dynamic treatment effects. In an event-study an observation's indicator for being in the treatment group is interacted with a dummy for its year relative to its treatment. This, in effect, incorporates leads and lags into the regression. I estimate the following

event-study model for individual i, CZ j, and year t using ordinary least squares

$$\text{Outcome}_{ijt} = \alpha + \sum_{\substack{k=-r \\ k\neq 0}}^{s} \sigma^k \text{EVERMERGE}_j^k + \mathbf{\Gamma X_{it}} + \eta_j + \delta_t + \varepsilon_{ijt}. \tag{6}$$

Here k is year relative to the merger and $k = 0$ is the year of merger, where r is the number of leads and s is the number of lags. Here the dummy for year of merger is excluded so that the point estimates of σ^{-r} to σ^s are relative to it. EVERMERGE^k is an indicator for being in the treatment group and in year k.

Results

Table 2 presents the difference-in-difference results for both log wage and log employment using the QCEW from 2001 to 2014. The results suggest that a hospital merger reduces the average wage in an exposed industry-commuting zone by 1.5 percent and this result is significant at the .01 level. The result for employment suggests mergers had no statistically significant impact on employment in exposed industry-commuting zones.

Table 2. Effect of Mergers on Labor Market Outcomes

	Est.	P-Value	N
Log(Average Wage)	-0.015	0.000	94102
Log(Employment)	0.001	0.954	94102

Note: Data from the 2001-2014 Quarterly Census of Employment and Wages (QCEW). Observation is the industry-commuting zone level.

Table 3 presents the difference-in-differences results using the 2005-2011 ACS. In the ACS analysis I control for a host of individual-level demographic characteristics including age, sex, race/ethnicity, education, citizenship status, and parental status. The results suggest that mergers resulted in a 0.8% reduction in wage and salary income on average, with a p-value of 0.019. The results for hours worked, employment, and part-time status were all statistically insignificant. In the

appendix, I also estimate a cumulative merger effect without dropping CZs with multiple mergers over the time period, finding similar results.

Table 3. Effect of Mergers on Labor Market Outcomes for ACS

	Est.	P-Value	N
Log(Wage and Salary Income)	-0.008	0.019	1696943
Log(Hours)	0.000	0.823	1812034
Employed	-0.002	0.158	2322008
Part Time	0.001	0.348	2322008
Full Time	-0.001	0.337	2322008

Note: Data from the 2005-2011 American Community Survey (ACS). Sample is limited to those ages 22-64. Full-time work defined as working at least 35 hours per week. Covariates include year and Commuting Zone fixed effects, marital status, sex, age, education, 2-digit industry number, citizenship status, race/ethnicity, and parental status. Regression are weighted by ACS weights.

In conjunction with this basic difference-in-differences design, I also estimate event-study specifications using the same data. Event-study designs have two primary benefits. First, they allow for testing whether treated and control labor markets had differential trends in the run-up to mergers that could bias the difference-in-difference result. Secondly, they allow testing whether there were dynamic effects as a result of treatment, including whether any effects fade or increase over time.

The event-study results for the 2001-2014 QCEW are presented in Figure 1. Here the coefficients for year relative to merger interacted with treatment status are plotted, along with their 95 percent confidence intervals. The year the merger occurred is the excluded year which the point estimates are relative to. I find no evidence of pre-trends in the years leading up to mergers in the QCEW data for either wages or for employment. Again, I find that mergers reduced wages and here it appears this effect grows over time. As in the difference-in-differences analysis, no effect on employment is detectable.

Figure 1. QCEW Event Study Estimates

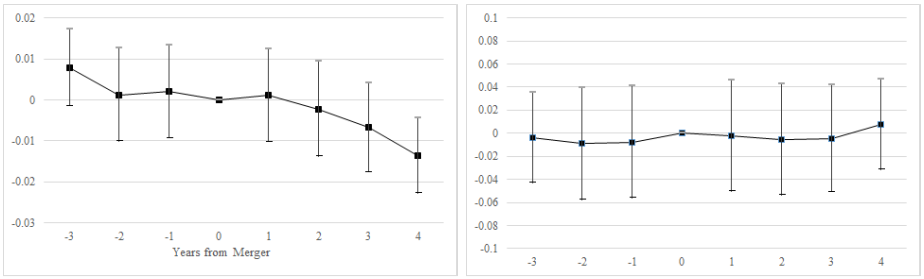

(a) Log(Average Wage) (b) Log(Employment)

Note:Data from the 2001-2014 Quarterly Census of Employment and Wages (QCEW). Point estimates and 95% Confidence Intervals are depicted for the three years prior to and four years subsequent to the merger. Observation is the industry-commuting zone level.

The event-study results for the 2005-2011 ACS are presented in Figure 2. Here again I find no evidence of pre-trends for wages, hours, or the work-status measures. Like the QCEW result for wages, the ACS results also suggest that the impacts of mergers on wages grew over time. While the difference-in-differences results found no overall average effect of mergers on non-wage outcomes in post-period years, the event-study results suggest some evidence that these measures were affected by mergers. In panel (b) of Figure 2, while coefficients on hours worked for one and two years post-merger are statistically insignificant, the coefficient for three years post-merger is marginally significant and negative, suggesting that hours were affected by mergers gradually. Similarly the event study results suggest mergers decreased full-time work and increased part-time work among adults in the sample, with the estimated effect on both measures statistically different from the base year after three years. These impacts on part-time and full-time work are consistent with the findings of Montgomery and Cosgrove (1993), Buchmueller (1999), and Baicker and Chandra (2006). It may be easier for firms to adjust the composition of their workforce between full- and part-time workers than deal with the costs through other potential channels.

Figure 2. ACS Event Study Estimates

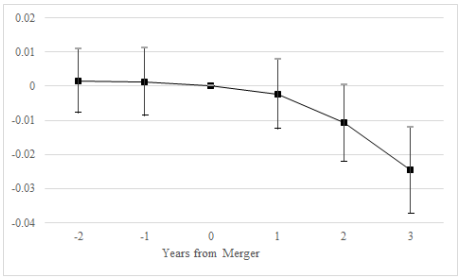

(a) Log(Wage and Salary Income)

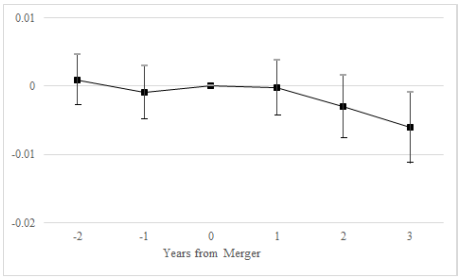

(b) Log(Hours)

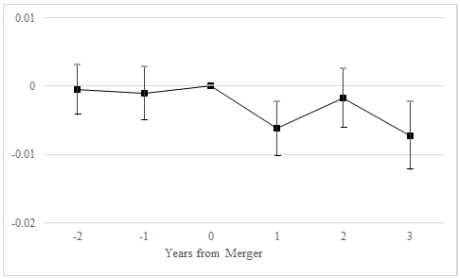

(c) Employed

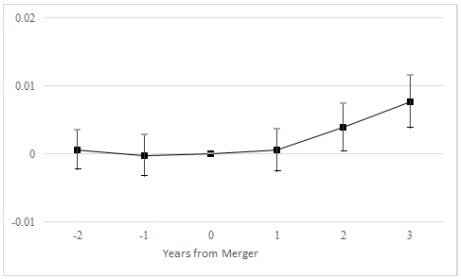

(d) Part Time

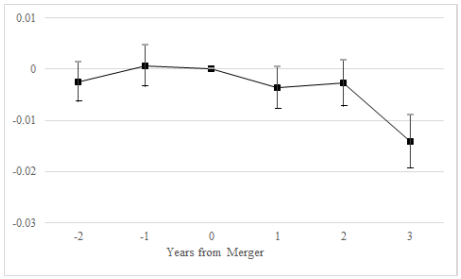

(e) Full Time

Notes: Data from the 2005-2011 American Community Survey (ACS). Point estimates and 95% Confidence Intervals are depicted for the two years prior to and three years subsequent to the merger. Sample is limited to those ages 22-64. Full-time work defined as working at least 35 hours per week. Covariates include year and commuting zone fixed effects, marital status, sex, age, education, 2-digit industry number, citizenship status, race/ethnicity, and parental status. Regressions are weighted by ACS person weights.

The delay in the impacts of mergers on labor markets for both the QCEW and ACS is not surprising in light of the findings of Cooper et al. (2019). The authors found that the effect of mergers on prices charged to insurers appear to grow over time, with results greatest 2+ years after mergers occurred. Therefore, the growth in effects over time is in line with the idea that prices charged by hospitals are driving these labor market changes.

These effects on hours and full-time status present in the event-study analysis are important for

interpreting the impact in both the QCEW and ACS analysis on annual wages. If hours worked, full-time status, and part-time status were unaffected by mergers, it would be likely that the decline in annual income was being driven by a decline in hourly wages, in line with a standard incidence story of the Summers (1989) model presented above. However with hours and work status changing as well, it could be either changes in hourly wages or the work arrangements themselves that are driving the income changes.

There may be some concern that the Great Recession is influencing the results presented here, say if it affected treated labor markets differently than control labor markets. In the appendix, I reestimate the QCEW event-study results on the years 2001 to 2007 and find similar results to the main analysis. Unfortunately the sensitivity of the ACS estimates cannot be similarly tested for since the Great Recession occurs in the middle of that sample period.

Conclusion

This paper examines the impact of hospital mergers on local labor markets using both micro-level data from the American Community Survey and aggregate payroll data from the Quarterly Census of Employment and Wages. The results suggest that consolidation in the hospital industry affects local labor markets in several ways. I find strong evidence that mergers have negative effects on annual wages. I also find evidence that mergers resulted in reductions in full-time employment and increases in part-time employment in commuting zones where they occurred. Due to data limitations I am not able to determine whether these reductions are from changes in hourly wages or whether they are solely being driven by the changes in hours worked. Thus whether individual workers bear the incidence of changes in the cost of their insurance through hourly wage reductions cannot be determined. It may be the case they simply substitute away from now-costlier full-time workers who are eligible for insurance coverage towards part-time workers usually ineligible for such

firm coverage. The findings of Cooper et al. (2019) suggest that the impact that hospital mergers have on prices increases over time, and here I find that the impact of the mergers on labor market outcomes also appears to grow over time.

These findings have important implications for policy-makers and regulators. Recent work on employer monopsony power such as Azar et al. (2017) and Benmelech et al. (2018) has suggested that the welfare implications of mergers extend beyond consumers to the employees in the industries being affected and that these impacts should be accounted for in merger reviews (see Prager and Schmitt (2019) for an example in the hospital industry). The findings in this paper suggest that, due to the unique relationship between health care costs and worker compensation in the United States, the implications of hospital mergers extend to workers outside of the directly relevant industry as well.

While beyond the direct scope of this paper, the findings in this study may also be relevant for understanding broader trends in slow wage growth experienced by the United States in recent years. As noted by Cooper et al. (2019), there were over 700 hospital mergers from 2001 to 2011 and the average Herfindahl-Hirschman Index (HHI) in the hospital industry increased by 19% over that period. If mergers do suppress wages through their impact on premiums, this rapid consolidation could be one factor contributing to slow wage growth seen in recent decades.

In light of this paper's findings, I believe there are several areas where future research would be best focused. This paper was only able to estimate the short run effects of these mergers and was not able to estimate an overall effect this concentration has had on the US economy. For example, if hospital concentration reduces local demand due to reductions in disposable income, there may be macroeconomic effects that cannot be explored in this partial equilibrium framework. Perhaps even more pressing is the need to incorporate the premium costs that employers face when studying labor markets. As repeatedly noted here, data limitations are a serious issue for understanding how

employers choose their compensation packages. Given that the surveys frequently used to study labor markets including the ACS, CPS and SIPP do not have this information, studies examining labor costs or workers' compensation using them are missing a large piece of the picture. Better data can potentially shed light on what is currently a black box. With private health insurance expenditures surpassing 6 percent of GDP in 2017 (Centers for Medicare & Medicaid Services 2017), the matter could not be more pressing.

References

Anand, P. (2017). Health insurance costs and employee compensation: Evidence from the national compensation survey. *Health Economics 26*(12), 1601–1616.

Azar, J., I. E. Marinescu, and M. Steinbaum (2017). Labor market concentration. *NBER Working Paper 24147*.

Baicker, K. and A. Chandra (2006). The labor market effects of rising health insurance premiums. *Journal of Labor Economics 24*(3), 609–634.

Benmelech, E., N. Bergman, and H. Kim (2018). Strong employers and weak employees: How does employer concentration affect wages? *NBER Working Paper 24307*.

Bhattacharya, J. and M. K. Bundorf (2009). The incidence of the healthcare costs of obesity. *Journal of Health Economics 28*(3), 649–658.

Buchmueller, T. (1999). Fringe benefits and the demand for part-time workers. *Applied Economics 31*(5), 551–563.

Buchmueller, T. and M. Lettau (1997). Estimating the wage-health insurance tradeoff: More data problems?

Bureau of Labor Statistics (2019). *Employer Costs for Employee Compensation-March 2019*. Bureau of Labor Statistics.

Centers for Medicare & Medicaid Services (2017). National health expenditure fact sheet. https://www.cms.gov/research-statistics-data-and-systems/statistics-trends-and-reports/nationalhealthexpenddata/nhe-fact-sheet.html. Accessed: 2019-09-01.

Clemens, J. and D. M. Cutler (2014). Who pays for public employee health costs? *Journal of Health Economics 38*, 65–76.

Cooper, Z., C. C. Stuart, M. Gaynor, and J. V. Reenan (2019). The price ain't right? hospital prices and health. *The Quarterly Journal of Economics 134*(1), 51–107.

Currie, J. and B. Madrian (1999). Health, health insurance and the labor market. *Handbook of Labor Economics 3*, 3309–3416.

Cutler, D. M. and B. Madrian (1998). Labor market response to rising health insurance costs: Evidence on hours worked. *RAND Journal of Economics 29*(3), 509–530.

Federal Trade Commission (2018). Hearings on competition and consumer protection in the 21st century. https://www.ftc.gov/news-events/events-calendar/2018/10/ftc-hearing-3-competition-consumer-protection-21st-century. Accessed: 2019-07-01.

Gaynor, M. and R. Town (2011). Competition in health care markets. *Handbook of Health Economics 2*, 499–637.

Gowrisankaran, G., A. Nevo, and R. Town (2015). Mergers when prices are negotiated: Evidence from the hospital industry. *American Economic Review 105*(1), 172–203.

Gruber, J. (1994). The incidence of mandated maternity benefits. *American Economic Review 84*(3), 622–641.

Gruber, J. and M. Hanratty (1995). The labor market effects of introducing national health insurance: Evidence from canada. *Journal of Business Economic Statistics 13*(2), 163–173.

Gruber, J. and A. B. Krueger (1991). The incidence of mandated employer-provided insurance: Lessons from workers' compensation insurance. In *Tax Policy and the Economy*, Volume 5.

Gruber, J. and R. McKnight (2003). Why did employee health insurance contributions rise? *Journal of Health Economics 22*(6), 1085–1104.

Kaiser Family Foundation (2018). State health facts: Health insurance coverage of adults 19-64. https://www.kff.org/other/state-indicator/adults-19-64. Accessed: 2019-07-01.

Kaiser Family Foundation (2019). Employer health benefits: 2018 annual survey.

Kolstad, J. T. and A. E. Kowalski (2016). Mandate-based health reform and the labor market: Evidence from the massachusetts reform. *Journal of Health Economics 47*, 81–106.

Lubotsky, D. and C. A. Olson (2015). Premium copayments and the trade-off between wages and employer-provided health insurance. *Journal of Health Economics 44*, 63–79.

Montgomery, M. and J. Cosgrove (1993). The effect of employee benefits on the demand for part-time workers. *ILR Review 47*(1), 87–98.

Moriya, A. S., W. B. Vogt, and M. Gaynor (2010). Hospital prices and market stucture in the hospital and insurance industries. *Health Economics, Policy, and Law 5*(4), 459–479.

Prager, E. and M. Schmitt (2019). Employer consolidation and wages: Evidence from hospitals. *Washington Center for Equitable Growth Working Paper*.

Qin, P. and M. Chernew (2014). Compensating wage differentials and the impact of health insurance in the public sector on wages and hours. *Journal of Health Economics 38*, 77–87.

Sommers, B. D. (2005). Who really pays for health insurance? the incidence of employer-provided health insurance with sticky nominal wages. *International Journal of Health Care Finance and Economics 5*(1), 89–118.

Summers, L. H. (1989). Some simple economics of mandated benefits. *American Economic Review 79*(2), 177–183.

Vistnes, J. and T. Selden (2011). Premium growth and its effect on employer-sponsored insurance. *International Journal of Health Care Finance and Economics 11*(1), 55–81.

Appendix

Cumulative Merger Effect

Here I present an alternative specification to the difference-in-differences analysis presented in the main text. Rather than restricting the sample only to commuting zones that experienced either no or one merger to isolate the impact of a single merger, I regress the outcome on a cumulative merger measure. The dependent variable of interest here is a count variable that increases by one for every year during the sample period that a merger occurs in a commuting zone. In this way no commuting zones are dropped from the sample. An issue with this method is that increases in the measure associated later mergers would capture any dynamic effects of earlier mergers in that same commuting zone, potentially biasing the estimates. This is particularly an issue for this analysis since I find strong evidence of dynamic post-merger effects.

Table A1 presents the results for the QCEW. Here the effect on wages remains negative and statistically significant and the effect on employment remains statistically insignificant, just as in the the main difference-in-differences analysis. Table A2 presents the results for the ACS. Here the impact of a merger on wage and salary income remains negative and statistically significant. Here the point estimates on employment and full-time employment become statistically significant. While this differs from the basic difference-in-differences analysis, these were findings that were present in the event-study model presented later in the paper.

Table A1 presents the results for the Quarterly Census of Employment and Wages from 2001 to 2014.

Table A1. Cumulative Effect of Mergers on Labor Market Outcomes

	Est.	P-Value	N
Log(Average Wage)	-0.004	0.000	119302
Log(Employment)	-0.003	0.483	119302

Note: Data from the 2001-2014 Quarterly Census of Employment and Wages (QCEW). Observation is the industry-commuting zone level.

Table A2. Cummulative Effect of Mergers on Labor Market Outcomes for ACS

	Est.	P-Value	N
Log(Wage and Salary Income)	-0.012	0.000	2261765
Log(Hours)	0.000	0.659	2413717
Employed	-0.002	0.004	3101002
Part Time	0.000	0.546	3101002
Full Time	-0.003	0.000	3101002

Note: Data from the 2005-2011 American Community Survey (ACS). Sample is limited to those ages 22-64. Full-time work defined as working at least 35 hours per week. Covariates include year and Commuting Zone fixed effects, marital status, sex, age, education, 2-digit industry number, citizenship status, race/ethnicity, and parental status. Regression are weighted by ACS weights.

Excluding the Great Recession

Here I present QCEW event-study results, restricting the data to 2001 to 2007. It may be possible that the results above are driven by a relationship between which commuting zones experienced a merger and the impact the Great Recession had at a local level. Again we find no evidence of pre-trends prior to the merger events occurring. We also find no evidence of an impact of the mergers on employment. There is evidence here again that mergers reduced average annual wages. The point estimate of -0.012 percent after 3 years is similar in magnitude to the earlier results and statistically significant at the 0.1 level.

Table A3. QCEW Event Study Estimates 2001-2007

	Est.	P-Value
Log(Average Annual Wage)		
Two years prior to merger	0.000	0.932
One year prior to merger	-0.001	0.858
One year after merger	-0.001	0.837
Two years after merger	-0.004	0.598
3+ years after merger	-0.012	0.094
Log(Employment)		
Two years prior to merger	0.001	0.962
One year prior to merger	-0.003	0.925
One year after merger	0.003	0.912
Two years after merger	0.003	0.922
3+ years after merger	0.006	0.868
Observations	56623	

Note: Data from the 2001-2007 Quarterly Census of Employment and Wages (QCEW). Observation is the industry-commuting zone level.