How Much do Firms Drive Healthcare Utilization: Evidence from Multi-Payer Claims Data*

Daniel Deibler
Federal Trade Commission

Daniel Hosken
Federal Trade Commission

Thomas Koch
Federal Trade Commission

Marshall Thomas
Federal Trade Commission

February 11, 2025

Abstract

We examine the extent to which firms are an important determinant of healthcare utilization in the United States. Using commercial healthcare claims from 14 states, we estimate a model which can separately identify the role of doctors, firms, facilities, and within-state geography in explaining healthcare utilization, specifically examining the role of each of these in whether a patient has a vaginal birth or a Caesarian section (C-section). We find that firms are important: 8% of the variation in C-section rates can be explained by firms, while 10% can be explained by doctors. We then show that doctors respond to firm characteristics, leveraging doctors' transitions across firms. Descriptively, we find that when doctors move from a high C-section firm to a low C-section firm, their C-section rate falls by at least 30 percentage points. To address endogeneity of firm changes and changing patient composition, we exploit quasi-experimental variation in doctors' movement across firms generated by acquisitions in a Stacked Difference-in-Difference design. We find that when a firm at the 50th percentile of the firm-level C-section distribution acquires a firm at the 25th percentile of the distribution, the acquired-doctors' C-Section rates increase by five percentage points. We find no evidence of changing patient composition, further evidence that doctors change their practice styles to match their firms. Aggregating the estimated firm effects we demonstrate two important facts. First, some of the geographic variation in healthcare utilization is explained by average differences across firms. And second, within a given state or county, there is wide variation in firm practice styles.

*The views expressed in this article are those of the authors. They do not necessarily represent those of the Federal Trade Commission or any of its Commissioners.

1 Introduction

Where people live has an enormous impact on the kind and amount of medical care that patients receive [Skinner, 2011]. For example, Caesarian section rates vary widely at the county-level and these differences are persistent [Baicker et al., 2006, Robinson et al., 2023]. What causes this variation by geography? Studies of other kinds of care have considered the role of different beliefs or training based on where they live [Cutler et al., 2019]. An alternative is that patients who live in different areas have different demand for healthcare [Badinski et al., 2023].

Another possible explanation is geographic differences in average physician practice groups - i.e. firms in some places behave differently than others. When doctors switch to practice groups with higher medicare spending per patient, their treatment intensity increases [Doyle and Staiger, 2024], and when dialysis centers are acquired by large chains, treatment styles change to match the acquiring firm [Eliason et al., 2020]. Part of this effect may be driven by changes in incentives [Gruber and Owings, 1996, Gruber et al., 1999, Singh et al., 2022, La Forgia, 2023, Durrance and Hankins, 2018]. However, it is difficult for researchers to isolate and quantify the role of the firm on physician treatment practices due to the endogeneity of physician movements. Physicians that change groups may be different than others, and when they move to different locations they are likely changing firms as well.

In this paper, we leverage a unique dataset to overcome these challenges and separately estimate the effect of geography, physicians, hospitals, and firms on one of the most common medical procedures: whether a patient has a vaginal birth or a C-section. We examine C-sections because of fixed patient demand for birth, the high rate of surgical intervention, and the extensive literature examining determinants of C-section rates. Our dataset consists of all private healthcare claims for multiple large payers from 14 states from 2015-2020, with unique identifiers for doctors, firms, and healthcare facilities. This allows us to quantify the

role of the firm relative to other factors, in a model identified by doctors' changing firms and treating patients across facilities. Because of the length and breadth of our data, we can also exploit quasi-experimental variation in physician movement across firms driven by mergers of physician practices. Mergers generally don't change a physician's location, meaning we can isolate effects that are driven by firms.

We estimate that the firm a doctor works for is nearly as important as physician practice style. Firms explain 8% of overall variation in C-section rates and doctors 10%. Patient characteristics (age, diagnoses) and state are the most important factors, explaining 30% of overall variation. This finding helps quantify the relative importance of doctors compared to firms. We then examine the role in which firms may change providers' individual behavior. First, we show that when physician change firms, it materially impacts their C-section rates - moving from firms in the fourth quartile to the first quartile of the estimated firm effect lowers physicians' C-section rates from approximately 70% to 40%.

Next, to address the fact that doctors switching firms are endogenous and correlated with geographic moves, we examine exogenous shifts of doctors across firms driven by mergers of physician practices.[1] Using a Stacked Difference-in-Differences design [Cengiz et al., 2019], we find that when a firm at the 50th percentile of the firm effect distribution acquires a firm at the 25th percentile, the doctors in the acquired firm increase their C-section rate by approximately five percentage points (pp). We find no change in patient characteristics or location fixed effect - suggesting the effect is not driven by a changing patient population or changing geography. This large change in C-section rate is credibly entirely driven by the change in firm.

Lastly, we provide estimates of the geographic variation in firm behavior, and demonstrate that cross-location differences in C-section rate is at least partly driven by cross-location

[1] As an example, if a doctor moves from firm A to firm B because the destination firm allows their doctors to perform more C-sections, that effect is driven by the doctor.

differences in the average firm. As a concrete example, Cook County, IL has a relatively low C-section rate, and relatively low average firm C-section rate while Dallas County, TX has a relatively high C-section rate and a relatively high average firm C-section rate. These averages mask an important fact - within a state or county there is a substantial variation in C-section intensity.

Our study contributes to several literatures. First, we add to literature on geographic variation in healthcare utilization. Robinson et al. [2023] examines the role of geography in C-section utilization, while Badinski et al. [2023] examine geographic variation leveraging a model based around doctors' movement over time. Our model builds on the work of Currie and MacLeod [2017] and Badinski et al. [2023] by utilizing a unique dataset to estimate the share of variation that firms (in addition to patients, doctors, hospitals, and location) play in healthcare provision.

Our work also adds to the new literature examining the role of the firm on physician practice styles. Eliason et al. [2020] demonstrate that acquisitions of dialysis providers resulted in acquired practices moving towards acquired groups. La Forgia [2023], finds that OB-GYN C-section rates changed when providers were acquired by management companies; the direction of that change (more or fewer C-sections) varied by the management company in ways that were consistent with their publicly declared corporate strategies. Chernew et al. [2021] finds that vertically integrated physicians are more likely to refer to hospital-based providers. Doyle and Staiger [2024] shows that physician group switchers change their practice styles, even within the same hospital. We add to this literature by quantifying the aggregate effects of firm practices, estimating how much the "role of the firm" explains overall variation, and using physician group mergers to show how doctor behavior changes.

We also add to the literature on firm and worker fixed effects from the labor literature, as in Abowd et al. [1999] and Card et al. [2013]. These fixed effect models are increasingly being used in healthcare. Badinski et al. [2023] and Ahammer and Schober [2023] use fixed

effects models to estimate a fixed physician component, while Marquardt [2023] estimates a model incorporating physician identifiers.

Lastly, we add to the work examining the role of mergers and firms on healthcare outcomes, such as Beaulieu et al. [2020] for hospital mergers, or Koch et al. [2021] for hospital acquisitions of physicians. We demonstrate that after a merger, doctors in acquired firms behave differently, strong evidence of different norms and practices across firms driving differences in healthcare utilization.

2 What Drives Variation in Healthcare Utilization?

Better understanding of whether a given medical procedure is appropriate is important for patients. For this reason, it is important to better understand the different determinants of overall healthcare utilization, especially in situations where the different levels of care include a surgical intervention.

Understanding the aggregate role of the firm in healthcare utilization is very difficult for several reasons. First, patient demand for certain medical services can be induced by doctors, and patients can seek out different doctors depending on their own varying demands for healthcare. Second, practice intensity can be difficult to identify, depending on the procedure, resulting in an analysis that is underpowered.

For these reasons, we believe that labor and delivery is the best setting to begin to examine the aggregate impact of the firm in explaining healthcare utilization. C-sections provide a unique setting in that aggregate demand for children is credibly unrelated to variation in average practice styles within a region. Additionally, labor and delivery has a very high rate of surgical intervention.[2] This high rate of surgical intervention, and the numerous "gray areas" in when to perform one [Currie and MacLeod, 2017] offer us sufficient

[2]In our data, similar to aggregate statistics, the overall C-section rate is approximately one-third.

power to be able to examine the role of different factors in whether or not a C-section is performed.

The fact that labor and delivery is so common also means that our analysis on the role of the firm is relevant regardless of whether the findings may be different for other specialties. For patients, a C-section is an intense abdominal procedure which (1) has a longer recovery time than a vaginal birth and (2) is substantially more expensive. Whether a C-section is medically appropriate, a la Baicker et al. [2006], and better ways of reducing medically unnecessary C-sections is an important consideration. Investigating the determinants of and the value created by more intensive healthcare has also been investigated by Abaluck et al. [2016] and Chandra and Staiger [2020].

A great deal of prior research has shown that overall healthcare utilization is driven by geography [Robinson et al., 2023]. Epstein and Nicholson [2009] documents substantial and persistent variation across doctors within several geographic areas in New York and Florida. Chyn and Shenhav [2022] find evidence of place-of-birth having an effect on childrens' outcomes.

However, geography is likely inclusive of several factors, most notably hospitals and firms, which can play an important role in outcomes as well. Hospitals are generally in fixed locations, and have their own policies about when it is reasonable to perform a C-section; hospitals indeed play an important role in determining the outcomes for many births [Kozhimannil et al., 2013, Fischer et al., 2022, Battaglia, 2023]. Some obstetricians have admitting privileges at several hospitals, and may change their decision-making based on the given hospital. Similarly, as discussed by Singh et al. [2022] new evidence suggests that firm organization can affect doctors' decision-making process.

Understanding whether firms are relevant is particularly difficult due to data availability. A model which separately measures doctor and firm care intensity will need enough movers between firms to be able to pin down any firm effect. Therefore, this model requires data

with enough information over time and across a sufficient geography.[3] Our unique dataset of all private healthcare claims from four large payers in 14 states allows us to examine each piece of the puzzle individually.

Differentiating between geography, firms, and doctors is important for understanding causes of varying C-section rates. Is it the case that Texan doctors perform C-sections at a higher rate or is it instead that firms (i.e. physician practices) in Texas are more likely to perform C-sections? The first explanation suggests that doctors' training or state laws may be the explanatory factor, while the second suggests that firms' incentive or practice rules may matter more. Each explanation has different policy implications.[4] Now we turn to describing our data and the model in more detail.

2.1 The Individual and the Firm

To separate out the effects of doctors, firms, hospitals, and geography, we use multiple fixed effect models. Abowd et al. [1999] and Card et al. [2013] both utilize these models to examine changes in "worker sorting" - the increasing tendency of "high-wage" workers to be located in "high-wage" firms. These models output a set of fixed effects at the relevant level of analysis. Additionally, because the model is assumed to be additively separable, it allows us to understand the source of variation of each component individually.

The labor literature, fixed-effects models have been used in analyses of wage inequality [Song et al., 2018], outsourcing [Goldschmidt and Schmieder, 2017], and rent sharing [Card et al., 2018]. This type of model has now begun to be used more widely - Badinski et al. [2023] utilize a similar model (along with a model of healthcare demand) to examine differences in

[3] As an example, an all-payer claims dataset from a single state would have doctors entering and leaving neighboring states.

[4] As an example, if training or state laws are more important, any policy change would be more likely to impact future doctors. If firm rules are an important determination of C-sections, that has the ability to reduce unnecessary C-sections among both future doctors and currently practicing doctors.

healthcare utilization across geographies.

Inasmuch the observable characteristics we use generate reasonable comparisons across different births, we can examine the role of doctors, firms, hospitals, and geography using these models. Similarly, because any individual doctor or procedure will have an extremely small contribution to the overall calculation of a firm-level or facility-level fixed effect, we can examine the role that a doctor moving from a "high fixed effect" to a "low fixed effect" firm, location, or hospital will have on their C-section rates.

This type of analysis extends the use of fixed-effect examining a firms' production function through the lens of a worker component (doctors) and a firm component (practices). This type of analysis, particularly the role that workers play in firm production, is not as widely understood. Even though it is particularly appropriate in our context, where workers are explicitly making decisions, these types of models could be expanded to better understand the role of workers in firms' individual production functions.

3 Data and Fixed Effects Model

This section describes the claims data we utilize to identify births and associated firm and facility information.

3.1 Claims Data

We utilize a set of healthcare claims data provided to the Federal Trade Commission by five large insurers: Aetna, Anthem, Cigna, Health Care Service Corporation (HCSC), and United Healthcare. These data were provided to the FTC as part of a study into Physician Group and Healthcare Facility Mergers.[5] These include all the healthcare claims for the payers'

[5]For more information see: https://www.ftc.gov/enforcement/competition-matters/2021/04/physician-group-healthcare-facility-merger-study. To the extent our findings relate to information obtained from this

beneficiaries in Colorado, Georgia, Indiana, Illinois, Kentucky, Maine, Missouri, Montana, Nevada, New Hampshire, New Mexico, Ohio, Oklahoma, and Texas from 2015-2020. The claims are separated into two types, which are outlined below.

Professional Claims

We restrict to professional claims related to birth based on the associated CPT codes which include separate identifiers for firms and doctors. We also restrict for patients above age 15 and below age 46. For more detail about these data and the cleaning procedure, see the Data Appendix.

After cleaning the data, we have a total of 2,934,002 births. As a comparison, we estimate approximately 3.4 million births performed under private insurance in 6 six years, meaning our data contains more than 85% of births covered by private insurance.[6] Table 1 provides summary statistics on the sample of births in our data. 34.4% of births are a C-Section, similar to the national rate where approximately 1/3rd of births are C-sections.[7] 10.0% of births are Vaginal Birth after C-Section (VBAC), and 0.5% of births are multiple birth. Examining geographic variation - El Paso County, CO is approximately at the 5th percentile of C-section rate, at 24.7%, while El Paso County, TX is approximately at the 95th percentile, with a C-section rate of 45.0%.[8].[9] Overall, this demonstrates that our data contains substantial geographic variation in C-Section rates as well.

6(b) study it has been aggregated or anonymized as necessary to protect trade secrets and confidential or privileged commercial or financial information.

[6]Based on an estimate of 2022 populations of 16-45 year old women of approximately 20.235 million, a birth rate of 56 births per thousand in 2022, and assuming half of births being covered by Medicaid in our study states.

[7]https://www.cdc.gov/nchs/data/vsrr/vsrr021.pdf

[8]Here, we limit to counties with at least 50 births in our data.

[9]The county with the lowest C-section rate is Baca County, CO at 10.7%, while the highest is Jim Hogg County, TX at 59.8%.

Table 1: Descriptive Statistics: Births

	(1) Vaginal Birth		(2) C-section		(3) All Births	
	Mean	SD	Mean	SD	Mean	SD
Age	29.198	5.468	30.746	5.493	29.731	5.526
Multiple Births	0.002	0.044	0.010	0.101	0.005	0.069
Large Infant	0.000	0.010	0.000	0.020	0.000	0.014
Previous C-section	0.013	0.112	0.267	0.442	0.100	0.300
Placenta Previa	0.000	0.008	0.003	0.055	0.001	0.033
Placental Abruption	0.001	0.036	0.006	0.075	0.003	0.053
Cord Prolapse	0.000	0.013	0.002	0.047	0.001	0.030
Breech	0.006	0.076	0.089	0.285	0.035	0.183
Cord Around Neck	0.051	0.219	0.022	0.146	0.041	0.198
Uterine Bleeding	0.000	0.005	0.000	0.008	0.000	0.006
Eclampsia	0.000	0.012	0.001	0.027	0.000	0.018
Hypertension	0.000	0.017	0.001	0.025	0.000	0.020
Hyper. Dur. Preg.	0.002	0.044	0.003	0.056	0.002	0.048
Lung Issues	0.001	0.024	0.000	0.022	0.001	0.023
Heart Issues	0.001	0.026	0.001	0.035	0.001	0.029
Diabetes	0.000	0.019	0.001	0.033	0.001	0.025
Anemia	0.001	0.037	0.002	0.044	0.002	0.040
Blood Disorder	0.000	0.013	0.000	0.013	0.000	0.013
Herpes	0.000	0.017	0.001	0.023	0.000	0.020
Other Disease	0.000	0.005	0.000	0.005	0.000	0.005
Polyhydramnios	0.003	0.055	0.007	0.085	0.005	0.067
Prev. Complications	0.001	0.024	0.002	0.040	0.001	0.031
Observations	1923872		1010130		2934002	

Facility Claims

Facility claims represent all claims provided by a hospital, and includes information on allowed amounts for the hospital, diagnosis codes, and services rendered by the hospital. These data also include Hospital TINs, which provide the relevant information about which hospital a given procedure is being performed at.

Linking Professional and Facility Claims

The claims data also contain facility and professional claims. We link these claims following a procedure outlined in the Data Appendix. After linking and subsequent restrictions, we are left with a total of 2,103,184 births.

4 Model and Sample Restrictions

In this section we outline the two-step model we use to estimate firm and facility effects. We do this by first estimating a model of C-section appropriateness based on patient-level observables. Then we use a multi-way fixed effects model estimation strategy.

Model of C-section appropriateness

In order to investigate the effect of firm and potentially facility-related effects, we use a two-stage model following Currie and MacLeod [2017] and Badinski et al. [2023].

First, using the sample of 2,934,002 births in Table 1, we estimate a discrete choice linear probability model to measure the appropriateness of a C-section for a given patient.

$$\rho_i = X'_{it}\beta \tag{1}$$

Where for a given patient i, ρ_i represents the predicted probability of having a C-section based on patient diagnoses. As in Currie and MacLeod [2017], we assume that there is an underlying state of the patient, which can be approximated by the predicted probability of a C-section $\hat{\rho}_i$ as a measure of the appropriateness of a C-section. As our explanatory variables we use patient diagnoses (multiple birth, breech, placenta previa, etc.) and state fixed effects. The state fixed effects account for potential variation across states in laws or regulations about appropriateness for a given procedure, or different medical malpractice rules.

Because we are estimating this model over the population of C-sections across the study states, we assume that the effect of any individual doctor, firm, or facility will be averaged out, and that after accounting for state-level variation, we will pick up a measure of the underlying appropriateness of a C-section. This model is the similar to the first step of Currie and MacLeod [2017].

The results of this estimation, and the set of variables used, are shown in Table 2. As we can see, the results suggest that certain diagnoses greatly increase the probability of a C-section, including having twins or having a previous C-section.

We then utilize these results to estimate $\hat{\rho}_i$, and test the accuracy of this model against observed C-sections versus vaginal births.[10] As we can see in Figure 1, we find that vaginal births are rarely performed when $\hat{\rho}_i$ >0.35. However, there are still a large number of C-sections performed when $\hat{\rho}_i$ <0.35, suggesting both that we have identified an appropriate model of appropriateness, but also that there is substantial room for discretion in whether to provide a C-section.

[10]Unlike Currie and MacLeod [2017] we use a linear probability model, which is less precise in the tails of the distribution and may have predicted probabilities outside of [0,1]. We find that our model has few predictions outside this range, and the majority of the distribution is away from the tail.

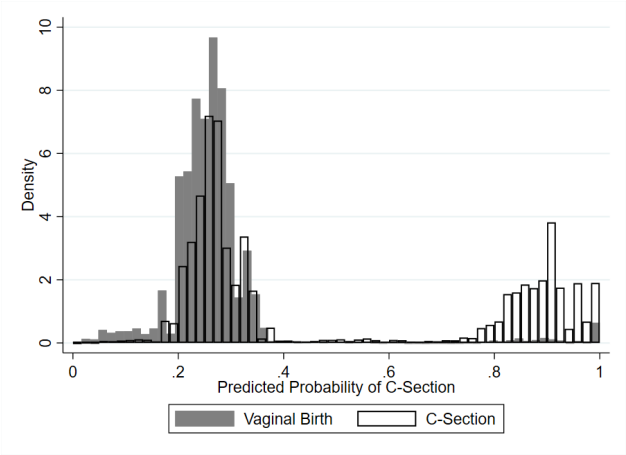
Table 2: Model of C-Section Appropriateness

	Patient Has C-section	
	Point Estimate	SE
Age 25-29	0.0276	(0.000727)
Age 30-34	0.0449	(0.000713)
Age 35+	0.0977	(0.000794)
Multiple Births	0.273	(0.00355)
Large Infant	0.343	(0.0173)
Prev. C-section	0.389	(0.00388)
Other Organ Issue	0.239	(0.00381)
Placentia Previa	0.553	(0.00743)
Placental Abruption	0.350	(0.00462)
Cord Prolapse	0.663	(0.00836)
Breech	0.571	(0.00134)
Cord Around Neck	-0.148	(0.00125)
Uterine Bleeding	0.200	(0.0397)
Eclampsia	0.375	(0.0133)
Hypertension	0.0194	(0.0166)
Hyper. Dur. Preg.	0.0848	(0.00507)
Lung Issues	-0.0435	(0.0105)
Heart Issues	0.112	(0.0113)
Diabetes	0.220	(0.00991)
Anemia	0.0291	(0.00613)
Blood Disorder	-0.00791	(0.0190)
Herpes	0.162	(0.0125)
Other Disease	-0.0162	(0.0520)
hydra	0.136	(0.00364)
Prev. Complications	0.220	(0.00801)
N	2904935	

Standard errors in parentheses

Coefficients on diagnoses in linear probability model. Estimates on state fixed effects not included.

Figure 1: C-sections by Predicted Probability



Represents distribution of predicted probabilities of C-section rates based on linear model of patient characteristics and state, and observed final outcome, i.e. whether or not the patient had a C-section or vaginal birth.

Multiple Fixed Effects Estimation Strategy

We then incorporate the individual $\hat{\rho}_i$ to estimate the following model:

$$y_{i,l,d,j,h,t} = \hat{\rho}_i + \phi_l + \delta_d + \pi_{j(d,t)} + \eta_h + P_{i,d} + P_{i,d,h} + T_t + \epsilon_{i,l,d,j,h,t} \tag{2}$$

Equation 2 describes the decision-making process about whether to have a C-section. On the left hand side is a binary indicator for whether patient i, who lives in three-digit zipcode l, seeing doctor d at firm j, in hospital h at time t will have a C-section. $j(d,t)$ is an assignment function, which denotes whether doctor d is in firm j at time t.[11]

$\hat{\rho}_i$ represents the effect of patient characteristics on the probability of having a C-section. ϕ_l represents the effect of geography alone, when separated from other factors. Here, ϕ_l is used as the three-digit zipcode where the patient lives. We believe this is analagous to geography of where the patient gives birth, as most patients give birth relatively close to

[11]Note that an extremely small fraction of doctors bill to multiple firms in the same period.

where they live.

δ_d is the "doctor fixed-effect", which can be thought of as a measure of physicians' overall propensity to perform a C-section.[12]

$\pi_{j(d,t)}$ is the "firm fixed-effect", which is the firm-level component. One way to interpret this measure is that certain firms have rules about when they start performing C-sections, or that firms have different training measures which impact physicians' decision-making.

η_h represents hospital fixed effects, based on the linked facility TIN. This can be thought of as the hospital-level component that is separate from geography. As an example, some hospitals may have rules about which patients they will or will-not perform C-sections on.

$P_{i,d}$ denotes the logarithm of the estimated cost of a C-section that will be paid to obstetricians relative to a regular birth for a given patient, while $P_{i,d,h}$ denotes the logarithm of the total estimated cost, inclusive of facility fees and fees paid to other professionals (i.e. anesthesiologists).

T_t denotes quarter-year fixed effects to account for variation driven by time, and $\epsilon_{i,l,d,j,h,t}$ is an error term.

This model treats the decision to have a C-section as a linear combination of patient's condition, where the patient lives, the doctors' surgical and diagnostic skill, the firm's decision and the hospital's decision, conditional on prices.

The model is similar to the mover-based studies of physicians [Badinski et al., 2023], but utilizes our unique dataset to examine the role that firms and facilities play as well, separate from geography. Because we have longitudinal claims data from a variety of payers, we can separately identify each of these components, and identify which is driving our variation.

The model is identified by several sources of variation. The doctor and firm fixed effects

[12] Unlike Currie and MacLeod [2017] we are not incorporating a physician-level diagnosis measure, primarily to focus on the role of firms versus doctors in explaining overall variation. When we have incorporated a physician-level diagnosis component all our results are similar.

are identified by doctors moving across firms, which allow us to identify doctors and firms in a given connected set [Card et al., 2013]. The hospital fixed effects are identified by doctors working across multiple hospitals, which happens both due to physicians moving, but also due to cross-hospital admitting privileges of obstetricians. Lasly, the location effects are identified by physicians seeing patients across multiple geographies, patients choosing between hospitals, and doctors moving across geographies as well.

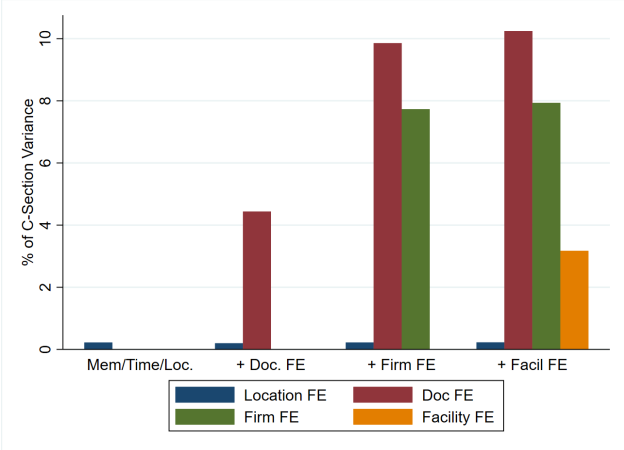
One important factor is that this model does not identify any factor prior to the C-Section decision, i.e. patients' choice of hospital and doctor. We assume that these factors, and important factors such as patient age, will affect the baseline decision-making process. The doctor, firm, and hospital fixed effects will instead pick up the role that these entities play in "gray areas" - ceretis paribus is that entity more or less likely to choose a C-Section when compared to other doctors, firms, or hospitals?

5 What Drives Variation in C-Section Rates?

One benefit of the linearity of this model is that it allows for separate addition of individual components to identify which factors are the most important determinants of C-section variation. This is particularly relevant to discussions about geographic variation in C-section rates across the United States - hospitals are obviously stationary, meaning that it is difficult to distinguish between geographic and hospital factors. By estimating the model using different components (i.e. including doctor FEs but not including firm FE) we can uncover what cross-correlation exists between these factors.

We outline these one-by-one versions of the model described by Equation 2. Based on that model, taking the variance of both sides allows us to identify the components that are driven by location, doctors, firms, and hospitals individually, and how the variance changes. The results of this analysis are shown in Figure 2.

Figure 2: C-Section Variation Explained under Various Models



Represents share of overall variation in C-section rates under different models, calculated by taking the variance of each object divided by the overall variation in C-section rates in the data.

As we can see in the first model, after conditioning on patients' predicted probability of having a C-section, geography only explains a small portion of overall variation in C-section rates. When we include doctor fixed effects, we find that a 5% of the variation in C-section rates driven by doctor fixed effects.

In the most-saturated model that includes location, doctor, firm, and facility fixed effects, we find that geography explains the least amount of variation. Patient characteristics (not shown) explain 25% of overall variation in C-section rates. Next most important are fixed physician components (doctor FE) - nearly 10% of the variation in C-section rates is explained by doctors alone. Firms are the next most important at 8%, and hospitals alone explain just under 5% of variation.

In Table 3, we provide tables illustrating the summary statistics on the various fixed effects estimates outputted from the model in Equation 2.

As we can see, the effects are small and centered around zero, but there is a substantial amount of variation - standard deviation for both effects is approximately 0.2 for doctors,

Table 3: Descriptive Statistics: Fixed Effects

	Mean	Std. Dev.
Location FE	0.006	0.109
Doctor FE	0.007	0.201
Firm FE	0.004	0.186
Hospital FE	-0.005	0.171

Represents descriptive statistics of estimated location, doctor, firm, and hospital fixed effects from multi-way fixed effect model.

firms and hospitals. Location alone on the other ahnd is a relatively small factor, with less overall variation than the other factors.

5.1 Do Firm Effects Matter?

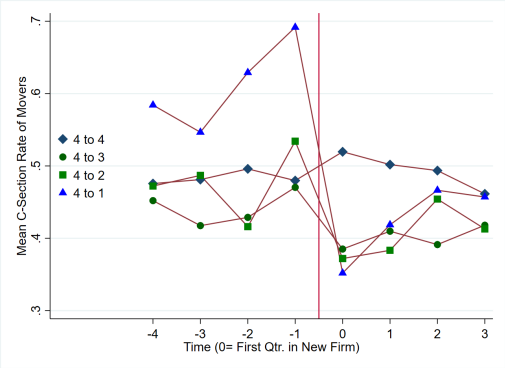
The results in Figure 2 provide evidence that firms explain are nearly as important as doctors in explaining the overall variation in C-section rates. We next examine whether these estimates have a direct relationship to doctors' behavior, by examining whether doctors' behavior changes after they change firms.

Previous studies have examined the role of doctors changing location [Badinski et al., 2023], our data allows us to examine the impact of doctors changing *firms*. This is particularly relevant from a policy perspective. If firms are an important part of the overall process of determining whether or not a patient receives a C-section, then more attention and research should be directed at understanding how firms contribute to this process, and whether they promote unnecessary C-sections.

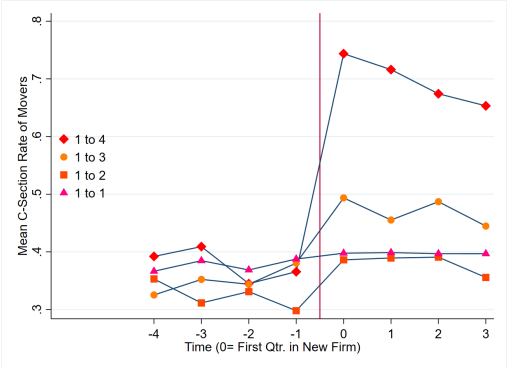
We next perform a movers' analysis, similar to Card et al. [2013]. First, we bin firms into quartiles based on their estimated fixed-effects. Then we examine the C-section rates before and after doctors move across firms. We identify "movers" based on whether doctors "main firm" changes in a given quarter, where their main firm is defined by the TIN at which the

Figure 3: Mover Analysis: Firms

Panel A: Doctors Leaving Firms in 4th Quartile of Firm FE



Panel B: Doctors Leaving Firms in 1st Quartile of Firm FE



Note: Represents average C-section rates of movers before and after a move. Physicians move from firms in the first quartile or fourth quartile of the firm fixed-effect distribution to other firms, broken out by quartile.

doctor performed the most births within a given quarter.

In Figure 3 we show the results of this analysis - Panel A shows the C-section rates for doctors leaving firms in the fourth quartile of the fixed-effects estimates, and Panel B shows the C-section rates for doctors leaving the 1st quartile C-section firms. Broadly, we find evidence doctors respond quite a lot to the firms they work for.

Doctors moving from the lowest quartile to the highest quartile see an 30pp increase in their C-section rate, from approximately 40% to approximately 70%. Doctors moving from the highest quartile to the lowest see their C-section rate decrease by 30pp, from 70% to 40%. We see similar patterns for doctors moving from the 4th or 1st quartile to the second or 3rd - when doctors move to firms with higher estimated fixed effects measures, they start to perform more C-sections.

This analysis shows that, descriptively, there is a change in C-section rates that occurs when doctors change to higher or lower-intensity C-section firms. This result demonstrates the importance of firms, and also illustrates the difficulty of analyses of geographic variation

in healthcare utilization. When doctors change location they are more than likely changing firms as well - when Badinski et al. [2023] finds limited effects of cardiologists changing geography, were the doctors changing to largely similar firms?

However, while these results are suggestive that firms play an important role in doctors' decision-making the transitions are endogenous. Firstly, doctors who want to perform more C-sections may seek out a high C-section firm - in the counterfactual world where they did not transition, it is possible they start to perform more C-sections regardless. This would violate any parallel trends assumption, and any assumption of exogeneity - the counterfactual outcomes are correlated with the unobservable decision to switch firms.

Additionally, a firm transition may be a combination of effects - the change in C-section rate observed due to changing firms may additionally represent a change in geography, facilities, or patient mix. Patient mix in particular is not accounted for in this descriptive transition model, meaning that we may be observing "appropriate" changes in C-section rates, if doctors are more likely to see higher-risk patients after they move. Lastly, it is always possible that even if the specific firms doctors move to is random, movers themselves may not be - if doctors who switch firms are more responsive to changes in firm effect, these results might overestimate the effects of firms, and not generalize when examining all doctors.

To address these concerns, we utilize an identification strategy that quasi-randomly sorts doctors across different firms. We do this by using mergers of physician practices, which shifts doctors but does not shift these other characteristics, especially patient mix. We describe our strategy in the next section.

6 Identification Strategy: Physician Practice Mergers

To address the endogeneity of both different patients and physician selection into movement between firms, we utilize a quasi-experimental source of physician movement across firms: mergers of practices.

These mergers provide useful variation for our analysis - individual doctors in acquired firms are not choosing to move individually. Instead, their firms were acquired by an outsider with potentially different rules.[13] This addresses the endogeneity of moves, particularly that doctors are deciding which firms to go to based on whether they perform more or fewer C-sections.

An additional benefit of this strategy is that it allows us to isolate the firm component. The merger of two practices may not broadly change the patient mix, facilities doctors work at, or their location. Instead, all that has changed is ownership. We test this assumption, and find no evidence of factors outside of ownership changing. This strategy allows us to isolate the effect of changing firm on C-section in a way that is independent of other factors. Using mergers as an identification for a change in practice styles has been done previously in the literature, particularly Eliason et al. [2020], who examined the effect of acquisitions on practice styles in dialysis centers.

We identify mergers in our data using large flows of doctors from "acquiring" to "acquired" firms, and the subsequent deaths of TINs. For more information on how we identify mergers, see the Data Appendix.

We identify 407 doctors in 91 firms who move from an acquired firm to an acquiring firm in our sample. There are approximately 19,820 obstetricians in the United States, meaning that this represents approximately 2% of all obstetricians in the country.[14]

[13] We acknowledge the possibility that firms may decide who acquires them. We address this possibility in the next section.

[14] See https://www.bls.gov/oes/current/oes291218.htm

With these doctors and treatment timings, we seek to estimate a differences-in-differences model which compares doctors whose firms are acquired to doctors in non-acquired firms, additionally comparing doctors who were acquired by practices with higher vs. lower firm fixed effects. Because there is variation in treatment timing, which can lead to negative weighting, we utilize a Stacked Difference-in-Differences design with an intensity of treatment measure, which addresses the timing concerns. [15]

First, for each unique k_d, we identify all doctors in k_d who did not move firms between k_d and $k_d + 1$, and who were never part of an acquisition in our sample, and never part of an acquiring firm. These control-group doctors are given the "acquired quarter" as a placebo treatment date. Following Cengiz et al. [2019], we then "stack" the data so there is no variation in treatment timing. [16]

6.1 Estimating Equation

In the Stacked design, denoting the treatment time m as $m = t - k$ we seek to estimate the following regression:

$$\begin{aligned} y_{i,l,d,j,h,t} = \alpha + \sum_{m=-8}^{7} \gamma_m \times \Delta\pi_{j(d,k)} \times 1\{t - k = m\} \\ + \tau_m \times \lambda_k + \tau_m \times \delta_d \\ \hat{\rho}_i + \phi_l + \eta_h + P_{i,d} + P_{i,d,h} + \epsilon_{i,l,d,j,h,t} \end{aligned} \tag{3}$$

Equation 3 is our estimating equation of the stacked difference in differences design.

[15] Because mergers occur at various times, using a simple two-way-fixed-effect estimator would result in the acquired doctors being used as a control-group for the treated doctors in later periods, which would bias our estimates. The "Stacking" design accounts for this by treating all mergers as though they occurred simultaneously.

[16] This implicitly allows for doctors in the control group to be in the data multiple times. These are treated as separate individuals.

Letting t be quarters in time, and k be the "acquired quarter". $y_{i,l,d,j,h,t}$ is whether or not patient i in time $m = t - k$ seeing doctor d, living in location l, at hospital h has a C-section.

$\Delta\pi_{j(d,k)}$ is the treatment effect, the change in the firm FE between the acquired firm and the acquiring firm. We use this intensity of treatment measure to account for the fact that acquisitions can result in doctors to higher or lower FE firms. Because the estimates of the FE measure need to be contextualized, we normalize the treatment effect by the estimated difference between the 25th and 50th percentile of the firm fixed-effect distribution.[17]

Our object of interest is γ_m which estimates the effect of a doctor in a firm at the 25th percentile of the firm-effect distribution being acquired by one at the 50th percentile of the distribution in period m. By using a dynamic setup we can examine pre-trends and dynamic effects over the two years before and after. By construction the treatment effect is zero for control-group doctors, since they do not change firms at time k. [18]

We include $\tau_m \times \lambda_k$ - acquired quarter fixed effects interacted with time fixed effects.[19] $\tau_m \times \delta_d$ are "acquired quarter" fixed effects interacted with doctor fixed effects. These are necessary for the stacking design.

Lastly, we include the predicted estimates of the other fixed effects (hospital, location, patient diagnoses, doctor diagnosis factor) from the model in Equation 2. We include these controls to make sure that we are not biasing the estimates due to other changes due to the merger. Not including these controls makes no difference in our estimates.[20]

This setup requires the standard difference-in-differences assumptions, which we outline

[17] We do not make restrictions based on mergers which occur close to the beginning or end of our data (i.e. in early 2015 or late 2020) because there will still be an estimate of the overall effect immediately before and after the merger.

[18] We allow both treatment and control group doctors to change firms in later quarters. We do not anticipate that this will bias our estimates, but may result in imprecise measures of the effect, especially if treated or control group doctors later change firms. Changing firms is an outcome variable, and conditioning on it would be inappropriate.

[19] These are equivalent to acquired quarter fixed effects interacted with calendar quarter fixed effects.

[20] For example, if the acquired doctors were in-network with one payor, but the acquiring firm was not, this might change the patient mix in some way. Including these controls attempts to account for those changes.

below.

6.2 Identification Assumptions

Assumption 1: Exogeneity

This assumption states that, in potential outcomes framework, counterfactual outcomes for the acquired doctors are mean independent of the merger-decision, i.e. that there is no relationship between being acquired and the change in counterfactual outcomes. Essentially, we are assuming that were it not for the acquisition, acquired doctors would have performed roughly the same amount of C-Sections.

This allows for the possibility that doctors have a voice in whether their practice is acquired, but that their behavior would have been similar in the counterfactual where the merger did not occur.

Assumption 2: SUTVA

This assumption states that for a given group of doctors in an acquired firm, there is no effect on the control group. Again, we believe this assumption is reasonable - we think it is unlikely that individual doctors would have their decision-making affected by a merger of other firms, conditional on seeing the same patients in the same locations. Additionally, because we remove doctors from the acquiring firm, there will be no spillovers.

Assumption 3: No Anticipation

This assumption states that the pre-merger outcomes for all doctors are unaffected by the merger before it occurs. We believe that this assumption is reasonable because of the staggered timing outlined above, where we only investigate when doctors change firms. Given the dynamic effects it is also possible to observe whether there is a violation of this assumption

Assumption 4: Parallel Trends

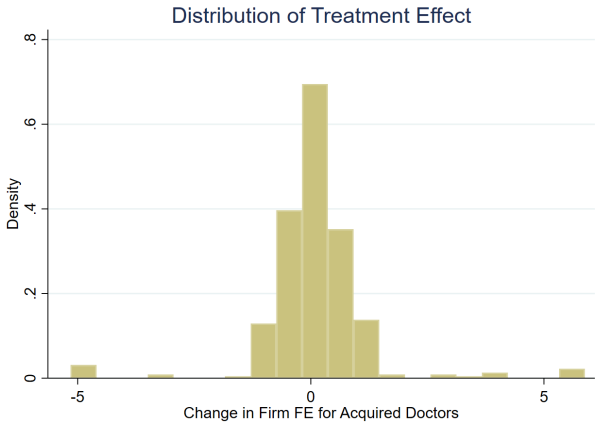
This assumption requires that the trends in C-Section rates for acquired and control group doctors would have been similar but-for the merger. Again, we believe this assumption is reasonable. Figure 3 suggests that C-Section rates are fairly stable over time prior to possibly endogenous job changes. Additionally, the dynamic specification allows us to visually examine whether there appears to be any violation of parallel trends, it appears that the parallel trend assumption holds.

Lastly, the specific setup assumes two facts. First, is linearity in effects, i.e. moving from the 25th to 35th percentile has half the effect of moving from the 25th to 45th percentile. The second is bi-directionality, moving from the 25th to the 35th percentile has the opposite effect of moving from the 25th to the 15th. We believe that the evidence in Figure 3 is consistent with the bi-directionality assumption. When doctors move to a higher (lower) FE firm they perform more (fewer) C-sections.

For linearity it is difficult to test, mainly due to lack of power. However, we find that median absolute value of the treatment effect is approximately 0.357, representing a change from the 25th to 34th percentile of the firm fixed effects distribution. The 90th percentile is approximately 1.167 - for nearly 90% of the acquired doctors, the change in the fixed effect measure is less than moving from the 25th to 50th percentiles (though this large change does occur for over 10% of doctors.) Figure Figure 4 shows the distribution of treatment effects among acquired doctors.

While overall linearity may not be true, the treatment effects are on average fairly small shifts, moving up approximately 10% of the distribution at the median, and we believe linearity for the smaller estimated treatment effects is a more reasonable assumption. Next, we turn to our results.

Figure 4: Distribution of Treatment Effect Among Acquired Doctors



Note: Represents distribution of treatment effect, calculated as the difference between firm fixed effect in acquiring and acquired practices for all acquired doctors. N=407. Standardized relative to a 1-quartile move based on the overall firm fixed-effect distribtution.

7 Results

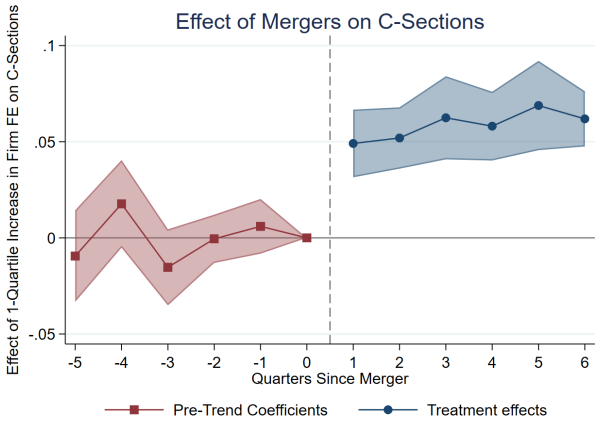
7.1 Effect of Mergers on C-Section Rates

This section provides our main estimates on the effect of doctors switching firms on C-section rates, using mergers as our source of identifying variation. We find evidence consistent with Figure 3, where changing firms dramatically alters physician practice style.

The overall results are shown in Figure 5 and shows our preferred estimation: the dynamic effects of doctors in at a firm in the 25th percentile being acquired by a firm at the 50th percentile of the firm-effects distribution.

Prior to acquisition, there is no evidence of any pre-trends or anticipation. After the doctors change firms, their C-Section rates increase by approximately 5 percentage points. This change is large, and is equivalent to the difference in C-section rates between Cook County, IL (Chicago) and Dallas County, TX (Houston), which are at approximately at the

Figure 5: Stacked Difference-in-Difference: Mergers



Note: Represents dynamic stacked difference-in-differences estimate of the effect of changing firm fixed effect on C-section rate among acquired doctors, relative to never-acquired doctors. The treatment effect represents the change in the estimated firm fixed effect between acquiring and acquired doctors. Clustered at the pre-treatment TIN level. Includes controls for provider interacted with cohort, and time interacted with cohort. Also includes controls for patient characteristics, facility fixed effect, location fixed effect, and prices.

25th percentile and 60th percentile of the geographic distribution respectively.

Most treatment effects are rather smaller than that. Using the median absolute value of the change in C-section rates is approximately 1.8 percentage points, corresponding to approximately an 5% increase in C-section rates relative to a baseline of 1/3rd of all births being C-sections.

There is no evidence of any dynamic effect; the effect is instantaneous and stable. This suggests that the change is a change in practice rules - rather than doctors becoming better diagnosticians or surgeons, which we would expect would increases over time due to learning, they abide by the rules at their new firm and behave accordingly.

These results demonstrate the important role that firms play in healthcare utilization, seemingly moreso than geography. When doctors change firms, their practice styles change substantially. When you only look at doctors who changed firms due to mergers, the effects are consistent with firms playing an important role in the provision of healthcare. In B we also explore whether there is heterogeneity in our results, and we find no evidence of heterogeneity. The fact that there is no heterogeneity for patients with a fairly low predicted probability of C-section is potentially evidence that these firms are performing more unnecessary C-sections.

7.2 Robustness Checks

Next, we examine our findings to see if there are potentially contemporaneous changes that could explain our results. One concern with our analysis is that mergers may change doctors' behavior in a way that the controls do not account for.

We test this possibility by examining whether there are changes in patient observables or other factors. We estimate Equation 3 using a binary treatment indicator - specifically, we estimate Equation 3 using a binary treatment indicator $\gamma_m \times Post$. Instead of using whether

a birth was a C-section as the outcome measure, we instead examine the relationship between the change in the firm fixed effect and variables measuring observed patient composition, facility fixed effects, or other factors that would suggest there has been a substantial change in patient mix.

We find little evidence of a change in patient characteristics or any other factor. In particular, while there is a small decline in the estimated facility effect, suggesting doctors are seeing patients at facilities which perform fewer C-sections, this may be (1) correlated with the overall firm effect, especially if the acquiring practice was a hospital system, and (2) is negative, suggesting that any increase in C-sections is in-spite of these changes. Additionally, the estimated change is only 0.031, representing a change of 17% of a standard deviation. Nevertheless, this the use of the facility effect as a control in our estimation.

There is no statistically significant estimated change on patient age, previous C-section rate, or the predicted probability of having a C-section. Lastly, there is no estimated change in the average relative cost, suggesting that there is not much change in the overall price of C-sections relative to vaginal births after acquisition. The lack of findings of any change in patient characteristics and price is

Lastly, we examine whether there is a change in number of births, in order to investigate whether these effects are being driven by changes in patient volume. For example, the model may be estimating effects that occur due to doctors having less time, and performing a C-section because they are quicker. This would change the interpretation of our results, doctors do more C-sections because they are seeing too many patients. We collapse our data to averages at the doctor-quarter level and run the same estimation as Equation 3, examining whether there was a change in total number of births in a quarter, conditional on quarterly-level averages of patient characteristics, facility and location fixed effects, and prices. The results of this analysis are shown in Figure 6.

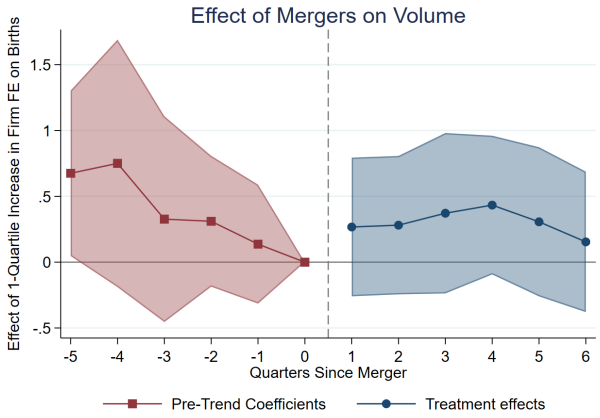
We find no evidence that the effects are driven by volume, suggesting that the changing

Table 4: Effect of Mergers on Patient Characteristics and Outcomes

Outcome	$\gamma_m \times Post$	SE
Location FE	-0.000	(0.000)
Facility FE	-0.031**	(0.010)
Multiple Birth	0.001	(0.001)
Prev. C-section	-0.009+	(0.005)
Breech	-0.001	(0.003)
Eclampsia	0.000	(0.000)
Patient Age	0.066	(0.055)
Patient Over 35	0.003	(0.004)
Predicted Prob. C-sec	-0.007	(0.005)
Csec. Rel. Cost	0.006	(0.012)

Standard errors in parentheses. Rows outline listed outcome variable. Column (1) represents point estimate of stacked difference-in-differences estimate of the effect of changing firm effect among acquired doctors, relative to never-acquired doctors. The treatment effect represents the change in the estimated firm fixed effect between acquiring and acquired doctors. Clustered at the pre-treatment TIN level. Includes controls for provider interacted with cohort, and time interacted with cohort. For regressions where outcome is patient characteristics, includes controls for patient characteristics, facility fixed effect, location fixed effect, and prices. For regressions where outcome is location and facility fixed effect includes controls for prices. + $p < 0.1$, * $p < 0.05$, ** $p < 0.01$

Figure 6: Difference-in-Difference: Mergers on Volume



Note: Represents dynamic stacked difference-in-differences estimate of the effect of changing firm fixed effect on total number of births by quarter among acquired doctors, relative to never-acquired doctors. The treatment effect represents the change in the estimated firm fixed effect between acquiring and acquired doctors. Clustered at the pre-treatment TIN level. Includes controls for provider interacted with cohort, and time interacted with cohort. Also includes controls for quarter-level averages of patient characteristics, facility fixed effect, location fixed effect, and prices.

C-section rates are driven by changing decisions about whether to perform a C-section, rather than doctors seeing substantially more patients.

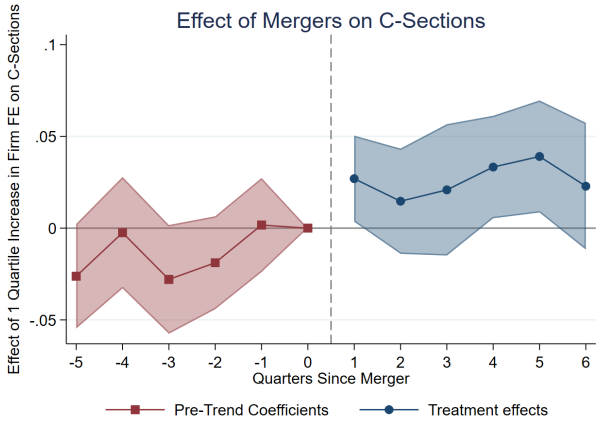
Figure 7 shows the estimated effects when using a binary estimator and restricting to doctors with positive changes in firm fixed effect. We find largely similar estimates when doctors are acquired by firms with a higher firm fixed effect - their C-section rate increases by 2.5 percentage points.

7.3 Alternative Mechanism: Is There a Learning Effect?

While we have focused on the role of the changing firm, another possible mechanism is that doctors learn from one another, and that could be explaining our effects.

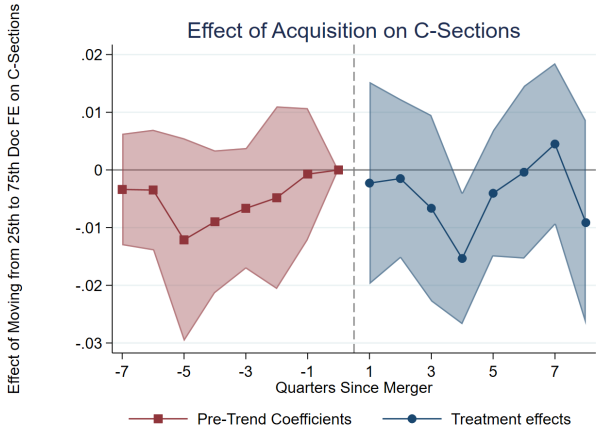
If it were the case that learning was happening, we would also expect learning on the side

Figure 7: Robustness: Binary Treatment



Note: Represents dynamic stacked difference-in-differences estimate of the effect of changing firm fixed effect on C-section rate among acquired doctors, relative to never-acquired doctors. The treatment effect is a binary indicator, and restricted to acqusitions with an increase in the firm fixed effect. Clustered at the pre-treatment TIN level. Includes controls for provider interacted with cohort, and time interacted with cohort. Also includes controls for patient characteristics, facility fixed effect, location fixed effect, and prices.

Figure 8: Difference-in-Difference: Learning



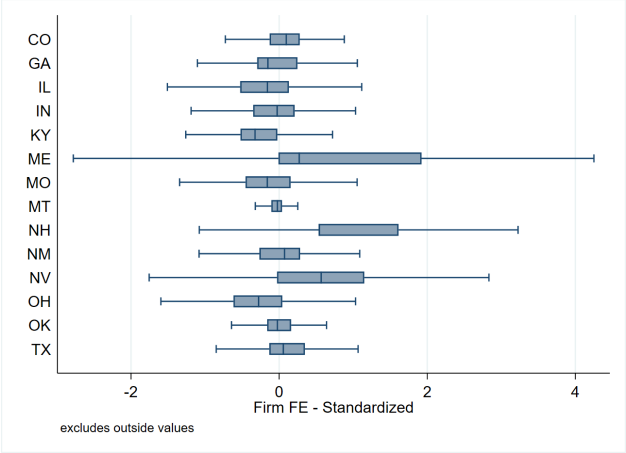
Note: Represents dynamic stacked difference-in-differences estimate of the effect of weighted average of acquired doctor fixed effects on C-section rate among doctors in acquiring firms, relative to doctors never in acquiring firms. Clustered at the pre-treatment TIN level. Includes controls for provider interacted with cohort, and time interacted with cohort. Also includes controls for patient characteristics, facility fixed effect, location fixed effect, and prices.

of the doctors in the acquiring firm, i.e. when practices merge, do the doctors in practice A and practice B converge to somewhere in between the two?

We test this using a variation of Equation 3, where instead of examining the changes for doctors at the acquired firm, we examine whether there are changes for doctors at the *acquiring* firm. We utilize the same Stacked Difference-in-Differences strategy, where the treated group is the doctors in the acquiring firm, and instead of $\Delta\pi_{J(d,t)}$, we estimate use $\Delta d_{J(d_j,t)}$, which is the difference in estimated fixed effect between doctor d in acquiring firm and the average of the doctor fixed effects for the doctors in the acquired firm. This represents the average change due to exposure to a new set of coworkers.

The results of this analysis are shown in Figure 8. We find no evidence of any learning effect, where doctors in acquiring firms seemingly do not change their behavior at all. There

Figure 9: Geographic Differences in Firm Effects

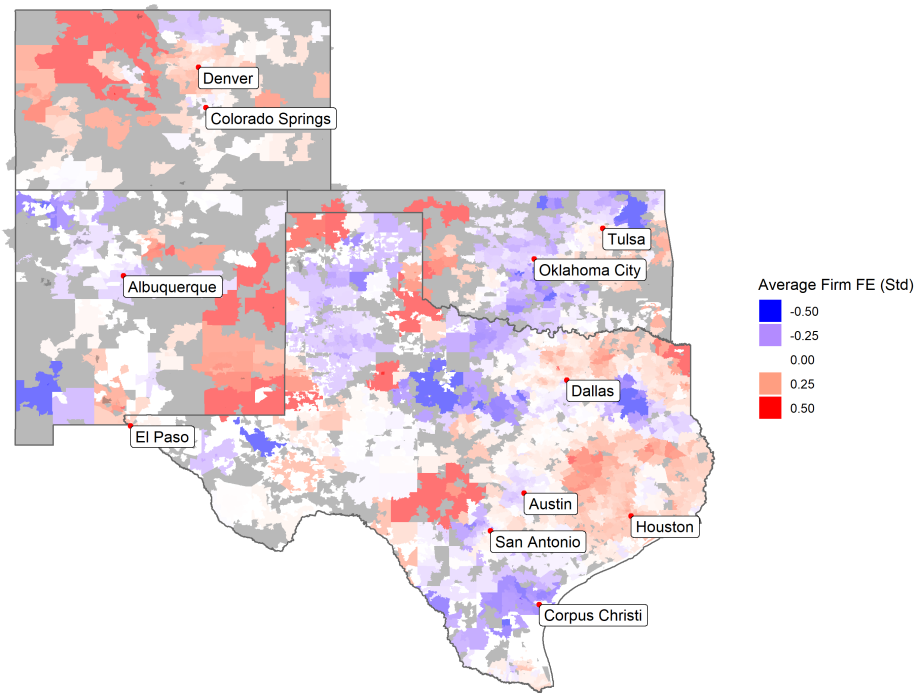


is little evidence of any change in C-Section rates. This result suggests that rather than just doctors, firms play an important role in healthcare outcomes. It also suggests that learning from new coworkers is potentially driving a very small part of overall variation. Instead, doctors appear to respond to the new rules at their firm.

8 Geographic Variation in Firm Fixed Effects

Following our findings that firms play an important role in the provision of healthcare, we now turn back to geography, to better understand what type of variation there might be *within* a given geography in the firm fixed effect. By only examining the role that geography plays, we might be missing underlying variation due to firms, which can confound our analyses. We attempt to show these effects in Figure 9, which shows the distribution of firm fixed effects within a given state. What we can see is that some states have higher average firm fixed effect than others, suggesting that state performs more C-sections overall. However, within every state the distribution is quite wide - there are high and low FE firms that perform a substantial number of births across the distribution.

Figure 10: Map of Average Firm Fixed Effect: South

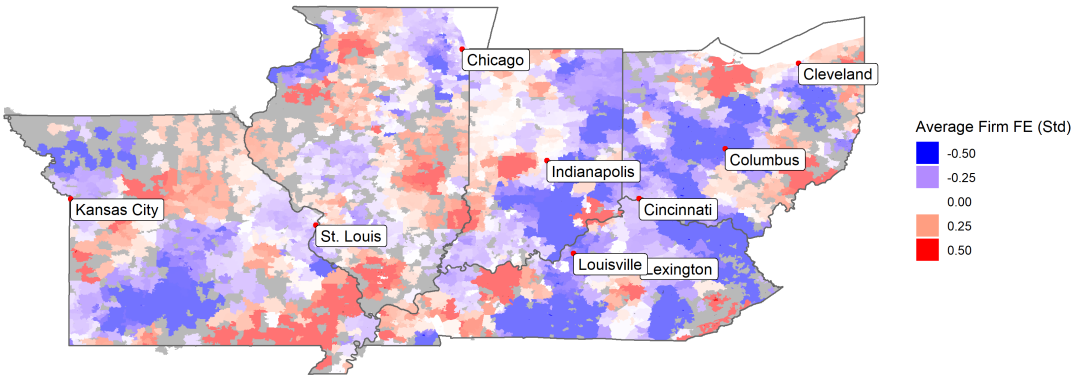


Note: Represents average firm fixed effect for all births in our final sample at the three digit zipcode. Restricted to zipcodes with more than 20 births. Average fixed effects values winsorized so the minimum value is -0.5 of the standardized value of the firm fixed effects, and 0.5 is the maximum value.

In Figures 10 and 11, we also provide maps of the average estimated firm fixed effect at the zipcode level - for collections of some of the states in our data.

The goal of this exercise was to demonstrate that by focusing on the role of just geography in the utilization of healthcare, researchers are (1) conflating firm changes with geographic changes, and (2) underestimating the amount of within-geography variation that may be occurring just due to firms. As we can see, certain cities have different average firm effects. In Texas, Houston and Dallas have higher average firm effects than San Antonio. At the

Figure 11: Maps of Average Firm Fixed Effect: Midwest



Note: Represents average firm fixed effect for all births in our final sample at the three digit zipcode. Restricted to zipcodes with more than 20 births. Average fixed effects values winsorized so the minimum value is -0.5 of the standardized value of the firm fixed effects, and 0.5 is the maximum value.

same time Chicago, Louisville, Columbus, and Cincinatti all have very low average firm effects. This does somewhat align with geographic differences - Cook County, IL has a lower average C-section rate than both Houston and Dallas for instance.

However, within each county, and even within zipcodes, there is substantial variation - the standard deviation of the firm effect within Cook and Harris Counties is approximately half that of the overall standard deviation, meaning that there is approximately half as much variation in Cook County, IL and Harris County, TX as across 14 states. Nevertheless, this mapping suggests that at least some of the geographic differences in healthcare utilization rates may be somewhat driven by differences across the average firm.

Overall, our results demonstrate that firm effects (and facility effects as well) seem to be an important determination of C-section rates. This fact suggests that analyses of geography are to some degree estimating changes across firms, but also potentially masking important within-geography variation. That does not suggest that geography is not important, but rather that what we think of as geography is partially a firm-related object. For example, we do not know *why* certain areas have higher average firm fixed effects. Understanding these objects and what they represent is an important avenue of future research. It is also important to note that a high average firm fixed effect is not necessarily perfectly correlated with a high C-section rate. Patient characteristics still drive an important share of variation in this model.

Our results have important policy implications. If variation in healthcare utilization is partially a firm object, that has different policy implications than a purely geography based variation in healthcare utilization. In-particular, the results demonstrate that doctors' behavior changes substantially when they change firms, which suggests that rules targeting practices can be effective in limiting unnecessary C-sections. These policies could benefit patients, both because a C-section is a major abdominal surgery with a long recovery time, but also because they are substantially more expensive to health plans. Our results also

suggest that it may be important to target policies depending on the distribution of firms, doctors, and facilities - if an area has a very high firm effect that may require a different policy response than an area which has a very high facility fixed effect.

9 Conclusion

In this paper, we seek to better understand the role of the firm in healthcare provision. Utilizing a unique dataset of all private healthcare claims from 14 states, we examine the role of the firm in whether or not patients have a vaginal birth or a C-section. Understanding the variation in C-section rates is important on its own terms - C-sections are extremely common and a major abdominal surgery.

We estimate a model which can separately identify the role of the firm in the provision of healthcare. We find that firms play an exceedingly important role - while doctors explain roughly 10% of overall variation in C-section rates, firms explain nearly 8%, and hospitals explain just under 5%. This suggests that the firm plays a very important role in the overall provision of healthcare.

We next examine how the firm might impact utilization, and find strong evidence that doctors' practice styles change when they switch firms. Using both descriptive analyses and an identification strategy based on practice mergers, we find that changing firms substantially alters how doctors perform C-Sections. When doctors change firms, their behavior changes as well. Using both a mover analysis and an analysis leveraging the quasi-experimental variation induced by mergers, we find that moving from the 25th to 50th percentile of the firm fixed-effect distribution increases C-section rates by approximately 5 percentage points, roughly equivalent to the difference in average C-section rates between Chicago and Dallas.

Lastly, we demonstrate the overlap between analyses of geography and firms. In particular, we find that certain localities have higher average firm effects, and those seem to

correspond to overall geographic variation in C-section rates. This suggests that the role of the firm is an important part of previously-measured differences in cross-geography utilization of healthcare. Nevertheless, we also show that within a given geography there is substantial variation in firm effects.

These findings demonstrate that, at least within the context of labor and delivery, that the firm is a very important component of which medical procedures patients are given. This finding can help inform future policy, particularly through changes to firm effects. Policies which affect firms naturally have a wider impact, and doctors appear to respond to firm-level rules and practice styles. In particular, if policymakers seek to reduce unnecessary C-sections, it is important to understand which rules and practices are most responsible for unnecessary C-Sections and target any changes at the firm level.

References

Jason Abaluck, Leila Agha, Chris Kabrhel, Ali Raja, and Arjun Venkatesh. The determinants of productivity in medical testing: Intensity and allocation of care. *American Economic Review*, 106(12):3730–3764, 2016.

John M. Abowd, Francis Kramarz, and David N. Margolis. High wage workers and high wage firms. *Econometrica*, 67(2):251–333, 1999.

Alexander Ahammer and Thomas Schober. Exploring variations in health-care expenditures–what is the role of practice styles. *Health Economics*, 29(6), 2023.

Ivan Badinski, Amy Finkelstein, Matthew Gentzkow, and Peter Hull. Geographic variation in healthcare utilization: The role of physicians. Technical report, National Bureau of Economic Research, 2023.

Katherine Baicker, Kasey S Buckles, and Amitabh Chandra. Geographic variation in the appropriate use of cesarean delivery: do higher usage rates reflect medically inappropriate use of this procedure? *Health Affairs*, 25(Suppl1):W355–W367, 2006.

Emily Battaglia. The effect of hospital maternity ward closures on maternal and infant health. 2023.

Nancy D Beaulieu, 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*, 382(1):51–59, 2020.

David Card, Jorg Heining, and Patrick Kline. Workplace heterogeneity and the rise of west german wage inequality. *The Quarterly Journal of Economics*, 128(3):967–1015, 2013.

David Card, Ana Rute Cardoso, Joerg Heining, and Patrick Kline. Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics*, 36(1), 01 2018.

Doruk Cengiz, Arindrajit Dube, Attila Lindner, and Ben Zipperer. The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics*, 134(3):1405–1454, 2019.

Amitabh Chandra and Douglas O Staiger. Identifying sources of inefficiency in healthcare. *The Quarterly Journal of Economics*, 135(2):785–843, 2020.

Michael Chernew, Zack Cooper, Eugene Larsen Hallock, and Fiona Scott Morton. Physician agency, consumerism, and the consumption of lower-limb mri scans. *Journal of Health Economics*, 76, 2021.

Eric Chyn and Na'ama Shenhav. Place effects and geographic inequality in health at birth. Technical report, National Bureau of Economic Research, 2022.

Janet Currie and W Bentley MacLeod. Diagnosing expertise: Human capital, decision making, and performance among physicians. *Journal of labor economics*, 35(1):1–43, 2017.

David Cutler, Jonathan S Skinner, Ariel Dora Stern, and David Wennberg. Physician beliefs and patient preferences: a new look at regional variation in health care spending. *American Economic Journal: Economic Policy*, 11(1):192–221, 2019.

Joseph J. Doyle and Becky Staiger. Physician group influences on treatment intensity and health: Evidence from physician switchers. Technical report, National Bureau of Economic Research, 2024.

Christine Piette Durrance and Scott Hankins. Medical malpractice liability exposure and ob/gyn physician delivery decisions. *Health Services Research*, 53(4):2633–2650, 2018.

Paul J. Eliason, Benjamin Heebsh, Ryan C. McDevitt, and James W. Roberts. How acquisitions affect firm behavior and performance: Evidence from the dialysis industry. *Quarterly Journal of Economics*, 135(1):221–267, 2020.

Andrew J Epstein and Sean Nicholson. The formation and evolution of physician treatment styles: an application to cesarean sections. *Journal of health economics*, 28(6):1126–1140, 2009.

Stefanie J Fischer, Heather Royer, and Corey D White. Health care centralization: The health impacts of obstetric unit closures in the us. Technical report, National Bureau of Economic Research, 2022.

Deborah Goldschmidt and Johannes F. Schmieder. The Rise of Domestic Outsourcing and the Evolution of the German Wage Structure*. *The Quarterly Journal of Economics*, 132 (3):1165–1217, 04 2017.

J Gruber and M Owings. Physician financial incentives and cesarean section delivery. *The Rand journal of economics*, 27(1):99–123, 1996.

Jon Gruber, John Kim, and Dina Mayzlin. Physician fees and procedure intensity: the case of cesarean delivery. *Journal of health economics*, 18(4):473–490, 1999.

Thomas G Koch, Brett W Wendling, and Nathan E Wilson. The effects of physician and hospital integration on medicare beneficiaries' health outcomes. *Review of Economics and Statistics*, 103(4):725–739, 2021.

Katy Backes Kozhimannil, Michael R Law, and Beth A Virnig. Cesarean delivery rates vary tenfold among us hospitals; reducing variation may address quality and cost issues. *Health Affairs*, 32(3):527–535, 2013.

Ambar La Forgia. The impact of management on clinical performance: evidence from physician practice management companies. *Management Science*, 69(8):4646–4667, 2023.

Kelli Marquardt. Mis(sed) diagnosis: Physician decision making and adhd. Technical report, R&R Journal of Political Economy, 2023.

Sarah Robinson, Heather Royer, and David Silver. Geographic variation in cesarean sections in the united states: Trends, correlates, and other interesting facts. Technical report, National Bureau of Economic Research, 2023.

Yashaswini Singh, Zirui Song, Daniel Polsky, Joseph D. Bruch, and Jane M. Zhu. Association of private equity acquisition of physician practices with changes in health care spending and utilization. *JAMA Health Forum*, 3(9), 2022.

Jonathan Skinner. Causes and consequences of regional variations in health care. In *Handbook of health economics*, volume 2, pages 45–93. Elsevier, 2011.

Jae Song, David J Price, Fatih Guvenen, Nicholas Bloom, and Till von Wachter. Firming Up Inequality*. *The Quarterly Journal of Economics*, 134(1):1–50, 10 2018. ISSN 0033-5533.

Elise G Valdes. Examining cesarean delivery rates by race: a population-based analysis using the robson ten-group classification system. *Journal of Racial and Ethnic Health Disparities*, 8(4):844–851, 2021.

Yanchao Yang and Maggie Mullen. Black-white disparities of cesarean section rate in the us: Trends and determinants. *Available at SSRN 4204876*, 2022.

A Data Appendix

A.1 Professional Claims

When a health care claim within a hospital is created, patients are separately billed for the services provided by the doctor, and for services provided by the hospital. Professional claims represent all claims provided by doctors, including obstetricians and anesthesiologists. These data include prices (allowed amounts) and CPT codes outlining the services rendered. The CPT codes provide the primary information about whether a birth was vaginal or a C-section. We denote a vaginal birth as CPT codes 59400, 59409, 59610, 59612, and 59614. A C-section is any claim with the codes 59510, 59514, 59515, 59618, 59620, and 59622. The professional claims data also include identifiers for doctors (National Provider Identifier - NPI) and identifiers for firms (Tax Identifier Number - TIN). Lastly, they include characteristics on patient age, diagnoses, and information on where the patient lives.[21]

Based on this, we restrict to all professional claims with a birth-related CPT in our data. We then remove any claims which are not missing a doctor identifier, patient identifier, or firm identifier. Because a given claim can have multiple lines and edits, we utilize the sum of all allowed charges for a given patient-doctor-firm in a given month. We remove any patients whose reported age is below 16 or above 45 years old, to avoid data errors and instances in which C-sections are almost always likely to happen due to patient age.

It is worth noting that overall number of births may seem somewhat low relative to overall birth rates. This is due to three factors: First, our data is constructed from only 14 states, including several smaller ones. Several large states, including New York, California, and Pennsylvania are not included. Second, we are only using larger payers. Medicaid is

[21] Unfortunately, a limitation of these data is that they do not include information on income or race. Valdes [2021] and Yang and Mullen [2022] illustrate the importance of race in determining whether a C-section is performed. We believe we account for this by including various levels of controls for geography and diagnosis, which should account for most of this variation. Nevertheless, this is a limitation of the study.

responsible for approximately 42% of births in the U.S.[22], while there are local insurers and other plans not in our dataset.

A.2 Linking Facility and Professional Claims

Because facility and professional claims are not inherently linked in claims data, it can be difficult to identify which facility a given procedure was performed at. However, this study is benefited from two underlying facts. First, the vast majority of all births are done on an inpatient basis with a single facility claim in a year. Second, people cannot give birth more than once a month. Because facility claims and professional claims both contain an identifier for the patient, we can link claims based on inpatient stays which occur in a given month. Generally, the admission date at a facility is the same as the birth, which leads to easy linkage by patient identifier and admission date. In instances where there are multiple facility claims that have different facility identifiers[23] , the birth was not coded as an inpatient claim, or there is an error with the dates, we utilize the following algorithm to link the claims.[24]

First, we load in all inpatient facility claims and remove duplicates of the same member, month, Facility TIN, and allowed amount, in order to remove duplicated data.

Next, we keep inpatient facility claims where the admission date was different from the discharge date, and had the highest allowed amount in a given month for a given member. We merge this dataset into our births dataset and identify the "unmatched" patients who are missing a facility claim.

For unmatched patients, we then identify any facility for which they were in an inpatient

[22]See https://www.medicaid.gov/medicaid/quality-of-care/quality-improvement-initiatives/maternal-infant-health-care-quality/index.html.

[23]e.g. if a woman went to an urgent care center in the weeks before giving birth

[24]Additionally, there may be some claims due to out-of-state care, i.e. when a patient with a Texas-based plan gives birth in New York state. That claim will be in our data, but we will fail to meaningfully identify enough births to credibly estimate any doctor or firm fixed effects.

claim and the admission date was the same as the discharge date in the same month as their birth. Lastly, for any instances in which the patient is still unmatched, we merge in any remaining facility claim in the same month.

This may still lead to duplicates if there are multiple facility claims, so we remove duplicates wherein there are multiple facility claims in the same month as the birth and one of the claims has a birth-related diagnosis and the other claims do not.

Similar to Currie and MacLeod [2017], we further restrict to remove instances where a firm or facility didn't perform a large enough number of C-sections or births generally. Specifically, we remove any firm whose highest number of C-sections or vaginal births in a given single quarter was less than 25.

We also remove any firm who performed less than 10 births or C-sections in a given quarter, and any doctor whose highest number of births participated in was less than five. This restriction is due to there being a large "left tail" of doctors who participate in very few births. These doctors are likely surgeons, medical students, or other providers who are involved in the birth but not the "main provider". Similarly, if a firm or facility does not participate in many births it may be due to it being emergency care, in which case the normal decision-making process is different in some way.

The individual-quarter restrictions are important to allow for stable estimates of prices. Gruber and Owings [1996], Gruber et al. [1999] find evidence that financial incentives are an important determinant of doctors' and patients' decision-making process. Because the observed claim does not include the counterfactual that the patient would have been charged, we calculate the relative price of a C-Section vs. a vaginal birth based on the total average amount charged for all patients. In instances where a firm participates in only vaginal births or only C-sections (or very few of one or the other) the prices will likely be misspecified.

This allows us to estimate appropriate provider-level and firm-level C-section rates. For example, if a doctor only performs one birth per quarter, and alternates between a C-Section

and a vaginal birth, that would show up as wild swings in C-Section rate, which caused by low volume of patients.

After these restrictions, we retain a total of 2,103,184 births. The drop is almost entirely due to removing firms which perform a very low number of C-sections.

It is not possible to separately identify Hospitals and TINs in some health systems. For example, if a health system uses a single TIN for all their hospitals, it is not possible to identify which hospital was used. Similarly, if a hospital changes TINs it will show up as a new identifier. This would not limit our study, but change the interpretation of the role of facilities into, instead, the role of health systems and ownership.

A.3 Identifying Mergers

We identify mergers in the following way. First, we limit to TINs and NPIs in the 14 states in our dataset. Next we remove small TINs by identifying the average number of doctors in the firm in a given quarter, and removed any TINs with fewer than two doctors.

Next, we generate a list of "possibly acquired" TINs and the associated NPIs. For every TIN we identify that last year and quarter for that TIN. Denote k as the last year and quarter for a given TIN, and let t be calendar time. If a TIN left the sample prior to Q4 2020, they were marked as "acquired". For every "acquired" firm, we then generate the list of NPIs in that firm in the last year and quarter $t = k$ and the prior quarter $t = k - 1$. We combine all NPIs in $t = k$ and $t = k - 1$.

Lastly, we identify acquisitions and acquiring groups based on the flow of doctors from the "possibly acquired" TIN to the "acquiring TIN". We identify an acquisition as any instance in which 80% of the doctors in the "possibly acquired" firm move to the "acquiring firm" by $t = k + 1$. Notably, this likely undercounts the overall number of mergers due to the high thresholds. There may be mergers where only 60% of the physicians change firms.

Some doctors leave the panel and do not flow to any firm, likely representing retirements or a move to a location not in our data. Lastly, we check to see if the acquiring firm existed prior to $t = k + 1$. If the acquiring firm did not exist, we remove those instances from our analysis.

We then identify the "treated" doctors as those whose "main job" based on number of births was in the acquired firm, and who moved to the acquiring firm at any point in k and the three quarters before. We use this procedure due to variation in movements and timing - i.e. if if a merger occurs in the middle of the quarter, we may observe some doctors moving at different times based on changing case volume. We defined the treatment quarter for a given doctor k_d as the quarter that doctor moved from the acquiring firm to the acquired firm.

B Heterogeneity Analysis: Mergers

In this section we provide an analysis of heterogeneity in the effect of mergers on doctor behavior based on observable patient outcomes. Overall, we find little evidence of heterogeneity - regardless of patient age or predicted probability of a C-section, the effect is remarkably similar. The fact that there is an increase in C-section rate is further evidence of the role of firm behavior - if doctors were becoming better diagnosticians due to access to the new firm's better equipment, we would expect different results. Additionally, the fact that there is no heterogeneity even among very low predicted C-section probability is evidence that the "high fixed-effect" firms are performing more unnecessary C-sections [Currie and MacLeod, 2017].

The results of this analysis are shown in Table 5.

Table 5: Heterogeneity Analysis

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Under 35	Over 35	$\hat{\rho}_i$ <.2	.2<$\hat{\rho}_i$ <.3	.3< $\hat{\rho}_i$ <.4	Neg. TE	Pos. TE
Post·TE	0.053**	0.047**	0.076**	0.066**	0.072**	0.052**	0.051**
	(0.008)	(0.012)	(0.012)	(0.007)	(0.020)	(0.012)	(0.008)
N	8590530	2342216	683717	6849093	1315495	10909310	10902037

Standard errors in parentheses
+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$