Labor and Product Market Power, Quality of Care, and the Consolidation of the US Hospital Industry*

Bradley Setzler†

June 25, 2025

Abstract

This paper analyzes, both theoretically and empirically, how hospitals leverage their simultaneous gains in market share over patients and workers through consolidation. I develop a framework for merger evaluation in which firms exploit both oligopoly power in the product market and oligopsony power in the labor market. The model allows firms to endogenously choose product quality, and shows that optimal quality provision depends on labor market power. From the model, I predict the effects of hospital mergers on patient and worker outcomes. Empirically, I combine several administrative panel data sources on the universe of hospitals over 1996-2022. Utilizing a propensity-score matched difference-in-differences design, I estimate the extent to which oligopoly and oligopsony power are exploited after merger events. Confirming the predictions of the model, merging hospitals not only raise prices and lower wages after consolidating, but also reduce the number of patients treated, the number of workers employed, and the quality of care provided to each patient. Patients face worse staffing ratios, report lower satisfaction with their care, and experience greater mortality. Merger spillovers lead competitors in the local market to cut wages even as they increase employment and increase patient volume. On aggregate, mergers increase the number of patients who fail to receive treatment and the number of workers who fail to attain employment. Finally, I use the estimated merger effects to recover the structural parameters, then use the estimated model to measure market power.

*Preliminary draft. Please do not distribute beyond the intended recipients. For valuable comments and suggestions, I thank seminar participants at ASSA, Brandeis, Leuven, MIT, Princeton, and Stanford. I thank Eric San Miguel for valuable contributions during an early stage of this project.

†Pennsylvania State University and NBER (bradley.setzler@gmail.com).

1 Introduction

The healthcare sector accounts for 17% of US GDP, which is a greater share than any other OECD country (OECD, 2023). The healthcare sector also employs more workers than any other sector in the US (Gottlieb, Mahoney, Rinz, and Udalova, 2025). Within this sector, hospitals are responsible for nearly all inpatient care, playing a fundamental role in Americans' lifetime welfare and mortality.[1] This market has been concentrated by a massive consolidation wave in the 21st century, such that today the median commuting zone has only three hospitals and only two distinct hospital systems.[2] This recent wave of mergers raises the concern that hospitals have important price-setting power. The concern is heightened by the fact that inpatient care is inherently urgent and geographically constrained, implying that patients have limited flexibility to shop for lower prices. Recent studies by Dafny, Ho, and Lee (2019), Cooper, Craig, Gaynor, and Van Reenen (2019), and Brand, Garmon, and Rosenbaum (2023) indeed document significant price increases following hospital mergers, suggesting that hospitals exercise local oligopolies over millions of vulnerable patients.

While increasing prices following hospital mergers have received attention, hospital concentration raises equally concerning issues in the labor market. Healthcare occupations typically require specialized training and certifications that have limited value outside of the hospital setting—both for patient care workers, such as nurses, and non-patient care workers, such as hospital administrators—implying inelastic local labor supply.[3] Mirroring patient markets, mergers have substantially lowered the number of distinct hospital systems competing for workers in local labor markets. A natural concern is that hospitals exploit these gains in labor market share to suppress wages and employment levels below competitive benchmarks. Furthermore, the staffing of workers per patient is a key determinant of the quality of care received by patients. If hospitals exploit labor market power by cutting employment, patients may be harmed not only by rising prices but also by reduced quality of care, leading to worsening medical outcomes.

This paper analyzes, both theoretically and empirically, how hospitals leverage their simultaneous gains in market share over patients and workers through consolidation. I

[1] Inpatient treatments are those that require patients to be hospitalized at least one night. Conditions that almost always require inpatient treatment include heart attack, stroke, serious infection, severe burn, major fracture, and conditions requiring a major surgery or transplant.

[2] I document that, out of fewer than 4,000 general-purpose hospitals in the US, there were more than 1,100 hospital consolidations between 2000 and 2018, nearly half of which were mergers among hospitals in the same local market. About 230 of those mergers met US DOJ and FTC (2023) criteria for "presumed anti-competitive" activity, and almost half of patients and hospital employees are located in hospital markets in which presumed anti-competitive mergers occurred.

[3] For example, using resume data, Schubert, Stansbury, and Taska (2024) find that workers in healthcare occupations are the least likley to move to jobs in other occupations.

develop a framework for merger evaluation in which firms exploit both oligopoly power in the product market and oligopsony power in the labor market. The model allows firms to endogenously choose product quality, and shows that optimal quality provision depends on labor market power. Empirically, I combine several administrative panel data sources on the universe of hospitals over 1996-2022. Utilizing a propensity-score matched difference-in-differences design, I estimate the extent to which oligopoly and oligopsony power are exploited after merger events.

The structural framework for merger evaluation has three key components. First, it takes as a starting point the modern framework for the evaluation of horizontal mergers for product market concentration in industrial organization.[4] Firms face a downward-sloping product demand curve which becomes more inelastic as the firm gains product market share, and firms strengthen markups and reduce output after horizontal mergers. Second, I add a rich model of the local labor market in which workers have horizontally-differentiated and vertically-differentiated preferences over employers.[5] Firms face an upward-sloping labor supply curve, implying that the marginal cost is increasing. Labor supply becomes more inelastic as the firm gains labor market share, and firms strengthen markdowns and reduce employment after horizontal mergers. Third, patient care is heterogeneous both in price and quality, and—through staffing per patient—workers are utilized to provide quality. Quality provision inherits an upward-sloping marginal cost curve from the labor market, so horizontal mergers shift the marginal cost curve for quality by concentrating labor.

From the equilibrium model, I derive several propositions that characterize how the outcomes of patients and workers are impacted by hospital mergers. First, I show that hospital mergers impact the merging hospitals through two channels: *labor market diversion*, in which

[4]For leading examples, see Houde (2012), Björnerstedt and Verboven (2016), and Miller and Weinberg (2017). This literature assumes the marginal cost is constant and product quality is predetermined. I introduce variable marginal cost curves and endogenous quality provision, both of which respond to mergers. For an alternative approach to study mergers in the dialysis industry, see Wollmann (2024), who permits endogenous product quality with increasing marginal costs, but assumes there is no price competition and the marginal cost curve does not shift in response to mergers. Related, Dingel, Gottlieb, Lozinski, and Mourot (2024) model hospital spatial competition on quality for Medicare patients, without price competition.

[5]For models of the labor market in which workers view employers as horizontally differentiated, see Card, Cardoso, Heining, and Kline (2018) and Lamadon, Mogstad, and Setzler (2022) for the monopsonistic competition case in which employers cannot exploit labor market concentration, and see Berger, Herkenhoff, and Mongey (2022), Roussille and Scuderi (2025), and Chan, Kroft, Mattana, and Mourifié (2024) for the oligopsonistic case in which employers may exploit labor market concentration. These papers do not combine product market power and labor market power and do not consider mergers, which are among my points of departure. Recent work by Kroft, Luo, Mogstad, and Setzler (2025) combines a monopsonistic competition labor market and monopolistic competition product market, but does not permit a role for concentration and therefore cannot be used for merger evaluation. Related, Yeh, Macaluso, and Hershbein (2022) and Rubens (2023) provide results on the measurement of wage markdowns under minimal assumptions about the labor and product markets, but do not fully specify economic models that could be used for ex ante predictions of the effects of mergers.

hospitals internalize the costs imposed on other hospitals in the same ownership system in the competition over local workers, and *product market diversion*, in which hospitals internalize the revenues lost by other hospitals in the same ownership system in the competition over local patients. Due to these two diversion effects, mergers lead hospitals to raise prices and cut wages, while reducing the number of patients treated and the number of workers employed. These two effects reinforce one another: the labor diversion effect amplifies the increase in prices and reduction in number of patients treated, while the product diversion effect amplifies the reductions in wages and employment.

Second, I show that hospital mergers have *spillovers* on other hospitals in the local market. From the perspective of local competitors, both the product demand curve and the labor supply curve shift outward, allowing competitors to raise prices and cut wages for any level of output. Local competitors re-optimize to treat more patients and employ more workers. However, since marginal costs fall and marginal revenues rise, it is ambiguous whether prices or wages rise or fall at local competitors. At the same time, the outside option shares—the share of patients who fail to receive treatment and workers who fail to attain employment—unambiguously rise.

Third, I show that a merger-induced increase in the *labor* market share impacts *patients* by incentivizing the hospital to reduce quality of care, providing another channel through which labor and product markets interact. As the hospital gains labor market power over workers, the labor diversion effect implies that it is more costly to provide high staffing for any given number of patients. However, the merger also reduces the number of patients treated, which may counteract the decline in labor in staffing by relieving patient care congestion. I provide a congestion threshold under which quality of care unambiguously falls after the merger, but above which quality of care may rise after the merger. An interesting implication is that the mergers that tend to improve quality are the ones involving hospitals operating inefficiently by prioritizing patient volume over quality of care.

Empirically, I investigate the impacts of hospital mergers using a comprehensive panel dataset on the universe of general-purpose hospitals in the US from 1996 to 2022. This dataset combines administrative records on hospital revenues and costs by broad patient category, employment hours and wages by broad occupational category, the ownership structure and consolidation events, patient satisfaction surveys, patient case mix index, risk-adjusted mortality rates, and geographic identifiers. Using this information, I construct market concentration for each hospital labor and product market as well as the changes in market concentration induced by mergers. For the purposes of the empirical analysis, I classify hospital mergers as "presumed anti-competitive" if they meet the current thresholds utilized by US courts (US DOJ and FTC, 2023). Roughly half of within-market hospital mergers

exceed the presumed anti-competitive thresholds, both for labor and product markets.

Utilizing a difference-in-differences design around hospital mergers matched on propensity score, I confirm—and put magnitudes on—the predictions of the model among presumed anti-competitive mergers. First, within the merging hospitals, I find that prices increase by approximately 8% while the number of patients treated decreases by 5%.[6] In the labor market, wages decrease by more than 1% while employment levels fall by 10%. I find strongly negative wage and employment effects for both patient care occupations—nurses, nursing aids, and hospitalists—and non-patient care occupations—administration, food services, sanitation, and maintenance.[7]

I also implement a spillover design, in which I compare the local competitors of merging hospitals to otherwise similar hospitals in other regions. I find that local competitors decrease wages by 1% yet increase employment by 4% and increase the number of patients treated by 6%, consistent with the model's predictions.[8] Aggregating across all hospitals in the market, I find that the total number of patients receiving care declines by 2-3% in the short-run and the total number of workers employed in the market also declines by 2-3%. Consistent with the predictions of the model, the number of patients who fail to receive treatment and workers who fail to attain employment in the local market rise.

Lastly, I find a substantial deterioration in the quality of care in the merging hospitals, whether using input-based or outcome-based measures. The staffing ratio declines by 6%, while patient satisfaction ratings decrease by more than 1 percentage points.[9] Decomposing the patient satisfaction ratings, patients note worsening cleanliness and quietness. Most concerning, risk-adjusted mortality rates within 30-days of being treated by the merging

[6]Dafny et al. (2019), Cooper et al. (2019), and Brand et al. (2023) estimate comparable price increases following hospital mergers using several data sources. These papers and the related empirical literature do not report a significant decrease in number of patients treated, which is an important effect for characterizing oligopoly power. A related literature studies the bargaining between insurers and hospitals that can explain large price increases without changes in the volume of patients (Gowrisankaran, Nevo, and Town, 2015; Ho and Lee, 2017). In our empirical model, we incorporate both classical oligopoly power of hospitals over patients and bargaining on prices with insurers, providing a decomposition of the total price increase into the contributions of oligopoly power versus bargaining power.

[7]Prager and Schmitt (2021) find a similar decline in wages after hospital mergers among certain non-patient care occupations—especially nursing administrators and pharmacists, which they refer to as "nurses and pharmacists" for short (see the definition on page 1 of their appendix). However, they do not report estimated wage effects for patient care workers, such as nurses, and do not detect significant effects on hospital employment, which are among my points of departure.

[8]Recent papers by Sharma (2023) for Brazilian textile manufacturing, Roussille and Scuderi (2025) for US tech workers, and Derenoncourt and Weil (2025) for large US retailers test for wage-setting spillovers in local labor markets. They do not find evidence of spillovers in those industries.

[9]Hackmann (2019) studies staffing ratios in the Pennsylvania nursing home industry, finding little responsiveness to changes in competition. Beaulieu, Dafny, Landon, Dalton, Kuye, and McWilliams (2020) find a similar decline in patient satisfaction following hospital mergers, but are unable to detect mortality effects of mergers using their short sample time frame that only comprises about one-fifth of my sample of mergers.

hospitals for heart failure or pneumonia increase by 0.4-0.8 percentage points, relative to national mortality rates of 12-13%. To place this finding in context, Cooper, Doyle, Graves, and Gruber (2022) estimate that a one standard deviation increase in the hospital price distribution is associated with a 0.5 percentage point decrease in the mortality rate, implying that the worsening mortality rates observed after hospital mergers are relatively large.

In the final section, I take the model to the data, developing a method of simulated moments estimator to recover all model parameters by matching the simulated model to the estimated merger effects. The empirical model includes rich additional sources of heterogeneity, including fixed administrative overhead costs that vary by hospital and hospital-specific coinsurance rates that parsimoniously allow for hospitals to leverage mergers to bargain for higher prices from insurers. The model yields several additional insights. First, patient and non-patient care workers are found to be gross complements in the production of hospital quality, and hospital quality exhibits increasing returns to scale, in line with the results using a distinct framework by Dingel et al. (2024). Second, about a quarter of the total observed log price increase is driven by hospitals leveraging increased bargaining power over insurers, with more than half explained by the exercise of classical oligopoly power. Third, across the range from the smallest to the largest hospitals in terms of market share, I find that the hospital-specific patient demand elasticity ranges from 2.7 at the smallest hospitals to 2.0 at the largest hospitals. Hospital-specific labor supply elasticities range from 4.7 in the smallest hospitals to 3.5 in the largest hospitals, with similar values for both patient care and non-patient care occupations. Fourth, estimating markups and markdowns, I find that prices exceed marginal costs by a range of 40-55% of the price, while wages fall short of marginal revenues by 17-23% as a share of marginal revenues.

This paper relates to a growing literature raising concerns about the anti-competitive effects on sellers of mergers among buyers, where employers and workers can be interpreted as buyers and sellers of labor, respectively.[10] Hemphill and Rose (2018) distinguish two types of harm that mergers may have on such sellers: "Some mergers of competing buyers harm sellers by increasing the merged firm's incentive to cut back on its purchases of an important input in order to drive down input prices—a classical exercise of *monopsony power*. [...] Mergers of competing buyers harm sellers alternatively, or in addition, by increasing the new firm's *bargaining leverage*." Recently, Hosken, Larson-Koester, and Taragin (2024) incorporate the latter channel, wage bargaining, into an otherwise standard model of mergers for product market power.[11] However, no prior research has incorporated the former channel, monopsony

[10]See Hemphill and Rose (2018), Naidu, Posner, and Weyl (2018), and Marinescu and Hovenkamp (2019). Buyer market power has recently been incorporated into US merger guidelines (US DOJ and FTC, 2023).

[11]The wage-bargaining model of Hosken et al. (2024) treats labor as a non-differentiated wholesale good sold to firms by an upstream supplier, i.e., a labor union. Such a model may be appropriate in an unionized

power, into an otherwise standard model of mergers for product market power, whether in the hospital context or any other setting.

2 A Model of Mergers for Labor and Product Market Power with Endogenous Quality and Overhead

In this section, I develop a model of mergers in the hospital industry. Hospitals simultaneously exploit product and labor market power by leveraging their market shares over patients and hospital workers, respectively. Hospitals hire for both patient care and non-patient care occupations, and utilize staffing ratios to endogenously provide quality of care to patients. Hospitals face heterogeneous overhead costs, which must be paid by hiring a fixed number of non-productive administrators. Given this environment, I derive comparative statics that characterize the effects of hospital mergers on the outcomes and competitiveness of hospital patient and labor markets.

2.1 Labor Supply, Product Demand, and Technology of Care

Product Demand. Patient i's utility from being treated by hospital h in market m and time t is,

$$u_{iht}^{Q} = -\beta_P P_{ht} + \beta_Y Y_{ht} + \xi_{ht}^{Q} + \varepsilon_{iht}^{Q}. \tag{1}$$

There are four components to patient utility. The first captures the price of treatment, P_{ht}, multiplied by the distaste for the price of treatment, $-\beta_P$.[12] The second captures the observable quality of care provided by the hospital, Y_{jt}, multiplied by the relative preference for quality, β_Y. I assume that Y_{jt} is endogenously chosen by the hospital, subject to constraints described below. The third, ξ_{ht}^{Q}, captures the residual component of hospital quality that is unobserved by the analyst but observed by patients. As is standard, residual quality is predetermined.[13] Combined, $\beta_Y Y_{ht} + \xi_{ht}^{Q}$ characterizes vertical differentiation in quality

labor market. However, only 13% of US hospital workers are unionized (Ahmed, Kadakia, Ahmed, Shultz, and Li, 2022). Hosken et al. (2024) emphasize that the bargaining approach is inconsistent with substantial changes in employment following mergers. I show that the monopsony approach is successful in explaining the large employment reductions that follow hospital mergers. Other related work includes Berger, Hasenzagl, Herkenhoff, Mongey, and Posner (2025), who consider labor market mergers in a model with no role for product markets, and Angerhofer, Collard-Wexler, and Weinberg (2025), who consider monopsonistic wage bargaining with teachers' unions when school districts have quality concerns.

[12] In the empirical model, I also incorporate a tractable role for insurers and distinguish the total price paid by insurers to hospitals from the coinsurance price paid by patients.

[13] I follow the standard terminology in industrial organization: a variable is referred to as "predetermined" at time t if it is taken as given (not a choice variable) when choices are made at time t. Predetermination is

across hospitals. The fourth, ε^Q_{iht}, represents patient i's match-specific taste for hospital h and is assumed to have the standard Gumbel distribution. These patient-hospital-specific tastes are unobserved by the analyst and hospital but observed by the patient, giving rise to horizontal product differentiation.

Aggregating across the distribution of match-specific tastes, the product demand of hospital h can be expressed in terms of the market share as,

$$s^Q_{ht} \equiv \frac{Q_{ht}}{\bar{Q}_{mt}} = \frac{\exp\left(-\beta_P P_{ht} + \beta_Y Y_{ht} + \xi^Q_{ht}\right)}{1 + \sum_{h'} \exp\left(-\beta_P P_{h't} + \beta_Y Y_{h't} + \xi^Q_{h't}\right)}, \tag{2}$$

where $\bar{Q}_{mt}$ denotes the total number of patients in market m at time t. The outside share of the product market, which is the share of patients who fail to attain treatment, is denoted by $s^Q_{0t} \equiv 1 - \sum_h s^Q_{ht}$.[14]

Labor Supply. I consider two broad occupational categories within hospitals: patient care and non-patient care. Patient care workers (nurses, nursing aides, and hospitalists) directly provide treatment to patients, while non-patient staff (administration, sanitation, food services, maintenance) support hospital operations. The number of patient care workers employed by the hospital is denoted by L_{ht} and the number of non-patient care workers employed by the hospital is denoted by N_{ht}. The wage offered to patient care workers is denoted by W^L_{ht} and the wage offered to non-patient care workers is denoted by W^N_{ht}.

In labor market m at time t, the preference of worker ι of occupation type E for working at hospital h is,

$$u^E_{\iota ht} = \gamma_E \log\left(W^E_{ht}\right) + \xi^E_{ht} + \varepsilon^E_{\iota ht}, \quad E = L, N. \tag{3}$$

This specification of labor preferences has three components. First, the term $\log\left(W^E_{ht}\right)$ captures diminishing marginal utility of income, and labor supply of occupation E is more elastic when the coefficient γ_E is greater. Second, hospital-occupation-specific job quality ξ^E_{ht} represents the value of amenities enjoyed by all workers in the hospital and occupation. These amenities are unobserved to the analyst but common knowledge among workers, giving rise to vertical employer differentiation. The value of amenities may differ between patient care and non-patient care occupations (i.e., $\xi^L_{ht} \neq \xi^N_{ht}$). Third, $\varepsilon^E_{\iota ht}$ represents worker ι's idiosyncratic and match-specific taste for the amenities of hospital h and is assumed to

not to be confused with statistical independence: the firm's optimal price and other choices are functions of the predetermined variables and thus depend upon them.

[14] As is standard, the outside option is normalized to have mean utility index of zero without loss of generality, and each patient draws a Gumbel-distributed taste for the outside option. Below, the analogous assumptions characterize the outside option in each labor market.

have the standard Gumbel distribution at time t.[15] These worker-hospital-specific tastes are unobserved to the analyst and hospital but known to the worker, giving rise to horizontal employer differentiation.

Aggregating across the distribution of match-specific tastes, the labor supply to hospital h in occupation E can be expressed in terms of the market share as,

$$s_{ht}^E \equiv \frac{E_{ht}}{\bar{E}_{mt}} = \frac{\exp\left(\gamma_E \log\left(W_{ht}^E\right) + \xi_{ht}^E\right)}{1 + \sum_{h'} \exp\left(\gamma_E \log\left(W_{h't}^E\right) + \xi_{h't}^E\right)}, \quad E = L, N, \tag{4}$$

where $\bar{E}_{mt}$ denotes the total number of workers of type E in market t. The outside share of each labor market, which is the share of workers who fail to attain employment, is denoted by $s_{0t}^E \equiv 1 - \sum_h s_{ht}^E$.[16]

Technology of Care. To treat Q_{ht} patients, the amount of patient care labor required is determined by the patient care technology,

$$Q_{ht} = T_{ht}(L_{ht}), \tag{5}$$

where $T_{ht}(\cdot)$ is the clinical treatment function, which is strictly increasing in the number of patient care workers. The treatment function may vary across hospitals and time, representing that some hospitals may be more efficient than others.[17]

The hospital also combines patient care labor, non-patient care labor, and the number of patients to determine quality,

$$Y_{ht} = \frac{F(L_{ht}, N_{ht} - \psi_{ht})}{Q_{ht}} \quad \text{if} \quad N_{ht} \geq \psi_{ht}, \quad \text{and} \quad Y_{ht} = -\infty \quad \text{if} \quad N_{ht} < \psi_{ht}. \tag{6}$$

We refer to $F(L_{ht}, N_{ht} - \psi_{ht})$ as *effective staffing*, so quality can be interpreted as effective staffing divided among patients. For example, if $F(L_{ht}, N_{ht}) = L_{ht} + N_{ht}$, then Y_{ht} is simply the staffing ratio.[18] The predetermined unobservable ψ_{ht} is the administrative *overhead cost*,

[15] Note that this structure of match-specific tastes in the cross-section does not restrict the correlational structure of match-specific tastes over time within individual (by Sklar's Theorem). That is, denoting $\vec{\varepsilon}_{\iota t}^E \equiv (\varepsilon_{\iota 1t}^E, \varepsilon_{\iota 2t}^E, ..., \varepsilon_{\iota \bar{h}t}^E)$, the model places no restrictions on the copula of $(\vec{\varepsilon}_{\iota t}^E, \vec{\varepsilon}_{\iota t'}^E)$ for $t' \neq t$.

[16] In their model of oligopsonistic labor markets, Berger et al. (2022) consider nested local labor markets with mobility across markets, but with no outside option. Here, we consider local labor markets with an outside option. Rubens, Setzler, and Yeh (2025) show that the outside option can be parameterized in terms of the cross-market mobility elasticity such that markdowns are identical between the two approaches.

[17] In the empirical application below, I consider the constant elasticity treatment function $T_{ht}(L_{ht}) = A_{ht} L_{ht}^{\alpha}$, where A_{ht} measures the relative productivity of h and α is the elasticity of patients to employment.

[18] In the empirical application below, I consider the constant elasticity of substitution functional form, $F(L, N) = (\delta L^{\rho} + (1 - \delta) N^{\rho})^{\phi/\rho}$. This allows that patient and non-patient care labor may be gross complements ($\rho < 0$) or gross substitutes ($\rho > 0$) in the provision of quality, and returns to scale in quality may be

that is, the minimum number of administrators required to keep the hospital operational by maintaining minimal standards ("keeping the lights on"). We normalize quality to be infinitely low if the overhead requirement is not satisfied, which is equivalent to assuming the hospital is shut down by regulators. Only once non-patient care labor exceeds the overhead cost do such workers contribute to effective staffing and thus to quality.

Before proceeding, it is useful to emphasize the differences in functional form between the two care technologies. Non-patient care workers cannot be substituted to provide patient care, such that N_{ht} is excluded from the patient care technology. This reflects that administrators and support staff typically cannot perform clinical duties. By contrast, patient care and non-patient care workers are combined in the provision of quality.

Multi-hospital System's Conduct. Hospital system H owning the set of hospitals $\mathcal{H}_H$ is assumed to maximize profits across its set of establishments in the market:

$$\max_{\{Q_{ht},Y_{ht},L_{ht},N_{ht}\}_{h\in\mathcal{H}_H}} \sum_{h\in\mathcal{H}_H} \left(P_{ht}Q_{ht} - W^L_{ht}L_{ht} - W^N_{ht}N_{ht}\right). \tag{7}$$

subject to product demand, the labor supply of each occupational category, the patient care technology, and the quality of care technology. The hospital system values profits equally across all hospitals, which is equivalent to interpreting a hospital system as a multi-product firm in which each hospital is a product.

Hospitals compete à la differentiated-Cournot in both the product and labor markets.[19] This means that, when the reference firm is choosing its price-quantity and wage-labor pairs, it perceives competitors as holding patient volume and employment fixed, with competitors

increasing ($\phi > 1$) or decreasing ($\phi < 1$). It includes the staffing ratio as a special case.

[19] In a differentiated market, the distinction between the common Nash conducts, Bertrand and Cournot, is subtle. In both cases, the firm can be interpreted as posting a price, from which the residual demand curve determines the quantity—mathematically, there is no difference between choosing price or quantity along an invertible residual demand curve. The salient difference is that they imply distinct residual demand curves, arising because a distinct choice variable is perceived to be held fixed *among competitors*. Under Bertrand, competitors are perceived to hold prices fixed, but quantities may adjust, as the reference firm varies its price-quantity pair. Under Cournot, competitors are perceived to hold quantities fixed, but prices may adjust, as the reference firm varies its price-quantity pair. To be clear, Cournot—not Bertrand—has the often-desirable property that each firm's first-order condition anticipates competitors' price responses. Differentiated-Cournot has been the standard conduct in models of oligopolistic product markets in international trade and macroeconomics since Atkeson and Burstein (2008) and a common specification in oligopsonistic labor markets since Berger et al. (2022). Differentiated-Bertrand is more common in empirical industrial organization, though differentiated-Cournot is drawing attention in the conduct testing literature (see, e.g., Magnolfi, Quint, Sullivan, and Waldfogel 2022). While comparative statics for horizontal mergers in the product market are well-known under differentiated-Bertrand competition and constant marginal costs (e.g., Nocke and Whinston 2022), I am not aware of prior work providing comparative statics for horizontal mergers under differentiated-Cournot competition, either with constant or increasing marginal costs, nor prior work providing comparative statics for mergers under simultaneous price and quality competition.

responding only by changing their prices and wages. Such a conduct seems natural in the labor market context, in which a firm attempting to poach workers likely anticipates that its competitors would raise wages to retain workers (Cournot) rather than keep wages fixed and respond through employment (Bertrand). Rubens et al. (2025) compare the theoretical properties of several conducts and build on Roussille and Scuderi (2025) to test the conduct of more than 100,000 local labor markets in the US, finding that differentiated-Cournot is one of the most common conducts and is the conduct that best fits the hospital labor market.

Welfare Measures. While we do not utilize the model to compute welfare measures in this paper, model-consistent welfare expressions are provided in Appendix A for completeness.

2.2 Comparative Statics for Mergers, given Quality

In this subsection, I characterize how mergers among hospitals in the same market affect the price-setting and wage-setting behavior of both the merging hospitals and their competitors. I do not yet account for endogenous quality provision for two reasons. First, it is rare to account for endogenous quality in merger evaluation—especially outside of the healthcare context—so these results may be applicable in a wider range of empirical settings. Second, comparative statics yield sharper predictions in the model without endogenous quality.

Let $\text{MP}^L_{ht} \equiv \frac{\partial T_{ht}(L_{ht})}{\partial L_{ht}} > 0$ denote the marginal product of patient care labor, $\theta^L_{ht} \equiv \frac{W^L_{ht}}{s^L_{ht}} \frac{\partial s^L_{ht}}{\partial W^L_{ht}} > 0$ denote the (residual) labor supply elasticity of patient care workers, and $\theta^Q_{ht} \equiv \frac{P_{ht}}{s^Q_{ht}} \frac{\partial s^Q_{ht}}{\partial P_{ht}} < 0$ denote the (residual) product demand elasticity of patients. Impose that $\beta_Y = 0$, such that patients do not value the endogenous component of quality, which in turn implies that $N_{ht} = 0$, as non-patient care workers produce no revenue if they cannot affect quality. In this case, the only quality of care component is the pre-determined factor, ξ^Q_{ht}.

The first-order condition for profit-maximization is as follows:

Lemma 1 (First-order condition in a single-establishment firm)**.** *Suppose quality of care is pre-determined. For a single-hospital system h, the first-order condition for profit-maximization is,*

$$\underbrace{\left(1 + 1/\theta^L_{ht}\right) \times W^L_{ht}}_{\equiv MC^L_{ht}} = \underbrace{\left(1 + 1/\theta^Q_{ht}\right) \times P_{ht} MP^L_{ht}}_{\equiv MR^L_{ht}}.$$

Furthermore,

$$\left|\frac{\partial \theta^Q_{ht}}{\partial s^Q_{ht}}\right| < 0 \quad \textit{and} \quad \left|\frac{\partial \theta^L_{ht}}{\partial s^L_{ht}}\right| < 0,$$

indicating that both product demand and labor supply become less elastic as the hospital gains market share in the product and labor market, respectively.

The following proposition follows immediately from Lemma 1:

Proposition 1 (Price-setting and wage-setting in a single-establishment firm)**.** *Suppose quality of care is pre-determined. For a single-hospital system h, (a) the optimal price satisfies,*

$$P_{ht} = \overbrace{\left(1+1/\theta^Q_{ht}\right)^{-1}}^{markup_{ht}} \times \frac{MC^L_{ht}}{MP^L_{ht}} = \underbrace{\overbrace{\left(1+1/\theta^Q_{ht}\right)^{-1}}^{markup_{ht}} \times \overbrace{\left(1+1/\theta^L_{ht}\right)}^{markdown^{-1}_{ht}}}_{double\ markup_{ht}} \times \frac{W^L_{ht}}{MP^L_{ht}},$$

and the optimal wage satisfies,

$$W^L_{ht} = \overbrace{\left(1+1/\theta^L_{ht}\right)^{-1}}^{markdown_{ht}} \times MR^L_{ht} = \underbrace{\overbrace{\left(1+1/\theta^L_{ht}\right)^{-1}}^{markdown_{ht}} \times \overbrace{\left(1+1/\theta^Q_{ht}\right)}^{markup^{-1}_{ht}}}_{double\ markdown_{ht}} \times P_{ht} MP^L_{ht}.$$

(b) Furthermore, the markup and double markup depend on product market share as,

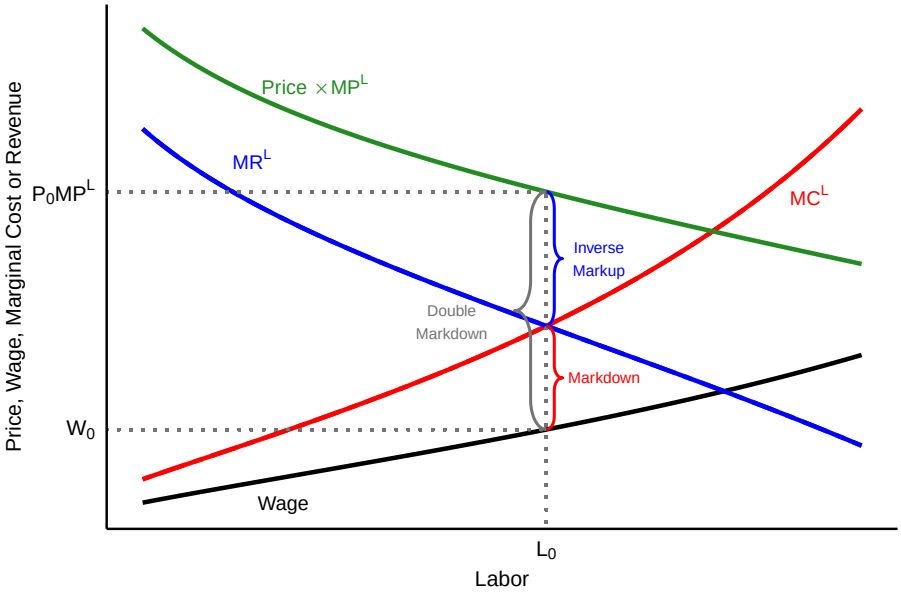
$$\frac{\partial(double\ markup_{ht})}{\partial s^Q_{ht}} > \frac{\partial(markup_{ht})}{\partial s^Q_{ht}} > 0,$$

while the markdown and double markdown depend on labor market share as,

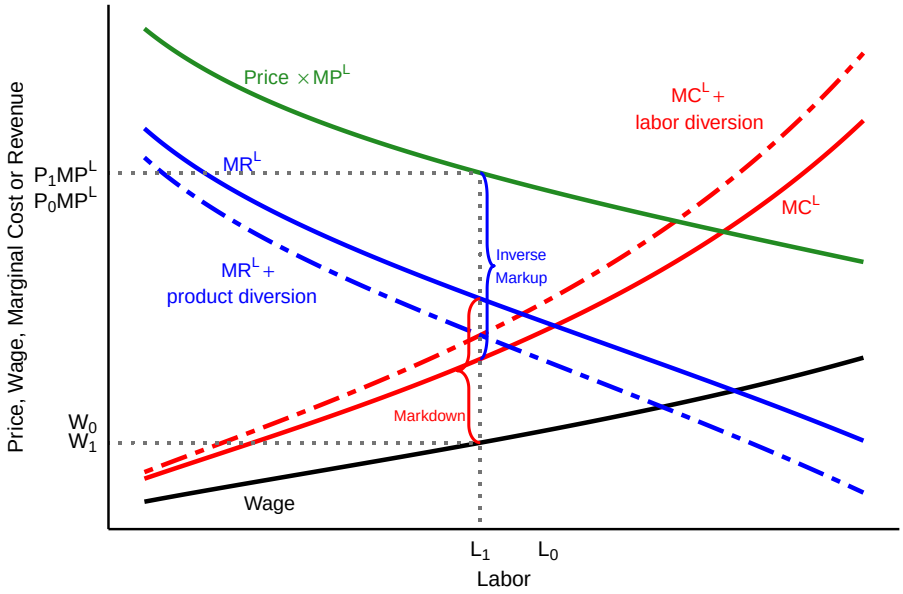
$$\frac{\partial(double\ markdown_{ht})}{\partial s^L_{ht}} < \frac{\partial(markdown_{ht})}{\partial s^L_{ht}} < 0.$$

The first insight from Proposition 1 is that product market power, as measured by the markup relative to the productivity-adjusted wage, W^L/MP^L, is determined not only by the elasticity of product demand, θ^Q, but also the elasticity of labor supply, θ^L. Similarly, not only upward-sloping labor supply but also downward-sloping product demand are relevant and distinct sources of labor market power, as measured by the markdown of wages relative to the value of the marginal product, $P \times \mathrm{MP}^L$. The upper panel of Figure 1 illustrates this proposition. This result is similar to the main insight of Kroft et al. (2025), but extended to allow for heterogeneous markups and markdowns that depend on market concentration.

The second insight from Proposition 1 is that, as the hospital gains market share in the product and labor market, its markup increases and its markdown decreases, indicating two



(a) Before Merger: Visualization of Proposition 1



(b) After Merger: Visualization of Proposition 2

Figure 1: The Merging Firm's Response to a Merger

Notes: This figure demonstrates how the first-order condition for profit-maximization changes from the perspective of hospital h in response to its merger with hospital g. Before the merger, the upper diagram visualizes hospital h's wage-setting and price-setting considerations, comparing its marginal revenue and cost curves. After the merger, the dashed curves in the lower diagram visualize how hospital h internalizes labor diversion from hospital g into its perceived marginal cost curve and product diversion from hospital g into its perceived marginal revenue curve. This figure is simulated from the actual model, setting $T_{ht}(L_{ht}) = A_{ht}L_{ht}$.

sources of greater market power.[20] Since labor and product market shares increase together through the patient care technology in equation (5), it follows that the double markup increases even more than the increase in the markup alone. This implies that, for the same change in hospital size, the effective gain in product market power is greater if the firm also exploits labor market power. Similarly, the double markdown suppresses wages to an even greater extent than the markdown alone if the hospital also exploits product market power.

I now consider how the choices of hospitals change in response to a merger:

Lemma 2 (First-order condition in a two-establishment firm)**.** *Suppose quality of care is pre-determined. If hospitals h and g merge to form a two-hospital system, the first-order condition for hospital h is,*

$$MC^{L}_{ht} + \underbrace{\frac{\partial W^{L}_{gt}}{\partial s^{L}_{ht}} s^{L}_{gt}}_{\text{labor diversion } (+)} = MR^{L}_{ht} + \underbrace{\frac{\partial P_{gt}}{\partial s^{Q}_{ht}} s^{Q}_{gt} MP^{L}_{ht}}_{\text{product diversion } (-)},$$

where MC^{L}_{ht} and MR^{L}_{ht} are defined in Lemma 1.[21]

The key result in Lemma 2 is that, from the perspective of hospital h, there is an increase in the marginal cost curve as well as a decrease in the marginal revenue curve when it merges with hospital g. I now provide intuition for these two channels.

The first channel, *labor market diversion*, effectively increases the marginal cost of providing patient care. It expresses that, for hospital h to increase employment, it must offer a higher wage. When hospital h offers a higher wage, some workers who would have chosen employment at hospital g are diverted to prefer hospital h. To maintain the same number of employees, hospital g must counter by increasing its wage for a given number of workers, which reduces the profits of hospital g. Since h values profits equally across h and g, it internalizes these lost profits at g as an increase in the marginal cost of hiring workers at h. The magnitude of the labor diversion effect is greater when s^{L}_{gt} is greater, as it becomes more likely that the workers hired by h were diverted from g rather than from other competitors.

[20] Given our parameterization, the markdown simplifies to $\left(1 + 1/\theta^{L}_{ht}\right)^{-1} = \frac{\gamma_L}{\gamma_L + 1 + s^{L}_{ht}/s^{L}_{0t}}$ and the markup simplifies to $\left(1 + 1/\theta^{Q}_{ht}\right)^{-1} = \frac{\beta_P P_{ht}}{\beta_P P_{ht} - 1 - s^{Q}_{ht}/s^{Q}_{0t}}$. The markup is greater for larger hospitals because they operate at an inelastic portion of the product demand curve, in accordance with Marshall's second law of demand. That labor supply also becomes more inelastic as hospital size increases can be viewed as the labor supply analogue to Marshall's second law. See the discussion by Autor, Dorn, Katz, Patterson, and van Reenen (2020).

[21] Given our parameterization, the product diversion term simplifies to $\frac{-s^{Q}_{gt}}{\beta_P s^{Q}_{0t}} \mathrm{MP}^{L}_{ht}$ and the labor diversion term simplifies to $\frac{W^{L}_{gt} s^{L}_{gt}}{\gamma_L s^{L}_{0t}}$.

The second channel, *product market diversion*, effectively decreases the marginal revenue of providing patient care. It expresses that, for hospital h to increase patients, it must offer a lower price. When hospital h offers a lower price, some patients who would have chosen treatment at hospital g are diverted to prefer hospital h. To maintain the same number of patients, hospital g must counter by lowering its price for a given number of patients, which reduces the profits of hospital g. Since h values profits equally across h and g, it internalizes these lost profits at g as a decrease in the marginal revenue from treating patients at h. The magnitude of the product diversion effect is greater when s_{gt}^{Q} is greater, as patients attracted to h are more likely to be diverted from g rather than from other competitors.

Combined, the product and labor market diversion effects imply the following results:

Proposition 2 (Responses to a merger in the merging firms)**.** *Suppose quality is predetermined. If hospitals h and g in market m merge at time t to form a two-hospital system H, the optimal product market choices of system H satisfy the following:*

(a) The price increases in both hospitals h and g.

(b) The wage decreases in both hospitals h and g.

(c) The number of patients treated decreases in system H.

(d) The number of workers employed decreases in system H.

(e) The price markup strengthens in hospital h.

(f) The wage markdown strengthens in hospital h.

(g) As the merging hospitals have greater market share, the magnitudes of the effects increase in parts (a-f).[22]

Proof. See Appendix A. □

The lower panel of Figure 1 provides the intuition for parts (a-f) of Proposition 2. In response to a merger, both labor diversion and product diversion effects shift the perceived marginal cost and perceived marginal revenue curves, resulting in lower employment and fewer patients. Furthermore, since the product demand curve is downward-sloping and the labor supply curve is upward-sloping, a decrease in number of patients and employment imply a higher price and a lower wage. With a lower wage and higher marginal revenue, the markdown decreases, and with a higher price and lower marginal cost, the markup increases.

Lastly, I characterize how spillovers from the merger affect other local competitors:

[22] Because $T_{ht}(\cdot)$ is monotonic, greater labor market share and greater product market share coincide. In the data, the within-market correlation between product and labor market shares is very high.

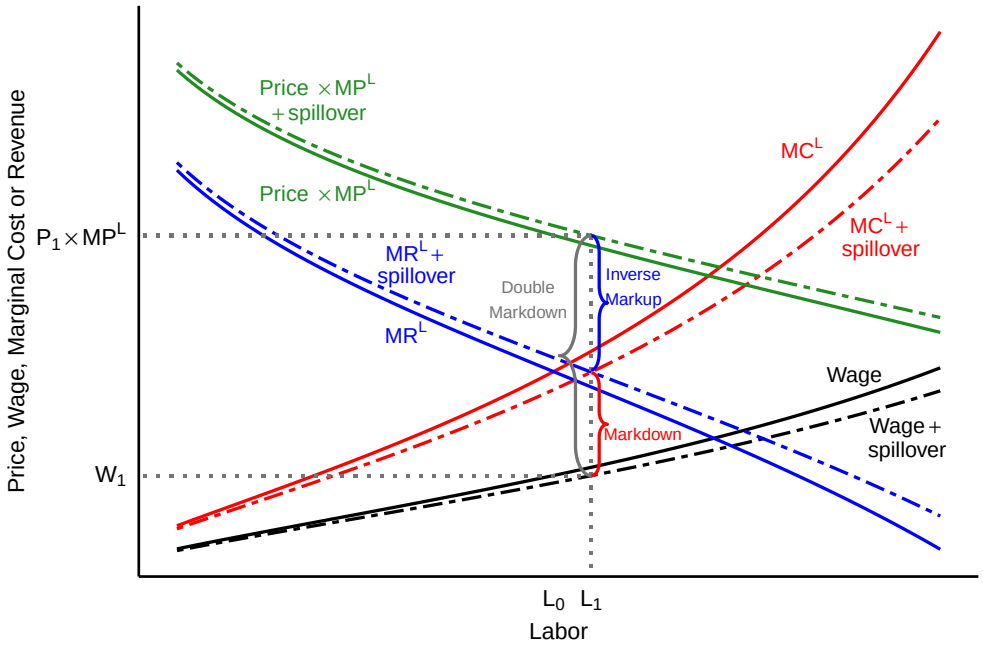


Figure 2: Competitor's Response to a Merger: Visualization of Proposition 3

Notes: This figure demonstrates how the first-order condition for profit-maximization changes from the perspective of local competitor j in response to a merger between two other hospitals. The solid curves represent the firm's considerations before the merger (similar to the upper panel of Figure 1), while the dashed curves incorporate spillovers from the merger. This figure is simulated from the actual model, setting $T_{ht}(L_{ht}) = A_{ht}L_{ht}$.

Proposition 3 (Aggregate and cross-firm spillovers in response to a merger)**.** *Suppose quality is pre-determined. If hospitals h and g in market m merge at time t to form a two-hospital system H, market m is affected as follows:*

(a) Each non-merging competitor j increases the number of patients treated.

(b) Each non-merging competitor j increases employment.

(c) The outside share of the product market increases.

(d) The outside share of the labor market increases.

(e) As the merging hospitals have greater market share, the magnitudes of the effects increase in parts (a-d).

Proof. See Appendix A. □

The intuition for Propositions 3(a-b) is provided by Figure 2, which characterizes the trade-offs faced by local competitors to the merging hospitals. Since the merging hospitals

optimally reduce the number of patients they treat and workers they employ, other hospitals in the same market experience an increase in patient demand—*product market spillovers*—as well as an increase in labor supply—*labor market spillovers*. Together, these spillovers lead to more patients being treated and more workers being employed in competitor hospitals. However, the decrease in patients and employment in the merging hospitals is not fully compensated by the increase in patients and employment in competitor hospitals, implying that some patients go without treatment and some workers go without employment.

It is important to observe that Proposition 3 does not provide unambiguous predictions about whether prices and wages will rise or fall in competitor hospitals. Consider the product market: from the perspective of a competitor hospital, the product demand curve has shifted outward after the merger, so the price is higher for any number of patients. This is a force that raises the price. However, the competitor's best response is to increase the number of patients until marginal revenue equals marginal cost, which is a force that lowers the price. If the marginal cost curve is steeper, a smaller increase in number of patients equates marginal revenue to marginal cost, so the price will not fall as far. On the other hand, the simultaneous outward shift in the labor supply curve implies that the price must fall further to equate marginal revenue and marginal cost. Similarly, whether the wage rises or falls in competitor hospitals after a local merger depends on several factors, including the slope and shift in the product demand curve.

2.3 Endogenous Quality of Care in Hospital Mergers

I now allow for the endogenous provision of quality. In particular, I permit $\beta_Y > 0$. There are two first-order conditions for profit-maximization:

Lemma 3 (First-order condition in a single-establishment firm with endogenous quality)**.** *Suppose quality of care is endogenous. For a single-hospital system h, the first-order condition with respect to non-patient care labor is,*

$$\underbrace{W_{ht}^N \times \left(1 + 1/\theta_{ht}^N\right)}_{\equiv MC_{ht}^N} = \overbrace{\underbrace{F_{ht}^N}_{\text{marginal quality (N)}} \times \underbrace{\left.\frac{\partial P_{ht}}{\partial Y_{ht}}\right|_{Q_{ht}}}_{=\beta_Y/\beta_P}}^{\text{returns from quality (N)}}.$$

where $F_{ht}^N \equiv \frac{\partial}{\partial N_{ht}} F\left(L_{ht}, N_{ht} - \psi_{ht}\right)$, and the first-order condition with respect to patient care

labor is,

$$MC^L_{ht} = MR^L_{ht} + \overbrace{\underbrace{\left(F^L_{ht} - Y_{ht}MP^L_{ht}\right)}_{\substack{\text{marginal quality (L)}\\ \text{with congestion}}} \times \underbrace{\frac{\partial P_{ht}}{\partial Y_{ht}}\Big|_{Q_{ht}}}_{=\beta_Y/\beta_P}}^{\text{returns from quality (L)}},$$

where MR^L_{ht} *and* MC^L_{ht} *are defined in Lemma 1 and* $F^L_{ht} \equiv \frac{\partial}{\partial L_{ht}} F\left(L_{ht}, N_{ht} - \psi_{ht}\right)$.

First, consider the first-order condition with respect to non-patient care labor, N. The term β_Y/β_P is the patients' marginal rate of substitution between quality and money. For each additional non-patient care worker hired, quality increases by F^N_{ht}/Q_{ht}, and the hospital can raise its price by $(\beta_Y/\beta_P) \times F^N_{ht}/Q_{ht}$ for each of its Q_{ht} patients, generating $\beta_Y/\beta_P \times F^N_{ht}$ in additional revenue through the quality channel. Since hiring non-patient care labor does not directly change the number of patients treated, the first-order condition simply equates the marginal cost of non-patient care labor, MC^N_{ht}, to the marginal revenue generated by the quality improvement these workers provide.

Also implicit in the first-order condition for non-patient care labor is the negative effect of overhead costs on quality. The marginal revenue from quality depends on workers not assigned to overhead, $N_{ht} - \psi_{ht}$, while the marginal cost of labor, MC^N_{ht}, depends on total hires, N_{ht}. An increase in the overhead requirement ψ_{ht} creates a larger wedge between these two quantities. As ψ_{ht} increases, the number of total staff N_{ht} required to achieve a given level of quality rises, and thus the marginal cost of quality rises. This increased effective cost of producing quality leads the hospital to reduce its optimal level of quality-producing staff, which unambiguously reduces the optimal quality of care, Y_{ht}.

Second, consider the first-order condition with respect to L. Hiring a patient care worker generates marginal revenue from attracting patients, MR^L_{ht}, and requires paying a marginal cost, MC^L_{ht}, as in Lemma 1. Furthermore, it generates two competing effects on quality. For a given number of patients, more patient care labor increases effective staffing and thus quality via the F^L_{ht} term. However, there is also a congestion effect: since patient care labor enables the treatment of more patients (by an amount MP^L_{ht}), the hospital's staff becomes diluted across a larger patient population. We say that the hospital operates in the *congested region* if $F^L_{ht} < Y_{ht}\text{MP}^L_{ht}$, which is the region over which the congestion effect dominates and the marginal revenues from the quality channel become negative. The congestion region can equivalently be expressed as $\theta^{F,L}_{ht} < \theta^{T,L}_{ht}$, where $\theta^{F,L}_{ht}$ is the elasticity of effective staffing, $F(\cdot)$, and $\theta^{T,L}_{ht} \equiv \frac{L_{ht}}{Q_{ht}}\text{MP}^L_{ht}$ is the elasticity of the patient production function, $T(\cdot)$, each with respect to L_{ht}. This expression makes clear that congestion occurs if an increase in patient care labor expands the number of patients treated proportionally more than it expands effective staffing, causing quality (effective staffing per patient) to fall.

Given endogenous quality provision, I now consider a merger between hospitals h and g:

Lemma 4 (First-order condition in a two-establishment firm with endogenous quality)**.** *Suppose quality of care is endogenous. If hospitals h and g in market m merge at time t to form a two-hospital system, the first-order condition for non-patient care labor is,*

$$MC_{ht}^{N} \quad + \quad \underbrace{\frac{\partial W_{gt}^{N}}{\partial s_{ht}^{N}} s_{gt}^{N}}_{\text{labor diversion in } N\ (+)} \quad = \quad \frac{\beta_Y}{\beta_P} F_{ht}^{N},$$

and the first-order condition for hospital h with respect to patient care labor is,

$$MC_{ht}^{L} + \underbrace{\frac{\partial W_{gt}^{L}}{\partial s_{ht}^{L}} s_{gt}^{L}}_{\text{labor diversion in } L\ (+)} \quad = \quad MR_{ht}^{L} + \frac{\beta_Y}{\beta_P}\left(F_{ht}^{L} - Y_{ht} MP_{ht}^{L}\right) + \underbrace{\frac{\partial P_{gt}}{\partial s_{ht}^{Q}} s_{gt}^{Q} MP_{ht}^{L}}_{\text{product diversion } (-)} .$$

The effect of the merger on non-patient care labor is straightforward: the marginal revenue from quality provision remains unchanged, $\beta_Y/\beta_P \times F_{ht}^{N}$, because the patients' valuation of quality and the quality of care technology are unaffected by market structure. However, labor diversion effectively increases the marginal cost of hiring these workers. With unchanged marginal revenue and higher marginal cost, non-patient care employment must fall at the merging hospitals. Figure 3 illustrates the hospital's quality of care decision in terms of choosing N for the simple case in which $F(L, N) = L + N$. Holding the number of patients fixed, the price-quality indifference curve is represented by an upward-sloping relationship between quality and price, with constant slope β_Y/β_P, and thus constant marginal revenue curve. Quality is chosen indirectly through the choice of non-patient care labor to equate marginal revenue and cost. After a merger, the marginal cost curve shifts upwards, while the marginal revenue curve is unchanged, implying that marginal revenue and marginal cost are equated by a lower choice of non-patient care labor and thus lower quality of care.

While Figure 3 makes clear why non-patient care labor is reduced by a merger and how this puts downward pressure on quality, the overall effect of a merger on quality is more nuanced because the number of patient care workers and patients adjust simultaneously. The following proposition summarizes the implications:

Proposition 4 (Responses to a merger with endogenous quality)**.** *Suppose quality of care is endogenous ($\beta_Y > 0$). If hospitals h and g in market m merge at time t to form a two-hospital system H, market m is affected as follows:*

(a) For patient care labor, the labor market responses to a merger are qualitatively the same as in Proposition 2.

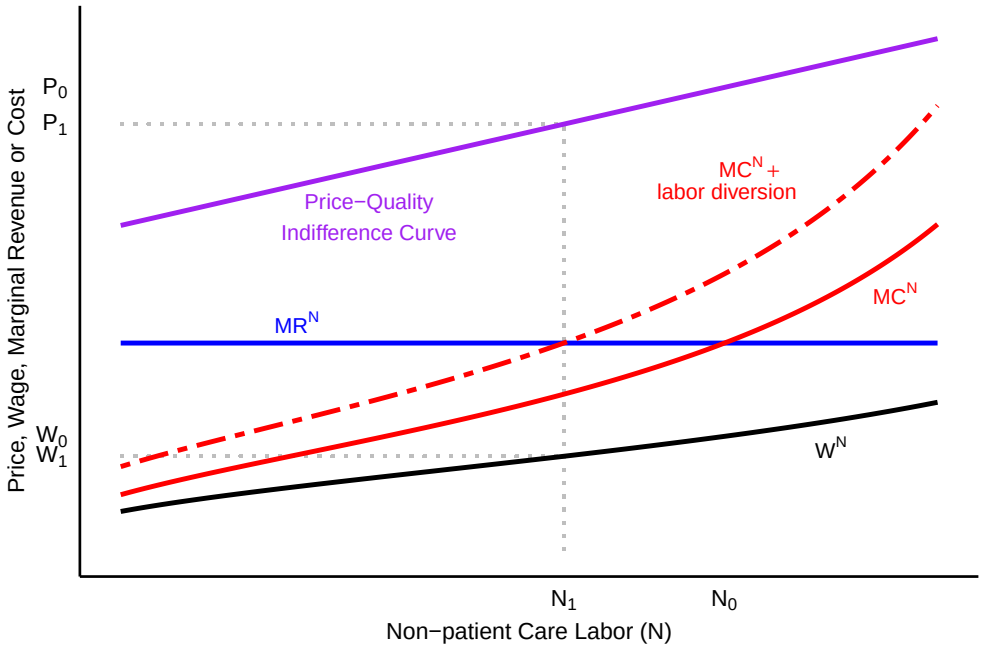


Figure 3: Quality-setting in Response to a Merger: Visualization of Lemma 4

Notes: This figure demonstrates how the first-order condition for profit-maximization with respect to non-patient care labor changes from the perspective of hospital h in response to its merger with hospital g. In response to the merger, the marginal cost of non-patient care labor shifts upwards due to labor diversion from hospital g, indicated by the difference between the solid and dashed MC^N lines. This figure is simulated from the actual model, setting $T_{ht}(L_{ht}) = A_{ht}L_{ht}$ and $F(L, N) = L + N$.

(b) The non-patient care wage and number of workers employed decrease in both hospitals h and g, and the non-patient care wage markdown strengthens in both.

(c) The number of patients treated decreases in system H.

(d) If the hospital operates below the congestion threshold, quality decreases and the sign of the price response is ambiguous.

(e) If the hospital operates above the congestion threshold, the price increases and the sign of the quality response is ambiguous.

Proof. See Appendix A. □

The endogeneity of quality does not overturn the quantity-related predictions of Proposition 2: merging hospitals reduce the number of patients treated and the number of workers employed—both for patient care and non-patient care workers. Wages fall and markdowns strengthen for both types of workers. However, in the presence of endogenous quality, price and quality responses are interdependent: a quality reduction places downward-pressure on

the price, so if the hospital operates in the non-congested region in which quality decreases in response to a merger, the price effect is mitigated and may be eliminated entirely. Conversely, if the hospital operates in a sufficiently congested region such that a merger increases quality, this reinforces the increase in price.

3 Data Sources and Descriptive Statistics

The administrative panel data for the US hospital industry provide a particularly apt empirical context for answering our research questions for several reasons. First, it is very rare for an industry-wide panel dataset in the US to include both product market outcomes—prices and quantities sold—and labor market outcomes—wages and employment. Our data include all four variables. Second, in other US administrative data on labor markets, it is unusual to observe: (i) hourly wages rather than total earnings—hours and hourly wages are not reported in IRS tax records or LEHD unemployment insurance records; (ii) employment by occupation within the firm—occupations are not formally reported in IRS or LEHD data; or (iii) wage and employment measures at the establishment-level rather than firm-wide level of aggregation—both IRS and LEHD records are firm-level. Our data provide hourly wages and employment by occupational category at the establishment-level. Third, it is unusual for an industry-wide dataset in the US to include measures of product quality, much less how quality changes over time. Our data provide both input-based and outcome-based measures of the quality of patient care over time. Fourth, the event study design in the next section requires a large number of mergers satisfying high-concentration thresholds in a large number of distinct markets. About one-fourth of all hospital markets in the US experienced such mergers, providing an unusually large sample.

3.1 Outcomes for Patients and Workers

Product and Labor Market Outcomes. My first data source is the Centers for Medicaid & Medicare Services (CMS) Hospital Cost Reports (HCRIS), spanning 1996 to 2022. This dataset comprises government-mandated reports from all Medicare-certified hospitals in the US. Following the literature, I exclude specialty and critical-access hospitals, which comprise a small share of overall inpatient treatments. The final sample contains about 3,400 unique hospitals and 81,000 annual observations.

In the product market, I define the quantity of patients, Q_{ht}, by the total number of inpatient discharges in the hospital-year. I measure prices by revenue-per-patient among non-Medicare inpatients, following Dafny (2009) and Dafny et al. (2019). Prices are adjusted

for composition such that it is as if all hospitals had the same payer and case mix index, following Brot, Cooper, Craig, and Klarnet (2024). Note that this adjustment is made over time within a hospital, so if the hospital's observed price were to fall over time due purely to a shift towards less-expensive cases (that is, a decrease in the case mix index), the composition-adjusted price would be unaffected.

I also use HCRIS to define the key labor market variables. At annual frequency, the reports include the total number of hours worked as well as the average wage per hour separately for several occupational categories. I follow the National Academy for State Health Policy (NASHP) in defining the patient care occupations as nurses, nursing aides, and the hospital's directly-employed physicians (commonly known as "hospitalists").[23] The remaining non-patient care occupations include administration, food services, sanitation, and maintenance. I measure the number of patient care workers, L_{ht}, and the number of non-patient care workers, N_{ht}, as the annualized number of full-time equivalent (FTE) workers in each occupational category (i.e. annual hours worked divided by 2,080). I measure the wage of patient care workers, W_{ht}^L, and the wage of non-patient care workers, W_{ht}^N, as the average wage per hour for the patient and non-patient care occupations, respectively.

Quality of Care Measures. To measure quality of care, I use three distinct approaches. First, I construct the simple staffing ratio as $(L_{ht} + N_{ht})/Q_{ht}$, which captures the total staff resources dedicated to each patient. The staffing ratio is commonly used as a measure of quality of care in the healthcare literature.

Second, I measure patient satisfaction using the Hospital Consumer Assessment of Healthcare Providers and Systems (HCAHPS), spanning 2008-2022 for the universe of hospitals. HCAHPS is a standardized national survey of a random sample of former patients. It includes an overall satisfaction rating, asks patients if they would recommend the hospital, and asks patients about their satisfaction with particular aspects of care, including cleanliness and quietness of the hospital environment.

Third, I measure mortality rates utilizing the Hospital Quality Initiative (HQI), spanning 2008-2021 for the universe of hospitals. These data include risk-adjusted 30-day all-cause mortality rates among patients originally treated at the hospital for specific conditions, particularly heart failure and pneumonia. These mortality rates, which are based on patient-level Medicare claims and eligibility information, are estimated using a statistical model to adjust for observable patient characteristics upon hospital arrival that predict mortality risk.

Table 1 demonstrates the relationship between a simple quality measure—the staffing

[23]Only 6% of physicians were directly employed by hospitals in 2012 (Kane, 2025), while more than 50% of all nurses are directly employed by hospitals (U.S. Bureau of Labor Statistics, 2023).

	Overall Patient Satisfaction Rating (%)			
	(1)	**(2)**	**(3)**	**(4)**
Staffing Ratio (log)	4.272***	2.824***	1.835***	1.475***
	(0.281)	(0.307)	(0.213)	(0.188)
Lagged Staffing Ratio (log)				0.709***
				(0.177)
Year FE	No	Yes	Yes	Yes
Market FE	No	Yes	No	No
Hospital FE	No	No	Yes	Yes
	Individual Item Satisfaction Ratings (%)			
	Cleanliness		**Quietness**	
	(5)		**(6)**	
Staffing Ratio (log)	1.457***		1.737***	
	(0.198)		(0.208)	
Year FE	Yes		Yes	
Hospital FE	Yes		Yes	

Table 1: Comparison between Staffing Ratio and Patient Satisfaction Ratings

Notes: This table presents fixed-effects regression estimates that relate the staffing ratio to patient satisfaction ratings. There are 3,149 hospitals satisfying the sample restrictions. Patient satisfaction ratings are available during 2008-2022. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels.

ratio—and patient satisfaction ratings. I regress overall patient satisfaction rating (the percent of top-box ratings) on log staffing ratio, with several control specifications. Column 1 provides the bivariate relationship without controls, while columns 2-4 progressively add year fixed effects, market fixed effects, hospital fixed effects, and the lagged value of the log staffing ratio. In all specifications, I find a strong positive relationship between staffing ratios and patient satisfaction, with a 1% increase in staffing ratio associated with a 1.5-4.3 percentage point increase in patient satisfaction, depending on the specification. Column 4 includes hospital and year fixed effects as well as the lagged staffing ratio, indicating that the contemporaneous relationship is strong even when controlling for hospital-specific factors and past staffing levels.

Furthermore, the bottom panel of Table 1 examines specific dimensions of care quality. I find that higher staffing ratios are associated with better ratings for hospital cleanliness (column 5) and quietness (column 6), with hospital and year fixed effects included. These results suggest that staffing ratios impact multiple dimensions of care quality as perceived by patients.

3.2 Market Concentration and Anti-competitive Mergers

Ownership and Concentration. To identify hospital ownership, I utilize the database developed by Cooper et al. (2019). They created and extensively validated a database on the universe of hospital mergers over 2001-2014. I supplement this database to include mergers over 1999-2018 by following their process and using the American Hospital Association (AHA) Annual Survey, Irving Levin Associates consulting reports, and digital newspaper indices to identify additional consolidation events and the year in which those consolidation events occurred.

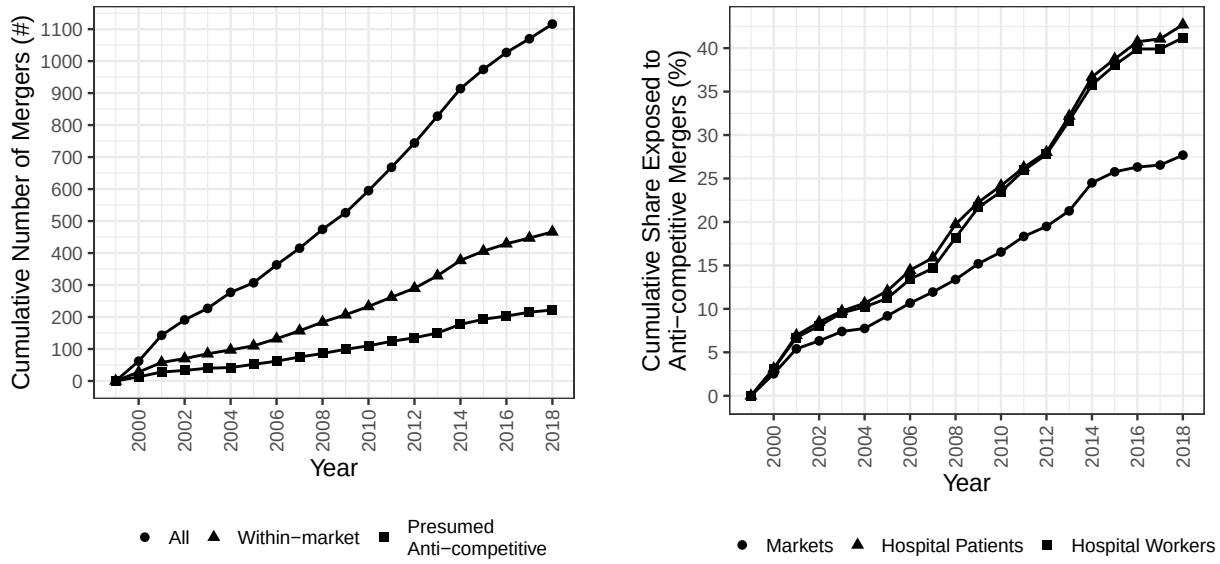
For the regression analysis, I follow Finkelstein, Gentzkow, and Williams (2021) for patient care markets and Prager and Schmitt (2021) for hospital labor markets by using commuting zones as the relevant geographic unit.[24] This results in 561 distinct markets across the US that include at least one hospital satisfying the sample criteria. To measure market concentration, I use the Herfindahl-Hirschman Index (HHI).[25] In the product market, denoting the inside market share by $\tilde{s}^Q_{ht}$, the HHI is defined by $\text{HHI} = \sum_h (\tilde{s}^Q_{ht})^2 \times 10,000$.[26] The change in the HHI induced by a merger between hospitals h and g at time t, based on pre-merger market shares, is $\Delta\text{HHI} = 2 \times \tilde{s}^Q_{ht} \times \tilde{s}^Q_{gt} \times 10,000$. Similarly, for labor markets, I calculate concentration using the inside labor market shares, $\tilde{s}^L_{ht}$ and $\tilde{s}^N_{ht}$, for patient and non-patient care workers, respectively.

Appendix Figure A1 provides several cross-sectional measures of hospital market concentration. In each subfigure, the x-axis displays quantiles across the distribution of markets, the solid line corresponds to year 2000, and the dashed line corresponds to year 2018. Subfigure (a) displays the number of hospitals per market, showing that the median market has only three hospitals. The bottom ventile of markets have only one hospital, and the top ventile of markets have at least seven hospitals. These counts have not changed over time. Subfigure (b) provides the analogous counts by number of hospital *systems* per market (i.e. commonly-owned groups of hospitals), revealing that the median market has decreased from

[24] Results are quite similar if defining a market as a circle of 30-mile radius around each hospital.

[25] Note that the HHI is utilized in this paper merely as a model-free descriptive statistic to narrow down a policy-relevant set of mergers in accordance with US DOJ and FTC (2023) guidelines. The HHI is taken as a pre-period covariate in the event study presented below, not as the driving variable. This is because the HHI is not a sufficient statistic for market power in a differentiated market. Instead, we use model-consistent markups and markdowns as the relevant measures of market power, as discussed in the previous section and estimated using the empirical model below.

[26] The inside market share is the share of *observed* patients in the market. It differs from the theoretical market share, which is the share of *potential* patients, inclusive of the outside option share. The inside market share is widely used in reduced-form studies of mergers because it is feasible to construct without a model. The HHI based on the internal market share is strictly greater than the HHI based on the potential market share if the outside option share is positive. In the model estimation below, we utilize the theoretical market share to ensure model consistency.



(a) Cumulative Number of Mergers

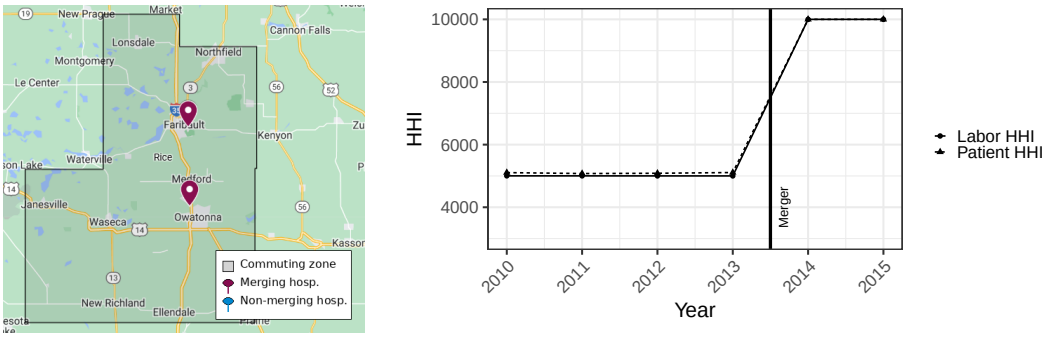
(b) Cumulative Exposure to Mergers

Figure 4: Summary of the Number of Mergers and Exposure to those Mergers

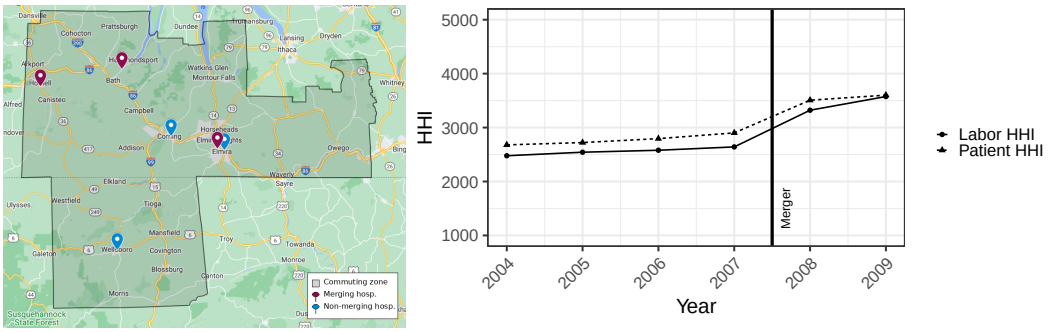
Notes: This figure presents the cumulative number of mergers (subfigure a) and the cumulative share of markets exposed to presumed anti-competitive mergers (subfigure b).

three systems to two systems. Subfigure (c) presents the distribution of HHI across markets, and subfigure (d) presents the market share of the largest hospital in each market. Both measures have shifted upwards over time as markets have become more concentrated. The market share of the largest hospital in the median market has risen from 60% to 65%.

Presumed Anti-competitive Mergers. US courts consider a merger to be "presumed anti-competitive" if $\Delta\text{HHI} > 100$ and post-merger $\text{HHI} > 1,800$ (US DOJ and FTC, 2023). Figure 4(a) presents the cumulative number of mergers over time relative to 1999, comparing all mergers, within-market mergers, and presumed anti-competitive mergers. The first result is that there are about 1,100 total merger events among the nearly 3,200 hospitals satisfying the sample definitions. The second result is that about 450 of those mergers were within the same market (i.e., $\Delta\text{HHI} > 0$). The third result is that just over half of these within-market mergers satisfy the presumed anti-competitive thresholds. To understand how exposed were patients and workers to presumed anti-competitive hospital mergers, Figure 4(b) presents the cumulative share of markets in which such a presumed anti-competitive merger occurred, demonstrating that more than one-fourth of all markets experienced such a merger. Nearly half of all hospital patients and workers belong to markets in which an anti-competitive hospital merger occurred.



(a) Merger between District One & Allina Health in the Commuting Zone of Faribault-Northfield-Otawonna, Minnesota



(b) Merger between Arnot Ogden, St. James Mercy, & Ira Davenport in the Commuting Zone of Corning-Elmira, New York

Figure 5: Examples of Mergers and Corresponding Changes in HHI

Notes: This figure provides two examples of mergers satisfying the "presumed anti-competitive" HHI thresholds. See the text for additional details.

In the next section, I utilize presumed anti-competitive mergers as the treated units in a difference-in-differences design. Before proceeding, I provide two examples of such treated units. Figure 5(a) considers the 2013 merger between District One Hospital and Allina Health in Minnesota. This merger occurred in a market with only two hospitals (Faribault-Northfield-Owatonna) and a population of approximately 123,000. Prior to the merger, the HHI was around 5,000 for both labor and product markets. After the merger, there was a single hospital system operating in the market, so the HHI rose to the maximum possible value of 10,000 in both markets. Figure 5(b) provides a much less extreme example. In 2008, Arnot Ogden Medical Center, St. James Mercy Health, and Ira Davenport Memorial hospitals merged in the Corning-Elmira region of New York, with a population of about 263,000. Before the merger, there were six distinct hospitals, and the HHIs were around 2,500 in the labor and product markets. Following the merger, there were only four distinct hospital systems, and HHIs increased to around 3,500.

4 Evidence from Merger Events

4.1 Empirical Design:

Staggered Difference-in-differences. In order to identify the causal effects of hospital consolidation, the empirical strategy employs a difference-in-differences (DiD) design. This approach compares changes in outcomes for hospitals involved in mergers against a matched-on-observables control group of hospitals that remain independent. The treated group is comprised of hospitals participating in mergers that satisfy the "presumed anti-competitive" threshold defined in the previous section. In cases where hospitals are involved in multiple mergers during the sample period, I focus only on the first merger event to protect against compounding effects. To implement this design, I follow Brot et al. (2024) by matching each merging entity to ten control hospitals based on pre-merger characteristics. These controls are selected from hospitals that are located in different markets and are not involved in any merger during the observation window (from four years before to seven years after the treatment hospital's merger).

The matching procedure uses a rich set of pre-merger covariates and lagged outcomes. This set includes hospital-specific characteristics in the product market (case mix index, percentage of Medicare patients, percentage of Medicaid patients, bed capacity, log number of patients, log price per patient), hospital-specific characteristics in the labor market (log number of workers in patient care, log number of workers in non-patient care, log wage in patient care, log wage in non-patient care), and regional characteristics (market unemployment rate, log market average income, and percentage of healthcare workers in the market). I estimate propensity scores using a logistic regression model and select with replacement the ten hospitals with the closest propensity scores to each merging entity.

The regression specification accounts for staggered treatment onset following the approach of Callaway and Sant'Anna (2021). Specifically, for each merged hospital h that consolidates in year t, I calculate the difference-in-differences estimator as,

$$\text{DiD}_{h,t,e} \equiv (Y_{h,t+e} - Y_{h,t-1}) - \underbrace{\mathbb{E}\left[Y_{h',t+e} - Y_{h',t-1} \mid h' \in C_h\right]}_{\substack{\text{change from } t-1 \text{ to } t+e, \\ \text{control hospitals matched to } h}}, \tag{8}$$

where $Y_{h,t+e}$ represents the outcome for hospital h at event time e relative to its merger year t, and C_h denotes the set of control hospitals matched to h. I then average these treatment effects across all mergers and across all merger cohorts:

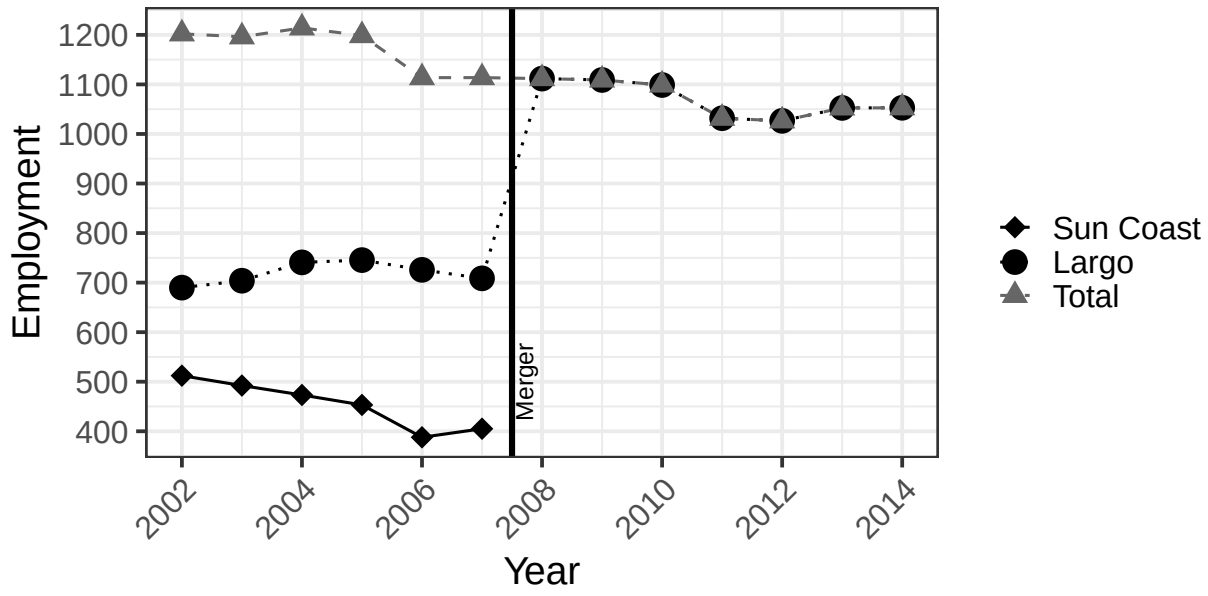


Figure 6: Example of Missing Data and Time-consistent Merging Firm Concept

Notes: This figure presents the employment reported to CMS by Largo Medical Center and Sun Coast Hospital. These two hospitals, located near Tampa, Florida, merged in 2008. After merging, Sun Coast stopped filing a separate report, instead reporting its employment and other outcomes under Largo's identification number.

$$\text{DiD}_e \equiv \sum_t \omega_{t,e} \times \frac{1}{|\mathcal{G}_t|} \sum_{h \in \mathcal{G}_t} \text{DiD}_{h,t,e}, \quad \omega_{t,e} \equiv \frac{|\mathcal{G}_t|}{\sum_t |\mathcal{G}_t|}, \tag{9}$$

where $\mathcal{G}_t$ represents the set of hospitals that merge in year t.

Time-consistent Firms. A data challenge in studying hospital mergers is that, in the HCRIS data, hospitals often jointly report their outcomes after merging, such that it is not possible to track each hospital separately after the merger. For example, as illustrated in Figure 6, when Largo Medical Center merged with Sun Coast Hospital in 2008 in Florida, Sun Coast began reporting its information under Largo's Medicare identification number. This gives the appearance that Sun Coast stopped operating, while Largo hired a massive number of workers in a single year, neither of which happened in reality. If one were to estimate DiD at the hospital-level only for hospitals consistently observed before and after the merger, one would only include Largo in the sample, and would mistakenly conclude that this merger caused a massive increase in employment among hospitals involved in mergers.

To address this issue, I construct time-consistent firm-by-market units of observation. For each merger, a time-consistent entity is defined by aggregating outcomes (through sums

or averages) across all hospitals that eventually consolidate, creating a consistent unit that can be followed over time. This approach allows me to determine whether total employment, patient volume, or quality metrics change after consolidation, even when individual hospital data are no longer separately reported. In Figure 6, this would mean that I estimate DiD using total employment (triangle symbols) rather than only using Largo's reported employment (circle symbols) as the outcome of interest.

4.2 Direct Effects of Mergers on the Merging Hospitals

The DiD estimation results are summarized in Table 2. Following consolidation, the number of patients treated by merged hospitals decreases substantially, with an average reduction of more than 6%. As demonstrated in Figure 7, the reduction in patients materializes by the second year after the merger. Composition-adjusted prices, which are measured with more noise, become statistically significant after five years with a price increase of about 8%. This price increase is in line with estimates from the literature using several data sources (Dafny et al., 2019; Cooper et al., 2019; Brand et al., 2023), while the prior literature has not previously reported a significant decrease in patient volume in response to mergers.

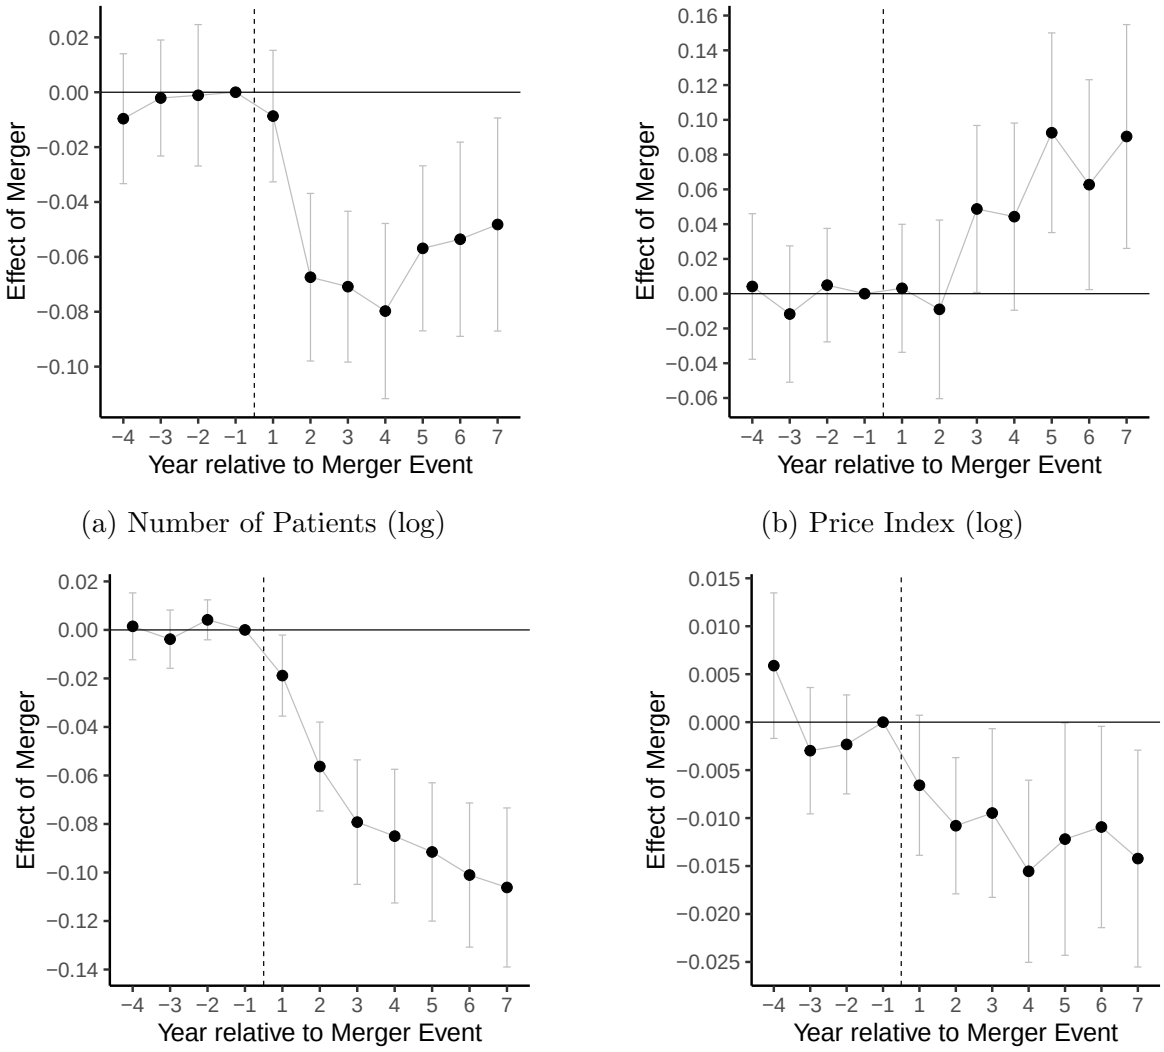
Turning to the labor market, wages decrease by more than 1%. These effects materialize quickly and persist over time. The employment effects are considerably larger, with hospitals reducing FTE employment by approximately 10% on average after five years. Using the AHA Survey to measure headcount employment for a slightly smaller sample of hospitals, the reduction in headcount employment is somewhat larger than the reduction in FTE hours of employment, exceeding 11%. Comparing panels in Figure 7, it is apparent that the employment reduction closely tracks both the wage reduction and the decline in patient volume over time. I also examine changes in the occupational mix of workers; the main estimates are reported in Table 2 with corresponding event studies in Appendix Figure A2. The reduction in employment is as large as 12% for non-patient care workers and 9% for patient care workers, while the reduction in wages reaches nearly 3% for non-patient care workers and 2% for patient care workers. Thus, the employment and wage reductions are greater for non-patient care workers, such as administrators, but are also economically and statistically significant for patient care workers, such as nurses.

Accounting for simultaneous changes in the number of patients and workers, the simple staffing ratio, $(L+N)/Q$, declines by approximately 6% after five years, as the reduction in labor is proportionally greater than the reduction in patients. As an alternative, I consider a Cobb-Douglas effective staffing technology, $L^{\delta}N^{1-\delta}/Q$, finding a reduction of 7% when calibrating $\delta = 1/3$ and 6% when calibrating $\delta = 2/3$. The decline in the staffing per patient

	Treatment Group Summary			DiD Effects of the Merger		
	Mergers	Hospitals	Outcome SD	Before Merger	After Merger	
Event Times:		{-1}		{-2,-3,-4}	{2,3,4}	{5,6,7}
Panel A.	**Product Market Outcomes**					
Number of Patients (log)	141	394	0.861	-0.004	-0.073***	-0.053***
				(0.013)	(0.016)	(0.019)
Price Index (log)	141	393	0.466	-0.001	0.028	0.082**
				(0.018)	(0.027)	(0.033)
Staffing Ratio (log)	141	394	0.261	0.010	-0.011	-0.060***
				(0.013)	(0.014)	(0.018)
Case Mix Index (log)	137	380	0.153	0.000	0.001	0.001
				(0.002)	(0.004)	(0.005)
Covered by Medicaid (share)	141	394	0.075	0.000	0.003	0.003
				(0.003)	(0.004)	(0.005)
Panel B.	**Labor Market Outcomes**					
Number of Workers, Headcount (log)	128	362	0.930	-0.007	-0.079***	-0.116***
				(0.008)	(0.021)	(0.026)
Number of Workers, FTE Hours (log)	141	394	0.887	0.001	-0.074***	-0.100***
				(0.006)	(0.013)	(0.017)
· · · Patient Care Occupations	140	392	0.935	-0.001	-0.068***	-0.087***
				(0.007)	(0.015)	(0.020)
· · · Non-patient Care Occupations	138	385	0.820	-0.003	-0.088***	-0.122***
				(0.009)	(0.016)	(0.022)
Hourly Wage (log)	141	394	0.204	0.000	-0.012***	-0.012**
				(0.003)	(0.004)	(0.006)
· · · Patient Care Occupations	140	392	0.196	0.000	-0.011**	-0.017**
				(0.005)	(0.005)	(0.008)
· · · Non-patient Care Occupations	138	385	0.231	-0.006	-0.021***	-0.027***
				(0.006)	(0.007)	(0.009)

Table 2: Direct Effects of Mergers on Outcomes in the Merging Hospitals

Notes: This table presents difference-in-differences (DiD) estimates of the effects of mergers on outcomes in the merging hospitals. Treated units are defined as time-consistent merging hospitals satisfying the "presumed anti-competitive" HHI thresholds, as defined in the text. Each treated unit is compared to a merger-specific control group of 10 hospitals from other markets, matched on a large set of pre-merger covariates using a propensity score approach. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels. Outcome SD refers to the standard deviation of the outcome variable.



(a) Number of Patients (log)

(b) Price Index (log)

(c) Number of Workers, FTE (log)

(d) Hourly Wage (log)

Figure 7: Direct Effects of Mergers on Outcomes in the Merging Hospitals over Time

Notes: This figure presents difference-in-differences estimates that compare treated hospitals, defined as merging hospitals satisfying the "presumed anti-competitive" HHI thresholds, to merger-specific control groups of 10 hospitals from other markets matched to treated units by propensity score. 90% confidence intervals are displayed as brackets.

suggests a decline in the quality of care in merging hospitals, which is investigated further in the next subsection.

4.3 Direct Effects of Mergers on Quality of Care

Figure 8 presents event study estimates for various measures of the quality of care provided to patients in the merging hospitals; a summary is provided in Appendix Table A1. A smaller set of mergers are studied in this subsection because the HCAHPS and HQI databases are only available in more recent years.

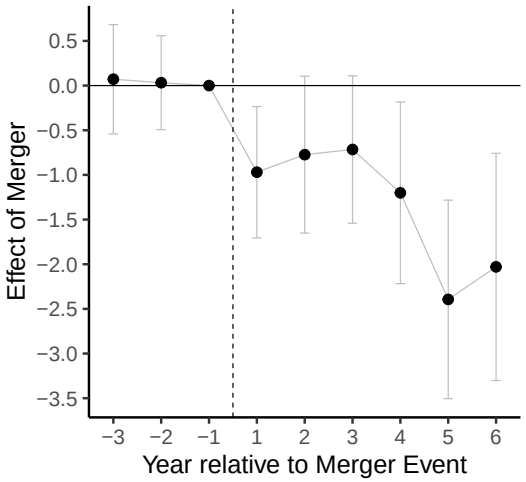
The top panels display the effect on patient-reported satisfaction from HCAHPS surveys. The percentage of patients reporting that the hospital provides the highest level of quality declines by approximately 2% post-merger, with this effect becoming stronger and statistically significant over time. A similar pattern is shown for the percentage who would recommend the hospital to others. When examining more specific ratings of hospital characteristics, patients report a decline in the cleanliness of the hospital environment. These consistent negative effects across multiple patient-reported measures provide compelling evidence that patients perceive deterioration in their care experience following hospital mergers.

The bottom panels use the HQI data to estimate effects on patient outcome measures—specifically, risk-adjusted 30-day mortality rates for patients admitted for heart failure or pneumonia. Both measures show concerning increases following mergers. The pneumonia mortality rate increases by approximately 0.8 percentage points (relative to the national average mortality rate of 13%), while heart failure mortality increases by as much as 0.6 percentage points (relative to the national average mortality rate of 12%). To place these estimates in context, Cooper et al. (2022) find that a one standard deviation increase in the hospital price distribution is associated with a 0.5 percentage point decrease in mortality.

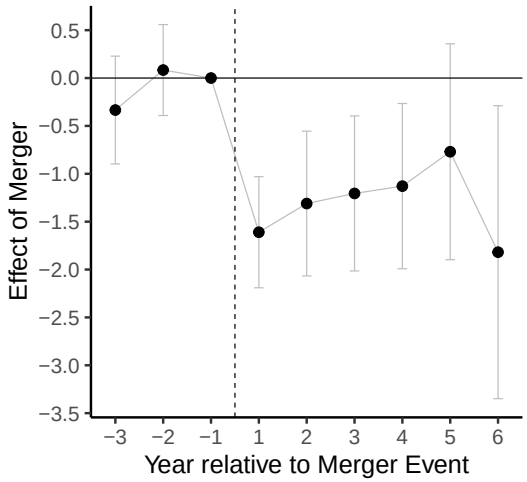
In sum, I find a decline in quality of care across multiple dimensions—staffing ratios, patient satisfaction, and medical outcomes—indicating that the quality reductions are perceived by patients and manifested in worsening health outcomes.

4.4 Aggregate and Spillover Effects of Local Mergers

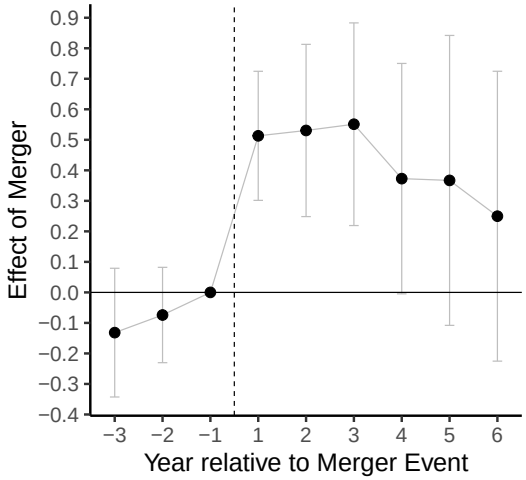
Beyond the direct effects on merging hospitals, the theory predicts that consolidation generates effects on the outside share of the market as well as spillovers on competitors in the same local market. These predictions are tested empirically utilizing two modified DiD designs. In order to estimate market-wide aggregate effects, I aggregate all firms in the market—including the merging firms—into a market-level aggregate outcome. The results are summarized in Panel A of Table 3, with corresponding event studies in Appendix Figure



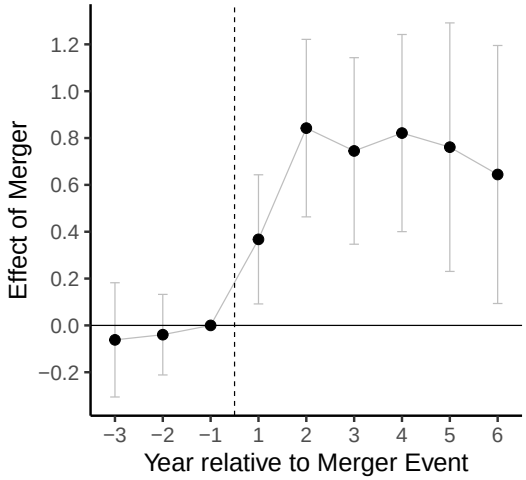
(a) Recommend Hospital (pp)



(b) Cleanliness Rating (pp)



(c) Heart Failure 30-Day Mortality (pp)



(d) Pneumonia 30-Day Mortality (pp)

Figure 8: Quality of Care Effects of Mergers in the Merging Hospitals over Time

Notes: This figure presents difference-in-differences estimates that compare treated hospitals, defined as merging hospitals satisfying the "presumed anti-competitive" HHI thresholds, to merger-specific control groups of 10 hospitals from other markets matched to treated units by propensity score. 90% confidence intervals are displayed as brackets.

A3. Similarly, in order to estimate spillover effects, I aggregate all non-merging competitors in the market into a time-consistent aggregate firm and estimate how it responds to the local merger. The results are summarized in Panel B of Table 3, with corresponding event studies in Appendix Figure A4.

The total number of patients treated in markets that experience a presumed anti-competitive merger declines by more than 3% in the shorter run (statistically significant) and 1% in the longer run (insignificant). By contrast, the aggregate price index does not demonstrate a statistically significant change. In the labor market, total market-wide employment decreases by 3-4% in both the shorter and longer run (significant). This reduction is similar and significant for both patient care and non-patient care occupations. The average hourly wage at the market level also experiences a statistically significant decrease of about 1%, with non-patient care occupations showing a more substantial reduction.

Regarding spillovers, non-merging hospitals significantly increase their volume of patients treated, by nearly 4% in the shorter term and 8% in the longer run. However, the price index does not exhibit statistically significant spillovers. In the labor market, competitors expand employment levels by 2% in the shorter run (statistically insignificant) and by 4% in the longer run (significant). This employment growth is strongest for non-patient care occupations. Notably, despite substantially increasing employment, competitor hospitals reduce hourly wages by about 2% in the longer run. This wage reduction is also stronger for non-patient care occupations.

In sum, after presumed anti-competitive mergers, local competitors increase patient volume and employment, despite decreasing wages and insignificantly changing prices. Aggregating across the merging and non-merging firms in the market, patient volume and employment decline, indicating that more patients go without treatment (especially in the shorter run) and more hospital workers go without jobs.

4.5 Alternative Interpretations of Findings

Our findings are consistent with the large set of predictions provided by Propositions 2, 3, and 4 for the labor market, patient care market, and quality of care outcomes, including both the predictions for the merging firms and the predictions for their competitors and outside shares. Nonetheless, certain findings may be consistent with alternative interpretations that are beyond the scope of the model. We now discuss some of these alternative interpretations and comment on their plausibility.

	Treatment Group Summary			DiD Effects of the Merger		
	Mergers	Hospitals	Outcome SD	Before Merger	After Merger	
Event Times:		{-1}		{-2,-3,-4}	{2,3,4}	{5,6,7}
Panel A.				**Market-wide Aggregate Effects**		
Number of Patients (log)	124	789	0.963	0.005	-0.033***	-0.011
				(0.008)	(0.012)	(0.016)
Price Index (log)	124	789	0.477	0.006	-0.006	-0.001
				(0.020)	(0.024)	(0.028)
Number of Workers (log)	124	789	0.978	0.001	-0.035***	-0.041***
				(0.006)	(0.009)	(0.013)
· · · Patient Care Occupations	124	789	1.008	0.006	-0.030**	-0.040**
				(0.007)	(0.012)	(0.016)
· · · Non-patient Care Occupations	124	789	0.938	0.000	-0.028**	-0.037**
				(0.009)	(0.011)	(0.018)
Hourly Wage (log)	124	789	0.210	-0.002	-0.008**	-0.008*
				(0.003)	(0.004)	(0.005)
· · · Patient Care Occupations	124	789	0.201	-0.002	-0.012**	-0.006
				(0.004)	(0.005)	(0.006)
· · · Non-patient Care Occupations	124	789	0.235	0.000	-0.014**	-0.023***
				(0.004)	(0.006)	(0.007)
Panel B.				**Within-Market Spillover Effects**		
Number of Patients (log)	82	222	1.168	-0.012	0.036**	0.083***
				(0.012)	(0.016)	(0.022)
Price Index (log)	82	222	0.531	-0.029	0.013	0.003
				(0.025)	(0.030)	(0.034)
Number of Workers (log)	82	222	1.065	-0.010	0.020	0.041**
				(0.008)	(0.012)	(0.019)
· · · Patient Care Occupations	82	222	1.159	-0.012	0.026*	0.033
				(0.009)	(0.015)	(0.023)
· · · Non-patient Care Occupations	82	222	0.965	0.008	0.041***	0.077***
				(0.011)	(0.014)	(0.022)
Hourly Wage (log)	82	222	0.220	0.007	-0.006	-0.022***
				(0.005)	(0.006)	(0.007)
· · · Patient Care Occupations	82	222	0.221	0.004	-0.011	-0.021*
				(0.007)	(0.008)	(0.010)
· · · Non-patient Care Occupations	82	222	0.243	-0.003	-0.019**	-0.027***
				(0.007)	(0.009)	(0.009)

Table 3: Aggregate and Spillover Effects of Mergers

Notes: This table presents difference-in-differences (DiD) estimates of the effects of mergers on outcomes aggregated to the CZ level. Treated CZs experienced a "presumed anti-competitive" merger, as defined in the text. Market-wide outcomes aggregate all hospitals in the treated CZ, while within-market spillover outcomes exclude hospitals involved in the merger. Each treated unit is compared to a merger-specific control group of 10 other CZs that did not experience a merger during the relevant time interval, matched on a large set of pre-merger covariates using a propensity score approach. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels. Outcome SD refers to the standard deviation of the outcome variable.

Patient Composition. The model suggests that the increase in price is a consequence of the large gain in product market concentration resulting from a presumed anti-competitive merger, amplified by the simultaneous gain in labor market concentration. However, an alternative interpretation is that the hospital shifted the composition of patients from low-price procedures to high-price procedures while keeping the price of each procedure constant, mechanically raising the average price at the hospital-level. Similarly, the increase in mortality rates among heart failure and pneumonia patients could arise from selecting a less-healthy pool of patients to treat.

There are several reasons to doubt that the price and mortality increases are due to changes in patient composition. First, as noted in Section 3, both the price and mortality indices are adjusted for time-varying patient composition prior to analysis. Second, as reported in Panel A of Table 2, presumed anti-competitive mergers have no impact on the hospital-level case mix index—an index of how costly to treat is the hospital's average procedure. Third, as also reported in Panel A of Table 2, the share of the hospital's patients covered by Medicaid—a proxy for the share of patients who are low-income—does not respond to presumed anti-competitive mergers. Fourth, because heart failure and pneumonia admissions are made on an emergency basis, patients have little opportunity to differentially select hospitals based on severity before versus after mergers. Fifth, Cooper et al. (2019) and Brand et al. (2023) utilize claims-level data to estimate changes in prices before and after mergers, controlling for patient characteristics; our estimates align with theirs.

Worker Composition. The model suggests that the decrease in the average wage of patient care workers is a consequence of the large gain in labor market concentration resulting from a presumed anti-competitive merger, amplified by the simultaneous gain in product market concentration and incentives to reduce quality of care. An alternative explanation is that the observed wage reduction results from a compositional shift within the hospital workforce. If the merger leads the hospital to lay off its more-skilled nurses, the calculated average wage for the remaining patient care workers would mechanically decrease.

To investigate skill composition changes directly, I utilize supplementary data from the AHA Annual Survey. This survey contains a finer breakdown of patient care employment by nurse occupations and thus allows us to examine changes in occupational composition. Two limitations of the survey are that it reports employment headcounts, rather than the labor hours reported in our primary HCRIS administrative dataset, and contains a slightly smaller sample. However, we found above that the surveyed headcount employment effects are similar, if slightly larger, than FTE employment effects. Using the AHA data, I compare the two available categories of nurse occupations: Registered Nurses (RN), who have a median

annual salary of $93,600, versus Licensed Practical Nurses (LPN) and Nursing Assistants (NA), who have median annual salaries of $62,340 and $49,500. The effect of mergers on the log ratio of the headcount employment of RN versus LPN and NA is 0.072 (with a standard error of 0.044), indicating a marginally significant shift towards higher-wage nurse occupations. Thus, to the extent that there is any composition bias for patient care workers, it actually biases us towards finding positive wage effects; the true wage effects may be even more strongly negative than reported above.

Administrative Efficiencies. The model suggests that the reduction in employment of non-patient care labor after a presumed anti-competitive merger is a consequence of the large gain in labor market concentration, amplified by the simultaneous gain in product market concentration and incentives to reduce quality of care. However, an alternative explanation for decreased employment after hospital mergers is a reduction in administrators that become redundant because the same administrator can frictionlessly perform the same task across multiple hospitals. If hospital mergers primarily change hospital operations by "trimming the fat" of excess administrators, this could explain a reduction in non-patient care workers unrelated to market power.[27]

Ex ante, there are several reasons to doubt that fixed cost reductions explain our findings. First, mergers lead to a large reduction in the number of patients treated and the number of patient care workers employed—the opposite of what one would predict from greater administrative efficiencies alone. Second, the primary administrative tasks performed in hospitals—scheduling patients, coordinating nurses, preparing claims, ordering medical supplies, and so on—are variable costs that scale with the number of patients treated and patient care workers employed. Third, nearly all of the acquired hospitals in our sample already belonged to a multi-hospital system prior to the merger, so any efficiencies from joining a multi-hospital system are likely to be in place in the pre-period.

Two recent studies use finer data on administrative costs to test directly for fixed cost reductions after hospital consolidation events. Gaynor, Sacarny, Sadun, Syverson, and Venkatesh (2023) provide a case study of a single hospital system merger using detailed financial, management, and clinical data. They find that the acquired system did not become more efficient, despite bringing in new management ostensibly to do so. Arnold, Gupta, Liu, and Olssen (2025) use detailed cost data from hospitals in California, which are more detailed than the cost data available for the national sample of hospitals, to estimate the effects of mergers on fixed and variable costs. Across several measures, they find no evidence that fixed costs are affected by mergers.

[27] In the model notation, this would mean that the overhead costs, ψ_{ht}, decrease in response to a merger.

Insurer Bargaining Power. In our model above, hospital systems increase prices after mergers by coordinating across hospitals to reduce patient volume—the classical exercise of oligopoly power—amplified by the simultaneous exercise of oligopsony power in labor markets. Such a model can rationalize large changes in the number of patients treated and cross-hospital patient diversion. However, oligopoly models struggle to generate log price changes as large as changes in log patient volume. Rationalizing such large price increases is made even more difficult when quality declines, since quality reduction dampens price increases. Motivated by this tension, an important literature considers the possibility that, as hospitals consolidate, they gain bargaining power relative to insurers, allowing hospitals to increase prices even while holding the number of patients fixed (Gowrisankaran et al., 2015; Ho and Lee, 2017; Dafny et al., 2019). Such models are successful in generating large price responses to hospital mergers, but struggle to generate large patient volume responses and cross-firm diversion. Rather than choosing one mechanism or the other, we propose a parsimonious and empirically-tractable model extension in the next section to allow that hospitals both exercise classical oligopoly power and increase bargaining power over insurers after a merger. As a result, our empirical model can rationalize large reductions in patient volume and quality while also rationalizing excessively large price increases.

5 Empirical Model and Estimation Results

Now that we have confirmed and quantified the predictions of the model, we structurally estimate the model parameters and examine the aptness of the model as an ex ante merger evaluation framework. There are three goals of this section. Our first goal is to provide a practical estimation strategy for recovering the model parameters and simulating realistic merger impacts, to which end, we also develop an extension with heterogeneous bargaining power over insurers that may increase in response to a merger. Our second goal is to show that the model is sufficiently flexible to reproduce the merger impacts presented in the previous section—both the direct effects of the mergers on the merging firms and the aggregate and spillover effects on the local market. Our third goal is to characterize market power—markdowns and markups—using the model-implied parameter distributions.

5.1 Model Estimation

Empirical Model In order to implement the model estimation, I must first parameterize the care technologies, which are nonparametric in Section 2. The patient care technology is specified as constant elasticity, $T_{ht}(L_{ht}) = A_{ht}L_{ht}^{\alpha}$, where A_{ht} is the relative productivity of h

and α is the elasticity of patients to employment. The quality of care technology is specified as constant elasticity of substitution, $F(L_{ht}, N_{ht} - \psi_{ht}) = (\delta (L_{ht})^{\rho} + (1-\delta)(N_{ht} - \psi_{ht})^{\rho})^{\phi/\rho}$. This allows that patient and non-patient care labor may be gross complements ($\rho < 0$) or gross substitutes ($\rho > 0$) in the provision of quality, and returns to scale in quality may be increasing ($\phi > 1$) or decreasing ($\phi < 1$). It includes the staffing ratio as a special case.

The empirical model also addresses the issue that insurers pay a large share of the hospital price, determined through hospital-insurer bargaining. Letting P_{ht}^{hos} denote the price received by the hospital from the insurer, and P_{ht}^{pat} denote the price paid by the patient ("coinsurance"), we have the accounting identity $P_{ht}^{\text{hos}} = \kappa_{ht} P_{ht}^{\text{pat}}$.[28] A higher value of κ_{ht} means that the hospital earns more from each dollar spent by a patient at hospital h, so κ_{ht} can be interpreted as the hospital's excess markup on insurers. From the hospital's perspective, P_{ht}^{hos} is the relevant price for measuring profits, but from the patient's perspective, P_{ht}^{pat} is the relevant price for determining demand. Holding P_{ht}^{pat} and thus patient demand fixed, increasing κ_{ht} can be interpreted as improving the hospital's bargaining power over insurers.

In order to allow for such bargaining power effects of mergers, I suppose that κ_{ht} may increase in response to the merger, such that hospitals extract greater prices from insurers even if holding the prices charged to patients fixed. For parsimony, this effect is parameterized as a proportional shift in κ_{ht}, that is, $\Delta \log \kappa_{ht} = \bar{\kappa}_{\Delta}$ is the change in κ_{ht} among the merging firms in response to a merger. Furthermore, by allowing for flexible heterogeneity in κ_{ht} at baseline, I solve an equilibrium in which hospital systems may be incentivized to steer patients towards high-κ_{ht} hospitals after merging, which is similar to the steering mechanism on predetermined coinsurance rates considered by Gowrisankaran et al. (2015). Appendix B provides the extended first-order conditions used in the estimation that allows the hospital to take into consideration κ_{ht}.

Estimation Strategy In industrial organization, it is common practice to estimate the components of the model, such as consumer preference parameters, using cross-sectional instruments. Such instruments—like the number of competitors in the market or the average quality of competitors—place strong implicit restrictions on cross-market sorting of competitors (see the discussion by Gandhi and Houde 2020). Such instruments can be difficult to rationalize. As an alternative, I propose to infer the model parameters from the ex post

[28] In the data, we observe P_{ht}^{hos} rather than P_{ht}^{pat}. The interpretation of the DiD estimate of the impact of mergers on the log price is not sensitive to the presence of heterogeneous κ_{ht}: assuming the coinsurance rate does not respond to the merger, $\Delta \log P_{ht}^{\text{hos}} = \Delta \log P_{ht}^{\text{pat}}$, so $\log P_{ht}^{\text{hos}}$ can be used as a proxy for $\log P_{ht}^{\text{pat}}$. However, a *change* in κ_{ht} in response to a merger alters the interpretation of this DiD estimate. Below, we provide a decomposition of the log price change in the total hospital price into the effect on patients' perceived prices versus the effect due to bargaining power changes.

merger evaluation presented in the previous section. By leveraging within-market changes over time, this approach places no restrictions on the cross-sectional sorting of competitors into markets. Instead, the identifying assumption is that mergers do not systematically induce changes in $\Lambda_{ht} \equiv (\xi^Q_{ht}, \xi^L_{ht}, \xi^N_{ht}, A_{ht}, \kappa_{ht}, \psi_{ht})$, except for proportional changes in κ_{ht}, as noted above. Appendix B develops the intuition for recovering the global model parameters, $\Xi \equiv (\beta_P, \beta_Y, \gamma_L, \gamma_N, \alpha, \delta, \rho, \phi, \bar{\kappa}_\Delta)$, using only the merger impacts. Given an estimate of Ξ, it is straightforward to recover the unobserved heterogeneity terms, Λ_{ht}, using model inversion.

In practice, I estimate the model using the method of simulated moments (MSM). Details are provided in Appendix B. Briefly, the estimation routine proceeds by (1) guessing a set of candidate global parameters, Ξ^*; (2) inferring the unobservable heterogeneity terms, Λ^*_{ht}, as a function of Ξ^*; (3) simulating the actual mergers that occurred in the treatment markets based on pre-period observables, Ξ^*, and Λ^*_{ht}; and (4) collecting the simulated average treatment effects on the merging firms and their competitors for various outcomes, denoted $\mathbf{M}^{sim}(\Xi^*)$. Given that the simulated merger effects can be obtained for any guess of Ξ^*, the MSM estimate, Ξ^{msm}, is the guess Ξ^* that minimizes the distance between these simulated moments and their observed counterparts, denoted $\mathbf{M}^{obs}$, which correspond to the merger effect estimates from the previous section. That is,

$$\Xi^{msm} \equiv \min_{\Xi^*} (\mathbf{M}^{obs} - \mathbf{M}^{sim}(\Xi^*))' \mathbf{W} (\mathbf{M}^{obs} - \mathbf{M}^{sim}(\Xi^*)),$$

where $\mathbf{W}$ is a weighting matrix. The MSM estimate of Λ_{ht} is the one inferred from Ξ^{msm}.

A final practical issue is that the model requires us to calibrate the baseline outside option shares, which is a standard challenge in merger evaluation. Since smaller outside shares imply greater market power, all else equal, we err on the side of larger outside shares to be conservative. In the product market, we suppose that 25% of inpatient treatments are diverted to the outside option at baseline, motivated by the estimate of Dingel et al. (2024) that 20% of all inpatient and outpatient Medicare treatments are out of region. In the labor market, we suppose that 50% of hospital workers are diverted to the outside option at baseline, motivated by the fact that more than 50% of all nurses are directly employed by hospitals (U.S. Bureau of Labor Statistics, 2023). The choice of 50% has also been used, for example, by Miller and Weinberg (2017) to calibrate the outside share in the beer industry.

5.2 Parameter Estimates and Simulated Merger Effects

The parameter estimates and fit of the simulated moments from the MSM estimation are presented in Table 4.

Parameter Estimates. The left side of the table provides the estimates of the global parameters. Regarding the patient preference parameters, we find a marginal rate of substitution, β_Y/β_P, of about 9.2. This translates to patients being willing to sacrifice about 0.30 standard deviations in the price distribution to improve one standard deviation in the quality distribution.[29] The estimated labor preference for the log-wage, γ_E, is estimated to be in the 4.5-4.8 range for both patient care and non-patient care labor.[30] To interpret these estimates, note that the weakest markdown permitted by the model is $\gamma_E/(1+\gamma_E)$, which is achieved as the hospital's labor market share approaches zero. From our estimates, the weakest possible markdown is 0.82 for patient care labor and 0.83 for non-patient care labor, such that workers are paid at least 17-18% less than their marginal revenue. We characterize the full empirical distribution of markdowns below.

Regarding the technology of care parameters, we find that the output elasticity of patient volume with respect to the number of patient care workers, α, is about 0.75, indicating substantially diminishing returns. In quality production, we find that the intensity of patient care labor, δ, is about 0.57, indicating that patient care and non-patient care labor are almost equally productive. The elasticity of substitution between patient and non-patient care labor, $1/(1-\rho)$, is estimated to be about 0.44, indicating that L and N are gross complements in production, working together to provide quality of care to patients. Finally, ϕ is estimated to be 1.06, indicating modestly increasing returns to scale, which aligns with the key mechanism of increasing returns to scale in quality from Dingel et al. (2024).

Simulated Merger Effects. The right side of the table presents the goodness of fit of the simulated moments relative to the estimated moments. In the product market, the model successfully replicates the qualitative patterns observed in the data: merging firms increase prices and reduce patient volume, competitors increase their patient volume, and the total number of patients treated in the market declines. Quantitatively, the model captures the majority of the direct price and quantity effects, though it slightly under-predicts the price increase (4.0% simulated vs. 5.5% estimated) and over-predicts the patient volume decrease (-8.3% vs. -6.3%). The model also provides a close fit for the aggregate reduction in patient volume (-2.6% vs. -2.2%). Since quality of care is not observable, we target the staffing ratio—which is in theory closely related to the quality measure—as an auxiliary moment for

[29] A standard deviation in the patient price distribution is 0.76 (measured in \$1,000 USD) and one standard deviation in the quality distribution is 0.024 (no natural units). The ratio of standard deviations is $0.024/0.760 \approx 0.032$. Then, $9.2 \times 0.032 \approx 0.30$ is the price of one standard deviation of quality along the indifference curve, measured in standard deviations of the patient price.

[30] Reassuringly, our merger-based instrument yields an estimate consistent with the 3-5 range recently found using several demand instruments in the US context (Kroft et al., 2025; Rubens et al., 2025).

MSM Parameter Estimates			Simulated Moment Fit		
Parameter	**Value**	**Description**	**Moment**	**Target**	**Simulated**
Patient Preferences			**Product Market: Patients**		
β_P	1.440	Disutility from price	Direct: $\Delta \log P_h$	0.055	0.040
β_Y	13.232	Utility from quality	Direct: $\Delta \log Q_h$	−0.063	−0.083
			Spillover: $\Delta \log \sum_{j \neq h} Q_j$	0.060	0.031
Labor Preferences			Aggregate: $\Delta \log \sum_j Q_j$	−0.022	−0.026
γ_L	4.552	Log-wage utility: Patient care			
γ_N	4.792	Log-wage utility: Non-patient care	**Labor Market: Patient Care**		
			Direct: $\Delta \log W_h^L$	−0.014	−0.028
Technology of Care			Direct: $\Delta \log L_h$	−0.078	−0.111
α	0.746	Quantity: Output elasticity	Spillover: $\Delta \log \sum_{j \neq h} L_j$	0.030	0.041
δ	0.574	Quality: Patient care intensity	Aggregate: $\Delta \log \sum_j L_j$	−0.035	−0.035
ρ	−1.262	Quality: Elasticity of substitution			
ϕ	1.063	Quality: Returns to scale	**Labor Market: Non-Patient Care**		
			Direct: $\Delta \log W_h^N$	−0.024	−0.023
Bargaining Power			Direct: $\Delta \log N_h$	−0.105	−0.098
$\bar{\kappa}_\Delta$	0.011	Change in log insurer markup	Spillover: $\Delta \log \sum_{j \neq h} N_j$	0.059	0.038
			Aggregate: $\Delta \log \sum_j N_j$	−0.033	−0.027
Calibrated Outside Shares					
s_0^Q	0.250	Patients	**Quality of Care**		
s_0^L	0.500	Labor: Patient care	Direct: $\Delta \log(\text{SR}_h)$	−0.036	−0.023
s_0^N	0.500	Labor: Non-patient care	Direct: $\Delta \log(Y_h)$		−0.027

Table 4: Estimates of Global Parameters and Fit of Simulated Moments

Notes: This table presents the estimates of the model parameters using the method of simulated moments. It also presents the goodness of fit of the simulated moments versus the targeted moments. The targeted moments are the unweighted average of the DiD estimates of the impacts of mergers over event times 2-7 from Tables 2 and 3. "SR" refers to the staffing ratio, $(L_h + N_h)/Q_h$. There are 13 simulated moments and 9 parameters to estimate. For convenience, we also report the calibrated outside shares and the simulated effect of mergers on the model-implied quality (Y). Prices are measured in thousands of dollars.

indirect inference. The model correctly predicts a decline in the staffing ratio, although the simulated reduction is slightly weaker than the targeted moment (-2.3% vs. -3.6%).

In the labor markets, the model again replicates the key qualitative patters in the data: merging firms cut wages and employment for both patient and non-patient care workers, competitors increase employment, and the total number of workers employed in the market declines. For patient care occupations, the model somewhat over-states the magnitudes of wage and employment reductions, and exactly matches the aggregate decline in employment (-3.5% vs. -3.5%). For non-patient care occupations, the model almost perfectly matches the estimated direct effects on wages (-2.3% vs. -2.4%) and employment (-9.8% vs. -10.5%), while also closely matching the aggregate reduction in employment (-2.7% vs. -3.3%). The spillover effects on competitor employment are also well-matched for both labor categories.

Overall, the model demonstrates a strong ability to replicate the complex set of direct, spillover, and aggregate effects of hospital mergers across product and labor dimensions. Furthermore, the model simulation yields two effects of mergers that cannot be produced directly from the data. First, the model yields the implied reduction in quality, Y, which is unobservable but implicitly inferred from revealed preferences in the model estimation. The estimated model implies that quality declines by nearly 3% among the merging hospitals in response to a merger—more than half as large as the percentage increase in prices, implying that prices alone substantially understate the welfare loss experienced by patients. Second, regarding the change in bargaining power over insurers, we find that the increase in the excess markup relative to the patient price, $\bar{\kappa}_\Delta$, is 0.011. In order to interpret this term, notice that we can use the estimated model to decompose the hospital price increase as,

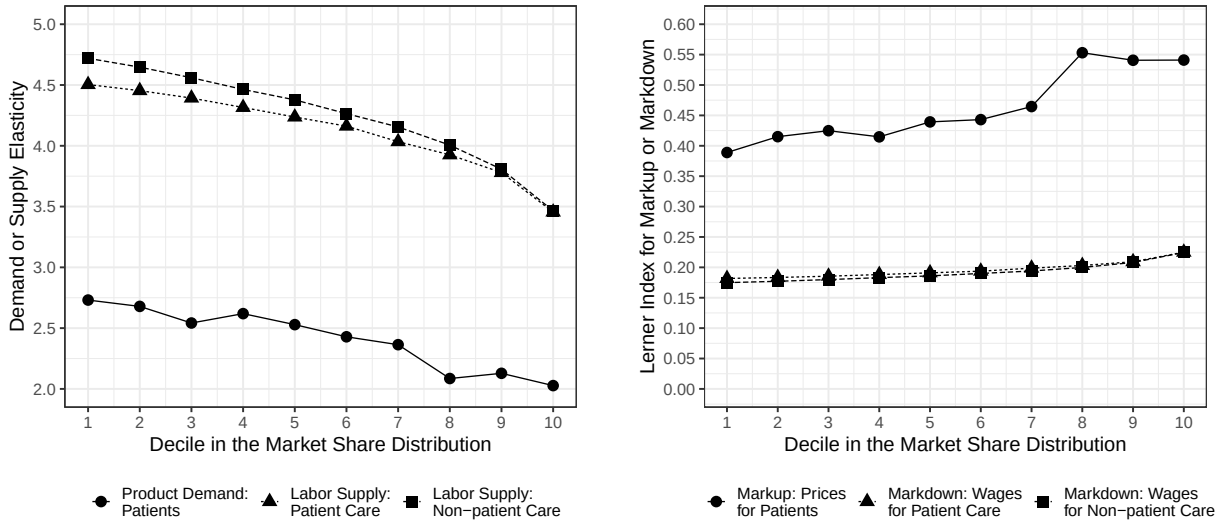
$$\underbrace{\mathbb{E}[\Delta \log P_h^{\text{hos}}]}_{5.5\%} = \underbrace{\mathbb{E}[\Delta \log P_h^{\text{pat}}]}_{2.9\%} + \underbrace{\bar{\kappa}_\Delta}_{1.1\%} + \underbrace{\text{residual}}_{1.5\%}$$

Thus, more than half of the observed price increase for hospitals is experienced by both patients and insurers, while nearly a fourth is only experienced by insurers, consistent with hospitals gaining bargaining power over insurers through mergers.

5.3 Market Power of Hospital Systems

Given our estimates of all global parameters, $\Xi \equiv (\beta_P, \beta_Y, \gamma_L, \gamma_N, \alpha, \delta, \rho, \phi, \bar{\kappa}_\Delta)$, and all firm-specific parameters, $\Lambda_{ht} \equiv (\xi^Q_{ht}, \xi^L_{ht}, \xi^N_{ht}, A_{ht}, \kappa_{ht}, \psi_{ht})$, we now possess all of the required components to characterize the market power of hospital systems.

The results are presented in Figure 9. We begin with the hospital-specific distribution of product demand and labor supply elasticities, which are presented in panel (a). As



(a) Demand and Supply Elasticities

(b) Lerner Index for Prices and Wages

Figure 9: The Empirical Distribution of Hospital Market Power

Notes: This figure presents the estimated distribution of hospital-specific product demand and labor supply elasticities (subfigure a) and of the Lerner index for price markups and wage markdowns (subfigure b). The x-axis is the decile of the market share distribution.

expected, the labor supply elasticities become smaller as the hospital moves up the market share distribution. The smallest hospitals face an average labor supply elasticity of about 4.5 for patient care labor and 4.7 for non-patient care labor. The largest hospitals face much less elastic labor supply, with a labor supply elasticity below 3.5 for both occupational categories. Hospital-specific product demand is substantially more inelastic.[31] The smallest hospitals face a product demand elasticity of about 2.7, while the largest hospitals face a product demand elasticity of about 2.0.[32] For comparison, Gaynor and Vogt (2003) find a hospital-specific patient demand elasticity of about 5.0.

In panel (b) of Figure 9, we present the estimates of markups on prices charged to patients and markdowns on wages for patient and non-patient care workers. For comparability, we convert both measures into a Lerner index. The Lerner index for prices is $1-1/\text{markup}_{ht}$ and the Lerner index for wages is $1-\text{markdown}_{ht}^{E}, E=L,N$. Regarding markups, we find that the Lerner index ranges from almost 0.40 for the smallest hospitals to 0.55 for the largest hospitals. This indicates that prices exceed marginal costs by at least 40% of the price and

[31]Note that this is the elasticity of patient demand with respect to the patient price, P_{ht}^{pat}, rather than the total price received by the hospital, P_{ht}^{hos}.

[32]Due to unobserved heterogeneity in marginal costs of patient treatment and quality of care, the product demand elasticity need not be monotonic in market share. Indeed, there is some evidence that the patient demand elasticity is not strictly decreasing in product market share.

at most 55% of the price across the empirical distribution. Regarding markdowns, we find that the Lerner index ranges from about 0.17 for the smallest hospitals to 0.23 for the largest hospitals, with nearly identical values for patient care and non-patient care workers. This indicates that wages fall short of marginal revenues by at least 17% of marginal revenues and at most 23% of marginal revenues across the empirical distribution.

6 Conclusions

This paper analyzed, both theoretically and empirically, how hospitals leverage their simultaneous gains in market share over patients and workers through consolidation. I developed a framework for merger evaluation in which firms exploit both oligopoly power in the product market and oligopsony power in the labor market. The model allowed firms to endogenously choose product quality, and showed that optimal quality provision depends on labor market power. From the model, I predicted the effects of hospital mergers on patient and worker outcomes. Empirically, I combined several administrative panel data sources on the universe of hospitals over 1996-2022. Utilizing a propensity-score matched difference-in-differences design, I estimated the extent to which oligopoly and oligopsony power are exploited after merger events. Confirming the predictions of the model, I found that merging hospitals not only raise prices and lower wages after consolidating, but also reduce the number of patients treated, the number of workers employed, and the quality of care provided to each patient. Patients face worse staffing ratios, report lower satisfaction with their care, and experience greater mortality. Merger spillovers lead competitors in the local market to cut wages even as they increase employment and increase patient volume. On aggregate, mergers increase the number of patients who fail to receive treatment and the number of workers who fail to attain employment. Finally, I used the estimated merger effects to recover the structural parameters, then used the estimated model to measure market power.

References

Ahmed, A. M., K. Kadakia, A. Ahmed, B. Shultz, and X. Li (2022). Trends in Labor Unionization Among US Health Care Workers, 2009-2021. *JAMA 328*(24), 2404–2411.

Angerhofer, T., A. Collard-Wexler, and M. Weinberg (2025). Monopsony and the Countervailing Power of Unions: Evidence from K-12 Teachers. *Unpublished manuscript*.

Arnold, J., A. Gupta, T. Liu, and A. Olssen (2025). Targeting Efficiencies? Reducing Costs through Hospital Consolidation. *Unpublished manuscript*.

Atkeson, A. and A. Burstein (2008). Pricing-to-Market, Trade Costs, and International Relative Prices. *American Economic Review 98*(5), 1998–2031.

Autor, D., D. Dorn, L. F. Katz, C. Patterson, and J. van Reenen (2020). The Fall of the Labor Share and the Rise of Superstar Firms. *The Quarterly Journal of Economics 135*(2), 645–709.

Beaulieu, N. D., L. S. Dafny, B. E. Landon, J. Dalton, I. Kuye, and J. M. McWilliams (2020). Changes in Quality of Care After Hospital Mergers and Acquisitions. *New England Journal of Medicine 382*(1), 51–59.

Berger, D., T. Hasenzagl, K. Herkenhoff, S. Mongey, and E. A. Posner (2025). Merger Guidelines for the Labor Market. *Journal of Monetary Economics*.

Berger, D., K. Herkenhoff, and S. Mongey (2022). Labor Market Power. *American Economic Review 112*(4), 1147–1193.

Björnerstedt, J. and F. Verboven (2016). Does Merger Simulation Work? Evidence from the Swedish Analgesics Market. *American Economic Journal: Applied Economics 8*(3), 125–164.

Brand, K., C. Garmon, and T. Rosenbaum (2023). In the Shadow of Antitrust Enforcement: Price Effects of Hospital Mergers from 2009 to 2016. *The Journal of Law and Economics 66*(4), 639–669.

Brot, Z., Z. Cooper, S. V. Craig, and L. Klarnet (2024). Is There Too Little Antitrust Enforcement in the US Hospital Sector? *American Economic Review: Insights 6*(4), 526–542.

Callaway, B. and P. H. Sant'Anna (2021). Difference-in-Differences with Multiple Time Periods. *Journal of Econometrics 225*(2), 200–230.

Card, D., A. R. Cardoso, J. Heining, and P. Kline (2018). Firms and Labor Market Inequality: Evidence and Some Theory. *Journal of Labor Economics 36*(S1), S13–S70.

Chan, M., K. Kroft, E. Mattana, and I. Mourifié (2024). An Empirical Framework for Matching with Imperfect Competition. *Unpublished manuscript*.

Cooper, Z., S. V. Craig, M. Gaynor, and J. Van Reenen (2019). The Price Ain't Right? Hospital Prices and Health Spending on the Privately Insured. *The Quarterly Journal of Economics 134*(1), 51–107.

Cooper, Z., J. J. Doyle, J. A. Graves, and J. Gruber (2022). Do Higher-Priced Hospitals Deliver Higher-Quality Care? *Unpublished manuscript*.

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

Dafny, L., K. Ho, and R. S. Lee (2019). The Price Effects of Cross-Market Mergers: Theory and Evidence from the Hospital Industry. *The RAND Journal of Economics 50*(2), 286–325.

Derenoncourt, E. and D. Weil (2025). Voluntary Minimum Wages: The Local Labor Market Effects of National Retailer Policies. *Quarterly Journal of Economics*.

Dingel, J., J. Gottlieb, M. Lozinski, and P. Mourot (2024). Market Size and Trade in Medical Services. *Unpublished manuscript*.

Finkelstein, A., M. Gentzkow, and H. Williams (2021). Place-Based Drivers of Mortality: Evidence from Migration. *American Economic Review 111*(8), 2697–2735.

Gandhi, A. and J.-F. Houde (2020). Measuring Substitution Patterns in Differentiated-Products Industries. *Unpublished manuscript*.

Gaynor, M., A. Sacarny, R. Sadun, C. Syverson, and S. Venkatesh (2023). The Anatomy of a Hospital System Merger: The Patient Did Not Respond Well to Treatment. *Review of Economics and Statistics*, 1–28.

Gaynor, M. and W. B. Vogt (2003). Competition among Hospitals. *The RAND Journal of Economics 34*(4), 764–785.

Gottlieb, J. D., N. Mahoney, K. Rinz, and V. Udalova (2025). The Rise of Healthcare Jobs. *Unpublished manuscript*.

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.

Hackmann, M. B. (2019). Incentivizing Better Quality of Care: The Role of Medicaid and Competition in the Nursing Home Industry. *American Economic Review 109*(5), 1684–1716.

Hemphill, C. S. and N. L. Rose (2018). Mergers that Harm Sellers. *The Yale Law Journal 127*(7), 2078–2203.

Ho, K. and R. S. Lee (2017). Insurer Competition in Health Care Markets. *Econometrica 85*(2), 379–417.

Hosken, D., M. Larson-Koester, and C. Taragin (2024). Labor and Product Market Effects of Mergers. *Unpublished manuscript*.

Houde, J.-F. (2012). Spatial Differentiation and Vertical Mergers in Retail Markets for Gasoline. *American Economic Review 102*(5), 2147–2182.

Kane, C. K. (2025). Physician Practice Characteristics in 2024. Policy Research Perspective 2025, American Medical Association.

Kroft, K., Y. Luo, M. Mogstad, and B. Setzler (2025). Imperfect Competition and Rents in Labor and Product Markets: The Case of the Construction Industry. *American Economic Review*.

Lamadon, T., M. Mogstad, and B. Setzler (2022). Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market. *American Economic Review 112*(1), 169–212.

Magnolfi, L., D. Quint, C. Sullivan, and S. Waldfogel (2022). Differentiated-Products Cournot Attributes Higher Markups than Bertrand–Nash. *Economics Letters 219*, 110804.

Marinescu, I. and H. J. Hovenkamp (2019). Anticompetitive Mergers in Labor Markets. *Indiana Law Journal 94*(3).

McFadden, D. (1999). Computing Willingness-to-Pay in Random Utility Models. In J. C. Moore, R. G. Riezman, and J. R. Melvin (Eds.), *Trade, Theory and Econometrics: Essays in Honour of John S. Chipman*, pp. 253–274. London: Routledge.

Miller, N. H. and M. C. Weinberg (2017). Understanding the Price Effects of the MillerCoors Joint Venture. *Econometrica 85*(6), 1763–1791.

Naidu, S., E. A. Posner, and G. Weyl (2018). Antitrust Remedies for Labor Market Power. *Harvard Law Review 132*(2), 536–601.

Nocke, V. and M. D. Whinston (2022). Concentration Thresholds for Horizontal Mergers. *American Economic Review 112*(6), 1915–1948.

OECD (2023). *Health at a Glance 2023: OECD Indicators*. Paris: OECD Publishing.

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

Roussille, N. and B. Scuderi (2025). Bidding for Talent: A Test of Conduct in a High-Wage Labor Market. *Unpublished manuscript*.

Rubens, M. (2023). Market Structure, Oligopsony Power, and Productivity. *American Economic Review 113*(9), 2382–2410.

Rubens, M., B. Setzler, and C. Yeh (2025). Conduct and Collusion in the US Labor Market. *Unpublished manuscript*.

Schubert, G., A. Stansbury, and B. Taska (2024). Employer Concentration and Outside Options. *Unpublished manuscript*.

Sharma, G. (2023). Monopsony and Gender. *Unpublished manuscript*.

Small, K. A. and H. S. Rosen (1981). Applied Welfare Economics with Discrete Choice Models. *Econometrica*, 105–130.

U.S. Bureau of Labor Statistics (2023, August). Why Hospitals Matter to the BLS Occupational Requirements Survey. Factsheet, U.S. Department of Labor.

US DOJ and FTC (2023). US DOJ and FTC Merger Guidelines. Technical report, United States Government.

Wollmann, T. G. (2024). How to Get Away with Merger: Stealth Consolidation and Its Real Effects on US Dialysis. *Unpublished manuscript*.

Yeh, C., C. Macaluso, and B. Hershbein (2022). Monopsony in the US Labor Market. *American Economic Review 112*(7), 2099–2138.

A Mathematical Details

A.1 Properties of Product Demand and Labor Supply

Details on residual product demand Recall,

$$s_{ht}^Q = \frac{\exp\left(-\beta_P P_{ht} + \beta_Y Y_{ht} + \xi_{ht}^Q\right)}{\mathcal{P}}, \quad \mathcal{P} \equiv 1 + \sum_{h'} \exp\left(-\beta_P P_{h't} + \beta_Y Y_{h't} + \xi_{h't}^Q\right).$$

Taking the log of both sides and rearranging, we can express the inverse product demand curve as,

$$P_{ht} = \frac{1}{\beta_P}\left(\beta_Y Y_{ht} + \xi_{ht} + \log s_{0t}^Q - \log s_{ht}^Q\right),$$

where we denote the outside option share by $s_{0t}^Q \equiv 1 - \sum_h s_{ht}^Q$ and use that $s_{0t} = 1/\mathcal{P}$.

Imposing Cournot conduct, such that competitors' quantity variables Q_{jt}, Y_{jt}, and L_{jt} are perceived as fixed for all $j \neq h$, the effect of a change in s_{ht} on inverse residual demand is,

$$\frac{\partial P_{ht}}{\partial s_{ht}^Q} = \frac{-1}{\beta_P}\left(\frac{s_{ht}^Q + s_{0t}^Q}{s_{0t}^Q s_{ht}^Q}\right) < 0.$$

It follows that the inverse of the residual product demand elasticity is,

$$1/\theta_{ht}^Q \equiv \frac{s_{ht}^Q}{P_{ht}}\frac{\partial P_{ht}}{\partial s_{ht}^Q} = \frac{-1}{\beta_P P_{ht}}\left(\frac{s_{ht}^Q}{s_{0t}^Q} + 1\right) \implies \theta_{ht}^Q = -\beta_P P_{ht}\frac{s_{0t}^Q}{s_{ht}^Q + s_{0t}^Q} < 0.$$

Taking the partial derivative of the inverse of the residual product demand elasticity with respect to s_{ht}^Q,

$$\frac{\partial(1/\theta_{ht}^Q)}{\partial s_{ht}^Q} = \frac{1}{\beta_P P_{ht}^2}\underbrace{\frac{\partial P_{ht}}{\partial s_{ht}^Q}}_{(-)}\left(\frac{s_{ht}^Q}{s_{0t}^Q} + 1\right) + \frac{-1}{\beta_P P_{ht}}\underbrace{\frac{\partial}{\partial s_{ht}^Q}\frac{s_{ht}^Q}{s_{0t}^Q}}_{(+)} < 0 \implies \frac{\partial \theta_{ht}^Q}{\partial s_{ht}^Q} > 0,$$

which implies the residual product demand elasticity becomes more inelastic as market share increases. The inverse cross-price effect of the market share at h on the price of competitor j is,

$$\frac{\partial P_{jt}}{\partial s_{ht}^Q} = \frac{\partial}{\partial s_{ht}^Q}\left(\frac{1}{\beta_P}\left(\beta_Y Y_{jt} + \xi_{jt} + \log s_{0t}^Q - \log s_{jt}^Q\right)\right) = \frac{1}{\beta_P}\frac{\partial \log s_{0t}^Q}{\partial s_{ht}^Q} = \frac{-1}{\beta_P s_{0t}^Q}.$$

Thus, we have,

$$\frac{\partial s_{jt}^{Q}}{\partial P_{ht}} = -\beta_P s_{0t}^{Q} < 0.$$

Similarly, the effect of a price increase on the outside option is,

$$\frac{\partial P_{ht}}{\partial s_{0t}^{Q}} = \frac{-1}{\beta_P s_{0t}^{Q}} \implies \frac{\partial s_{0t}^{Q}}{\partial P_{ht}} = -\beta_P s_{0t}^{Q}.$$

Details on residual labor supply Recall that,

$$s_{ht}^{E} = \frac{\exp\left(\gamma_E \log\left(W_{ht}^{E}\right) + \xi_{ht}^{E}\right)}{\mathcal{W}}, \quad \mathcal{W}^{E} \equiv 1 + \sum_{h'} \exp\left(\gamma_E \log\left(W_{h't}^{E}\right) + \xi_{h't}^{E}\right).$$

Taking the log of both sides and rearranging, we can express the inverse labor supply curve as,

$$\log\left(W_{ht}^{E}\right) = \frac{1}{\gamma_E}\left(\log s_{ht}^{E} - \log s_{0t}^{E} - \xi_{ht}^{E}\right),$$

where we denote the outside option share by $s_{0t}^{E} \equiv 1 - \sum_h s_{ht}^{E}$ and use that $s_{0t}^{E} = 1/\mathcal{W}^{E}$.

Imposing Cournot conduct, such that competitors' quantity variables Q_{jt}, Y_{jt}, and L_{jt} are perceived as fixed for all $j \neq h$, the effect of a change in s_{ht}^{E} on inverse residual labor supply is,

$$\frac{\partial W_{ht}^{E}}{\partial s_{ht}^{E}} = \frac{W_{ht}^{E}}{\gamma_E}\left(\frac{1}{s_{ht}^{E}} + \frac{1}{s_{0t}^{E}}\right),$$

which implies that the residual labor supply effect and elasticity are,

$$\frac{\partial s_{ht}^{E}}{\partial W_{ht}^{E}} = \frac{\gamma_E}{W_{ht}^{E}}\left(\frac{s_{0t}^{E} s_{ht}^{E}}{s_{0t}^{E} + s_{ht}^{E}}\right) > 0, \quad \theta_{ht}^{E} = \frac{W_{ht}^{E}}{s_{ht}^{E}}\frac{\partial s_{ht}^{E}}{\partial W_{ht}^{E}} = \gamma_E\left(\frac{s_{0t}^{E}}{s_{ht}^{E} + s_{0t}^{E}}\right) > 0.$$

Taking the partial derivative of the residual labor supply elasticity with respect to the wage,

$$\frac{\partial \theta_{ht}^{E}}{\partial W_{ht}^{E}} = \gamma_E \frac{\left(s_{ht}^{E} + s_{0t}^{E}\right)\frac{\partial s_{0t}^{E}}{\partial W_{ht}^{E}} - s_{0t}^{E}\frac{\partial\left(s_{ht}^{E} + s_{0t}^{E}\right)}{\partial W_{ht}^{E}}}{\left(s_{ht}^{E} + s_{0t}^{E}\right)^2} = \gamma_E \frac{s_{ht}^{E}\overbrace{\frac{\partial s_{0t}^{E}}{\partial W_{ht}^{E}}}^{(-)} - s_{0t}^{E}\overbrace{\frac{\partial s_{ht}^{E}}{\partial W_{ht}^{E}}}^{(+)}}{\left(s_{ht}^{E} + s_{0t}^{E}\right)^2} < 0,$$

which implies supply becomes more inelastic as wages increase. Given quality, higher wages imply higher employment, so supply is more inelastic in higher employment firms. The inverse cross-wage effect of the wage at h on competitor j is,

$$\frac{1}{W_{jt}^{E}}\frac{\partial W_{jt}^{E}}{\partial s_{ht}^{E}} = \frac{1}{\gamma_E s_{0t}^{E}} \implies \frac{\partial W_{jt}^{E}}{\partial s_{ht}^{E}} = \frac{W_{jt}^{E}}{\gamma_E s_{0t}^{E}}.$$

A.2 Model Proofs for Subsection 2.2

Details on Lemma 1 and Proposition 1 For the single-establishment firm with pre-determined quality, the first-order condition is,

$$\underbrace{\frac{\partial P_{ht}}{\partial L_{ht}}}_{\frac{\partial P_{ht}}{\partial Q_{ht}}\frac{\partial Q_{ht}}{\partial L_{ht}}} Q_{ht} + P_{ht}\frac{\partial Q_{ht}}{\partial L_{ht}} - \frac{\partial W_{ht}^L}{\partial L_{ht}}L_{ht} - W_{ht}^L = 0.$$

Recall that $\frac{\partial Q_{ht}}{\partial L_{ht}} = \text{MP}_{ht}^L$. Rearranging,

$$\left(\frac{\partial P_{ht}}{\partial Q_{ht}}\frac{Q_{ht}}{P_{ht}} + 1\right) P_{ht}\text{MP}_{ht}^L = \left(\frac{\partial W_{ht}^L}{\partial L_{ht}}\frac{L_{ht}}{W_{ht}^L} + 1\right) W_{ht}^L,$$

which completes the proof of Lemma 1. Proposition 1(a) follows from rearranging Lemma 1. Proposition 1(b) follows from the own-price product demand and own-wage labor supply becoming more inelastic as the price and wage increase, respectively, all else equal, which were shown above.

Details on Lemma 2 For the two-establishment firm with pre-determined quality, the first-order condition with respect to L_{ht} is,

$$\text{MR}_{ht}^L + \underbrace{\frac{\partial P_{gt}}{\partial L_{ht}}}_{\frac{\partial P_{gt}}{\partial Q_{ht}}\frac{\partial Q_{ht}}{\partial L_{ht}}} Q_{gt} + P_{gt} \underbrace{\frac{\partial Q_{gt}}{\partial L_{ht}}}_{0\text{ (Cournot)}} - \text{MC}_{ht}^L - \frac{\partial W_{gt}^L}{\partial L_{ht}}L_{gt} - W_{gt}^L \underbrace{\frac{\partial L_{gt}}{\partial L_{ht}}}_{0\text{ (Cournot)}} = 0,$$

where we used that, by definition of Cournot competition, competitors are perceived to respond with price and wage changes, not with changes in quantities of patients and labor.

Details on Propositions 2(a-b) We can use the demand and labor supply parameterizations to express Lemma 2 as,

$$\underbrace{\left(P_{ht} - \frac{1}{\beta_P}\left[1 + \frac{s_{ht}^Q}{s_{0t}^Q}\right]\right)\text{MP}_{ht}^L}_{\text{MR}_{ht}^L} + \underbrace{\frac{-s_{gt}^Q}{\beta_P s_{0t}^Q}\text{MP}_{ht}^L}_{\text{product diversion}} = \underbrace{W_{ht}^L + \frac{W_{ht}^L}{\gamma_L}\left[1 + \frac{s_{ht}^L}{s_{0t}^L}\right]}_{\text{MC}_{ht}^L} + \underbrace{\frac{W_{gt}^L s_{gt}^L}{\gamma_L s_{0t}^L}}_{\text{labor diversion}}.$$

Before the merger (i.e. without the diversion effects),

$$P_{ht}^{\text{before}} = \frac{1}{\beta_P}\left[1 + \frac{s_{ht}^Q}{s_{0t}^Q}\right] + \frac{W_{ht}^L}{\text{MP}_{ht}^L}\left(1 + \frac{1}{\gamma_L}\left[1 + \frac{s_{ht}^L}{s_{0t}^L}\right]\right),$$

and after the merger (by including the diversion effects),

$$P_{ht}^{\text{after}} = P_{ht}^{\text{before}} + \left(\underbrace{\frac{s_{gt}^{Q}}{\beta_P s_{0t}^{Q}}}_{\text{product diversion}} + \frac{1}{\text{MP}_{ht}^{L}} \underbrace{\frac{W_{gt}^{L} s_{gt}^{L}}{\gamma_L s_{0t}^{L}}}_{\text{labor diversion}} \right) > P_{ht}^{\text{before}}.$$

Similarly, the wage before the merger is,

$$W_{ht}^{L,\text{before}} = \left(1 + \frac{1}{\gamma_L}\left[1 + \frac{s_{ht}^{L}}{s_{0t}^{L}}\right]\right)^{-1} \text{MP}_{ht}^{L} \left(P_{ht} - \frac{1}{\beta_P}\left[1 + \frac{s_{ht}^{Q}}{s_{0t}^{Q}}\right]\right),$$

and the wage after the merger is,

$$W_{ht}^{L,\text{after}} = W_{ht}^{L,\text{before}} - \left(1 + \frac{1}{\gamma_L}\left[1 + \frac{s_{ht}^{L}}{s_{0t}^{L}}\right]\right)^{-1} \left(\text{MP}_{ht}^{L} \underbrace{\frac{s_{gt}^{Q}}{\beta_P s_{0t}^{Q}}}_{\substack{\text{product} \\ \text{diversion}}} + \underbrace{\frac{W_{gt}^{L} s_{gt}^{L}}{\gamma_L s_{0t}^{L}}}_{\substack{\text{labor} \\ \text{diversion}}} \right).$$

Details on Propositions 2(c-d) We utilize a proof by contradiction. Suppose first that $s_{ht}^{Q,\text{after}} + s_{gt}^{Q,\text{after}} \geq s_{ht}^{Q,\text{before}} + s_{gt}^{Q,\text{before}}$. Recall that we have already shown that $P_{ht}^{\text{before}} < P_{ht}^{\text{after}}$ and $P_{gt}^{\text{before}} < P_{gt}^{\text{after}}$. From the demand curve, we can write,

$$1 < \frac{s_{ht}^{Q,\text{after}} + s_{gt}^{Q,\text{after}}}{s_{ht}^{Q,\text{before}} + s_{gt}^{Q,\text{before}}} = \frac{s_{0t}^{Q,\text{after}}}{s_{0t}^{Q,\text{before}}} \underbrace{\frac{\exp\left(-\beta_P P_{ht}^{\text{after}} + \xi_{ht}^{Q}\right) + \exp\left(-\beta_P P_{gt}^{\text{after}} + \xi_{gt}^{Q}\right)}{\exp\left(-\beta_P P_{ht}^{\text{before}} + \xi_{ht}^{Q}\right) + \exp\left(-\beta_P P_{gt}^{\text{before}} + \xi_{gt}^{Q}\right)}}_{<1},$$

which implies that $s_{0t}^{Q,\text{after}} > s_{0t}^{Q,\text{before}}$. However, holding quantities fixed at all competitor firms, the difference in outside shares $s_{0t}^{Q} = 1 - \sum_{j'} s_{j't}^{Q}$ due to the merger is, $s_{0t}^{Q,\text{after}} - s_{0t}^{Q,\text{before}} = (s_{ht}^{Q,\text{before}} + s_{gt}^{Q,\text{before}}) - (s_{ht}^{Q,\text{after}} + s_{gt}^{Q,\text{after}})$ so $s_{0t}^{Q,\text{after}} > s_{0t}^{Q,\text{before}}$ is only possible if $(s_{ht}^{Q,\text{before}} + s_{gt}^{Q,\text{before}}) > (s_{ht}^{Q,\text{after}} + s_{gt}^{Q,\text{after}})$, which is a contradiction. Thus, it must be the case that $s_{ht}^{Q,\text{after}} + s_{gt}^{Q,\text{after}} < s_{ht}^{Q,\text{before}} + s_{gt}^{Q,\text{before}}$.

Next, suppose that $s_{ht}^{L,\text{after}} + s_{gt}^{L,\text{after}} \geq s_{ht}^{L,\text{before}} + s_{gt}^{L,\text{before}}$. Recall that we have already

shown that $W_{ht}^{\text{after}} < W_{ht}^{\text{before}}$ and $W_{gt}^{\text{after}} < W_{gt}^{\text{before}}$. From the supply curve, we can write,

$$1 < \frac{s_{ht}^{L,\text{after}} + s_{gt}^{L,\text{after}}}{s_{ht}^{L,\text{before}} + s_{gt}^{L,\text{before}}} = \frac{s_{0t}^{L,\text{after}}}{s_{0t}^{L,\text{before}}} \underbrace{\frac{\exp\left(\gamma_L W_{ht}^{L,\text{after}} + \xi_{ht}^L\right) + \exp\left(\gamma_L W_{gt}^{L,\text{after}} + \xi_{gt}^L\right)}{\exp\left(\gamma_L W_{ht}^{L,\text{before}} + \xi_{ht}^L\right) + \exp\left(\gamma_L W_{gt}^{L,\text{before}} + \xi_{gt}^L\right)}}_{<1},$$

which implies that $s_{0t}^{L,\text{after}} > s_{0t}^{L,\text{before}}$, leading to the same contradiction we saw previously for product market shares. Thus, it must be the case that $s_{ht}^{L,\text{after}} + s_{gt}^{L,\text{after}} < s_{ht}^{L,\text{before}} + s_{gt}^{L,\text{before}}$.

Details on Propositions 2(e-f) The markup is $P_{ht}/(\text{MC}_{ht}^L/\text{MP}_{ht}^L)$. We know that P_{ht} increases. Furthermore, MP_{ht}^L is monotonically decreasing in L_{ht}, while MC_{ht}^L is monotonically increasing in L_{ht}. Since L_{ht} decreases, we have that the numerator rises and the denominator falls, so the markup rises. The markdown is $W_{ht}^L/\text{MR}_{ht}^L$. We know that W_{ht} decreases, and MR_{ht}^L is monotonically decreasing in L_{ht}. Since L_{ht} decreases, we have that the numerator falls and the denominator rises, so the markdown strengthens.

Details on Propositions 2(g) The proof of Proposition 2(a) above demonstrates that $P_{ht}^{\text{after}} - P_{ht}^{\text{before}}$ is increasing in the size of hospital g (as measured by s_{gt}^Q or s_{gt}^L). Furthermore, the proof of Proposition 2(b) above demonstrates that $W_{ht}^{L,\text{after}}/W_{ht}^{L,\text{before}}$ is decreasing in the size of hospital g. Greater price and wage responses lead to greater quantity and employment responses from the proof of Propositions 2(c-d). By the same reasoning as the proof of Propositions 2(e-f), the markup and markdown strengthen when the employment response is stronger.

Details on Proposition 3(a-b) Since L_{jt} is determined by the intersection of the downward-sloping curve MR_{jt}^L and the upward-sloping curve MC_{jt}^L, and Q_{jt} is monotonic in L_{jt}, it suffices to show that MR_{jt}^L is strictly greater and MC_{jt}^L is strictly lesser for all choices of L_{jt} after the merger.

Marginal revenue: Recall that $MR_{jt}^L = \left(1 + 1/\theta_{jt}^Q\right) P_{jt} MP_{jt}^L$, where MP_{jt}^L is unaffected by the merger, given L_{jt}. Since $s_{0t}^{Q,\text{after}} > s_{0t}^{Q,\text{before}}$, it follows that $P_{jt}^{\text{before}} < P_{jt}^{\text{after}}$ for any $s_{jt}^Q = Q_{jt}/\bar{Q}_{mt}$ to satisfy the inverse product demand curve to firm j. In other words, to hold output constant after the merging parties raise their prices, competitor j must increase its price. Furthermore, $\theta_{jt}^Q = -\beta_P P_{jt} \frac{s_{0t}^Q}{s_{jt}^Q + s_{0t}^Q}$ becomes more negative, both due to the increase in P_{jt} and the increase in s_{0t}^Q, given Q_{jt}, which implies that $\left(1 + 1/\theta_{jt}^Q\right)$ becomes greater. Thus, the marginal revenue curve for competitor j strictly increases in response to the merger.

Marginal cost: Recall that $MC_{jt}^L = W_{jt}\left(1 + 1/\theta_{jt}^L\right)$. Since $s_{0t}^{L,\text{after}} > s_{0t}^{L,\text{before}}$, it follows that $W_{jt}^{\text{before}} > W_{jt}^{\text{after}}$ for any $s_{jt}^L = L_{jt}/\bar{L}_{mt}$ to satisfy the inverse labor supply curve to firm j. In other words, to employ a constant amount of labor after the merging parties reduce their wages, competitor j must offer a lower wage. Furthermore, $\theta_{jt}^L = \gamma_L \frac{s_{0t}^L}{s_{jt}^L + s_{0t}^L}$ becomes greater, due to the increase in s_{0t}^L, given L_{jt}, which implies that $\left(1 + 1/\theta_{jt}^L\right)$ becomes smaller. Thus, the marginal cost curve for competitor j strictly decreases in response to the merger.

Details on Proposition 3(c-d) Suppose $s_{0t}^{Q,\text{after}} = s_{0t}^{Q,\text{before}}$ and $s_{0t}^{L,\text{after}} = s_{0t}^{L,\text{before}}$. We will show that this leads to a contradiction.

For any non-merging firm j, recall that the FOC is,

$$\left(P_{jt} - \frac{1}{\beta_P}\left(1 + \frac{s_{jt}^Q}{\bar{s}_{0t}^Q}\right)\right) MP_{jt}^L = W_{jt}^L\left(1 + \frac{1}{\gamma_L}\left(1 + \frac{s_{jt}^L}{\bar{s}_{0t}^L}\right)\right).$$

Given that $\bar{s}_{0t}^Q$ and $\bar{s}_{0t}^L$ are constant, the FOC is otherwise unchanged and must thus have the same solution, $L_{jt}^{\text{after}} = L_{jt}^{\text{before}}$ and $Q_{jt}^{\text{after}} = Q_{jt}^{\text{before}}$. Since all market shares (including the outside option) must sum to one, and shares remain unchanged for each non-merging j as well as the outside shares, it follows that the combined market share of the merging firms must be unchanged: $s_{ht}^{Q,\text{after}} + s_{gt}^{Q,\text{after}} = s_{ht}^{Q,\text{before}} + s_{gt}^{Q,\text{before}}$, and similarly for the labor market shares. This contradicts our earlier proof that the merging firms must reduce output and labor in response to a merger because the MRPL strictly decreases and the MCL strictly increases (due to labor and product market diversion effects).

A similar contradiction holds if we assume the outside shares decrease in response to the merger: the FOCs imply that each non-merging j must decrease output and employment, so the sum of market shares can only equal one if the merging parties strictly increase their combined market share.

Details on Proposition 3(e) First, we show that if merging firm g is larger, the increase in the outside shares, s_{0t}^Q and s_{0t}^L, is greater. Since the shares must add to one, we have that the changes induced by the merger must satisfy $\Delta s_{0t}^Q + \Delta s_{Ht}^Q + \sum_{j \neq H} \Delta s_{jt}^Q = 0$. Rearranging this gives $-\Delta s_{Ht}^Q = \Delta s_{0t}^Q + \sum_{j \neq H} \Delta s_{jt}^Q$. It follows that $-\frac{d\Delta s_{Ht}^Q}{d\Delta s_{0t}^Q} = 1 + \sum_{j \neq H} \frac{d\Delta s_{jt}^Q}{d\Delta s_{0t}^Q}$. Since $\frac{d\Delta s_{jt}^Q}{d\Delta s_{0t}^Q} > 0$, it follows that $\frac{d\Delta s_{Ht}^Q}{d\Delta s_{0t}^Q} < -1$ and thus $\frac{d\Delta s_{0t}^Q}{d\Delta s_{Ht}^Q} \in (-1, 0)$. Thus, for a larger decrease in Δs_{Ht}^Q induced by the merger (i.e., Δs_{Ht}^Q becomes more negative, meaning $-\Delta s_{Ht}^Q$, the positive reduction in the merged firm's share, is greater because firm g is larger), there is a larger increase in s_{0t}^Q, which is what we wanted to show. An analogous argument applies to s_{0t}^L.

Second, we show that if merging firm g is larger, the response in each competitor j is also larger. Recall that the increases in L_{jt} and Q_{jt} for these firms are driven by shifts in their MR^L_{jt} and MC^L_{jt} curves, which are induced by the increases in s^Q_{0t} and s^L_{0t}. As established in the first part of this proof for Proposition 3(c), the increases in the outside shares, Δs^Q_{0t} and Δs^L_{0t}, are indeed greater when firm g is larger. Given that the magnitudes of the MR^L_{jt} and MC^L_{jt} shifts for non-merging firms j are monotonically increasing with the magnitudes of these changes in s^Q_{0t} and s^L_{0t} respectively (from the mechanisms detailed in Proposition 3(a)), it follows that a larger g induces larger employment and quantity effects in competitor j.

A.3 Model Proofs for Subsection 2.3

Details on Lemma 3 The first-order condition with respect to non-patient care labor N_{ht} is:

$$\frac{\partial P_{ht}}{\partial Y_{ht}} \frac{\partial Y_{ht}}{\partial N_{ht}} Q_{ht} - \frac{\partial (W^N_{ht} N_{ht})}{\partial N_{ht}} = 0.$$

From the inverse product demand $P_{ht} = \frac{1}{\beta_P}(\beta_Y Y_{ht} + \xi_{ht} + \log s^Q_{0t} - \log s^Q_{ht})$, we have $\frac{\partial P_{ht}}{\partial Y_{ht}} = \beta_Y/\beta_P$. The marginal cost of non-patient care labor is $MC^N_{ht} \equiv W^N_{ht}(1 + 1/\theta^N_{ht})$, where θ^N_{ht} is the residual labor supply elasticity for non-patient care workers. Thus, the first-order condition becomes:

$$\underbrace{W^N_{ht}(1 + 1/\theta^N_{ht})}_{\equiv MC^N_{ht}} = Q_{ht} \frac{F_N(L_{ht}, N_{ht} - \psi_{ht})}{Q_{ht}} \frac{\beta_Y}{\beta_P}.$$

The first-order condition with respect to patient care labor L_{ht} is:

$$\left(\frac{\partial P_{ht}}{\partial Q_{ht}} \frac{\partial Q_{ht}}{\partial L_{ht}} + \frac{\partial P_{ht}}{\partial Y_{ht}} \frac{\partial Y_{ht}}{\partial L_{ht}} \right) Q_{ht} + P_{ht} \frac{\partial Q_{ht}}{\partial L_{ht}} - \frac{d(W^L_{ht} L_{ht})}{dL_{ht}} = 0,$$

Rearranging,

$$\underbrace{\left(\left(\frac{\partial P_{ht}}{\partial Q_{ht}} \frac{Q_{ht}}{P_{ht}} \right) + 1 \right) P_{ht} MP^L_{ht}}_{\equiv MR^L_{ht}} + \frac{\partial P_{ht}}{\partial Y_{ht}} \frac{\partial Y_{ht}}{\partial L_{ht}} Q_{ht} = \underbrace{W^L_{ht}(1 + 1/\theta^L_{ht})}_{\equiv MC^L_{ht}},$$

where $\frac{\partial P_{ht}}{\partial Y_{ht}} = \beta_Y/\beta_P$. Expanding $\frac{\partial Y_{ht}}{\partial L_{ht}}$ completes the proof.

Details on Proposition 4(a-b) Many of the proofs are essentially identical when allowing for endogenous quality:

Rearranging Lemma 3, before the merger, the wage is,

$$W_{ht}^{L,\text{before}} = \left(1 + \frac{1}{\gamma_L}\left[1 + \frac{s_{ht}^L}{s_{0t}^L}\right]\right)^{-1}\left[\text{MP}_{ht}^L\left(P_{ht} - \frac{1}{\beta_P}\left[1 + \frac{s_{ht}^Q}{s_{0t}^Q}\right]\right) + \frac{\beta_Y}{\beta_P}\left(1 - \frac{L_{ht}+N_{ht}}{L_{ht}}\theta_{ht}^T\right)\right].$$

Rearranging Lemma 4, after the merger, the wage is,

$$W_{ht}^{L,\text{after}} = W_{ht}^{L,\text{before}} - \left(1 + \frac{1}{\gamma_L}\left[1 + \frac{s_{ht}^L}{s_{0t}^L}\right]\right)^{-1}\left(\underbrace{\text{MP}_{ht}^L\frac{s_{gt}^Q}{\beta_P s_{0t}^Q}}_{\substack{\text{product}\\ \text{diversion}}} + \underbrace{\frac{W_{gt}^L s_{gt}^L}{\gamma_L s_{0t}^L}}_{\substack{\text{labor}\\ \text{diversion}}}\right) < W_{ht}^{L,\text{before}}.$$

Since $W_{ht}^{\text{after}} < W_{ht}^{\text{before}}$ and $W_{gt}^{\text{after}} < W_{gt}^{\text{before}}$, and the labor supply structure is unchanged, the proof of Proposition 2(d) still holds, so $s_{ht}^{L,\text{after}} + s_{gt}^{L,\text{after}} < s_{ht}^{L,\text{before}} + s_{gt}^{L,\text{before}}$.

The markdown is the ratio between the wage and the total marginal revenue product. With endogenous quality, the total marginal revenue from patient care labor includes both the direct revenue from treating patients and the indirect revenue from quality improvements:

$$\text{Total MR}_{ht}^L = \text{MP}_{ht}^L\left(P_{ht} - \frac{1}{\beta_P}\left[1 + \frac{s_{ht}^Q}{s_{0t}^Q}\right]\right) + \frac{\beta_Y}{\beta_P}\left(1 - \frac{L_{ht}+N_{ht}}{L_{ht}}\theta_{ht}^T\right).$$

The markdown is therefore:

$$\text{markdown}_{ht} = \frac{W_{ht}^L}{\text{Total MR}_{ht}^L} = \left(1 + \frac{1}{\gamma_L}\left[1 + \frac{s_{ht}^L}{s_{0t}^L}\right]\right)^{-1}.$$

Thus, among hospitals for which s_{ht}^L decreases, then $\left(1 + \frac{1}{\gamma_L}\left[1 + \frac{s_{ht}^L}{s_{0t}^L}\right]\right)^{-1}$ becomes smaller, indicating that the markdown strengthens.

From Lemma 4, it follows immediately that s_{ht}^N and s_{gt}^N must strictly decrease to satisfy the first-order condition, which also implies that W_{ht}^N and W_{gt}^N strictly decrease and the total marginal revenue terms increase, which implies non-patient care markdowns strengthen.

Details on Propositions 4(c-e) Warning: The rest of this subsection is out of date but will be fixed soon.

Returning to the product market, the endogeneity of quality requires more than cosmetic changes to the proof structure. We now seek to show that $s_{ht}^{Q,\text{after}} + s_{gt}^{Q,\text{after}} < s_{ht}^{Q,\text{before}} + s_{gt}^{Q,\text{before}}$.

Clearly, if

$$\frac{\exp\left(-\beta_P P_{ht}^{\text{after}} + \beta_Y Y_{ht}^{\text{before}} + \xi_{ht}^Q\right) + \exp\left(-\beta_P P_{gt}^{\text{after}} + \beta_Y Y_{gt}^{\text{after}} + \xi_{gt}^Q\right)}{\exp\left(-\beta_P P_{ht}^{\text{before}} + \beta_Y Y_{ht}^{\text{before}} + \xi_{ht}^Q\right) + \exp\left(-\beta_P P_{gt}^{\text{before}} + \beta_Y Y_{gt}^{\text{before}} + \xi_{gt}^Q\right)} < 1,$$

then the same proof from Proposition 2(c) yields the desired result. Since $P_{ht}^{\text{after}} > P_{ht}^{\text{before}}$ and $P_{gt}^{\text{after}} > P_{gt}^{\text{before}}$, a sufficient condition for this to hold is that $Y_{ht}^{\text{after}} \leq Y_{ht}^{\text{before}}$ and $Y_{gt}^{\text{after}} \leq Y_{gt}^{\text{before}}$.

Suppose instead that $Y_{ht}^{\text{after}} > Y_{ht}^{\text{before}}$ by such a large margin that $-\beta_P P_{ht} + \beta_Y Y_{ht} + \xi_{ht}^Q$ increases, which would be required to overturn the condition from Proposition 2(c). We established that N_{ht} strictly decreases, and that L_{ht} decreases (except in special cases). Then, it is only possible for $Y_{ht} = (L_{ht} + N_{ht})/Q_{ht}$ to increase substantially if Q_{ht} decreases substantially. But if Q_{ht} and thus s_{ht} decreases substantially, then it cannot be the case that $\exp(-\beta_P P_{ht} + \beta_Y Y_{ht} + \xi_{ht}^Q)$ increases, contradicting the assumption. Intuitively, the only way to *increase* the number of patients treated—given that the price increases—is to substantially increase quality, but if staffing is reduced, the only way to increase quality is to *decrease* the number of patients treated, which is a contradiction. Thus, the condition required to overturn the proof in Proposition 2(c) cannot be met by increasing quality.

Totally differentiating the staffing ratio,

$$dY_{ht} = \frac{d(L_{ht} + N_{ht})}{Q_{ht}} - \frac{(L_{ht} + N_{ht})}{Q_{ht}^2} dQ_{ht}.$$

Dividing both sides by $Y_{ht} = \frac{L_{ht}+N_{ht}}{Q_{ht}}$ and using the fact that $dQ_{ht} = \theta_{ht}^T \frac{dL_{ht}}{L_{ht}} Q_{ht}$,

$$\frac{dY_{ht}}{Y_{ht}} = \frac{d(L_{ht} + N_{ht})}{L_{ht} + N_{ht}} - \frac{dQ_{ht}}{Q_{ht}}.$$

Since $\frac{d(L_{ht}+N_{ht})}{L_{ht}+N_{ht}} = \ell_{ht} \frac{dL_{ht}}{L_{ht}} + (1 - \ell_{ht}) \frac{dN_{ht}}{N_{ht}}$ and $\frac{dQ_{ht}}{Q_{ht}} = \theta_{ht}^T \frac{dL_{ht}}{L_{ht}}$,

$$\frac{dY_{ht}}{Y_{ht}} = \frac{dL_{ht}}{L_{ht}} \left(\ell_{ht} - \theta_{ht}^T\right) + (1 - \ell_{ht}) \frac{dN_{ht}}{N_{ht}}.$$

Propositions 4(b) established that $dN_{ht} < 0$ and $dL_{ht} < 0$ (except in special cases). Therefore, if $\ell_{ht} \geq \theta_{ht}^T$, each term in $\frac{dY_{ht}}{Y_{ht}}$ is negative (and the first term is zero if the condition holds with equality while the second term remains negative), so quality declines in response to the merger. However, if $\ell_{ht} < \theta_{ht}^T$, the first term becomes positive while the second remains negative. Quality increases if and only if the positive contribution from reduced congestion

outweighs the negative contribution from reduced non-patient care staffing:

$$\left|\frac{dL_{ht}}{L_{ht}}\right|\left(\theta_{ht}^{T}-\ell_{ht}\right)>(1-\ell_{ht})\left|\frac{dN_{ht}}{N_{ht}}\right|.$$

Details on First-Order Conditions with General Staffing $Y_{ht}=f(L_{ht},N_{ht})/Q_{ht}$: We consider the case where quality is determined by the more general functional form $Y_{ht}=f(L_{ht},N_{ht})/Q_{ht}(L_{ht})$, with $\frac{\partial f}{\partial L_{ht}}\equiv f_L>0$ and $\frac{\partial f}{\partial N_{ht}}\equiv f_N>0$. The hospital system maximizes profits $\Pi_{ht}=P_{ht}Q_{ht}-W_{ht}^{L}L_{ht}-W_{ht}^{N}N_{ht}$, where $Q_{ht}=T_{ht}(L_{ht})$, and P_{ht} is a function of Q_{ht} and Y_{ht} such that $\frac{\partial P_{ht}}{\partial Y_{ht}}=\beta_Y/\beta_P$.

The first-order condition (FOC) with respect to non-patient care labor N_{ht} is:

$$MC_{ht}^{N}=\frac{\partial P_{ht}}{\partial Y_{ht}}\frac{\partial Y_{ht}}{\partial N_{ht}}Q_{ht}.$$

Given $Y_{ht}=f(L_{ht},N_{ht})/Q_{ht}$, the partial derivative of Y_{ht} with respect to N_{ht} is $\frac{\partial Y_{ht}}{\partial N_{ht}}=\frac{f_N}{Q_{ht}}$ (since Q_{ht} is not a direct function of N_{ht}). Thus, $\frac{\partial Y_{ht}}{\partial N_{ht}}Q_{ht}=f_N$. Substituting this into the FOC gives:

$$MC_{ht}^{N}=\left(\frac{\beta_Y}{\beta_P}\right)f_N.$$

This is similar to the FOC for N_{ht} in Lemma 3, where $f_N=1$. The term $(\beta_Y/\beta_P)f_N$ is the marginal revenue from non-patient care labor via its contribution to quality.

The first-order condition with respect to patient care labor L_{ht} is:

$$MC_{ht}^{L}=\left(\frac{\partial P_{ht}}{\partial Q_{ht}}\frac{\partial Q_{ht}}{\partial L_{ht}}+\frac{\partial P_{ht}}{\partial Y_{ht}}\frac{\partial Y_{ht}}{\partial L_{ht}}\right)Q_{ht}+P_{ht}\frac{\partial Q_{ht}}{\partial L_{ht}}.$$

This can be written as:

$$MC_{ht}^{L}=MR_{ht}^{L}+\left(\frac{\beta_Y}{\beta_P}\right)\left(\frac{\partial Y_{ht}}{\partial L_{ht}}Q_{ht}\right).$$

Since $Y_{ht}=f(L_{ht},N_{ht})/Q_{ht}$:

$$\frac{\partial Y_{ht}}{\partial L_{ht}}=\frac{f_L\cdot Q_{ht}(L_{ht})-f(L_{ht},N_{ht})\cdot\frac{\partial Q_{ht}}{\partial L_{ht}}}{Q_{ht}^{2}}=\frac{f_LQ_{ht}-f\cdot MP_{ht}^{L}}{Q_{ht}^{2}}.$$

Therefore,

$$\frac{\partial Y_{ht}}{\partial L_{ht}}Q_{ht}=\frac{f_LQ_{ht}-f\cdot MP_{ht}^{L}}{Q_{ht}}=f_L-Y_{ht}MP_{ht}^{L}=\frac{f}{L_{ht}}\left(\theta_{ht}^{f,L}-\theta_{ht}^{T}\right).$$

where $\theta_{ht}^{f,L} \equiv \frac{L_{ht}}{f} f_L$ be the elasticity of the quality input function f with respect to patient care labor L_{ht}. Substituting this into the FOC for L_{ht}:

$$MC_{ht}^L = MR_{ht}^L + \left(\frac{\beta_Y}{\beta_P}\right)(f_L - Y_{ht} MP_{ht}^L) = MR_{ht}^L + \left(\frac{\beta_Y}{\beta_P}\right) Y_{ht} \frac{Q_{ht}}{L_{ht}} \left(\theta_{ht}^{f,L} - \theta_{ht}^T\right).$$

Thus, the firm is operating at a congested interval if $\theta_{ht}^{f,L} < \theta_{ht}^T$. Note that, if $f(L,N) = L^\rho N^{1-\rho}/Q$ is Cobb-Douglas, then $\theta_{ht}^{f,L} = \rho$ is constant.

A.4 Welfare Effects of Mergers

The final task for which I use the model is to define welfare measures. I define the welfare effects of a merger using changes induced in compensating variation (CV) by each merger, separately for patients, workers, and hospital owners.

Patient welfare is defined as

$$\text{CV}_{mt}^{Q,D} \equiv \bar{Q}_{mt} \times \frac{1}{\beta_P} \left[\log\left(\sum_{ht} \exp(\beta_P P_{ht}^D + \beta_Y Y_{ht}^D + \xi_{ht}^Q) + 1\right)\right], \quad D = 0, 1,$$

where $D = 1$ denotes the outcome with the merger and $D = 0$ denotes the counterfactual outcome without the merger. This is the familiar expression from Small and Rosen (1981) which makes use of the linearity of (indirect) utility in prices for a discrete good. The effect on patient welfare is then,

$$\text{CV}_{mt}^Q \equiv \text{CV}_{mt}^{Q,1} - \text{CV}_{mt}^{Q,0}.$$

The worker welfare effect is measured by first defining worker-specific change in the CV induced by the merger,

$$\max_h \left\{U(W_h^{E,0}, \xi_h^E, \varepsilon_{hi}^E)\right\} = \max_h \left\{U(W_h^{E,1} + \text{CV}_i^E, \xi_h^E, \varepsilon_{hi}^E)\right\}.$$

This expression accounts for the diminishing returns of (indirect) utility in income. The effect on worker welfare is then,

$$\text{CV}_{mt}^E \equiv \bar{E}_{mt} \mathbb{E}[\text{CV}_i^E], \quad E = L, N.$$

Since CV_i^E does not have a closed-form representation, one may follow McFadden (1999) to approximate CV_{mt}^E numerically.

Finally, hospital owner welfare is defined as profits:

$$\text{CV}^{\pi,D}_{mt} \equiv \sum_h \left(P^D_{ht} Q^D_{ht} - W^{L,D}_{ht} L^D_{ht} - W^{N,D}_{ht} N^D_{ht} \right).$$

The effect on hospital welfare is then,

$$\text{CV}^{\pi}_{mt} \equiv \text{CV}^{\pi,1}_{mt} - \text{CV}^{\pi,0}_{mt}.$$

B Estimation Details

B.1 Motivation for the Identifying Restrictions

Below, we utilize the method of simulated moments (MSM) to recover the model parameters of interest. Before proceeding, we develop a constructive argument for identification that makes clear the exogeneity conditions upon which the MSM estimator implicitly relies. We focus on the simpler model of Section 2.2 such that identification arguments have closed-form representations.

Consider the recovery of the labor supply parameter, γ_L. From the inverse labor supply curve for patient care labor L, we have,

$$\mathbb{E}[\Delta \log W^L_h] = \frac{1}{\gamma_L} \left(\mathbb{E}[\Delta \log s^L_h] - \mathbb{E}[\Delta \log s^L_0] + \mathbb{E}[\Delta \xi^L_h] \right).$$

where Δ denotes the change induced by the merger. Using that $\Delta \log s^L_h = \Delta \log \frac{L_h}{L} = \Delta \log L_h$ and $\Delta \log s^L_0 = \Delta \log(1 - \sum s^L_j) \approx -\Delta \log(\sum s^L_j) = -\Delta \log \sum L_j$,

$$\gamma_L \approx \frac{\overbrace{\mathbb{E}[\Delta \log L_h]}^{\text{direct DiD for } L} + \overbrace{\mathbb{E}[\Delta \log(\sum L_j)]}^{\text{aggregate DiD for } L}}{\underbrace{\mathbb{E}[\Delta \log W^L_h]}_{\text{direct DiD for } W^L}} + \frac{\overbrace{\mathbb{E}[\Delta \xi^L_h]}^{\text{amenity bias for } L}}{\mathbb{E}[\Delta \log W^L_h]}.$$

Thus, the merger-based DiD provides a valid moment to recover γ_L if it does not shift amenities, i.e, $\mathbb{E}[\Delta \xi^L_h] = 0$. The same argument applies to γ_N for non-patient care labor, since the labor supply structure is symmetric.

Next, from the treatment technology, we have,

$$\mathbb{E}[\Delta \log Q_h] = \alpha \mathbb{E}[\Delta \log L_h] + \mathbb{E}[\Delta \log A_h].$$

Rearranging,

$$\alpha = \frac{\overbrace{\mathbb{E}[\Delta \log Q_h]}^{\text{direct DiD for } Q}}{\underbrace{\mathbb{E}[\Delta \log L_h]}_{\text{direct DiD for } L}} + \frac{\overbrace{\mathbb{E}[\Delta \log A_h]}^{\text{productivity bias}}}{\mathbb{E}[\Delta \log L_h]} .$$

Thus, the merger-based DiD provides a valid moment to recover α if it does not shift productivity, i.e, $\mathbb{E}[\Delta \log A_h] = 0$.

Lastly, consider the recovery of the distaste for price parameter, β_P. From the inverse product demand curve,

$$\mathbb{E}[\Delta P_h] = \frac{1}{\beta_P} \left(\mathbb{E}[\Delta \log s_0^Q] - \mathbb{E}[\Delta \log s_h^Q] + \mathbb{E}[\Delta \xi_h^Q] \right),$$

Using the first-order Taylor expansion $\mathbb{E}[\Delta P_h] \approx \mathbb{E}[P_h] \times \mathbb{E}[\Delta \log P_h]$,

$$\beta_P \approx \frac{\overbrace{\mathbb{E}[\Delta \log Q_h]}^{\text{direct DiD for } Q} + \overbrace{\mathbb{E}[\Delta \log(\sum Q_j)]}^{\text{aggregate DiD for } Q}}{-\underbrace{\mathbb{E}[P_h]}_{\text{data}} \cdot \underbrace{\mathbb{E}[\Delta \log P_h]}_{\text{direct DiD for } P}} + \frac{\overbrace{\mathbb{E}[\Delta \xi_h^Q]}^{\text{unobs. quality bias}}}{\mathbb{E}[\Delta P_h]} ,$$

Thus, to a first-order approximation, the merger-based DiD provides a valid moment to recover β_P if it does not shift unobserved quality, i.e, $\mathbb{E}[\Delta \xi_h^Q] = 0$.

In sum, we have shown that the ex post merger effects from the DiD design, both the direct effects on the merging firms and the aggregate effects on the market, can be utilized to recover the key structural parameters if the merger does not induce changes in (i) the unobserved labor amenities ξ_h^L and ξ_h^N, (ii) the unobserved product quality ξ_h^Q, or (iii) unobserved productivity A_h.

B.2 First-order Conditions with Insurer Markups

Let P_{ht}^{hos} be the price received by the hospital from the insurer and P_{ht}^{pat} be the price paid by the patient. The relationship between them is given by $P_{ht}^{\text{hos}} = \kappa_{ht} P_{ht}^{\text{pat}}$, where κ_{ht} is the inverse of the coinsurance rate. The hospital's profit is a function of P_{ht}^{hos}, while the patient's utility and demand are functions of P_{ht}^{pat}.

The hospital system's profit maximization problem is:

$$\max_{\{Q_{ht}, Y_{ht}, L_{ht}, N_{ht}\}_{h \in \mathcal{H}_H}} \sum_{h \in \mathcal{H}_H} \left(P_{ht}^{\text{hos}} Q_{ht} - W_{ht}^L L_{ht} - W_{ht}^N N_{ht} \right)$$

From the patient's perspective, the inverse demand function is

$$P_{ht}^{\text{pat}} = \frac{1}{\beta_P}\left(\beta_Y Y_{ht} + \xi_{ht} + \log s_{0t}^Q - \log s_{ht}^Q\right).$$

Therefore, the inverse demand in terms of the hospital's price is,

$$P_{ht}^{\text{hos}} = \frac{\kappa_{ht}}{\beta_P}\left(\beta_Y Y_{ht} + \xi_{ht} + \log s_{0t}^Q - \log s_{ht}^Q\right).$$

The first-order condition with respect to non-patient care labor for a single-hospital system is:

$$\text{MC}_{ht}^N = \underbrace{\frac{\partial P_{ht}^{\text{hos}}}{\partial Y_{ht}}\frac{\partial Y_{ht}}{\partial N_{ht}}Q_{ht}}_{\text{returns from quality } (N)} = \underbrace{F_{ht}^N \times \frac{\kappa_{ht}}{\beta_P}\beta_Y}_{\text{returns from quality } (N)}$$

where $F_{ht}^N \equiv \frac{\partial}{\partial N_{ht}} F\left(L_{ht}, N_{ht} - \psi_{ht}\right)$. The first-order condition with respect to patient care labor is:

$$\text{MC}_{ht}^L = \underbrace{\left(1 + 1/\theta_{ht}^Q\right) \times P_{ht}^{\text{hos}}\text{MP}_{ht}^L}_{\equiv \text{MR}_{ht}^L} + \underbrace{\left(F_{ht}^L - Y_{ht}\text{MP}_{ht}^L\right) \times \frac{\kappa_{ht}}{\beta_P}\beta_Y}_{\text{returns from quality } (L)}$$

where MR_{ht}^L and MC_{ht}^L are the marginal revenue and marginal cost of patient care labor, respectively, and $F_{ht}^L \equiv \frac{\partial}{\partial L_{ht}} F\left(L_{ht}, N_{ht} - \psi_{ht}\right)$. The term $\left(1 + 1/\theta_{ht}^Q\right)$ now incorporates the patient-paid price in the elasticity:

$$\theta_{ht}^Q = -P_{ht}^{\text{hos}} \frac{s_{0t}^Q}{s_{ht}^Q + s_{0t}^Q}\frac{\beta_P}{\kappa_{ht}}.$$

B.3 Method of Simulated Moments Estimator

The method of simulated moments (MSM) algorithm proceeds as follows:

1. Guess parameters $\Xi^* \equiv (\alpha^*, \beta_Y^*, \gamma_L^*, \gamma_N^*, \delta^*, \rho^*, \phi^*, \bar{\kappa}_\Delta^*)$. Calibrate outside shares $s_0^{L,*}, s_0^{N,*}, s_0^{Q,*}$.
2. Infer market shares as $s_h^{L,*} = L_h/\bar{L}_m^*$, $s_h^{N,*} = N_h/\bar{N}_m^*$, and $s_h^{Q,*} = Q_h/\bar{Q}_m^*$, where $\bar{L}_m^* = (\sum_{j\in m} L_j)/(1 - s_0^{L,*})$, $\bar{N}_m^* = (\sum_{j\in m} N_j)/(1 - s_0^{N,*})$, and $\bar{Q}_m^* = (\sum_{j\in m} Q_j)/(1 - s_0^{Q,*})$.
3. Infer the labor supply elasticities, $\theta_h^{E,*} = \gamma_E^*\left(\frac{s_0^{E,*}}{s_h^{E,*} + s_0^{E,*}}\right)$, and the corresponding marginal costs of labor, $MC_h^{E,*} = W_h^E(1 + 1/\theta_h^{E,*})$, for each labor type $E \in \{L, N\}$.

4. Invert the labor supply curves to recover $\xi_h^{L,*}, \xi_h^{N,*}$:

$$\xi_h^{L,*} = \log(s_h^{L,*}/s_0^{L,*}) - \gamma_L^* \log(W_h^L)$$
$$\xi_h^{N,*} = \log(s_h^{N,*}/s_0^{N,*}) - \gamma_N^* \log(W_h^N)$$

5. Infer $A_h^* = Q_h / L_h^{\alpha^*}$ and thus $\text{MP}_h^{L,*} = \alpha^* A_h^* L_h^{\alpha^*-1}$.

6. Define the composite parameter $\tilde{\kappa}_h \equiv \kappa_h / \beta_P$. Infer ψ_h^* and the composite parameter $\tilde{\kappa}_h^*$ by numerically solving the following system of two first-order conditions simultaneously for the two firm-specific unobservables:

$$MC_h^{N,*} = \tilde{\kappa}_h^* \beta_Y^* F_N(L_h, N_h - \psi_h^*)$$
$$MC_h^{L,*} = (1 + 1/\theta_h^{Q,*}) P_h^{\text{hos}} \text{MP}_h^{L,*} + \tilde{\kappa}_h^* \beta_Y^* \left(F_L(L_h, N_h - \psi_h^*) - Y_h^* \text{MP}_h^{L,*} \right)$$

   where the inferred quality of care, $Y_h^* = F(L_h, N_h - \psi_h^*)/Q_h$, and the residual product demand elasticity, $\theta_h^{Q,*} = -\frac{P_h^{\text{hos}}}{\tilde{\kappa}_h^*} \frac{s_0^{Q,*}}{s_h^{Q,*} + s_0^{Q,*}}$, are themselves functions of these unknowns.

7. Infer $\xi_h^{Q,*}$ using the previously recovered $\tilde{\kappa}_h^*$:

$$\xi_h^{Q,*} = \log(s_h^{Q,*}/s_0^{Q,*}) + \frac{P_h^{\text{hos}}}{\tilde{\kappa}_h^*} - \beta_Y^* Y_h^*$$

8. Given Ξ^*, all model primitives have been recovered. Given these recovered primitives $\Lambda_h^* \equiv (\xi_h^{L,*}, \xi_h^{N,*}, \xi_h^{Q,*}, A_h^*, \psi_h^*, \tilde{\kappa}_h^*)$, impose the actual merger within each relevant market m and simulate the post-merger equilibrium. In the post-merger simulation, the composite parameter for each firm is shifted by the guessed global parameter, such that the new rate is $\tilde{\kappa}_h' = \tilde{\kappa}_h^* \exp(\bar{\kappa}_\Delta^*)$. The simulation solves for the new equilibrium outcomes (new $P_h', Q_h', \dots$) given this shift and the change in market structure. Calculate the log change in each outcome (e.g., $\log(P_h') - \log(P_h)$). Average these log changes across relevant firms and markets to form simulated merger effects.

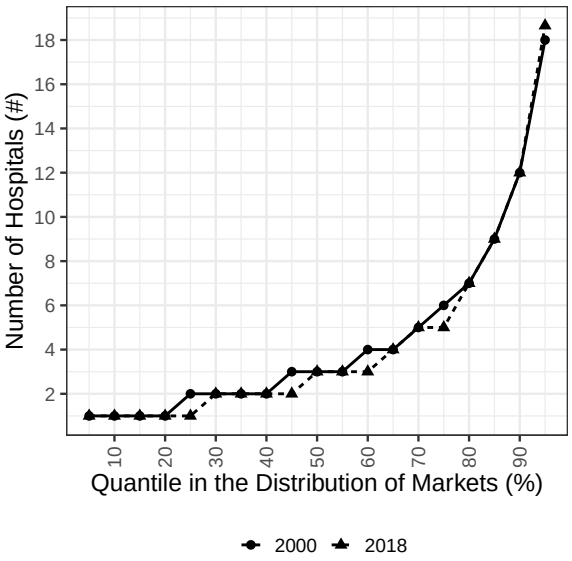
This algorithm returns the simulated moments $\mathbf{M}^{sim}(\Xi^*)$. We compare them to the observed moments $\mathbf{M}^{obs}$, which were motivated in the constructive identification discussion above. We then solve,

$$\Xi^{msm} = \min_{\Xi^*} (\mathbf{M}^{obs} - \mathbf{M}^{sim}(\Xi^*))' \mathbf{W} (\mathbf{M}^{obs} - \mathbf{M}^{sim}(\Xi^*)).$$

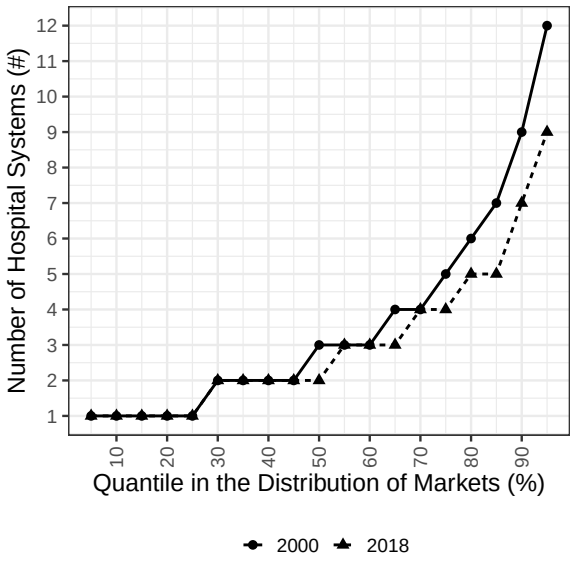
where in practice we use the diagonal weighting matrix in place of $\mathbf{W}$. The MSM estimate of Λ_h is the one that results from inverting the model evaluated at Ξ^{msm}.

Lastly, note that the estimation procedure recovers the composite parameter $\tilde{\kappa}_h \equiv \kappa_h/\beta_P$, and thus does not separately recover κ_h and β_P. Since the first-order conditions and markups and markdowns can be expressed only in terms of $\tilde{\kappa}_h$ without loss of generality, we do not need these parameters for our results. Nonetheless, we can recover the underlying parameters β_P and κ_h to scale if using an external measure of the average insurer markup, $E[\kappa_h]$. The true β_P is then identified by $\beta_P = \bar{\kappa}/E[\tilde{\kappa}_h]$. Once β_P is known, the true distribution of κ_h is recovered as $\kappa_h = \tilde{\kappa}_h \cdot \beta_P$. Since the average coinsurance rate in the US is 20%, a natural calibration value is $\bar{\kappa} = 5$.

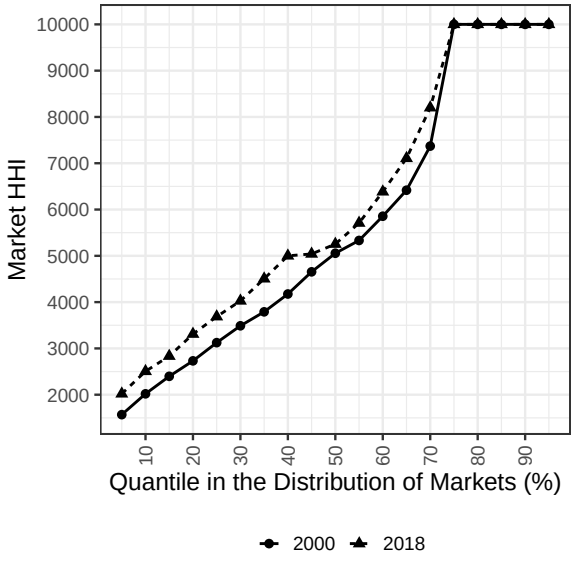
C Additional Figures and Tables



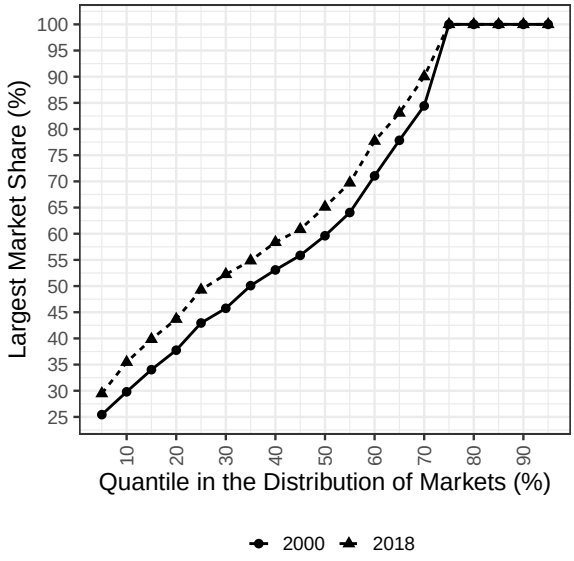
(a) Number of Hospitals per Market



(b) Number of Hospital Systems per Market



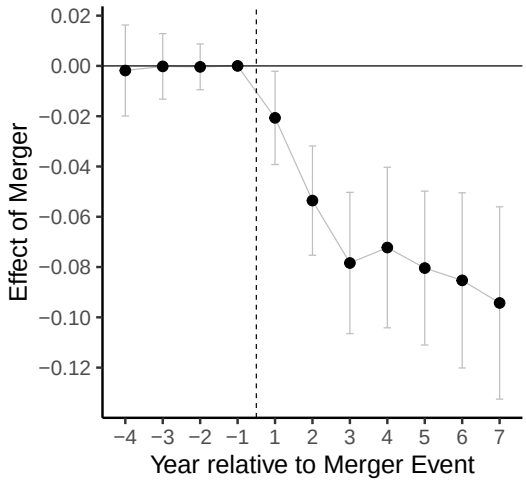
(c) HHI



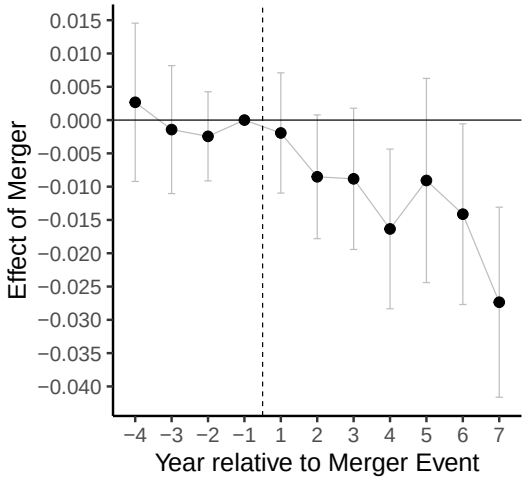
(d) Largest Market Share

Figure A1: Distribution of HHIs and Largest Market Shares across Markets

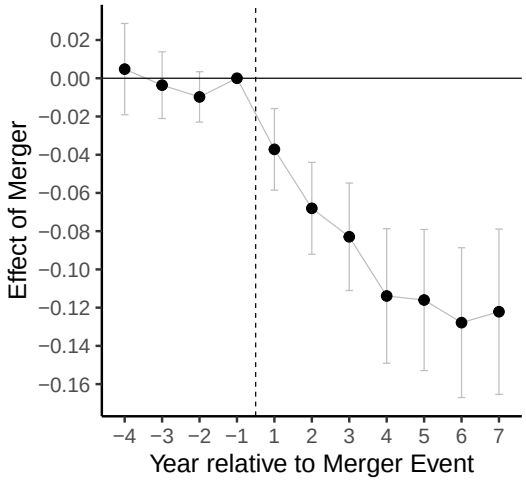
Notes: This figure presents deciles in the distributions of HHIs and largest market shares across markets in 2000 and 2018.



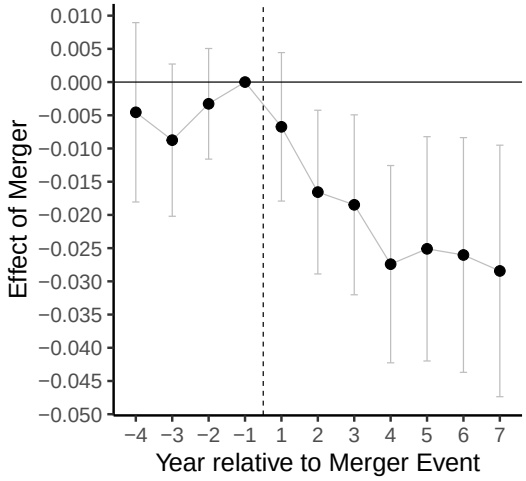
(a) Patient Care: Number of Workers, FTE (log)



(b) Patient Care: Hourly Wage (log)



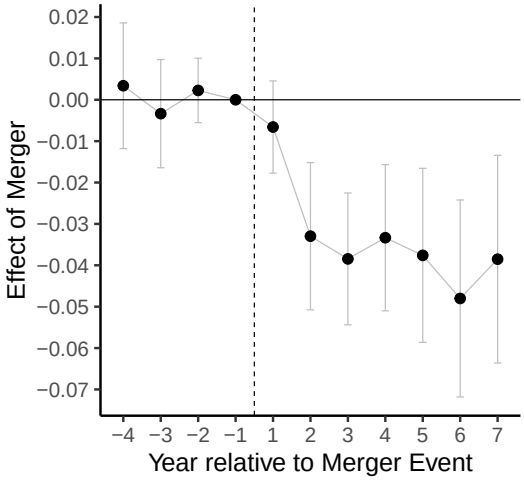
(c) Non-patient Care: Number of Workers, FTE (log)



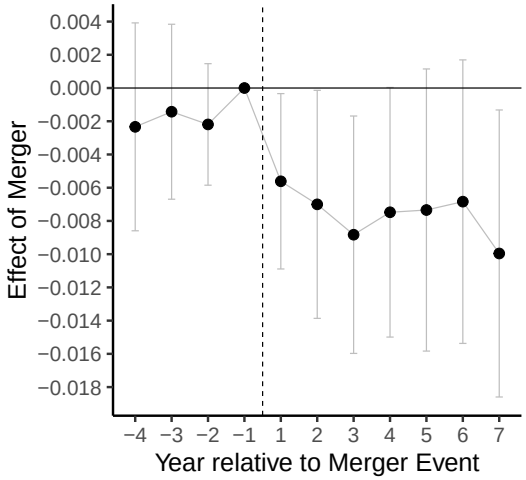
(d) Non-patient Care: Hourly Wage (log)

Figure A2: Direct Effects of Mergers on the Merging Hospitals: Labor Market Outcomes by Occupational Category

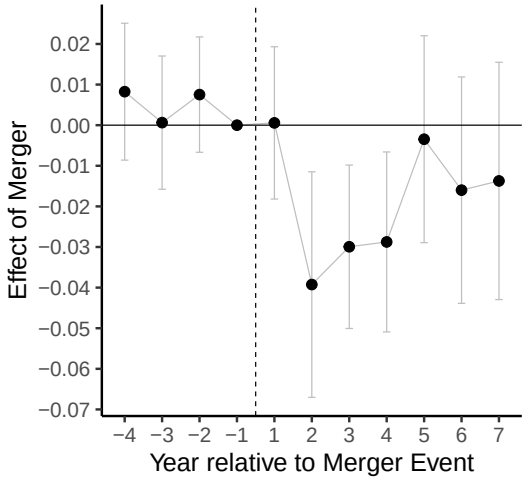
Notes: This figure presents difference-in-differences estimates that compare treated hospitals, defined as merging hospitals satisfying the "presumed anti-competitive" HHI thresholds, to merger-specific control groups of 10 hospitals from other markets matched to treated units by propensity score. The number of workers refers to the sum of workers employed among the merging hospitals, measured consistently across event times to account for reporting changes. The log wage refers to the employment-weighted average of the log wage among the merging hospitals, also measured consistently across event times.



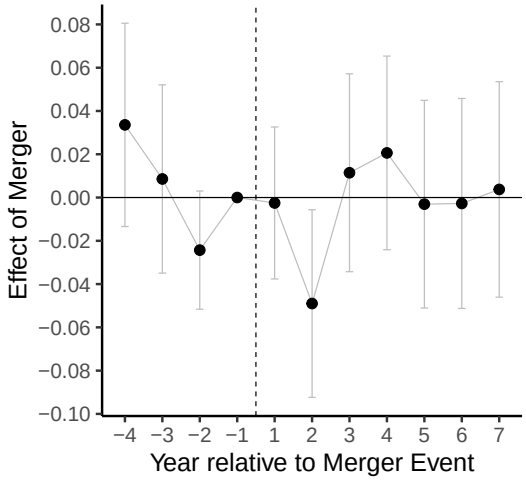
(a) Aggregate: Number of Workers (log)



(b) Aggregate: Hourly Wage (log)



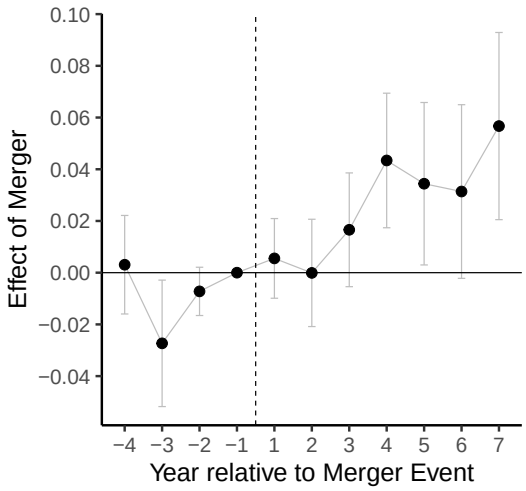
(c) Aggregate: Number of Patients (log)



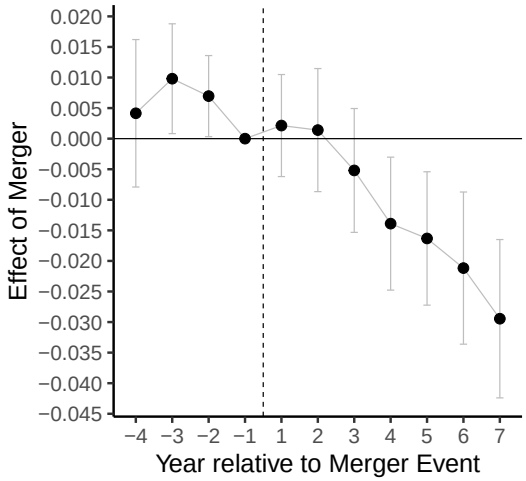
(d) Aggregate: Price Index (log)

Figure A3: Aggregate Effects of Mergers

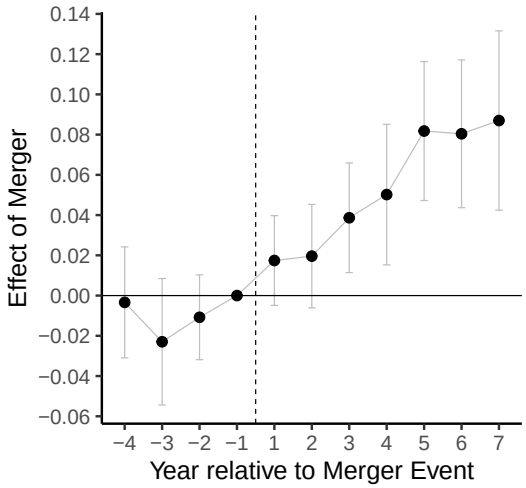
Notes: This figure displays difference-in-differences estimates of the aggregate effects of mergers by event time. Outcomes are aggregated across all hospitals within the commuting zone. The sum is used to aggregate the number of workers and patients, while employment-weighted and patient-weighted means are used for the hourly wage and price index, respectively. Treated units are the hospitals within commuting zones experiencing a presumed anti-competitive merger, and event time is relative to that merger. These treated units are matched to 10 control groups of similarly aggregated hospitals in other markets that did not experience mergers, using aggregates of the covariates specified in Section 4.1. 90% confidence intervals are indicated by brackets.



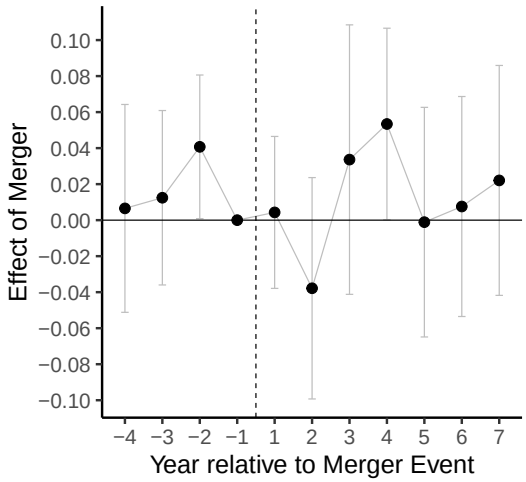
(a) Spillover: Number of Workers (log)



(b) Spillover: Hourly Wage (log)



(c) Spillover: Number of Patients (log)



(d) Spillover: Price Index (log)

Figure A4: Spillover Effects of Mergers

Notes: Notes: This figure displays difference-in-differences estimates of the spillover effects of mergers on non-merging competitor hospitals by event time. Outcomes are aggregated across all non-merging competitor hospitals within the commuting zone. The sum is used to aggregate the number of workers and patients, while employment-weighted and patient-weighted means are used for the hourly wage and price index, respectively. Treated units are the aggregated non-merging competitor hospitals within commuting zones experiencing a presumed anti-competitive merger, and event time is relative to that merger. These treated units are matched to 10 control groups of similarly aggregated hospitals in other markets that did not experience mergers, using aggregates of the covariates specified in Section 4.1. 90% confidence intervals are indicated by brackets.

	Treatment Group Summary			DiD Effects of the Merger	
	Mergers	Hospitals	Outcome SD	Before Merger	After Merger
Event Times:		{-1}		{-2,-3,-4}	{2,3,4,5}
Patient Satisfaction Survey:					
· · · Overall rating (%)	71	219	5.368	-0.253	-1.061**
				(0.276)	(0.460)
· · · Recommend the hospital (%)	71	219	6.563	0.052	-1.256**
				(0.305)	(0.488)
· · · Cleanliness (%)	71	219	4.463	-0.125	-1.081**
				(0.261)	(0.441)
· · · Quietness (%)	71	219	7.943	-0.181	-1.192**
				(0.320)	(0.496)
30-day All-cause Mortality Rates:					
· · · Admitted for heart failure (%)	69	213	1.392	-0.103	0.456**
				(0.102)	(0.195)
· · · Admitted for pneumonia (%)	72	222	2.224	-0.051	0.796***
				(0.114)	(0.234)

Table A1: Direct Effects of Mergers on Quality of Care in the Merging Hospitals

Notes: This table presents difference-in-differences estimates that compare treated hospitals, defined as merging hospitals satisfying the "presumed anti-competitive" HHI thresholds, to merger-specific control groups of 10 hospitals from other markets matched to treated units by propensity score. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels. Outcome SD refers to the standard deviation of the outcome variable.