NBER WORKING PAPER SERIES

LABOR AND PRODUCT MARKET POWER, ENDOGENOUS QUALITY, AND THE CONSOLIDATION OF THE US HOSPITAL INDUSTRY

Bradley Setzler

Working Paper 34180
http://www.nber.org/papers/w34180

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2025

There are no sources of funding or relevant financial relationships to disclose. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Bradley Setzler. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

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

Bradley Setzler

NBER Working Paper No. 34180
August 2025
JEL No. I11, J42, L41

ABSTRACT

Existing structural analyses of the harmful effects of market consolidation focus on either product or labor markets in isolation, ignoring that product market competitors often compete for workers as well. This paper develops a unified framework for merger evaluation, finding that firms' simultaneous exercise of oligopoly power in the product market and oligopsony power in the labor market amplifies the harm from mergers to both consumers and workers. The model also demonstrates how merger-induced gains in labor market power incentivize firms to reduce product quality, highlighting an additional channel for consumer harm. The model's predictions are tested and quantified in the context of the recent consolidation of the US hospital industry. Linking panel data from several sources on all US hospitals from 1996-2022, a difference-in-differences design is estimated for nearly 150 high-concentration within-market mergers. Hospital mergers significantly reduce patient volume, increase prices, reduce employment, lower wages, and deteriorate quality of care, resulting in higher patient mortality. After recovering the structural parameters, the estimated model replicates observed merger impacts. Counterfactual exercises reveal that ignoring increased labor (product) concentration would lead one to under-predict the harm of mergers to consumers (workers).

Bradley Setzler
Pennsylvania State University
Department of Economics
and NBER
bradley.setzler@gmail.com

1 Introduction

Recent literature documents rising market concentration in US product and labor markets, prompting concerns about increasing market power (Autor, Dorn, Katz, Patterson, and van Reenen, 2020; de Loecker, Eeckhout, and Unger, 2020). Mergers are a natural avenue through which markets become concentrated, with strong policy-relevance for economists: Antitrust authorities utilize structural economic models to predict which mergers may harm the public (Farrell and Shapiro, 2010; Hovenkamp and Shapiro, 2018). Merger evaluation in antitrust has traditionally focused on product market consolidation and the resulting harm to consumers through higher prices and reduced output. A developing literature raises concerns about the anti-competitive effects on workers of mergers among employers (Hemphill and Rose 2018; Naidu, Posner, and Weyl 2018; Marinescu and Hovenkamp 2019), and employer market power has recently been incorporated into US merger guidelines (US DOJ and FTC, 2023). Yet, despite the recent interest in employer market power and the fact that product market competitors often compete for workers as well, existing structural analyses of market consolidation focus on either labor or product markets in isolation.

What new insights about the anti-competitive effects of market consolidation can be learned from examining the interactions between labor and product market power? In this paper, 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 how labor market power distorts quality provision. I test and quantify the predictions of the model for the recent consolidation of the US hospital industry. The empirical analysis combines 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. I use the estimated merger effects to identify the parameters of the model, then use the estimated model to analyze the importance of accounting for both labor and product market power in merger evaluation.

The structural framework for merger evaluation has three key components. First, it takes as a starting point the modern differentiated-products framework for the evaluation of horizontal mergers for product market concentration in industrial organization. Second, I add a rich model of the local labor market in which workers have horizontally-differentiated and vertically-differentiated preferences over local employers. 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, products differ both in price and quality, and 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 characterizing how mergers impact consumers and workers. Mergers affect the merging parties through two channels: *labor market diversion*, where producers internalize costs imposed on other commonly-owned local producers when competing for workers, and *product market diversion*, where producers internalize revenues lost by other commonly-owned local producers when competing for consumers. When a merger expands the set of commonly-owned local producers, these diversion effects lead the merging parties to raise prices and cut wages while reducing output and employment, with each effect amplifying the other. Mergers also create *spillovers* on local competitors, shifting both product demand and labor supply curves outward and allowing competitors to increase output and employment without increasing costs. The outside shares—consumers who fail to consume a product and workers who fail to attain employment—unambiguously rise. Finally, a merger-induced increase in the *labor* market share impacts *consumers* by incentivizing the merging parties to reduce quality, providing another channel through which labor and product markets interact.

In order to test the model predictions in the real world, I study the consolidation of the US hospital industry. After a massive wave of mergers in the 21st century, the median commuting zone has only three hospitals and only two distinct hospital systems. Rising concentration of hospital markets raises the concern that hospitals have important price-setting power, heightened by the inherently urgent and local nature of much patient care. 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, implying inelastic local labor supply.[1] A natural concern is that hospitals exploit labor market consolidation to suppress wages below competitive benchmarks.[2] Furthermore, staffing levels are a key determinant of quality of care. If hospitals exercise labor market power by reducing employment, patients may be harmed not only by rising prices but also by worsening medical care. Thus, US hospital mergers provide an ideal environment for testing the model's predictions.

The empirical analysis combines administrative data on all US hospitals over 1996-2022, covering 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. Market concentration is measured for each hospital labor and product market as well as the changes in local market concentration induced by mergers. For the purposes of the empirical analysis, I focus on "presumed anti-competitive" mergers that meet the current thresholds utilized by US

[1] For example, using resume data, Schubert, Stansbury, and Taska (2024) find that workers in healthcare occupations are the least likely to move to jobs in other occupations. US workers are also known to be highly immobile across geographic markets; see the review by Autor, Dorn, Hanson, Jones, and Setzler (2025).

[2] See Prager and Schmitt (2021) for difference-in-differences evidence that high-concentration mergers reduce wages in the hospital industry and Arnold (2021) for other industries.

courts (US DOJ and FTC, 2023). One-fourth of commuting zones have experienced a presumed anti-competitive hospital merger so far in the 21st century.

I confirm the predictions of the model among presumed anti-competitive mergers using a propensity score-matched difference-in-differences design. This design compares merging hospitals to similar non-merging hospitals in other markets. Within the merging hospitals, the price increases by about 7% while the number of patients treated decreases by at least 4%. In the labor market, wages decrease by 2-4% across patient care and non-patient care occupations while employment levels fall by 9-13%. Using a spillover design, I find that local competitors of merging hospitals decrease wages by 3% yet increase employment by 6% and patient volume by 5%.[3] Aggregating across the market, patient volume and employment decline by as much as 3%, with rising numbers of patients who fail to receive treatment and workers who fail to attain employment. I also find substantial deterioration in quality of care: The staffing ratio declines by nearly 7%, patient satisfaction ratings decrease by more than one percentage point, and risk-adjusted all-cause mortality rates for heart failure and pneumonia patients increase by 0.5-0.8 percentage points—relatively large effects compared to baseline mortality rates of 12-13%.

In the final section, I take the model to the hospital data, developing a method of simulated moments estimator to recover all model parameters by matching the simulated equilibrium to the estimated merger effects. To better represent this industry's institutional features, the model is extended to include markups that hospitals charge to insurers. Hospitals may raise their markups on insurers in response to a merger, providing a reduced-form representation of the hospital-insurer bargaining outcome. The model yields several insights. First, patient and non-patient care workers are gross complements in hospital quality production, and quality provision exhibits increasing returns to scale in labor. Second, two-thirds of the merger-induced price increase is driven by gains in bargaining power over insurers, with another one-third driven by classical oligopoly power over patients. Third, comparing the smallest to largest hospitals in terms of market share, demand elasticities range from 3.4 to 2.5 while labor supply elasticities range from 5.5 to 2.6. Fourth, estimated markups and markdowns show that prices exceed marginal costs by 32% in the smallest hospitals to 40% in the largest, while wages fall short of marginal revenues by 18% to 27%.

Lastly, I use the estimated model to provide ex ante merger evaluations in counterfactual scenarios. I find that, if we ignore labor market competition, we understate consumer harm, and if we ignore product market competition, we understate worker harm. The extent to which harm is understated depends in practice on the magnitude of labor and product diversion effects. In turn, the diversion effects depend quantitatively on the degree of concentration and the elasticity of demand

[3]Recent papers by Sharma (2023), Roussille and Scuderi (2025), and Derenoncourt and Weil (2025) test for wage-setting spillovers in other industries but do not find evidence of spillovers, suggesting that US hospitals are a particularly extreme case of oligopsony.

and supply in the product and labor markets. Thus, if cross-market diversion effects are substantial, we must account for *both* product and labor market power to correctly predict merger-induced harm to workers *or* consumers

Related literature. This paper relates to three active literatures. The first is the literature on labor market power. In this literature, workers view employers as horizontally-differentiated, which generates an upward-sloping labor supply curve and thus an increasing marginal cost curve from the perspective of the firm. In the first strand of this literature, employers are strategically small and therefore would not take advantage of labor market concentration after a merger (Card, Cardoso, Heining, and Kline 2018, Lamadon, Mogstad, and Setzler 2022, Kroft, Luo, Mogstad, and Setzler 2025). In the second strand, employers are oligopsonistic such that larger employers can exploit market share to place stronger markdowns on wages (Berger, Herkenhoff, and Mongey 2022, Roussille and Scuderi 2025, Chan, Kroft, Mattana, and Mourifié 2024). This paper falls within the second strand. In the context of mergers, Berger, Hasenzagl, Herkenhoff, Mongey, and Posner (2025) characterize the effects of labor market mergers on labor market outcomes under oligopsony but without product market power, while Hosken, Larson-Koester, and Taragin (2024) incorporate wage bargaining into a model of mergers with product market power.[4] However, no prior research has incorporated oligopsony power into a model of mergers with product market power.

The second related literature provides the modern framework for the ex ante evaluation of horizontal mergers for product market concentration in industrial organization. Leading examples are Nevo (2000), Björnerstedt and Verboven (2016), and Miller and Weinberg (2017). By incorporating labor market power, I extend this framework in two ways. First, I introduce increasing marginal cost curves that shift upwards in response to a merger-induced gain in labor market concentration. This mechanism amplifies the price markup and the reduction in output. Conversely, a merger-induced gain in product market concentration shifts downwards the marginal revenue product of labor, which amplifies wage markdowns and the reduction in employment. Thus, if both oligopoly and oligopsony are present, ex ante merger evaluation may substantially understate the harm to consumers (workers) if it only accounts for gains in product (labor) market concentration. Second, I endogenize product quality as a function of labor, demonstrating that a merger-induced gain in labor market concentration can distort the provision of quality. Thus, mergers may harm consumers not only through higher prices and less output, but also through lower-quality products.

Third, this paper relates to the literature specifically on the effects of recent mergers in the US

[4] As Hosken et al. (2024) discuss, the wage-bargaining approach is inconsistent with substantial employment effects of mergers, which are present in my empirical context. A wage-bargaining approach can be motivated by unionized labor markets (see, e.g., Angerhofer, Collard-Wexler, and Weinberg 2025 for teachers unions). However, unionized labor markets are rare in the US: only 6% of the private sector (BLS, 2025) and only 13% of hospital workers (Ahmed, Kadakia, Ahmed, Shultz, and Li, 2022) are unionized.

hospital industry. There are two strands of this literature. The first strand provides difference-in-differences evidence that hospital mergers lead to a substantial increase in the price (Dafny, 2009; Cooper, Craig, Gaynor, and Van Reenen, 2019; Brand, Garmon, and Rosenbaum, 2023), decrease in the wage (Prager and Schmitt 2021), and decrease in patient satisfaction (Beaulieu, Dafny, Landon, Dalton, Kuye, and McWilliams 2020) among the merging parties. I extend this literature in three ways. First, I show that mergers have quantity effects: merging hospitals not only raise prices and reduce wages, but also decrease the volume of patients and the number of workers employed. Second, I show that mergers negatively impact several other measures of quality, including worsening mortality. Third, I show that mergers have spillover effects on non-merging competitors. The second strand of this literature analyzes the role of hospital-insurer bargaining in the setting of higher prices after mergers (Gowrisankaran, Nevo, and Town, 2015; Ho and Lee, 2017, 2019; Dafny, Ho, and Lee, 2019). While the empirical application focuses on markets for patients and hospital workers, I account for this institutional detail by incorporating a reduced-form representation of hospital-insurer bargaining effects of mergers in the quantitative model.

2 A Model of Oligopoly and Oligopsony Power in Mergers

In this section, I develop a model of mergers in which firms simultaneously exploit product and labor market power by leveraging their market shares over consumers and workers, respectively. Firms hire for both production and support occupations, and choose staffing to endogenously provide product quality to consumers. In this environment, I derive comparative statics that characterize the effects of mergers on the outcomes and competitiveness of consumer and labor markets. Appendix A provides mathematical details.

2.1 Labor Supply, Product Demand, and Technology

We denote a product by h, a consumer by i, a worker by ι, a market by m, and a time period by t. We interchangeably let h index the product and the producer of that product. A producer could refer specifically to a production line within a multi-product plant, an establishment within a multi-establishment firm, and so on. A worker is employed by one producer and a consumer consumes one product, per time and market. The owner of the system of production lines or establishments is referred to as the firm.

Product Demand. In product market m at time t, consumer i's utility from consuming product h is,

$$u^{Q}_{iht} = -\beta_P P_{ht} + \beta_Y Y_{ht} + \xi^{Q}_{ht} + \varepsilon^{Q}_{iht}. \tag{1}$$

There are four components to consumer utility. The first captures the price, P_{ht}, multiplied by the marginal disutility of lost income, $-\beta_P$.[5] The second captures the endogenous component of quality, Y_{ht}, multiplied by the relative preference for quality, β_Y. I assume that Y_{ht} is chosen or produced, subject to constraints described below. The third, ξ^Q_{ht}, captures the residual component of product quality. As is standard, residual quality is predetermined.[6] Combined, $\beta_Y Y_{ht} + \xi^Q_{ht}$ characterizes vertical differentiation in quality across products, which is unobserved by the analyst but observed by consumers. The fourth, ε^Q_{iht}, represents consumer i's taste for product h and is assumed to have the standard Gumbel distribution. These consumer-product-specific tastes are unobserved by the analyst and producer but observed by the consumer, giving rise to horizontal product differentiation.

Aggregating across the distribution of match-specific tastes in a market, the product demand for 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 consumers in market m at time t. The outside share of the product market—the share of consumers who fail to consume the product in the local market—is $s^Q_{0t} \equiv 1 - \sum_h s^Q_{ht}$.[7]

Labor Supply. I consider two broad occupational categories within firms: production and support. The number of production workers employed in the production of product h is denoted by L_{ht} and the number of support workers is denoted by N_{ht}. The wage offered to production workers is denoted by W^L_{ht} and the wage offered to support workers is W^N_{ht}. In the application to the hospital industry, the production occupations refer to patient care workers and support occupations refer to non-patient care workers. 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.

In labor market m at time t, the preference of worker ι of occupation type $E \in \{L, N\}$ for

[5] It is computationally feasible, but muddles the analytical comparative statics, to allow for more flexible utility functions, such as random coefficients or non-linear transformations of price.

[6] 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 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.

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

working on product 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, product-occupation-specific job quality ξ^E_{ht} represents the value of amenities enjoyed by all workers employed towards that product and occupation. Amenities are unobserved to the analyst but common knowledge among workers, giving rise to vertical employer differentiation. Amenities may differ between production and support 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 product h and is assumed to have the standard Gumbel distribution at time t.[8] These worker-product-specific tastes are unobserved to the analyst and firm but known to the worker, giving rise to horizontal employer differentiation.

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

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

where $\bar{E}_{mt}$ denotes the total number of workers of type E in market t. The outside share of each labor market—the share of workers who fail to attain employment—is $s^E_{0t} \equiv 1 - \sum_h s^E_{ht}$.[9]

Technology. To supply Q_{ht} consumers, the amount of production labor required is determined by the technology,

$$Q_{ht} = T_{ht}(L_{ht}), \quad \frac{\partial T_{ht}}{\partial L_{ht}} > 0, \quad \frac{\partial^2 T_{ht}}{\partial L^2_{ht}} \leq 0. \tag{5}$$

The patient treatment technology T_{ht} may vary across products and time, representing that some products may be more efficiently produced than others. In the empirical application to the US hospital industry, I consider the constant elasticity patient treatment function $T_{ht}(L_{ht}) = A_{ht} L^{\alpha}_{ht}$, where $A_{ht} > 0$ measures the relative productivity of h and $0 < \alpha \leq 1$ is the elasticity of patients

[8] 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 $\vec{\varepsilon}^E_{\iota t} \equiv (\varepsilon^E_{\iota 1 t}, \varepsilon^E_{\iota 2 t}, ..., \varepsilon^E_{\iota \bar{h} t})$, the model places no restrictions on the copula of $(\vec{\varepsilon}^E_{\iota t}, \vec{\varepsilon}^E_{\iota t'})$ for $t' \neq t$.

[9] 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.

treated with respect to employment.

The firm also combines production labor, support labor, and the number of consumers to determine the endogenous component of product quality,

$$Y_{ht} = \frac{F(L_{ht}, N_{ht})}{Q_{ht}}, \quad \frac{\partial F}{\partial E_{ht}} > 0, \quad \frac{\partial^2 F}{\partial E_{ht}^2} \leq 0, \quad E = L, N. \tag{6}$$

The term $F(L_{ht}, N_{ht})$ is referred to as *effective staffing*, and quality can be interpreted as effective staffing divided among consumers. For example, if $F(L_{ht}, N_{ht}) = L_{ht} + N_{ht}$, then Y_{ht} is simply the staffing ratio—employment per consumer. In the empirical application to the US hospital industry, I consider the constant elasticity of substitution functional form, $F(L_{ht}, N_{ht}) = (\delta L_{ht}^{\rho} + (1-\delta) N_{ht}^{\rho})^{\phi/\rho}$, which includes the staffing ratio as a special case.

Before proceeding, it is useful to emphasize the differences in functional form between the two technologies. Support workers cannot be substituted to produce output, such that N_{ht} is excluded from the production technology. In the hospital context, this reflects that administrators, maintenance workers, and other support staff cannot perform clinical duties. By contrast, production and support workers are combined in the provision of quality.

Multi-product Firm's Conduct. Firm H owning the set of producers $\mathcal{H}_H$ is assumed to maximize profits across its set of products in the market:

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

subject to product demand, the labor supply of each occupational category, the production technology, and the quality technology. The firm values profits equally across all products.

Firms compete à la differentiated-Cournot in both the product and labor markets.[10] This means that, when the reference firm is choosing its price-quantity and wage-labor pairs, it perceives competitors as holding output and employment fixed, with competitors responding via 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 (as in the

[10] 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 standard in empirical industrial organization, though differentiated-Cournot is drawing attention in the conduct testing literature, in part because it can rationalize larger markups (see Magnolfi, Quint, Sullivan, and Waldfogel 2022). While comparative statics for horizontal mergers in the product market are well-known under differentiated-Bertrand or *non*-differentiated-Cournot (see, e.g., Nocke and Whinston 2022), I am not aware of prior work providing comparative statics for horizontal mergers in the product market 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.

Cournot first-order condition) rather than keep wages fixed as employment adjusts (Bertrand).[11] 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 the conduct that best fits the hospital labor market.

Welfare Measures. While the model is not utilized to compute welfare measures in this paper, model-consistent welfare expressions are provided in Appendix B.

2.2 Comparative Statics for Mergers, given Quality

This subsection characterizes how mergers among producers in the same market affect the price-setting and wage-setting behavior of both the merging producers and their competitors. It does 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 quality responses.

Let $\text{MP}^L_{ht} \equiv \frac{\partial T_{ht}(L_{ht})}{\partial L_{ht}} > 0$ denote the marginal product of production 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 production 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 consumers. Impose that $\beta_Y = 0$, such that consumers do not value the endogenous component of quality, which in turn implies that $N_{ht} = 0$, as support workers produce no revenue if not via returns to quality. In this case, the only product quality component is the predetermined factor, ξ^Q_{ht}.

Prior to a merger, the first-order condition for profit-maximization is as follows:

Lemma 1 (First-order condition in a single-product firm)**.** *Suppose quality is predetermined. For a single-product firm h, the first-order condition for profit-maximization is,*

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

Furthermore, $\left|\frac{\partial \theta^Q_{ht}}{\partial s^Q_{ht}}\right| < 0$ *and* $\left|\frac{\partial \theta^L_{ht}}{\partial s^L_{ht}}\right| < 0$*, indicating that both product demand and labor supply*

[11] The results of Section 2.2 are qualitatively the same under differentiated-Bertrand or differentiated-Cournot competition. However, when allowing for endogenous quality responses in Section 2.3, the first-order conditions are much more tractable in the Cournot case. The reason is that the endogenous component of quality is a function of quantities (Q, L, N), and under Cournot, competitors' quantities are perceived as fixed, which implies competitors' quality (Y) is also perceived as fixed, yielding tractable diversion expressions.

become less elastic as the firm gains market share in the respective market.

The following proposition follows immediately from Lemma 1:

Proposition 1 (Price-setting and wage-setting in a single-product firm)**.** *Suppose quality of care is predetermined. For a single-product firm 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,

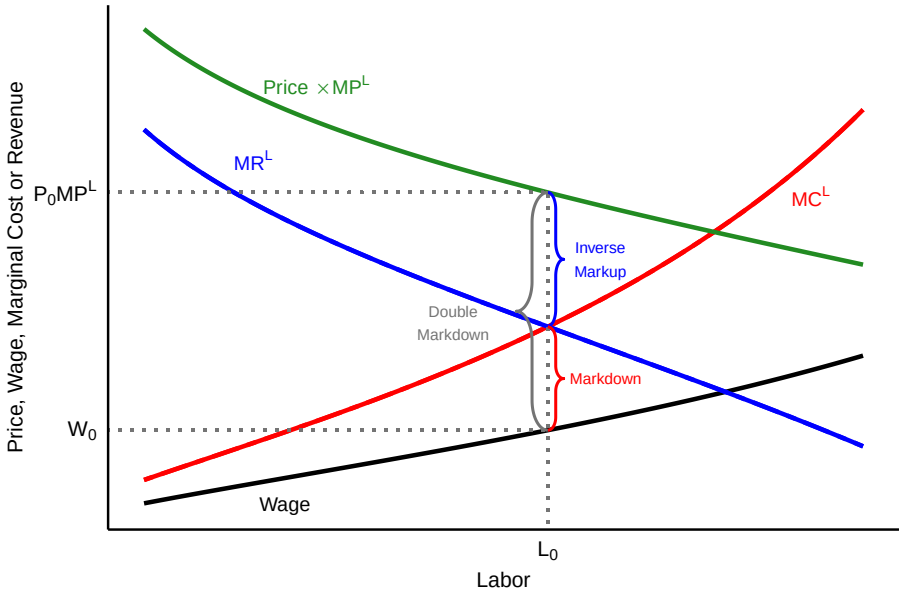
$$\frac{\partial(double\ markup_{ht})}{\partial s^{Q}_{ht}} > \frac{\partial(markup_{ht})}{\partial s^{Q}_{ht}} > 0,$$

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

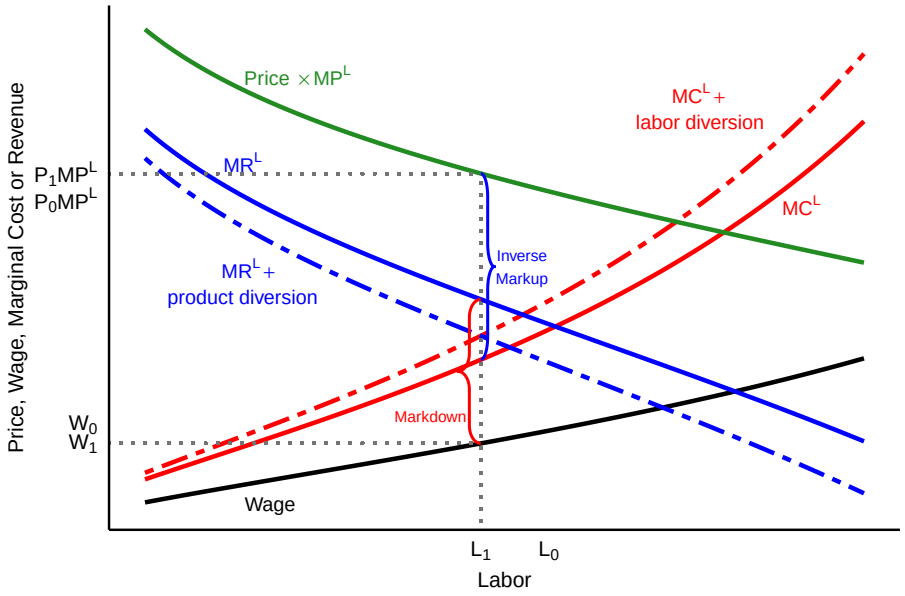
$$\frac{\partial(double\ markdown_{ht})}{\partial s^{L}_{ht}} < \frac{\partial(markdown_{ht})}{\partial s^{L}_{ht}} < 0.$$

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 firm gains market share in the product and labor market, its markup increases and its markdown decreases, indicating two sources of greater



(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 producer h in response to its merger with producer g. Before the merger, the upper diagram visualizes producer 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 producer h internalizes labor diversion from producer g into its perceived marginal cost curve and product diversion from producer g into its perceived marginal revenue curve. It is simulated from the actual model, setting $T_{ht}(L_{ht}) = A_{ht}L_{ht}$.

market power.[12] Since labor and product market shares increase together through the production technology in equation (5), the double markup increases even more than the increase in the markup alone. This implies that, for the same change in the firm's market share, 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 firm also exploits product market power.

Consider how the choices of firms change in response to a merger:

Lemma 2 (First-order condition in a two-product firm)**.** *Suppose quality is predetermined. If producers h and g are commonly-owned in a two-product firm, the first-order condition for 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.[13]

The key result in Lemma 2 is that, from the perspective of producer h, the marginal cost curve is greater and the marginal revenue curve is lesser when it is commonly-owned with local producer g. I now provide intuition for these two channels.

The first channel, *labor market diversion*, increases the perceived marginal cost of production. It expresses that, for producer h to increase employment, it must offer a higher wage. When producer h offers a higher wage, some workers who would have chosen employment at producer g are diverted to prefer producer h. To maintain the same number of employees, producer g must counter by increasing its wage for a given number of workers, which reduces the profits of producer 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.

The second channel, *product market diversion*, decreases the perceived marginal revenue of production. It expresses that, for producer h to increase output sold, it must offer a lower price.

[12] 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 firms 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 firm size increases can be viewed as the labor supply analogue to Marshall's second law. See the related discussion by Autor et al. (2020).

[13] The product diversion term simplifies to $\frac{-s^Q_{gt}}{\beta_P s^Q_{0t}} \text{MP}^L_{ht}$ and the labor diversion term simplifies to $\frac{s^L_{gt}}{\gamma_L s^L_{0t}} W^L_{gt}$.

When producer h offers a lower price, some consumers who would have chosen product g are diverted to prefer product h. To maintain the same number of consumers, producer g must counter by lowering its price, which reduces the profits of producer g. Since h values profits equally across h and g, it internalizes these lost profits at g as a decrease in the marginal revenue at h. The magnitude of the product diversion effect is greater when s^Q_{gt} is greater, as consumers attracted to h are more likely to be diverted from g rather than from other competitors.

The product and labor market diversion effects imply the following in equilibrium:

Proposition 2 (Responses to a merger in the merging firms)**.** *Suppose quality is predetermined. If producers h and g in market m merge at time t to form a two-producer firm H, equilibrium outcomes of the merging firm change as follows:*

(a) The output sold and workers employed decrease for firm H.

(b) The price increases and the wage decreases for both producers h and g.[14]

(c) The price markup and wage markdown strengthen for both producers h and g.

(d) As the merging producers have greater market share, the magnitudes of the effects increase in parts (a-c).[15]

The lower panel of Figure 1 provides the intuition for Proposition 2; additional mathematical details are provided in Appendix A.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 less output sold. Furthermore, since the product demand curve is downward-sloping and the labor supply curve is upward-sloping, a decrease in output 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, we characterize the equilibrium effects of a merger on local competitors.

Proposition 3 (Aggregate and cross-firm spillovers in response to a merger)**.** *Suppose quality is predetermined. If producers h and g in market m merge at time t to form a two-producer firm H, equilibrium outcomes of market m change as follows:*

(a) Each non-merging competitor j increases output sold and workers employed.

(b) The outside shares of the product market and the labor market increase.[16]

[14] There exists a special case, in which several improbable conditions must be met, such that at least one but not necessarily both merging producers increases price and decreases wage, which is a weaker result than in part (b). An analysis of this special case is provided by Appendix A.2.

[15] Because $T_{ht}(\cdot)$ is monotonic, a greater labor market share coincides with a greater product market share. In the empirical application, the correlation between product and labor market shares is 0.989 for patient care workers and 0.980 for non-patient care workers.

[16] There exists a special case, in which the larger and smaller merging establishments change output in opposite directions and several other improbable conditions are met, such that at least one but not necessarily both outside shares increase, which is a weaker result than in part (b). An analysis of this special case is provided by Appendix A.2.

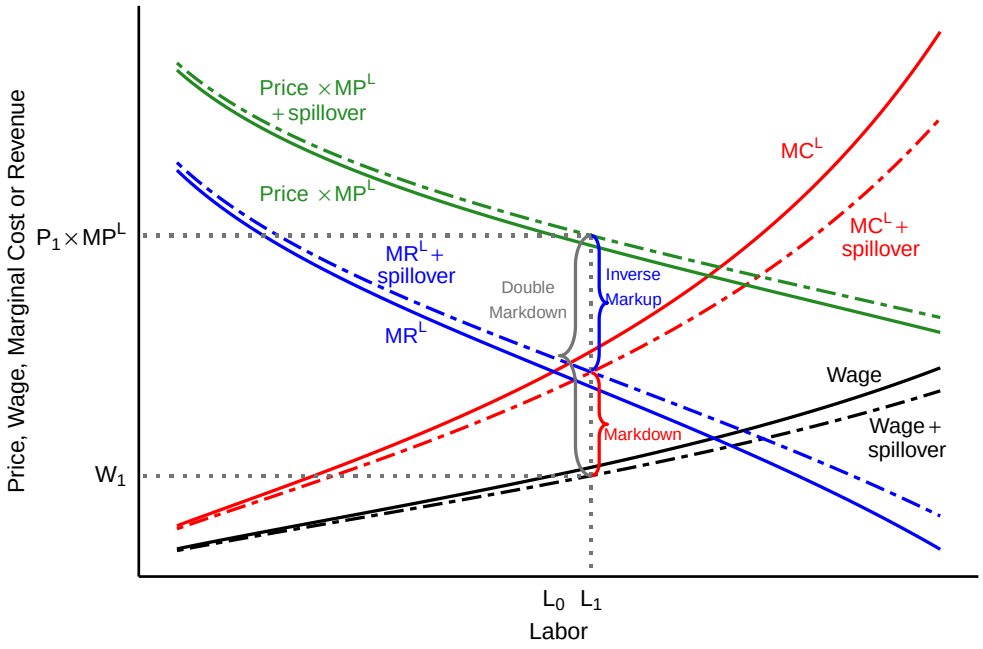


Figure 2: Competitor's Response to a Merger: Visualization of Proposition 3

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 producers. The solid curves represent the firm's considerations before the merger (similar to Figure 1a), while the dashed curves incorporate spillovers from the merger. It is simulated from the actual model, setting $T_{ht}(L_{ht}) = A_{ht}L_{ht}$.

(c) As the merging producers have greater market share, the magnitudes of the effects increase in parts (a-b).

The intuition for Propositions 3(a) is provided by Figure 2, which characterizes the trade-offs faced by local competitors to the merging producers. Since the merging producers optimally reduce output sold and workers employed, other producers in the same market experience an increase in demand—*product market spillovers*—as well as an increase in labor supply—*labor market spillovers*. Together, these spillovers lead to more output sold and more workers employed by competitors. However, the decrease in output and employment in the merging producers is not fully compensated by the increase in output and employment in competitor producers, implying that some consumers go without consumption 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 for competitors. Consider the product market: from the perspective of a competitor, the product demand curve has shifted outward after the merger, so the price is higher for any level of output. This is a force that raises the price. However, the competitor's best response is to increase output 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 output equates marginal

revenue to marginal cost, so the price does 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 competitors after a local merger depends on several factors, including the slope and shift in the product demand curve.

2.3 Comparative Statics for Mergers under Endogenous Quality

I now allow for the endogenous determination 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-producer firm h, the first-order condition with respect to support 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 $F_{ht}^N \equiv \frac{\partial}{\partial N_{ht}} F\left(L_{ht}, N_{ht}\right)$*, and the first-order condition with respect to production labor is,*

$$MC_{ht}^L = MR_{ht}^L + \overbrace{\underbrace{\left(F_{ht}^L - Y_{ht} MP_{ht}^L\right)}_{\substack{\text{marginal quality (L)} \\ \text{with congestion}}} \times \underbrace{\left.\frac{\partial P_{ht}}{\partial Y_{ht}}\right|_{Q_{ht}}}_{=\beta_Y/\beta_P}}^{\text{returns from quality (L)}},$$

where MR_{ht}^L *and* MC_{ht}^L *are defined in Lemma 1 and* $F_{ht}^L \equiv \frac{\partial}{\partial L_{ht}} F\left(L_{ht}, N_{ht}\right)$.

Consider the first-order condition with respect to support labor, N. The term β_Y/β_P is the consumers' marginal rate of substitution between quality and money. For each additional support worker hired, quality increases by F_{ht}^N/Q_{ht}, and the producer can raise its price by $(\beta_Y/\beta_P) \times F_{ht}^N/Q_{ht}$ for each of its Q_{ht} consumers, generating $\beta_Y/\beta_P \times F_{ht}^N$ in additional revenue through the quality channel. Since hiring support labor does not directly change output sold, the first-order condition simply equates the marginal cost of support labor, MC_{ht}^N, to the marginal revenue from the quality improvement these workers generate.

Next, consider the first-order condition with respect to L. Hiring a production worker generates marginal revenue from attracting consumers, MR_{ht}^L, and requires paying a marginal cost, MC_{ht}^L, as in Lemma 1. Furthermore, it generates two competing effects on quality. For a given output level, more production labor increases effective staffing and thus quality via the F_{ht}^L term. However, there

is also a *congestion* effect: since production labor enables more output (by an amount MP^L_{ht}), the producer's staff becomes diluted across a larger consumer population. We say that the producer 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 L via 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 $\theta^{T,L}_{ht}$ is the elasticity of production, $T(\cdot)$, each with respect to L. This expression makes clear that congestion occurs if an increase in production labor expands output proportionally more than it expands effective staffing, causing quality (effective staffing per consumer) to fall.

Given endogenous quality provision, consider a merger between producers h and g:

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

$$MC^N_{ht} \quad + \quad \underbrace{\frac{\partial W^N_{gt}}{\partial s^N_{ht}} s^N_{gt}}_{\text{labor diversion in N (+)}} \quad = \quad \frac{\beta_Y}{\beta_P} F^N_{ht},$$

and the first-order condition for producer h with respect to production labor for h is,

$$MC^L_{ht} + \underbrace{\frac{\partial W^L_{gt}}{\partial s^L_{ht}} s^L_{gt}}_{\text{labor diversion in L (+)}} \quad = \quad MR^L_{ht} + \frac{\beta_Y}{\beta_P}\left(F^L_{ht} - Y_{ht} MP^L_{ht}\right) + \underbrace{\frac{\partial P_{gt}}{\partial s^Q_{ht}} s^Q_{gt} MP^L_{ht}}_{\text{product diversion (−)}} .$$

The effect of the merger on support labor is straightforward: the marginal revenue from quality provision remains unchanged, $\beta_Y/\beta_P \times F^N_{ht}$, because the consumers' valuation of quality and the quality 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, support employment must fall at the merging producers. Figure 3 illustrates the producer's quality of care decision in terms of choosing N for the simple case in which $F(L, N) = L + N$. Holding the output level 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 support 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 support labor and thus lower product quality.

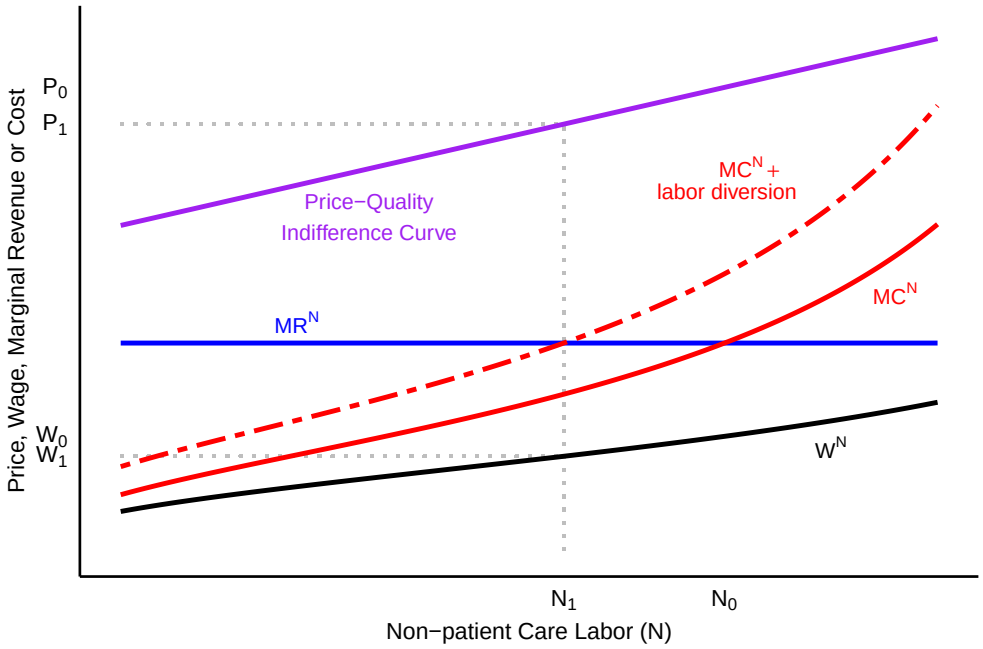


Figure 3: Quality-setting in Response to a Merger: Visualization of Lemma 4

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

While Figure 3 makes clear why support labor is reduced by a merger and how this places downward pressure on quality, the overall effect of a merger on quality is more nuanced because the number of production workers and the output level 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 producers h and g in market m merge at time t to form a two-producer firm H, equilibrium outcomes of market m are as follows:*

(a) Output sold and production labor employment decrease in firm H.

(b) For support labor, the wage decreases and the markdown strengthens for both producers h and g, while employment decreases in firm H.

(c) Each non-merging competitor j increases output, the number of production workers employed, and the number of support workers employed.

(d) The outside shares increase in the product market, the production labor market, and the support labor market.

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

The endogeneity of quality does not overturn the quantity-related predictions of Proposition 2: merging producers reduce output sold and the number of workers employed—both for production and support workers. Competitors respond by increasing output and employment of each occupation, while the outside shares increase for consumers and both types of workers. The key difference is that, in the presence of endogenous quality, price and quality responses are interdependent: a quality reduction places downward-pressure on the price, so if the producer operates in the non-congested region in which quality decreases in response to a merger, the price effect is attenuated and, in extreme cases, may even be eliminated entirely. Conversely, if the producer operates in a sufficiently congested region such that a merger increases quality, this reinforces the increase in price. Furthermore, since marginal revenues from quality may increase or decrease in response to a merger depending on the degree of congestion, it is possible for the marginal revenue of production labor to increase sufficiently that the wage is no longer guaranteed to decrease. Additional mathematical details are provided in Appendix A.3.

3 Data Sources and Descriptive Statistics

I now test the model predictions in the context of the US hospital industry. The administrative panel data from the US hospital industry provide a particularly apt empirical context for several reasons. First, the event study design requires a large number of mergers satisfying high-concentration thresholds in many distinct markets. About one-fourth of all hospital markets in the US experienced such mergers, providing an unusually large sample. Second, while it is rare for an industry-wide panel dataset in the US to include both product market outcomes—prices and output—and labor market outcomes—wages and employment, our data include all four variables. Third, our data provide several features that are unusual in other US administrative labor market datasets: hourly wages rather than annual earnings, employment by occupational category, and wage and employment measures at the establishment level rather than firm-wide (none of which are available in IRS W-2 or Census LEHD data). Fourth, while 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.

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, P_{ht}, 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").[17] 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^L_{ht}, and the wage of non-patient care workers, W^N_{ht}, as the average wage per hour for the patient and non-patient care occupations, respectively.

All monetary variables are inflation-adjusted to 2018 USD using the PCE-PI index (BEA, 2025) and winsorized above and below to protect against influential outliers.

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

[17] 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 (BLS, 2023).

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.

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 consolidations and associated event years.

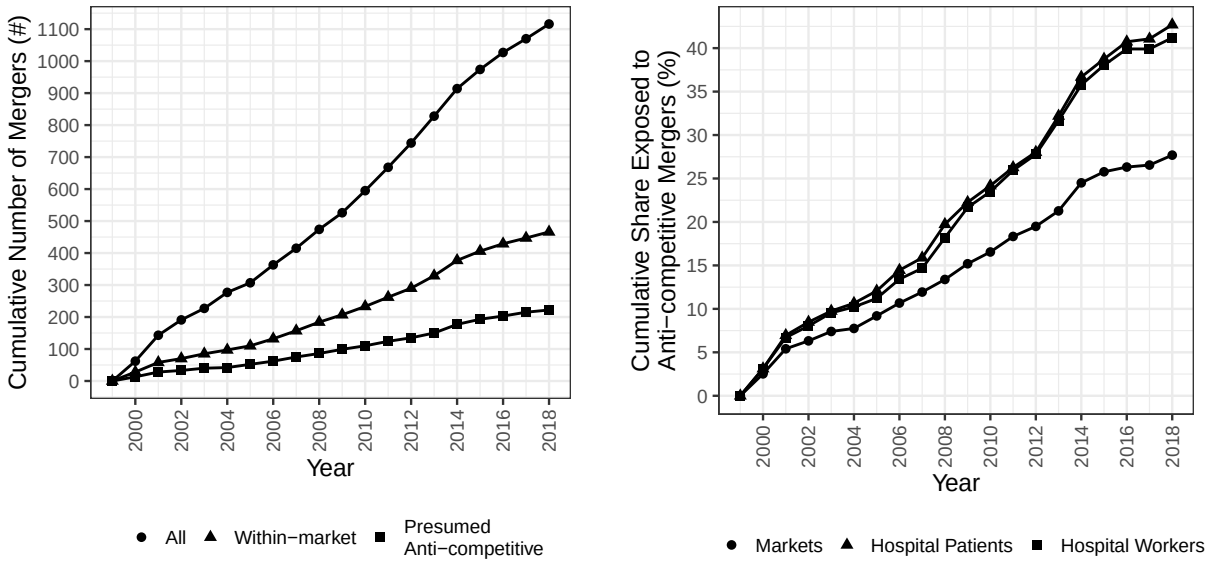
For the regression analysis, I follow Finkelstein, Gentzkow, and Williams (2021) for patient care markets and Prager and Schmitt (2021) for hospital labor markets by using commuting zones as the relevant geographic unit.[18] This results in 561 distinct markets across the US that include at least one hospital satisfying the sample criteria. To measure market concentration, I use the Herfindahl-Hirschman Index (HHI).[19] In the product market, denoting the inside market share by $\tilde{s}^{Q}_{ht}$, the HHI is defined by $\text{HHI} = \sum_h (\tilde{s}^{Q}_{ht})^2 \times 10,000$.[20] 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 panel, 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. Panel (a) displays the number of hospitals per market, showing that the median market has only three hospitals. The

[18] Results are quite similar if defining a market as a circle of 30-mile radius around each hospital, which is similar to the market concept considered by Brot et al. (2024).

[19] 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 courts (US DOJ and FTC, 2023). 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 to measure market power, as discussed above and estimated using the empirical model below.

[20] The inside market share is the share of *observed* patients. It differs from the theoretical market share, which is the share of *potential* patients, inclusive of the outside share. Formally, $\tilde{s}^{Q}_{ht} = s^{Q}_{ht}/(1 - s^{Q}_{0t})$. The inside market share is widely used in reduced-form studies of mergers, as it does not require a model.



(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).

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. Panel (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 three systems to two systems. Panel (c) presents the distribution of HHI across markets, and panel (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., $\Delta\text{HHI} \neq 0$). 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 such mergers occurred.

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 control units 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,

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

where $Y_{h,t+e}$ represents the outcome for hospital h at event time e relative to its merger year t, and

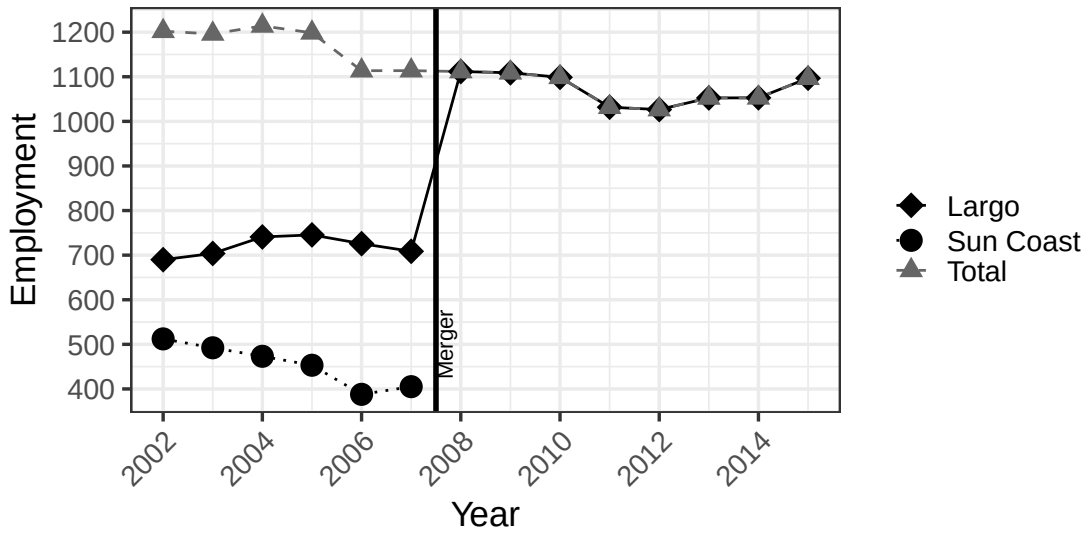


Figure 5: Example of Missing Data and Time-consistent Merging Firm Concept

Notes: This figure presents the employment reported to CMS by Largo Medical Center (diamond symbols) and Sun Coast Hospital (circles). These two hospitals, located near Tampa, Florida, completed their merger by 2008. After merging, Sun Coast stopped filing a separate report, instead consolidating its employment and other outcomes under Largo's identification number. The triangle symbols represent the time-consistent merging firm concept.

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:

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

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. Figure 5 provides an example: when Largo Medical Center (diamond symbols) merged with Sun Coast Hospital (circle symbols) 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 after 2008, 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.[21] For each merger, a time-consistent entity is defined by aggregating outcomes (through sums or averages, as appropriate) across all hospitals that eventually consolidate, creating a consistent unit that can be followed over time. This approach allows me to determine whether outcomes such as patient volume, employment, prices, and wages change after consolidation, even when individual hospital data are no longer separately reported. In Figure 5, for example, I would estimate DiD using total employment (triangle symbols) rather than only using Largo's reported employment (diamonds) as the outcome of interest.

4.2 Direct Effects of Mergers on the Merging Hospitals

The DiD estimation results are summarized in Table 1. Following consolidation, the number of patients treated by merged hospitals decreases substantially by more than 4%. As demonstrated in Figure 6, 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 7%. 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 hospital consolidation.

Turning to the labor market, wages decrease by nearly 2% overall. 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 12%. Comparing panels in Figure 6, it is apparent that employment and wage reductions strengthen over time. I also examine changes in the occupational mix of workers; the main estimates are reported in Table 1 with corresponding event studies in Appendix Figure A2. The reduction in employment is nearly 13% for non-patient care workers and nearly 9% for patient care workers, while the reduction in wages is about 4% 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 nearly 7% after five years, as the reduction in labor is proportionally

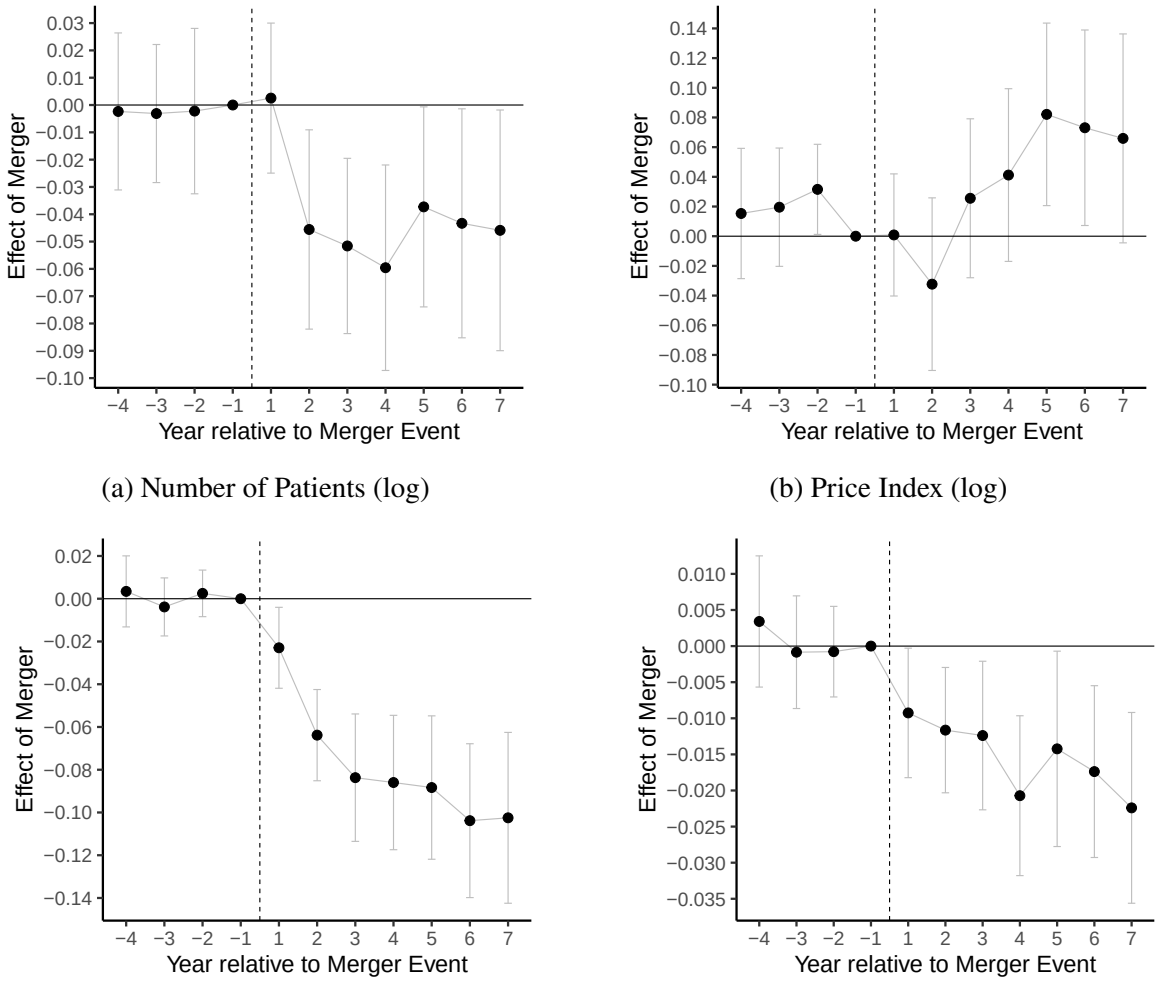
[21] The issue is prevalent in our sample: out of the 147 mergers utilized in the DiD analysis, 30 are subject to post-merger consolidated reporting among the merging hospitals.

	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)	147	411	0.848	-0.003	-0.052	-0.042
				(0.013)	(0.016)	(0.019)
Price Index (log)	146	405	0.377	0.022	0.011	0.074
				(0.015)	(0.025)	(0.030)
Staffing Ratio (log)	147	411	0.265	0.012	-0.021	-0.066
				(0.012)	(0.013)	(0.016)
Case Mix Index (log)	143	397	0.151	0.000	0.001	0.001
				(0.002)	(0.004)	(0.005)
Covered by Medicaid (share)	147	411	0.076	-0.001	0.000	-0.002
				(0.003)	(0.004)	(0.005)
Panel B.	**Labor Market Outcomes**					
Number of Workers, Headcount (log)	133	373	0.912	0.006	-0.080	-0.126
				(0.009)	(0.022)	(0.027)
Number of Workers, FTE Hours (log)	147	411	0.874	0.001	-0.078	-0.098
				(0.006)	(0.013)	(0.017)
· · · Patient Care Occupations	146	409	0.920	-0.004	-0.061	-0.085
				(0.007)	(0.014)	(0.018)
· · · Non-patient Care Occupations	144	402	0.808	0.001	-0.101	-0.129
				(0.009)	(0.015)	(0.022)
Hourly Wage (log)	147	411	0.128	0.001	-0.015	-0.018
				(0.003)	(0.004)	(0.006)
· · · Patient Care Occupations	146	409	0.127	-0.001	-0.011	-0.018
				(0.005)	(0.005)	(0.007)
· · · Non-patient Care Occupations	144	402	0.154	-0.008	-0.032	-0.043
				(0.005)	(0.007)	(0.009)

Table 1: 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 on propensity score. Outcome SD refers to the standard deviation of the outcome variable.

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 $\delta = 2/3$ (more patient care intensity) and 9% when calibrating $\delta = 1/3$ (more non-patient care intensity). The decline in the staffing per patient suggests a decline in the quality of care in merging hospitals, which we investigate below.



(a) Number of Patients (log)

(b) Price Index (log)

(c) Number of Workers, FTE (log)

(d) Hourly Wage (log)

Figure 6: 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. 95% confidence intervals are displayed as brackets.

4.3 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 local competitors. 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 2, with corresponding event studies in Appendix Figure 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 2, 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 about 3% in the shorter run (statistically significant) and 1.5% 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 2-3% in both the shorter and longer run (significant). This reduction is stronger for non-patient care workers (more than 4%) than patient care workers (about 3%), and both are statistically significant. The average hourly wage at the market level also experiences a decrease of about 1%, though the aggregate wage reduction is only significant for non-patient care workers.

Regarding spillovers, non-merging hospitals significantly increase their volume of patients by nearly 5%. However, the price index does not exhibit statistically significant spillovers. In the labor market, competitors expand employment levels by nearly 6%. This employment growth is much stronger for non-patient care occupations (9%) than patient care occupations (3%). Notably, despite substantially increasing employment, competitor hospitals reduce hourly wages by around 3% for patient and non-patient care occupations. Employment gains and wage reductions are statistically significant for both occupations among local competitors.

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.

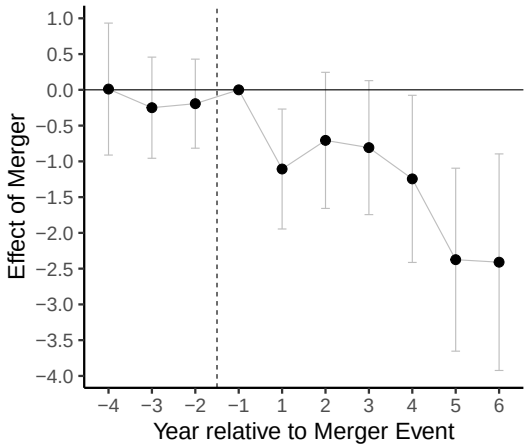
4.4 Direct Effects of Mergers on Quality of Care

Figure 7 presents estimated effects of mergers for several measures of the quality of care in the merging hospitals. A smaller set of mergers are studied because the HCAHPS and HQI databases are only available in more recent years. The top panel of the figure displays the effect for two key measures of patient-reported satisfaction from HCAHPS surveys. The percentage of patients who would recommend the hospital to others declines by more than 1-2 percentage points (pp) post-merger, with this effect becoming stronger and statistically significant over time. A similar pattern is observed for the percentage of patients who report that the hospital provides the highest level of satisfaction (not shown for brevity). When examining more specific ratings of hospital characteristics, patients report a significant decline in the cleanliness of the hospital environment, with a comparable effect for quietness (omitted for brevity). These results imply that patients perceive a deterioration in their care following hospital mergers.

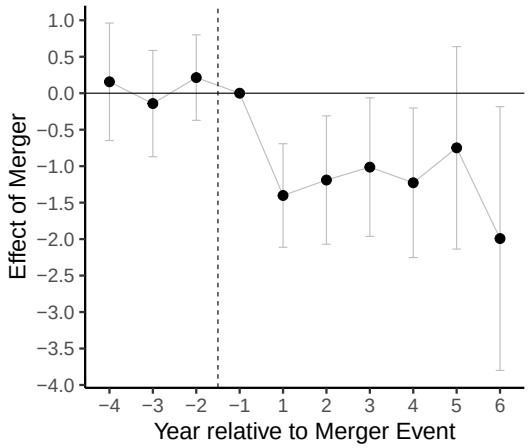
	Treatment Group Summary			DiD Effects of the Merger		
	Mergers	Hospitals	Outcome SD	Before Merger	After Merger	
Event Times:	{-1}			{-2,-3,-4}	{2,3,4}	{5,6,7}
Panel A.	**Market-wide Aggregate Effects**					
Number of Patients (log)	128	845	0.978	-0.004	-0.029	-0.015
				(0.008)	(0.011)	(0.014)
Price Index (log)	128	845	0.367	0.014	-0.007	0.001
				(0.015)	(0.021)	(0.025)
Number of Workers (log)	128	845	0.993	-0.001	-0.025	-0.027
				(0.006)	(0.009)	(0.014)
· · · Patient Care Occupations	128	845	1.023	-0.001	-0.022	-0.031
				(0.007)	(0.011)	(0.015)
· · · Non-patient Care Occupations	128	845	0.953	-0.004	-0.036	-0.043
				(0.008)	(0.011)	(0.017)
Hourly Wage (log)	128	845	0.207	-0.002	-0.007	-0.008
				(0.003)	(0.004)	(0.005)
· · · Patient Care Occupations	128	845	0.198	0.000	-0.005	0.000
				(0.004)	(0.005)	(0.006)
· · · Non-patient Care Occupations	128	845	0.233	-0.001	-0.014	-0.024
				(0.004)	(0.006)	(0.008)
Panel B.	**Within-Market Spillover Effects**					
Number of Patients (log)	87	245	1.188	-0.009	0.011	0.047
				(0.013)	(0.015)	(0.021)
Price Index (log)	87	245	0.412	0.003	0.003	0.010
				(0.021)	(0.027)	(0.032)
Number of Workers (log)	87	245	1.094	-0.002	0.039	0.057
				(0.009)	(0.013)	(0.020)
· · · Patient Care Occupations	87	245	1.184	-0.003	0.028	0.031
				(0.010)	(0.015)	(0.022)
· · · Non-patient Care Occupations	87	245	0.997	-0.001	0.043	0.089
				(0.012)	(0.015)	(0.023)
Hourly Wage (log)	87	245	0.218	-0.003	-0.014	-0.031
				(0.005)	(0.006)	(0.007)
· · · Patient Care Occupations	87	245	0.216	-0.001	-0.016	-0.026
				(0.006)	(0.008)	(0.010)
· · · Non-patient Care Occupations	87	245	0.240	-0.007	-0.022	-0.032
				(0.008)	(0.009)	(0.010)

Table 2: Aggregate and Spillover Effects of Mergers

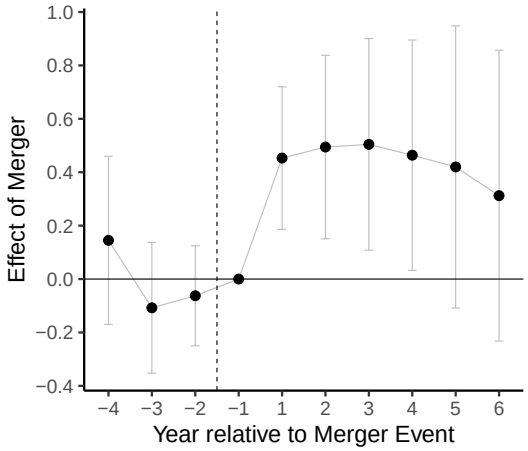
Notes: This table presents difference-in-differences (DiD) estimates of the effects of mergers on outcomes aggregated to the CZ level. Treated CZs experienced a "presumed anti-competitive" merger, as defined in the text. Market-wide outcomes aggregate all hospitals in the treated CZ, while within-market 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 on propensity score. Outcome SD refers to the standard deviation of the outcome variable.



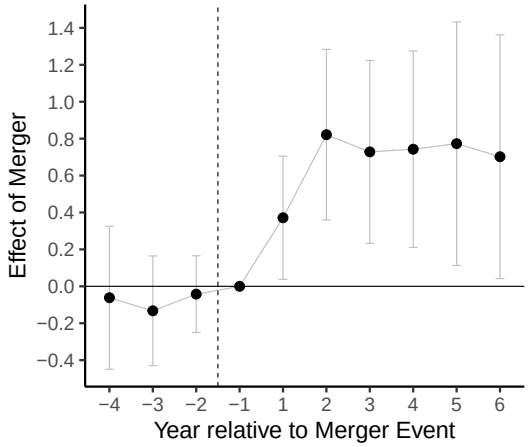
(a) Recommend Hospital (pp)



(b) Cleanliness Rating (pp)



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



(d) Pneumonia 30-Day Mortality (pp)

Figure 7: 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. 95% confidence intervals are displayed as brackets.

The bottom panel of the figure uses 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 0.8pp (relative to the national average mortality rate of 13%), while heart failure mortality increases by as much as 0.5pp (relative to the national average mortality rate of 12%). To place these estimates in context, Cooper, Doyle, Graves, and Gruber (2022) find that a one standard deviation increase in the hospital price distribution is associated with a 0.5pp 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.

4.5 Alternative Interpretations of Findings

Our findings are consistent with the large set of predictions provided by Propositions 2-4 for the patient care market, the two labor markets, 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.

Patient Composition. The model suggests that the increase in price is a consequence of product market power, amplified by labor market power. 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 hospital's average price. Similarly, the increase in mortality rates among heart failure and pneumonia patients could arise from selecting less-healthy patients.

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 1, 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 1, 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, 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 labor market power, amplified by product market power. 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 nurses could mechanically decrease.

To investigate skill composition changes directly, I utilize supplementary data from the AHA Annual Survey. This survey contains a finer breakdown of 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 verified in Panel B of Table 1 that the surveyed headcount employment effects are similar, if slightly larger, than FTE employment effects. Using the headcount data, I compare the two categories of nurses that are available across the full sample time frame: higher-skilled Registered Nurses versus lower-skilled Licensed Practical Nurses and Nursing Assistants. The effect of mergers on the share of nurse employment in the high-skilled category is 0.0132 (standard error 0.0074), indicating a marginally significant shift towards higher-skilled nurse occupations. Thus, to the extent that there is any composition bias, it actually biases us towards finding positive wage effects for patient care workers; 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 is a consequence of labor market power and incentives to reduce quality. However, an alternative explanation for decreased employment is that administrators become redundant after mergers because the same administrator can frictionlessly perform the same task across multiple hospitals. Of course, the fact that mergers lead to a large reduction in the number of patients treated and the number of patient care workers employed is the opposite of what one would predict from greater administrative efficiencies alone. Still, one may wish to test directly for reductions in fixed costs.

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 population 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, while variable labor costs are affected.

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 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 a quality reduction dampens the price increase.

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 when holding the number of patients fixed (Gowrisankaran et al., 2015; Ho and Lee, 2017). Such models are successful in generating large price responses to hospital mergers, but struggle to generate patient volume responses and cross-firm diversion. Indeed, as shown by Ho and Lee (2019), it would be counter to the bargaining model to observe large changes in patient volume after mergers, since both insurers and hospitals are better off renegotiating prices than dropping hospitals from insurer networks.

Rather than choosing only the oligopoly mechanism (which cannot rationalize disproportionately large price increases) or the bargaining mechanism (which cannot rationalize large reductions in patient volume), we propose in the next section to account for both mechanisms in an empirically-tractable model extension. Nonetheless, the oligopoly mechanism remains our primary focus, as it is relevant both in the hospital context and other contexts.

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 four goals of this section. The first goal is to provide a practical estimation strategy for recovering the model parameters. The 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. The third goal is to characterize product demand and labor supply elasticities as well as market power—markdowns and markups—using the model-implied parameter distributions. The fourth goal is to examine how biased ex ante merger evaluation would become if labor (product) market power were ignored when evaluating consumer (worker) harm.

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, $F(L_{ht}, N_{ht}) = (\delta\,(L_{ht})^{\rho} + (1-\delta)\,(N_{ht})^{\rho})^{\phi/\rho}$. 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.

In order to better represent the institutional features of the US hospital industry, the empirical model includes an extension: markups that hospitals charge to insurers. 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, the insurer markup κ_{ht} satisfies the accounting identity $P_{ht}^{\text{hos}} = \kappa_{ht} P_{ht}^{\text{pat}}$.[22] 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 a reduced-form representation of the hospital gaining bargaining power over insurers, in the spirit of the hospital-insurer bargaining literature (Gowrisankaran et al., 2015; Ho and Lee, 2017, 2019).

To permit reduced-form bargaining power effects, I allow κ_{ht} to 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. The baseline value of κ_{ht} is obtained by inverting the first-order condition, and the proportional shift parameter $\bar{\kappa}_{\Delta}$ is chosen to best fit the simulated merger impacts, as described below. Appendix C.1 provides the extended first-order conditions used in the estimation that allow the hospital to take into consideration κ_{ht} and $\bar{\kappa}_{\Delta}$.

Estimation Strategy In industrial organization, it is common to estimate structural parameters, such as consumer preferences, using cross-sectional instruments. These 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 restrictions can be difficult to rationalize. As an alternative, I propose to infer the model parameters from the ex post 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 $\Lambda_{ht} \equiv (\xi_{ht}^{Q}, \xi_{ht}^{L}, \xi_{ht}^{N}, A_{ht}, \kappa_{ht})$, except for proportional changes in κ_{ht}, as noted above. Appendix C.2 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

[22] 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 existence of heterogeneous κ_{ht}: if κ_{ht} is constant, $\Delta \log P_{ht}^{\text{hos}} = \Delta \log P_{ht}^{\text{pat}}$, 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 on insurer markups.

estimate of Ξ, it is straightforward to recover the unobserved heterogeneity terms, Λ_{ht}, using model inversion; the model inversion equations are provided in Appendix C.3.

In practice, I estimate the model using the method of simulated moments (MSM). Details are provided in Appendix C.3. 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 shares—a standard challenge in merger evaluation.[23] 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 25% of inpatient treatments are diverted outside at baseline, motivated by the estimate of Dingel, Gottlieb, Lozinski, and Mourot (2024) that one-fifth of inpatient and outpatient Medicare treatments are out of region. In the labor market, we suppose 40% of hospital workers are diverted outside at baseline, motivated by more than half of nurses being directly employed by hospitals (BLS, 2023).

5.2 Parameter Estimates and Simulated Merger Effects

The parameter estimates and fit of the MSM estimation are presented in Table 3.

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 2.9. This translates to patients being willing to sacrifice about 0.44 standard deviations in the price distribution to improve one standard deviation in the quality distribution.[24] The estimated labor preference for the log-wage, γ_E, is estimated to be 5.6 for patient care labor and 4.5 for

[23] For example, Miller and Weinberg (2017) calibrate the outside share in the beer industry to 50%.

[24] A standard deviation in the patient price distribution is 0.71 (measured in $1,000 USD) and one standard deviation in the quality distribution is 0.11 (no natural units). The ratio of standard deviations is $0.11/0.71 \approx 0.15$. Then, $2.9 \times 0.15 \approx 0.44$ is the price of one standard deviation of quality along the indifference curve, measured in standard deviations of the patient price.

Panel A. MSM Parameter Estimates			Panel B. Simulated Moment Fit		
Parameter	**Value**	**Description**	**Moment**	**Target**	**Simulated**
Patient Preferences			**Product Market: Patients**		
β_P	1.932	Disutility from price	Direct: $\Delta \log P_h$	0.042	0.035
β_Y	5.558	Utility from quality	Direct: $\Delta \log Q_h$	−0.047	−0.058
			Spillover: $\Delta \log \sum_{j \neq h} Q_j$	0.029	0.009
Labor Preferences			Aggregate: $\Delta \log \sum_j Q_j$	−0.022	−0.018
γ_L	5.606	Log-wage utility: Patient care			
γ_N	4.511	Log-wage utility: Non-patient care	**Labor Market: Patient Care**		
			Direct: $\Delta \log W_h^L$	−0.014	−0.023
Technology of Care			Direct: $\Delta \log L_h$	−0.073	−0.110
α	0.530	Quantity: Output elasticity	Spillover: $\Delta \log \sum_{j \neq h} L_j$	0.030	0.017
δ	0.384	Quality: Patient care intensity	Aggregate: $\Delta \log \sum_j L_j$	−0.027	−0.030
ρ	−1.575	Quality: Elasticity of substitution			
ϕ	1.223	Quality: Returns to scale	**Labor Market: Non-Patient Care**		
			Direct: $\Delta \log W_h^N$	−0.038	−0.028
Bargaining Power			Direct: $\Delta \log N_h$	−0.115	−0.113
$\bar{\kappa}_\Delta$	0.022	Change in log insurer markup	Spillover: $\Delta \log \sum_{j \neq h} N_j$	0.066	0.018
			Aggregate: $\Delta \log \sum_j N_j$	−0.039	−0.020
Calibrated Outside Shares					
s_0^Q	0.250	Patients	**Quality of Care**		
s_0^L	0.400	Labor: Patient care	Direct: $\Delta \log(\text{SR}_h)$	−0.044	−0.053
s_0^N	0.400	Labor: Non-patient care	Direct: $\Delta \log(Y_h)$		−0.079

Table 3: 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 1 and 2. "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 in 2018 USD.

non-patient care labor.[25] 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 0.85 for patient care labor and 0.82 for non-patient care labor, such that workers are paid at least 12-15% 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 0.53, indicating substantially diminishing returns. In quality production, we find that the intensity of patient care labor, δ, is about 0.38, indicating that non-patient care labor is moderately more productive than patient care labor in quality provision. The elasticity of substitution between patient and non-patient care labor, $1/(1-\rho)$, is estimated to be about 0.39, indicating that L and N are gross complements in the

[25] Our merger-based instrument yields an estimate consistent with the 3-7 range recently found using several demand-side instruments in the US context (Lamadon et al., 2022; Kroft et al., 2025).

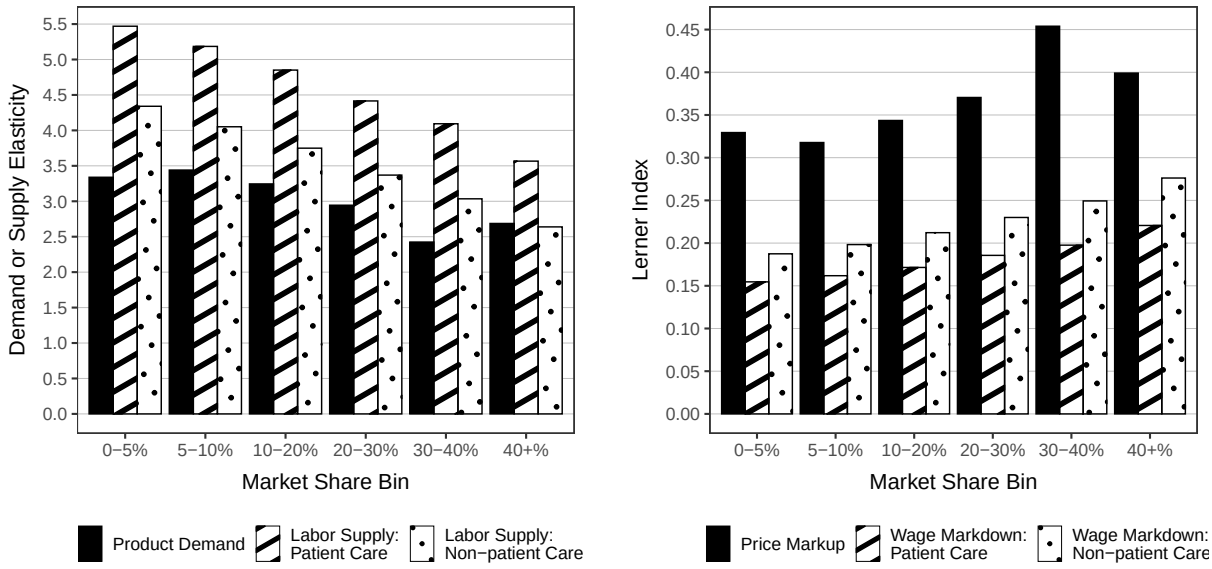
provision of quality of care. Finally, ϕ is estimated to be 1.22, indicating substantially increasing returns to scale in quality, which aligns with the increasing returns to scale in quality mechanism from the spatial trade model of Dingel et al. (2024).

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 (3.5% simulated vs. 4.2% estimated) and over-predicts the patient volume decrease (-5.8% vs. -4.7%). The model also provides a close fit for the aggregate reduction in patient volume (-1.8% vs. -2.2%). Since quality of care is not observable, we target the staffing ratio—closely related to the quality production—as an auxiliary moment for indirect inference. The model correctly predicts a staffing ratio decrease (-5.3% vs. -4.4%).

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 (-2.7% vs. -3.0%). For non-patient care labor, the model successfully predicts the direct effects on wages (-2.8% vs. -3.8%) and employment (-11.3% vs. -11.5%). The spillover effects on competitor employment are matched qualitatively though understated.

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 with only 9 global parameters. 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 about 8% among the merging hospitals in response to a merger—more than 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, the increase in the excess markup relative to the patient price, $\bar{\kappa}_\Delta$, is 0.022. In order to interpret this term, we can use the estimated model to decompose the hospital price increase,

$$\underbrace{\mathbb{E}[\Delta \log P_h^{\text{hos}}]}_{4.2\%} = \underbrace{\mathbb{E}[\Delta \log P_h^{\text{pat}}]}_{1.3\%} + \underbrace{\bar{\kappa}_\Delta}_{2.2\%} + \underbrace{\text{residual}}_{0.7\%}.$$



(a) Demand and Supply Elasticities

(b) Lerner Index for Prices and Wages

Figure 8: 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 theoretical market share distribution. In panel b, we ignore pre-merger common ownership for interpretability of the x-axis variable.

Thus, roughly two-thirds of the explained price increase for hospitals is borne only by insurers, consistent with hospitals gaining bargaining power over insurers through mergers, while the other one-third is attributable to a price increase for patients.

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 unobserved heterogeneity, $\Lambda_{ht} \equiv (\xi^Q_{ht}, \xi^L_{ht}, \xi^N_{ht}, A_{ht}, \kappa_{ht})$, we now possess all of the required components to characterize the market power of hospital systems.

The results are presented in Figure 8. We begin with the hospital-specific distribution of product demand and labor supply elasticities, which are presented in panel (a). As 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 5.5 for patient care labor and 4.5 for non-patient care labor. The largest hospitals face much less elastic labor supply, with a labor supply elasticity around 3.5 for patient care and 2.6 for non-patient care occupational categories. Hospital-specific product demand is substantially more inelastic.[26] The smallest hospitals face

[26]Note that this is the elasticity of patient demand with respect to the patient price, P^{pat}_{ht}.

a product demand elasticity of about 3.4, while hospitals with high market share face a product demand elasticity of around 2.5.[27]

In panel (b) of Figure 8, 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. In the notation of Section 2, 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 for prices—the extent to which price exceeds marginal cost, as a share of price—ranges from around 0.32 for the smaller hospitals to 0.40 for the larger hospitals. For patient care workers, we find that the Lerner index for wages—the extent to which the wage falls short of marginal revenue, as a share of marginal revenue—ranges from about 0.15 for the smallest hospitals to 0.22 for the largest hospitals. For non-patient care workers, the Lerner index for wages ranges from about 0.18 for the smallest hospitals to 0.27 for the largest hospitals.

5.4 Implications for Merger Evaluation

In this final section, I take the estimated structural primitives as given and project the impacts of mergers from the quantitative model, as in ex ante merger evaluation in antitrust (Farrell and Shapiro, 2010; Hovenkamp and Shapiro, 2018). I consider two counterfactual experiments which highlight interactions between labor and product markets. In the first scenario, the role of labor diversion effects is shut down. This is implemented by assuming that the merging hospitals remain labor market competitors after the merger, even though they collaborate in the product market. In the second scenario, the role of product diversion effects is shut down by assuming that the merging hospitals remain competitors for patients after the merger. I use the estimated model from Table 3, ignoring gains in insurer bargaining power for now to simplify the focus to oligopoly and oligopsony mechanisms.

The results are presented in Table 4. Panel A examines how the outcomes for patients would be different if we were to shut down labor diversion effects of the merger. We find that the reduction in patient volume would be 19% weaker, which is a non-trivial change. More dramatically, the reduction in quality of care would be weakened by about 45%. However, there is little effect of removing the labor diversion effect on the price level; there is less upward pressure from the oligopoly mechanism, but also less downward pressure from quality provision, such that the price effect slightly increases on net. The increase in the markup (as measured by the Lerner index) due to the merger is about 11% weaker, and the increase in the outside share of the product market is about

[27]For comparison, Gaynor and Vogt (2003) find a hospital-specific patient demand elasticity in the 4-6 range. Due to unobserved heterogeneity in marginal costs of patient treatment and quality of care, the product demand elasticity need not be monotonic in observed market share. Furthermore, the price markups account for heterogeneous insurer markups, which may be non-monotonic in market share.

Table 4: Ex Ante Predicted Merger Effects under Counterfactual Scenarios

Panel A. Outcomes for Patients			**Panel B. Outcomes for Labor (Patient Care)**		
	Baseline	**No Labor Diversion**		**Baseline**	**No Product Diversion**
Quantity (log)	-0.071	-0.057	Employment (log)	-0.134	-0.028
	(100.0%)	(80.8%)		(100.0%)	(21.2%)
Price (log)	0.011	0.014	Wage (log)	-0.028	-0.006
	(100.0%)	(122.7%)		(100.0%)	(21.1%)
Markup (Lerner)	0.054	0.048	Markdown (Lerner)	0.095	0.013
	(100.0%)	(88.8%)		(100.0%)	(13.8%)
Quality of Care (log)	-0.118	-0.065			
	(100.0%)	(55.4%)			
Outside share (log)	0.031	0.028	Outside share (log)	0.024	0.005
	(100.0%)	(87.8%)		(100.0%)	(20.8%)

Notes: This table presents the predicted changes in key outcomes following a merger under the estimated baseline model and various counterfactual scenarios. Gains in insurer bargaining power are ignored to simplify the focus to oligopoly and oligopsony mechanisms. Markup and markdown effects are reported as changes in Lerner indices (defined in the text), while other outcomes show log changes. Percentages in parentheses represent the relative magnitude of the counterfactual effect compared to the baseline effect.

12% weaker. Together, these results tell us that labor diversion effects of mergers meaningfully amplify the reductions in patient quantity and quality, while also somewhat amplifying increases in price markups and patients pushed outside of the local market. Thus, we would understate several channels through which mergers harm consumers if we ignored changes in labor market power.

Panel B examines how the outcomes for patient care workers, such as nurses, would be different if we were to shut down product diversion effects of the merger. We find that, if we only account for labor market power when evaluating the effects of a merger on labor outcomes, we predict merger effects that are only about 19% as large for employment and wage outcomes. The increase in the markdown would be only 14% as large, and the increase in the outside share would be only 19% as large. Results are similar for non-patient care labor, omitted for brevity. This means that product diversion is much more important than labor diversion in determining the harm of mergers for workers. The reason for this is that patients are much more inelastic than workers, leading to diversion effects that are larger in the product market than the labor market. On average, labor diversion effects are valued around $1,100 while product diversion effects are valued around $9,500, despite nearly identical concentration on the labor and product sides. Thus, even though workers are substantially harmed by hospital mergers, that harm is largely a byproduct of the firm's efforts to exploit gains in concentration over inelastic patients, passed as harm to workers through

the technological relationship between output and labor.

6 Concluding Remarks

In sum, this paper has demonstrated that cross-market diversion is theoretically and empirically important when predicting merger-induced harm to workers *or* consumers. Firms that compete in the product market often compete in the labor market as well, and employment and output are fundamentally linked by production. If we ignore labor market competition, we understate consumer harm; if we ignore product market competition, we understate worker harm. The extent to which harm is understated depends in practice on the magnitude of labor and product diversion effects. In turn, the diversion effects depend quantitatively on the degree of concentration and the elasticity of demand and supply in the product and labor markets. This paper underscores the need—and provides analytical tools—to integrate labor-side concentration into merger evaluations and antitrust policy.

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, D. (2021). Mergers and Acquisitions, Local Labor Market Concentration, and Worker Outcomes. *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, G. Hanson, M. R. Jones, and B. Setzler (2025). Places versus People: The Ins and Outs of Labor Market Adjustment to Globalization. In *Handbook of Labor Economics*, Volume 6. Elsevier.

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. *Quarterly Journal of Economics 135*(2).

BEA (2025). U.S. Bureau of Economic Analysis. Personal Consumption Expenditures: Chain-type Price Index. Federal Reserve Bank of St. Louis. Accessed August 8, 2025.

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).

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

BLS (2025). Union Members. `https://www.bls.gov/news.release/union2.htm`.

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).

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. *RAND Journal of Economics 50*(2), 286–325.

de Loecker, J., J. Eeckhout, and G. Unger (2020). The Rise of Market Power and the Macroeconomic Implications. *The Quarterly Journal of Economics 135*(2), 561–644.

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

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

Farrell, J. and C. Shapiro (2010). Antitrust Evaluation of Horizontal Mergers: An Economic Alternative to Market Definition. *The B.E. Journal of Theoretical Economics 10*(1), 1–41.

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. *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.

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.

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.

Ho, K. and R. S. Lee (2019). Equilibrium provider networks: Bargaining and exclusion in health care markets. *American Economic Review 109*(2), 473–522.

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

Hovenkamp, H. and C. Shapiro (2018). Horizontal Mergers, Market Structure, and Burdens of Proof. *The Yale Law Journal 127*, 1996.

Kane, C. (2025). Physician Practice Characteristics in 2024. *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 115*(9).

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.

Nevo, A. (2000). Mergers with differentiated products: The case of the ready-to-eat cereal industry. *The RAND Journal of Economics 31*(3), 395–421.

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

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. *Unpublished manuscript*.

Rubens, M., B. Setzler, and C. Yeh (2025). Conduct in US Labor Markets. *Unpublished*.

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.

US DOJ and FTC (2023). Merger Guidelines. *Technical Report*.

Appendix

This is the appendix to "Labor and Product Market Power, Endogenous Quality, and the Consolidation of the US Hospital Industry," by Bradley Setzler, dated August, 2025.

A Mathematical Details

A.1 Properties of Product Demand and Labor Supply

A.1.1 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,

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

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

By Cournot conduct, competitors' quantity variables Q_{jt}, L_{jt}, and N_{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 (closer to zero from below) as market share increases. The inverse cross-price effect of the market share at h on the price of competitor j is,

$$\frac{\partial P_{jt}}{\partial s^Q_{ht}} = \frac{\partial}{\partial s^Q_{ht}}\left(\frac{1}{\beta_P}\left(\beta_Y Y_{jt} + \xi^Q_{jt} + \log s^Q_{0t} - \log s^Q_{jt}\right)\right) = \frac{1}{\beta_P}\frac{\partial \log s^Q_{0t}}{\partial s^Q_{ht}} = \frac{-1}{\beta_P s^Q_{0t}} < 0.$$

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

$$\frac{\partial P_{ht}}{\partial s^Q_{0t}} = \frac{1}{\beta_P s^Q_{0t}} > 0$$

A.1.2 Details on residual labor supply

Recall that,

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

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

$$\log\left(W^E_{ht}\right) = \frac{1}{\gamma_E}\left(\log s^E_{ht} - \log s^E_{0t} - \xi^E_{ht}\right),$$

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

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

$$\frac{\partial W^E_{ht}}{\partial s^E_{ht}} = \frac{W^E_{ht}}{\gamma_E}\left(\frac{1}{s^E_{ht}} + \frac{1}{s^E_{0t}}\right),$$

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

$$\frac{\partial s^E_{ht}}{\partial W^E_{ht}} = \frac{\gamma_E}{W^E_{ht}}\left(\frac{s^E_{0t}s^E_{ht}}{s^E_{0t} + s^E_{ht}}\right) > 0, \quad \theta^E_{ht} = \frac{W^E_{ht}}{s^E_{ht}}\frac{\partial s^E_{ht}}{\partial W^E_{ht}} = \gamma_E\left(\frac{s^E_{0t}}{s^E_{ht} + s^E_{0t}}\right) > 0.$$

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 market share at h on competitor j is,

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

A.2 Proofs for Subsection 2.2

This subsection provides the proofs for the comparative statics of mergers presented in Section 2.2. We denote the pre-merger equilibrium with superscript before and the post-merger equilibrium with superscript after. Without loss of generality, we normalize the total labor force $\bar{L}_{mt} = 1$, such that a hospital's employment level is equal to its labor market share, $L_{jt} = s_{jt}^L$, with a similar normalization for patients. Consequently, the patient care production function is written as $s_{jt}^Q = T_{jt}(s_{jt}^L)$, and the marginal product of labor is a function of the labor share, $\text{MP}_{jt}^L(s_{jt}^L)$.

A.2.1 Proof of Lemma 2

After hospitals h and g merge to form system H, the system's total profit is $\pi_H = \pi_h + \pi_g$. The system chooses the labor share for each of its establishments to maximize total profit. The first-order condition (FOC) with respect to the labor share of establishment h, s_{ht}^L, is:

$$\frac{\partial \pi_H}{\partial s_{ht}^L} = \frac{\partial \pi_h}{\partial s_{ht}^L} + \frac{\partial \pi_g}{\partial s_{ht}^L} = 0$$

The first term is the FOC for a single-establishment firm, $\text{MR}_{ht}^L - \text{MC}_{ht}^L$. The second term captures the externalities. Under Cournot conduct, this is:

$$\frac{\partial \pi_g}{\partial s_{ht}^L} = s_{gt}^Q \frac{\partial P_{gt}}{\partial s_{ht}^L} - s_{gt}^L \frac{\partial W_{gt}^L}{\partial s_{ht}^L}$$

Using the chain rule, $\frac{\partial P_{gt}}{\partial s^L_{ht}} = \frac{\partial P_{gt}}{\partial s^Q_{ht}} \frac{ds^Q_{ht}}{ds^L_{ht}} = \frac{\partial P_{gt}}{\partial s^Q_{ht}} \text{MP}^L_{ht}$. The FOC becomes:

$$(\text{MR}^L_{ht} - \text{MC}^L_{ht}) + \left(s^Q_{gt} \frac{\partial P_{gt}}{\partial s^Q_{ht}} \text{MP}^L_{ht} - s^L_{gt} \frac{\partial W^L_{gt}}{\partial s^L_{ht}} \right) = 0$$

Rearranging this expression gives:

$$\text{MR}^L_{ht} + s^Q_{gt} \frac{\partial P_{gt}}{\partial s^Q_{ht}} \text{MP}^L_{ht} = \text{MC}^L_{ht} + s^L_{gt} \frac{\partial W^L_{gt}}{\partial s^L_{ht}}$$

The term added to the right-hand side is the labor diversion effect. Substituting the expression for the inverse cross-wage effect, $\frac{\partial W^L_{gt}}{\partial s^L_{ht}} = \frac{W^L_{gt}}{\gamma_L s^L_{0t}}$, this term is:

$$\text{labor diversion} = s^L_{gt} \frac{\partial W^L_{gt}}{\partial s^L_{ht}} = s^L_{gt} \left(\frac{W^L_{gt}}{\gamma_L s^L_{0t}} \right) > 0$$

The term added to the left-hand side is the product diversion effect. Substituting the expression for the inverse cross-price effect, $\frac{\partial P_{gt}}{\partial s^Q_{ht}} = \frac{-1}{\beta_P s^Q_{0t}}$, this term is:

$$\text{product diversion} = s^Q_{gt} \frac{\partial P_{gt}}{\partial s^Q_{ht}} \text{MP}^L_{ht} = s^Q_{gt} \left(\frac{-1}{\beta_P s^Q_{0t}} \right) \text{MP}^L_{ht} < 0$$

This completes the proof.

A.2.2 Best Response Functions, given Quality

Before proving the main propositions, we establish how a non-merging competitor responds to changes in the aggregate market conditions, as summarized by the outside option shares.

Lemma 5 (Competitor Best-Response Functions)**.** *For any non-merging competitor hospital j, its profit-maximizing choice of employment and patient volume is an increasing function of the outside option shares. Specifically:*

(a) The optimal labor share, s^L_{jt}, is strictly increasing in the product market outside share, s^Q_{0t}. Consequently, the optimal patient share s^Q_{jt} is also strictly increasing in s^Q_{0t}.

(b) The optimal labor share, s^L_{jt}, is strictly increasing in the labor market outside share, s^L_{0t}. Consequently, the optimal patient share s^Q_{jt} is also strictly increasing in s^L_{0t}.

Proof. The first-order condition (FOC) for a non-merging hospital j sets the marginal cost of labor equal to the marginal revenue product of labor: $\text{MC}^L_{jt} = \text{MR}^L_{jt}$. In equilibrium, the firm's optimal

choice of labor share, s^L_{jt}, is a function of the outside shares s^Q_{0t} and s^L_{0t}. We can therefore write the FOC as an identity that holds for all values of the parameters:

$$\text{MC}^L_{jt}(s^L_{jt}(s^Q_{0t}, s^L_{0t}), s^L_{0t}) \equiv \text{MR}^L_{jt}(s^L_{jt}(s^Q_{0t}, s^L_{0t}), s^Q_{0t})$$

To find how the optimal labor share s^L_{jt} changes with an outside share, we totally differentiate this identity.

For part (a), consider totally differentiate the FOC identity with respect to s^Q_{0t}:

$$\frac{d(\text{MC}^L_{jt})}{ds^Q_{0t}} = \frac{d(\text{MR}^L_{jt})}{ds^Q_{0t}}$$

Using the chain rule on both sides:

$$\frac{\partial \text{MC}^L_{jt}}{\partial s^L_{jt}} \frac{ds^L_{jt}}{ds^Q_{0t}} + \frac{\partial \text{MC}^L_{jt}}{\partial s^Q_{0t}} = \frac{\partial \text{MR}^L_{jt}}{\partial s^L_{jt}} \frac{ds^L_{jt}}{ds^Q_{0t}} + \frac{\partial \text{MR}^L_{jt}}{\partial s^Q_{0t}}$$

The term $\frac{\partial \text{MC}^L_{jt}}{\partial s^Q_{0t}}$ is zero because the marginal cost of labor does not depend directly on the product market outside share. Rearranging the expression to solve for $\frac{ds^L_{jt}}{ds^Q_{0t}}$:

$$\frac{ds^L_{jt}}{ds^Q_{0t}} \left(\frac{\partial \text{MC}^L_{jt}}{\partial s^L_{jt}} - \frac{\partial \text{MR}^L_{jt}}{\partial s^L_{jt}} \right) = \frac{\partial \text{MR}^L_{jt}}{\partial s^Q_{0t}} \implies \frac{ds^L_{jt}}{ds^Q_{0t}} = \frac{\frac{\partial \text{MR}^L_{jt}}{\partial s^Q_{0t}}}{\frac{\partial \text{MC}^L_{jt}}{\partial s^L_{jt}} - \frac{\partial \text{MR}^L_{jt}}{\partial s^L_{jt}}}$$

The denominator is the slope of the marginal cost of labor curve minus the slope of the marginal revenue product of labor curve. The second-order condition requires this to be positive. The numerator is the partial derivative of MR^L_{jt} with respect to s^Q_{0t}, holding s^L_{jt} constant. An increase in s^Q_{0t} raises product prices for any given quantity, which in turn raises the value of the marginal revenue product of labor. Thus, the numerator is positive.

$$\frac{ds^L_{jt}}{ds^Q_{0t}} = \frac{(+)}{(+)} > 0$$

This proves that the optimal labor share increases with the product market outside share. Since $s^Q_{jt} = T_{jt}(s^L_{jt})$, the patient share increases as well.

For part (b), we totally differentiate the FOC identity with respect to s^L_{0t}:

$$\frac{\partial \text{MC}^L_{jt}}{\partial s^L_{jt}} \frac{ds^L_{jt}}{ds^L_{0t}} + \frac{\partial \text{MC}^L_{jt}}{\partial s^L_{0t}} = \frac{\partial \text{MR}^L_{jt}}{\partial s^L_{jt}} \frac{ds^L_{jt}}{ds^L_{0t}} + \frac{\partial \text{MR}^L_{jt}}{\partial s^L_{0t}}$$

Here, the direct effect of s^L_{0t} on marginal revenue product of labor is zero. Rearranging gives:

$$\frac{ds^L_{jt}}{ds^L_{0t}} \left(\frac{\partial \text{MC}^L_{jt}}{\partial s^L_{jt}} - \frac{\partial \text{MR}^L_{jt}}{\partial s^L_{jt}} \right) = -\frac{\partial \text{MC}^L_{jt}}{\partial s^L_{0t}} \implies \frac{ds^L_{jt}}{ds^L_{0t}} = \frac{-\frac{\partial \text{MC}^L_{jt}}{\partial s^L_{0t}}}{\frac{\partial \text{MC}^L_{jt}}{\partial s^L_{jt}} - \frac{\partial \text{MR}^L_{jt}}{\partial s^L_{jt}}}$$

The denominator is again positive. The numerator is the negative of the partial derivative of MC^L_{jt} with respect to the labor outside share. An increase in s^L_{0t} makes labor more available and shifts the marginal cost of labor curve down. Therefore, $\frac{\partial \text{MC}^L_{jt}}{\partial s^L_{0t}} < 0$, which makes the numerator positive.

$$\frac{ds^L_{jt}}{ds^L_{0t}} = \frac{(-)(-)}{(+)} > 0$$

Thus, the optimal labor share also increases with the labor market outside share. □

A.2.3 Proof of Proposition 3

The proof proceeds by contradiction. We assume that the outside share in the product market does not increase after the merger ($\Delta s^Q_{0t} \leq 0$), then show that this violates the FOC for the merged firm.

Assume $\Delta s^Q_{0t} \leq 0$. By Lemma 5(a), a non-increasing outside share implies that the optimal share for every non-merging competitor j also does not increase, $\Delta s^Q_{jt} \leq 0$. This means the change in the total share of all non-merging competitors is non-positive, $\Delta S^Q_C = \sum_{j \notin H} \Delta s^Q_j \leq 0$. The market share identity requires that the sum of all changes in shares is zero: $\Delta S^Q_H + \Delta S^Q_C + \Delta s^Q_{0t} = 0$. Rearranging, $\Delta S^Q_H = -(\Delta S^Q_C + \Delta s^Q_{0t})$. Since both ΔS^Q_C and Δs^Q_{0t} are non-positive, their sum is also non-positive. Therefore, the market share identity requires that the merged firm's total share must not decrease: $\Delta S^Q_H \geq 0$.

We now show that $\Delta S^Q_H \geq 0$ and $\Delta s^Q_{0t} \leq 0$ contradicts the FOC of the merged firm. Due to the diversion effects, the FOC for establishment h changes from $\text{MR}^L_{ht} - \text{MC}^L_{ht} = 0$ before the merger to

$\text{MR}^L_{ht} - \text{MC}^L_{ht} > 0$ after the merger evaluated at the initial choices of the merged firm. Expanding,

$$d(\text{MR}^L_{ht} - \text{MC}^L_{ht}) = \underbrace{\frac{\partial(\text{MR}^L_{ht} - \text{MC}^L_{ht})}{\partial s^L_{ht}}}_{(-)} \Delta s^L_{ht} + \underbrace{\frac{\partial(\text{MR}^L_{ht} - \text{MC}^L_{ht})}{\partial s^Q_{0t}}}_{(+)} \Delta s^Q_{0t} + \underbrace{\frac{\partial(\text{MR}^L_{ht} - \text{MC}^L_{ht})}{\partial s^L_{0t}}}_{(+)} \Delta s^L_{0t} > 0$$

The first coefficient is negative by the second-order condition. The second and third coefficients are positive, as established in the proof of Lemma 5. Except in a special case discussed below, $\Delta S^Q_H \geq 0$ implies $\Delta S^L_H \geq 0$, which in turn implies $\Delta s^L_{0t} \leq 0$ by the market share identity and Lemma 5(b). Since $\Delta s^Q_{0t} \leq 0$ and $\Delta s^L_{0t} \leq 0$, for the entire $d(\text{MR}^L_{ht} - \text{MC}^L_{ht})$ expression to be positive, the first term—which is the only potentially positive term—must be strictly positive:

$$\underbrace{\frac{\partial(\text{MR}^L_{ht} - \text{MC}^L_{ht})}{\partial s^L_{ht}}}_{(-)} \Delta s^L_{ht} > 0 \implies \Delta s^L_{ht} < 0$$

A symmetric argument holds for establishment g. Thus, the assumption that outside shares do not increase leads to the implication that the merged firm must reduce its labor share, and consequently its patient share, $\Delta S^Q_H < 0$, contradicting the earlier result that $\Delta S^Q_H \geq 0$. Thus, it must instead be that $\Delta s^Q_{0t} > 0$. An identical line of reasoning proves that $\Delta s^L_{0t} > 0$. Thus, we have proven Proposition 3(b).

Proposition 3(a) follows directly from Proposition 3(b) and the best-response functions in Lemma 5: since we have proven that $\Delta s^Q_{0t} > 0$ and $\Delta s^L_{0t} > 0$, the lemma immediately implies that $\Delta s^Q_{jt} > 0$ and $\Delta s^L_{jt} > 0$ for all non-merging competitors j. Regarding Proposition 3(c), the magnitudes of the labor and product diversion terms derived in the proof of Lemma 2 are proportional to the merging partner's market shares, s^L_{gt} and s^Q_{gt}. All else equal, a merger between larger establishments creates a larger initial shock to the system, as the internalization effects are stronger. This larger initial shock leads to larger equilibrium adjustments for all firms in the market.

Discussion of special case: The structure of the proof of Proposition 3(b) was to (i) assume $\Delta s^Q_0 \leq 0$, (ii) use the accounting identity and best-response functions to show that this implies $\Delta S^Q_H \geq 0$, (iii) utilize that $\Delta S^Q_H \geq 0$ implies $\Delta S^L_H \geq 0$, (iv) repeat the logic of step (ii) to infer $\Delta s^L_0 \leq 0$ from $\Delta S^L_H \geq 0$, and (v) imposing $\Delta s^Q_0 \leq 0$ and $\Delta s^L_0 \leq 0$ in the expression for the merger-induced change, the contradiction follows from the FOC. Alternatively, we could have assumed *both* $\Delta s^Q_0 \leq 0$ and $\Delta s^L_0 \leq 0$ initially, immediately contradicting the FOC expression, avoiding step (iii) entirely. However, this would only show that at least one of $\Delta s^Q_0 > 0$ and $\Delta s^L_0 > 0$ must be true, which is a weaker result than Proposition 3(b) when step (iii) cannot be utilized.

Under which conditions does $\Delta S_H^Q \geq 0$ imply $\Delta S_H^L \geq 0$, such that step (iii) can be utilized? Suppose that the two merging producers adjust output in the same direction. Then, $\Delta S_H^Q \geq 0$ occurs when both $\Delta s_h^Q \geq 0$ and $\Delta s_g^Q \geq 0$. From the monotonicity of T, this implies $\Delta s_h^L \geq 0$ and $\Delta s_g^L \geq 0$, so step (iii) is valid. Thus, the only cases in which the stronger version of Proposition 3(b) may not hold are those in which the two merging producers adjust output in opposite directions. The opposite-signed response scenario is a well-known special case that arises in theory, but is often dismissed as improbable in practice, in the merger evaluation literature; see related discussions by Farrell and Shapiro (2010) and Nocke and Whinston (2022).

Even in the special case of opposite-signed responses, step (iii) of the proof above is valid under several standard scenarios. First, suppose linearity, $T(s_h^L) = s_h^L$. Then, $\Delta S_H^Q = \Delta S_H^L$, so step (iii) is valid. Second, suppose $T(s_h^L)$ is concave, so $s_h^L = T^{-1}(s_h^Q)$ is convex. Without loss of generality, suppose $s_h^Q > s_g^Q$. If $\Delta s_h^Q \geq 0$, $\Delta S_H^L \geq 0$ follows from Jensen's inequality, so step (iii) is valid. Lastly, note that we do not necessarily need step (iii) to complete the stronger version of the proof: as long as the combined term $\frac{\partial(\mathrm{MR}_{ht}^L - \mathrm{MC}_{ht}^L)}{\partial s_{0t}^L}\Delta s_{0t}^L$ is not larger in magnitude than $\frac{\partial(\mathrm{MR}_{ht}^L - \mathrm{MC}_{ht}^L)}{\partial s_{0t}^Q}\Delta s_{0t}^Q$, $\Delta s_h^L < 0$ is still required to ensure $d(\mathrm{MR}_{ht}^L - \mathrm{MC}_{ht}^L) > 0$ in step (v), yielding the desired contradiction.

A.2.4 Proof of Proposition 2

Regarding Proposition 2(a), this follows from Proposition 3. Recall that the patient share satisfies $\Delta S_H^Q = -\Delta S_C^Q - \Delta s_{0t}^Q$. Since we proved $\Delta S_C^Q > 0$ and $\Delta s_{0t}^Q > 0$, it must be that $\Delta S_H^Q < 0$. The same logic applies to the labor market, yielding $\Delta S_H^L < 0$.

Regarding Proposition 2(b), we first prove that wages must decrease. Assume for contradiction that $W_h^{L,\text{after}} \geq W_h^{L,\text{before}}$. Recall that the change in the FOC requires that $\mathrm{MR}_{ht}^L > \mathrm{MC}_{ht}^L$, since the term (Labor Diversion - Product Diversion) is positive (Lemma 2). Recall that $MC_h^L = W_h^L(1 + 1/\theta_h^L)$, where $W_h^L \propto (s_h^L/s_0^L)^{1/\gamma_L}$ and $(1 + 1/\theta_h^L) = 1 + (1/\gamma_L)(s_h^L/s_0^L + 1)$ are strictly increasing functions of the ratio s_h^L/s_0^L. Since s_0^L is increasing, s_h^L must increase by at least as much as s_0^L to satisfy $W_h^{L,\text{after}} \geq W_h^{L,\text{before}}$. This also implies that MC_h^L increases. Recall also that $MR_h^L = P_h(1 + 1/\theta_h^Q)MP_h^L$. Notice that MP_h^L is weakly decreasing in s_h^L, $P_h(1 + 1/\theta_h^Q)$ is strictly decreasing in s_h^Q/s_0^Q, and $s_h^Q = T_h(s_h^L)$ is strictly increasing in s_h^L. Therefore, except in a special case in which $\Delta(s_h^Q/s_0^Q) < 0$ (discussed below), MR_h^L must decrease if $W_h^{\text{after}} \geq W_h^{L,\text{before}}$. The post-merger first-order condition requires that $MR_h^{L,\text{after}} - MC_h^{\text{after}} \geq$ (Labor Diversion − Product Diversion) > 0. However, our finding that MR_h decreases while MC_h^L increases implies that $MR_h^{L,\text{after}} - MC_h^{\text{after}} < MR_h^{L,\text{before}} - MC_h^{L,\text{before}} = 0$. This result, that $MR_h^{L,\text{after}} - MC_h^{L,\text{after}}$ must be negative, directly contradicts the first-order condition from Lemma 2. The proof that price must increase follows the analogous steps, starting from the assumption that $P_h^{\text{after}} \leq P_h^{\text{before}}$, leading to

the contradiction that $MR_h^{L,\text{after}} < MC_h^{L,\text{after}}$.

Regarding Proposition 2(c), we define the realized price markup as the ratio of price to productivity-adjusted marginal cost of labor, $P_{ht}/(W_{ht}^L/\text{MP}_{ht}^L)$, and the realized wage markdown as the ratio of the wage to the marginal revenue product of labor, $W_{ht}^L/(P_{ht}\text{MP}_{ht}^L)$. For the markup, we proved that price increases ($P_h^{\text{after}} > P_h^{\text{before}}$) and the wage decreases ($W_h^{\text{after}} < W_h^{\text{before}}$). The reduction in labor share at establishment h means $\text{MP}_{ht}^L(s_h^{L,\text{after}}) \geq \text{MP}_{ht}^L(s_h^{L,\text{before}})$ by weakly diminishing returns. In the ratio $P_{ht}/(W_{ht}^L/\text{MP}_{ht}^L)$, the numerator increases while the denominator decreases (as W_{ht}^L falls and MP_{ht}^L rises). Thus, the price markup strengthens. For the markdown, in the ratio $W_{ht}^L/(P_{ht}\text{MP}_{ht}^L)$, the numerator decreases while the denominator increases (as both P_{ht} and MP_{ht}^L rise). Thus, the entire ratio must fall, meaning the wage markdown strengthens.

Finally, part (d) of Proposition 2 follows from the same logic as Proposition 3(c). The diversion effects that drive all subsequent results are larger when the merging firms have larger pre-merger market shares, all else equal.

Discussion of special case: The structure of the proof of Proposition 2(b) was to (i) assume $\Delta W_h^L > 0$; (ii) using the labor supply structure and the earlier result that $\Delta s_0^L > 0$, infer that $\Delta s_h^L > 0$ from $\Delta W_h^L > 0$ which in turn implies that the marginal cost has increased; (iii) using the marginal revenue structure, $\Delta s_h^Q > 0$, and $\Delta(s_h^Q/s_0^Q) > 0$, infer that the marginal revenue has decreased; and (iv) having established that marginal cost increases and marginal revenue decreases, the contradiction follows from the FOC. However, there may be special cases in which step (iii) is invalid because $\Delta(s_h^Q/s_0^Q) < 0$. Alternatively, we could have assumed both $\Delta W_h^L \geq 0$ and $\Delta W_g^L \geq 0$, which immediately leads to a contradiction: both $\Delta s_h^L \geq 0$ and $\Delta s_g^L \geq 0$, so $\Delta S_H^L \geq 0$ and thus $\Delta s_0^L \leq 0$. However, this would only show that at least one of $\Delta W_h^L < 0$ and $\Delta W_g^L < 0$ must be true, which is a weaker result than Proposition 2(b) when step (iii) cannot be utilized.

Under which conditions can we utilize step (iii)? Suppose that the two merging producers adjust labor in the same direction. Then, since $\Delta s_h^L > 0$ from step (ii), and assuming $\Delta s_g^L > 0$, we contradict that $\Delta S_H^L > 0$ and thus $\Delta s_0^L < 0$. Thus, the only cases in which the stronger version of Proposition 2(b) may not hold are those in which the two merging producers adjust labor in opposite directions. Even with opposite-signed responses, we can establish the contradiction under standard assumptions. In particular, consider the constant elasticity production function, $T(s_h^L) = (s_h^L)^\alpha$. Then, using the result in step (ii) that $\Delta \log s_h^L > \Delta \log s_0^L$, it follows that $\Delta \log s_h^Q \geq \Delta \log s_0^Q$ if $\alpha \geq \frac{\Delta \log s_0^Q}{\Delta \log s_h^L}$, from which step (iii) is valid, even with opposite-signed responses among the two merging producers. This condition generalizes: letting θ_h^T denote the elasticity of production with respect to labor, $\theta_h^T \geq \frac{\Delta \log s_0^Q}{\Delta \log s_h^L}$ is a sufficient condition for step (iii) to be valid. Lastly, note that we do not necessarily need step (iii): even if s_0^Q increased so much that MR_h^L increased, as long as

MC_h^L increased even more due to the assumed increase in W_h^L and corresponding increase in s_h^L/s_0^L, step (iv) would still satisfy $d(\text{MR}_{ht}^L - \text{MC}_{ht}^L) < 0$, yielding the desired contradiction.

A.3 Proofs for Subsection 2.3

A.3.1 Proofs of Lemma 3-4

Proof of Lemma 3. For a single-establishment firm, the first-order condition (FOC) with respect to non-patient care labor N_{ht} is:

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

From the inverse product demand $P_{ht} = \frac{1}{\beta_P}(\beta_Y Y_{ht} + \xi_{ht}^Q + \log s_{0t}^Q - \log s_{ht}^Q)$, we have $\frac{\partial P_{ht}}{\partial Y_{ht}} = \beta_Y/\beta_P$. The derivative of quality with respect to non-patient care labor is $\frac{\partial Y_{ht}}{\partial N_{ht}} = \frac{F_{ht}^N}{Q_{ht}}$, where $F_{ht}^N \equiv \frac{\partial}{\partial N_{ht}} F(L_{ht}, N_{ht})$. The marginal cost of non-patient care labor is $\text{MC}_{ht}^N \equiv W_{ht}^N(1 + 1/\theta_{ht}^N)$, where θ_{ht}^N is the residual labor supply elasticity. Substituting these into the FOC gives:

$$\frac{\beta_Y}{\beta_P} \frac{F_{ht}^N}{Q_{ht}} Q_{ht} - \text{MC}_{ht}^N = 0 \implies \text{MC}_{ht}^N = \frac{\beta_Y}{\beta_P} F_{ht}^N \equiv \text{MR}_{ht}^{Y,N}.$$

The FOC with respect to patient care labor L_{ht} is:

$$\left(\frac{\partial P_{ht}}{\partial s_{ht}^Q} \frac{ds_{ht}^Q}{ds_{ht}^L} + \frac{\partial P_{ht}}{\partial Y_{ht}} \frac{\partial Y_{ht}}{\partial s_{ht}^L} \right) s_{ht}^Q + P_{ht} \frac{ds_{ht}^Q}{ds_{ht}^L} - \frac{d(W_{ht}^L s_{ht}^L)}{ds_{ht}^L} = 0.$$

Rearranging terms yields:

$$\underbrace{\left(\left(\frac{\partial P_{ht}}{\partial s_{ht}^Q} \frac{s_{ht}^Q}{P_{ht}} \right) + 1 \right) P_{ht} \text{MP}_{ht}^L}_{\equiv \text{MR}_{ht}^L} + \underbrace{\frac{\partial P_{ht}}{\partial Y_{ht}} \frac{\partial Y_{ht}}{\partial s_{ht}^L} s_{ht}^Q}_{\text{MR}_{ht}^{Y,L}} = \underbrace{W_{ht}^L (1 + 1/\theta_{ht}^L)}_{\equiv \text{MC}_{ht}^L},$$

where $\frac{\partial P_{ht}}{\partial Y_{ht}} = \beta_Y/\beta_P$. The final term, $\frac{\partial Y_{ht}}{\partial s_{ht}^L} s_{ht}^Q$, captures the marginal effect of patient care labor on quality, multiplied by total patient volume. Expanding this term using the quotient rule on the definition of quality, $Y_{ht} = F(\cdot)/s_{ht}^Q$:

$$\frac{\partial Y_{ht}}{\partial s_{ht}^L} = \frac{(\frac{\partial F}{\partial s_{ht}^L}) s_{ht}^Q - F(\cdot)(\frac{ds_{ht}^Q}{ds_{ht}^L})}{(s_{ht}^Q)^2} = \frac{F_{ht}^L s_{ht}^Q - F(\cdot)\text{MP}_{ht}^L}{(s_{ht}^Q)^2} = \frac{F_{ht}^L}{s_{ht}^Q} - Y_{ht} \frac{\text{MP}_{ht}^L}{s_{ht}^Q}$$

Multiplying by s^Q_{ht} gives $\text{MR}^{Y,L}_{ht} = \frac{\partial Y_{ht}}{\partial s^L_{ht}} s^Q_{ht} = F^L_{ht} - Y_{ht}\text{MP}^L_{ht}$. Substituting this into the FOC for L_{ht} completes the proof.

Proof of Lemma 4. After hospitals h and g merge to form system H, the FOCs are derived from the system's total profit, $\pi_H = \pi_h + \pi_g$. The FOC with respect to non-patient care labor N_{ht} is $\frac{\partial \pi_H}{\partial N_{ht}} = \frac{\partial \pi_h}{\partial N_{ht}} + \frac{\partial \pi_g}{\partial N_{ht}} = 0$. The first term is the single-firm FOC derived above, $\frac{\beta_Y}{\beta_P} F^N_{ht} - \text{MC}^N_{ht}$. The second term captures the diversions with respect to hospital g. Under Cournot conduct, competitor quantities are fixed. A change in N_{ht} affects Y_{ht}, which changes P_{ht}. However, since competitor quantities s^Q_{gt} and the outside share s^Q_{0t} are held fixed, there is no diversion effect of N_{ht} on P_{gt}. The only externality is in the market for non-patient care labor:

$$\text{(labor diversion for } N) = \frac{\partial \pi_g}{\partial N_{ht}} = -s^N_{gt} \frac{\partial W^N_{gt}}{\partial N_{ht}}$$

Combining these terms, the full FOC is $(\frac{\beta_Y}{\beta_P} F^N_{ht} - \text{MC}^N_{ht}) - s^N_{gt} \frac{\partial W^N_{gt}}{\partial N_{ht}} = 0$.

The FOC with respect to patient care labor L_{ht} is $\frac{\partial \pi_H}{\partial L_{ht}} = \frac{\partial \pi_h}{\partial L_{ht}} + \frac{\partial \pi_g}{\partial L_{ht}} = 0$. The first term is the single-firm FOC with quality effects derived in Lemma 3. The second term captures the diversions with respect to hospital g. A change in L_{ht} does not affect the market for non-patient care labor. The diversions are therefore identical to those derived in Lemma 2:

$$\frac{\partial \pi_g}{\partial s^L_{ht}} = s^Q_{gt} \frac{\partial P_{gt}}{\partial s^L_{ht}} - s^L_{gt} \frac{\partial W^L_{gt}}{\partial s^L_{ht}} = \text{(product diversion)} - \text{(labor diversion for } L)$$

Combining these terms and rearranging gives the expression in the lemma.

A.3.2 Best Response Functions with Endogenous Quality

With quality as a choice variable, firms choose both patient care labor (s^L_{jt}) and non-patient care labor (s^N_{jt}). Before proving Proposition 4, we must establish how a non-merging competitor responds to changes in the aggregate market conditions in this more complex environment. This new lemma replaces Lemma 5 for the endogenous quality case.

Lemma 6 (Competitor Best-Response Functions with Endogenous Quality)**.** *For any non-merging competitor hospital j in the model with endogenous quality, its profit-maximizing choice of both patient care and non-patient care labor is an increasing function of the outside option shares in all markets. Specifically, for $E \in \{L, N\}$ and $k \in \{Q, L, N\}$:*

$$\frac{ds^E_{jt}}{ds^k_{0t}} > 0$$

Proof. A competitor j chooses its labor shares, s_{jt}^L and s_{jt}^N, to satisfy a system of two first-order conditions (FOCs) derived in Lemma 3:

$$\mathcal{F}_L(s_{jt}^L, s_{jt}^N, s_{0t}^Q, s_{0t}^L) \equiv \text{MC}_{jt}^L - \text{MR}_{jt}^L - \frac{\beta_Y}{\beta_P}(F_{jt}^L - Y_{jt}\text{MP}_{jt}^L) = 0$$
$$\mathcal{F}_N(s_{jt}^L, s_{jt}^N, s_{0t}^N) \equiv \text{MC}_{jt}^N - \frac{\beta_Y}{\beta_P}F_{jt}^N = 0$$

To find the response of the optimal choices (s_{jt}^L, s_{jt}^N) to a change in a parameter (e.g., s_{0t}^Q), we use the implicit function theorem for a system of equations.

The solution is given by:

$$\begin{pmatrix} ds_{jt}^L/ds_{0t}^k \\ ds_{jt}^N/ds_{0t}^k \end{pmatrix} = -[J(\mathcal{F})]^{-1} \begin{pmatrix} \partial\mathcal{F}_L/\partial s_{0t}^k \\ \partial\mathcal{F}_N/\partial s_{0t}^k \end{pmatrix}$$

where $J(\mathcal{F})$ is the Jacobian matrix of the FOC system with respect to the choice variables:

$$J(\mathcal{F}) = \begin{pmatrix} \frac{\partial\mathcal{F}_L}{\partial s_{jt}^L} & \frac{\partial\mathcal{F}_L}{\partial s_{jt}^N} \\ \frac{\partial\mathcal{F}_N}{\partial s_{jt}^L} & \frac{\partial\mathcal{F}_N}{\partial s_{jt}^N} \end{pmatrix}$$

The diagonal terms of the Jacobian, $\frac{\partial\mathcal{F}_L}{\partial s_{jt}^L}$ and $\frac{\partial\mathcal{F}_N}{\partial s_{jt}^N}$, are positive by the second-order conditions for profit maximization. The off-diagonal terms are equal by Young's Theorem. The sign of these off-diagonal terms depends on the cross-partial derivative of the effective staffing function, F_{LN}. We assume patient care and non-patient care labor are not strong substitutes in the production of quality, such that $F_{LN} \geq 0$, which implies the off-diagonal terms of the Jacobian are non-positive. For the equilibrium to be stable, the determinant of the Jacobian, $|J(\mathcal{F})| = \frac{\partial\mathcal{F}_L}{\partial s_{jt}^L}\frac{\partial\mathcal{F}_N}{\partial s_{jt}^N} - (\frac{\partial\mathcal{F}_L}{\partial s_{jt}^N})^2$, must be positive. This is a standard condition requiring the direct effects of the FOCs to be stronger than the cross-effects. The inverse of the Jacobian is $[J(\mathcal{F})]^{-1} = \frac{1}{|J(\mathcal{F})|}\text{adj}(J(\mathcal{F}))$. Under these conditions, all elements of the adjugate matrix, and thus all elements of the inverse Jacobian $[J(\mathcal{F})]^{-1}$, are non-negative.

We now evaluate the vector of partial derivatives with respect to each outside share.

Response to a change in s_{0t}^Q. The FOC for non-patient care labor, $\mathcal{F}_N$, has no direct dependence on s_{0t}^Q, so $\partial\mathcal{F}_N/\partial s_{0t}^Q = 0$. Specifically, neither the MC_{jt}^N term nor the quality-related term $-\frac{\beta_Y}{\beta_P}F_{jt}^N$ directly depend on s_{0t}^Q, as the former is a function of s_{jt}^N and s_{0t}^N, and the latter is a function of only the choice variables L and N. The FOC for patient care labor, $\mathcal{F}_L$, depends directly on s_{0t}^Q through the marginal revenue product term, MR_{jt}^L. Neither MC_{jt}^L nor the marginal revenue from quality, $-\frac{\beta_Y}{\beta_P}(F_{jt}^L - Y_{jt}\text{MP}_{jt}^L)$, have s_{0t}^Q as a direct argument, and the partial derivative of these terms with respect to s_{0t}^Q is therefore zero when holding the choice variables constant. An increase in s_{0t}^Q raises

MR^L_{jt}, so $\partial \mathcal{F}_L / \partial s^Q_{0t} = -\partial \text{MR}^L_{jt} / \partial s^Q_{0t} < 0$. Thus,

$$\begin{pmatrix} ds^L_{jt}/ds^Q_{0t} \\ ds^N_{jt}/ds^Q_{0t} \end{pmatrix} = -\underbrace{[J(\mathcal{F})]^{-1}}_{(+)} \underbrace{\begin{pmatrix} < 0 \\ 0 \end{pmatrix}}_{(-)} = \begin{pmatrix} > 0 \\ > 0 \end{pmatrix}$$

We see that an increase in the product market outside share leads the firm to hire more of both types of labor.

Response to a change in s^L_{0t}. The FOC $\mathcal{F}_N$ does not depend on s^L_{0t} directly, so $\partial \mathcal{F}_N / \partial s^L_{0t} = 0$. This is because neither MC^N_{jt} nor the quality term $-\frac{\beta_Y}{\beta_P} F^N_{jt}$ are direct functions of s^L_{0t}. The FOC $\mathcal{F}_L$ depends directly on s^L_{0t} through MC^L_{jt}. The quality-related component of $\mathcal{F}_L$, which is $-\frac{\beta_Y}{\beta_P}(F^L_{jt} - Y_{jt}\text{MP}^L_{jt})$, does not directly depend on the outside share s^L_{0t}, and the same holds for the MR^L_{jt} term. Thus, the only direct effect on $\mathcal{F}_L$ comes from the change in MC^L_{jt}. An increase in s^L_{0t} lowers the marginal cost of labor, so $\partial \mathcal{F}_L / \partial s^L_{0t} = \partial \text{MC}^L_{jt} / \partial s^L_{0t} < 0$. The resulting best response is:

$$\begin{pmatrix} ds^L_{jt}/ds^L_{0t} \\ ds^N_{jt}/ds^L_{0t} \end{pmatrix} = -\underbrace{[J(\mathcal{F})]^{-1}}_{(+)} \underbrace{\begin{pmatrix} < 0 \\ 0 \end{pmatrix}}_{(-)} = \begin{pmatrix} > 0 \\ > 0 \end{pmatrix}$$

Response to a change in s^N_{0t}. The FOC $\mathcal{F}_L$ does not depend directly on s^N_{0t}, so $\partial \mathcal{F}_L / \partial s^N_{0t} = 0$. This is because none of its components, MC^L_{jt}, MR^L_{jt}, or the quality-related term, are direct functions of s^N_{0t}. The FOC $\mathcal{F}_N$ depends on s^N_{0t} through MC^N_{jt}. The quality term in $\mathcal{F}_N$, which is $-\frac{\beta_Y}{\beta_P} F^N_{jt}$, is a function of the choice variables but does not directly depend on the outside share s^N_{0t}, so the entire direct effect is contained in the marginal cost term. An increase in s^N_{0t} lowers the marginal cost of non-patient care labor, so $\partial \mathcal{F}_N / \partial s^N_{0t} = \partial \text{MC}^N_{jt} / \partial s^N_{0t} < 0$. Thus,

$$\begin{pmatrix} ds^L_{jt}/ds^N_{0t} \\ ds^N_{jt}/ds^N_{0t} \end{pmatrix} = -\underbrace{[J(\mathcal{F})]^{-1}}_{(+)} \underbrace{\begin{pmatrix} 0 \\ < 0 \end{pmatrix}}_{(-)} = \begin{pmatrix} > 0 \\ > 0 \end{pmatrix}$$

Thus, for all three outside option shares, the best response to an increase in any of the three outside option shares is to increase employment of *both* labor types. This completes the proof. □

A.3.3 Proof of Proposition 4

First, we establish that a merger must increase the outside shares in all markets using the best response functions and a proof by contradiction. The remaining parts of the proposition then

follow from this key intermediate result.

Assume for contradiction that at least one of the outside shares does not increase post-merger. For instance, assume the patient outside share does not increase: $\Delta s^Q_{0t} \leq 0$. A parallel argument holds if we assume $\Delta s^L_{0t} \leq 0$ or $\Delta s^N_{0t} \leq 0$. By Lemma 6, each competitor's optimal choice of both patient care and non-patient care labor is an increasing function of all outside option shares. Thus, if all outside shares are non-increasing, every non-merging competitor j will not increase employment of either labor type, which implies patient volume will also not increase ($\Delta s^Q_{jt} \leq 0$). This means the change in the total patient market share of all non-merging competitors is non-positive: $\Delta S^Q_C = \sum_{j \notin H} \Delta s^Q_j \leq 0$. The market share identity for the patient market is $\Delta S^Q_H + \Delta S^Q_C + \Delta s^Q_{0t} = 0$. Given the assumption that $\Delta s^Q_{0t} \leq 0$ and its implication that $\Delta S^Q_C \leq 0$, the merged firm's total share must not decrease: $\Delta S^Q_H \geq 0$. We next show that the outcome $\Delta S^Q_H \geq 0$ contradicts the merged firm FOCs. Note that we require the outside shares to shift in the same direction, which is true except in the special case discussed in the context of Proposition 3 above.

Let us define the pre-merger FOC expressions as $\mathcal{G}_L \equiv \text{MR}^L_{ht} + \text{MR}^{Y,L}_{ht} - \text{MC}^L_{ht} = 0$ and $\mathcal{G}_N \equiv \text{MR}^{Y,N}_{ht} - \text{MC}^N_{ht} = 0$. From Lemma 4, the post-merger equilibrium must satisfy a new set of conditions where the MR versus MC gap equals the net diversion effects. This means the change in the value of the FOC expressions from the pre-merger to the post-merger equilibrium must be strictly positive. We can analyze this required change by taking the total differential of the system:

$$d\mathcal{G}_L = \frac{\partial \mathcal{G}_L}{\partial s^L_{ht}} \Delta s^L_{ht} + \frac{\partial \mathcal{G}_L}{\partial s^N_{ht}} \Delta s^N_{ht} + \frac{\partial \mathcal{G}_L}{\partial s^Q_{0t}} \Delta s^Q_{0t} + \frac{\partial \mathcal{G}_L}{\partial s^L_{0t}} \Delta s^L_{0t} > 0$$
$$d\mathcal{G}_N = \frac{\partial \mathcal{G}_N}{\partial s^L_{ht}} \Delta s^L_{ht} + \frac{\partial \mathcal{G}_N}{\partial s^N_{ht}} \Delta s^N_{ht} + \frac{\partial \mathcal{G}_N}{\partial s^N_{0t}} \Delta s^N_{0t} > 0$$

The partial derivatives $\partial \mathcal{G} / \partial s$ represent the slopes of the marginal profit curves. By the second-order conditions for profit maximization, the own-derivatives are negative: $\partial \mathcal{G}_L / \partial s^L_{ht} < 0$ and $\partial \mathcal{G}_N / \partial s^N_{ht} < 0$. The partial derivatives with respect to the outside shares, $\partial \mathcal{G} / \partial s^k_{0t}$, are positive, as a larger outside market makes conditions more favorable.

We can express this system in matrix form. Let $\Delta \mathbf{s} = [\Delta s^L_{ht}, \Delta s^N_{ht}]^T$. Let $\mathbf{J}_\mathcal{G}$ be the Jacobian matrix of the system with respect to the choice variables. We have,

$$\mathbf{J}_\mathcal{G} \Delta \mathbf{s} + \begin{pmatrix} \frac{\partial \mathcal{G}_L}{\partial s^Q_{0t}} \Delta s^Q_{0t} + \frac{\partial \mathcal{G}_L}{\partial s^L_{0t}} \Delta s^L_{0t} \\ \frac{\partial \mathcal{G}_N}{\partial s^N_{0t}} \Delta s^N_{0t} \end{pmatrix} > \begin{pmatrix} 0 \\ 0 \end{pmatrix}$$

Using $\Delta s^k_{0t} \leq 0$ for all k, and since the coefficients $\partial \mathcal{G} / \partial s^k_{0t}$ are positive, the vector term involving

the outside shares is non-positive. The inequality therefore requires:

$$\mathbf{J}_{\mathcal{G}}\Delta\mathbf{s} > -\begin{pmatrix}\text{non-positive}\\ \text{non-positive}\end{pmatrix} \implies \mathbf{J}_{\mathcal{G}}\Delta\mathbf{s} \geq \begin{pmatrix}0\\0\end{pmatrix}$$

The Jacobian $\mathbf{J}_{\mathcal{G}}$ of the marginal profit functions has negative diagonal elements and non-negative off-diagonal elements (by $F_{LN} \geq 0$). For stability, the matrix $-\mathbf{J}_{\mathcal{G}}$ must have an inverse with all non-negative elements. This property implies that if $\mathbf{J}_{\mathcal{G}}\Delta\mathbf{s}$ is a non-negative vector, then the vector $\Delta\mathbf{s}$ must be non-positive. Therefore, for the FOCs to be satisfied under the assumption of non-increasing outside shares, we must have $\Delta s^L_{ht} \leq 0$ and $\Delta s^N_{ht} \leq 0$. Since the diversion effects that drive the inequality are strictly positive, this must be a strict inequality, which implies that the merged system's total market share must decrease, $\Delta S^Q_H < 0$, which directly contradicts the result from the market share identity that $\Delta S^Q_H \geq 0$. The contradiction is achieved, establishing Proposition 4(c-d).

Since $\Delta s^Q_{0t} > 0$ and $\Delta S^Q_C = \sum_{j\notin H} \Delta s^Q_{jt} > 0$ from the best response function, the market clearing condition requires $\Delta S^Q_H + \Delta S^Q_C + \Delta s^Q_{0t} = 0$, from which it follows that $\Delta S^Q_H < 0$. Therefore, the total number of patients treated by the merged system decreases. The same market clearing argument holds in the two labor markets, establishing Proposition 4(a).

The proof by contradiction of Proposition 4(b) follows the same logic as Proposition 2(b). Assume for contradiction that $W^{N,\text{after}}_{ht} \geq W^{N,\text{before}}_{ht}$. Since s^N_{0t} increases and $W^N_{ht} \propto (s^N_{ht}/s^N_{0t})^{1/\gamma_N}$, the wage increase requires s^N_{ht} to increase proportionally more than s^N_{0t}. This increases MC^N_{ht}. On the other hand, marginal revenue diminishes by concavity of the quality production function. Greater MC and lesser MR violates Lemma 4 with positive labor diversion, achieving the contradiction. The wage markdown $W^N_h/(\frac{\beta_Y}{\beta_Y}F^N_h)$ becomes smaller (i.e. stronger) since W^N_h decreases and F^N_h increases.

Finally, Proposition 4(e) follows from the same logic as Proposition 4(c-d). The diversion effects that drive all subsequent results are larger when the merging firms have larger pre-merger market shares, all else equal.

B 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,

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

Note also that,

$$\log \text{CV}_{mt}^{Q,1} - \log \text{CV}_{mt}^{Q,0} = \log(1/s_0^{Q,1}) - \log(1/s_0^{Q,0}) = -\Delta \log s_0^Q$$

Thus, the increase in the log outside share is equal to the decrease in the log welfare of consumers.

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:

$$\text{CV}_{mt}^{\pi,D} \equiv \sum_h \left(P_{ht}^D Q_{ht}^D - W_{ht}^{L,D} L_{ht}^D - W_{ht}^{N,D} N_{ht}^D \right).$$

The effect on hospital welfare is then,

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

C Estimation Details

C.1 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} can be interpreted as the additional markup on insurers. There is also a markup in P_{ht}^{pat}, so there are effectively two markups on insurers. 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 $F_{ht}^{N} \equiv \frac{\partial}{\partial N_{ht}} F\left(L_{ht}, N_{ht}\right)$ and marginal cost expressions are not affected by the introduction of κ_{ht}. 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 $F_{ht}^{L} \equiv \frac{\partial}{\partial L_{ht}} F\left(L_{ht}, N_{ht}\right)$. The term $\left(1 + 1/\theta_{ht}^{Q}\right)$ now incorporates the patient-paid price in the

elasticity:

$$\theta^{Q}_{ht} = -P^{\text{hos}}_{ht} \frac{s^{Q}_{0t}}{s^{Q}_{ht} + s^{Q}_{0t}} \frac{\beta_P}{\kappa_{ht}}.$$

Lastly, regarding the multi-hospital system, the labor diversion terms are unaffected by the presence of κ_{ht}. The product diversion term becomes $\frac{-s^{Q}_{gt}}{s^{Q}_{0t}} \frac{\kappa_{ht}}{\beta_P} \text{MP}^{L}_{ht}$.

C.2 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 without endogenous quality from 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 $\Delta \log s^{L}_{h} = \Delta \log \frac{L_h}{\bar{L}} = \Delta \log L_h$ 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,

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

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 $\mathbb{E}[\Delta P_h] \approx \mathbb{E}[P_h] \times \mathbb{E}[\Delta \log P_h]$,

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

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, for the model of Section 2.2, 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—are sufficient 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.

C.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 $\tilde{\kappa}_h^*$ by inverting the first-order condition for N:

$$MC_h^{N,*} + \sum_g \frac{W_g^N s_g^{N,*}}{\gamma_N^* s_0^{N,*}} = \tilde{\kappa}_h^* \beta_Y^* F_N^*(L_h, N_h).$$

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, $\Lambda_h^* \equiv (\xi_h^{L,*}, \xi_h^{N,*}, \xi_h^{Q,*}, A_h^*, \tilde{\kappa}_h^*)$, have been recovered. Impose the actual merger within each relevant market m and simulate the post-merger equilibrium. This requires simultaneously solving the full system of $2 \times N_m$ stacked FOCs across all N_m hospitals in the market. The labor diversion and product diversion terms are assigned to each FOC based on post-merger ownership. In the post-merger simulation, the composite insurer markup for each firm ($\tilde{\kappa}_h$) is shifted proportionally by the guessed global parameter, i.e., $\tilde{\kappa}_h' = \tilde{\kappa}_h^* \exp(\bar{\kappa}_\Delta^*)$. 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 appropriate hospitals.

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

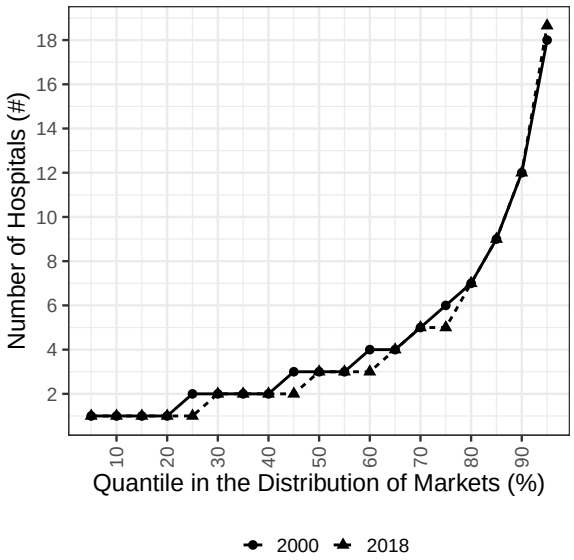
$$\Xi^{msm} = \min_{\Xi^*} (\mathbf{M}^{obs} - \mathbf{M}^{sim}(\Xi^*))' \mathbf{W} (\mathbf{M}^{obs} - \mathbf{M}^{sim}(\Xi^*)).$$

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}.

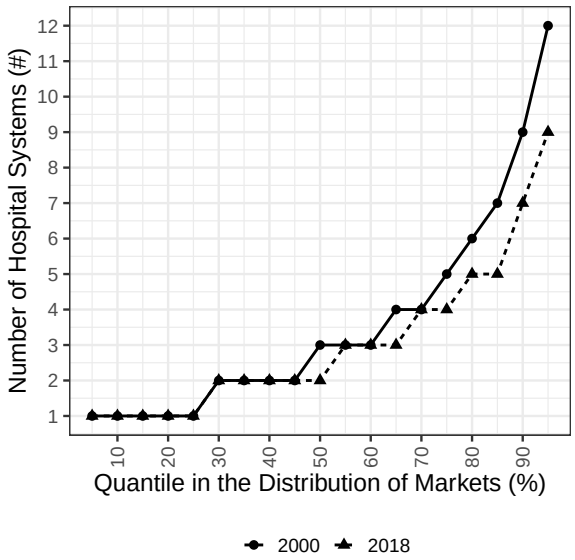
Lastly, note that the estimation procedure recovers the composite parameter $\tilde{\kappa}_h \equiv \kappa_h/\beta_P$, and thus does not separately recover κ_h and β_P. Since the first-order conditions and markups and markdowns can be expressed only in terms of $\tilde{\kappa}_h$ without loss of generality, we do not need to

separate 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 $\bar{\kappa} = 5$, so we use this value when interpreting the magnitude of β_P.

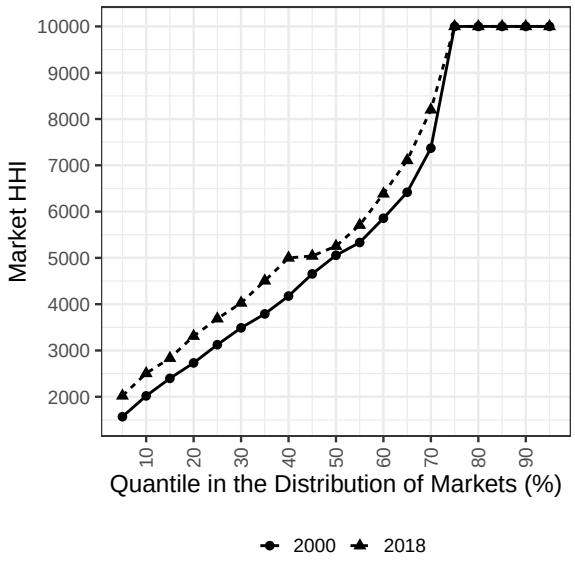
D Additional Figures



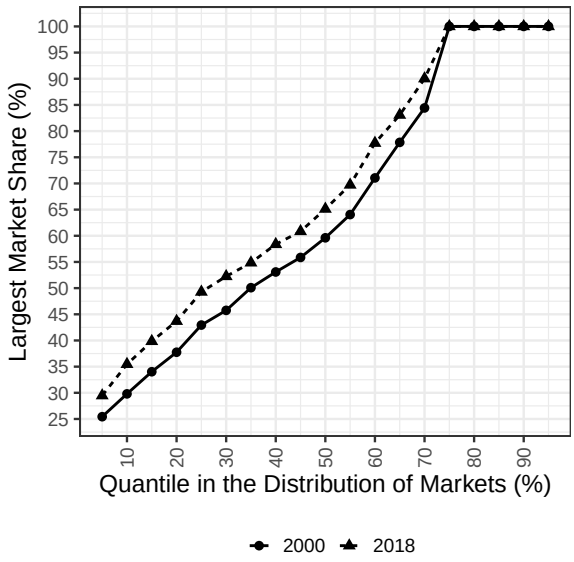
(a) Number of Hospitals per Market



(b) Number of Hospital Systems per Market



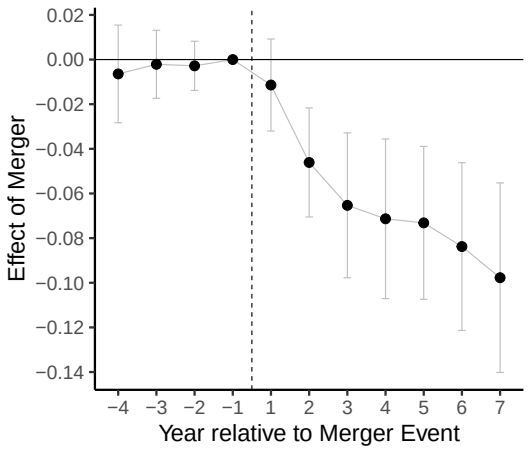
(c) HHI



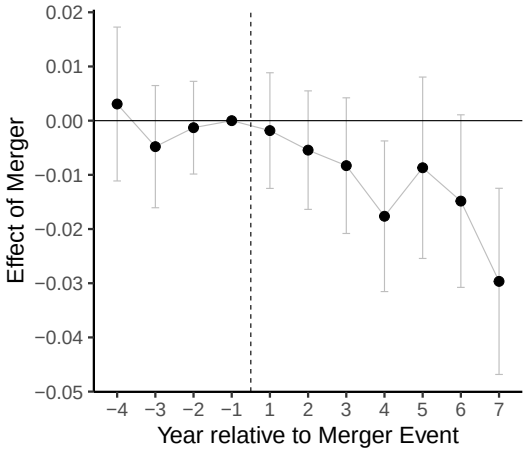
(d) Largest Market Share

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

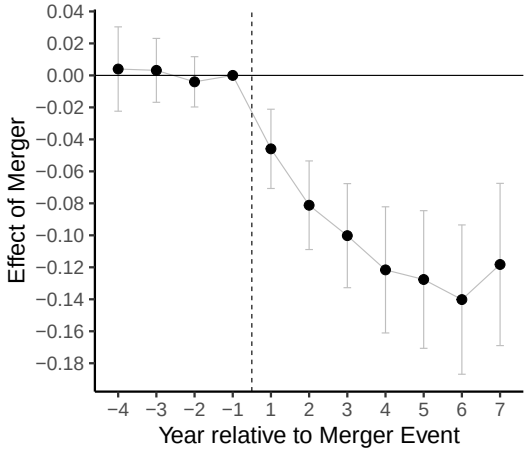
Notes: This figure presents quantiles in the distributions of hospitals per market, hospital systems per market, HHIs, and largest market shares across markets in 2000 and 2018.



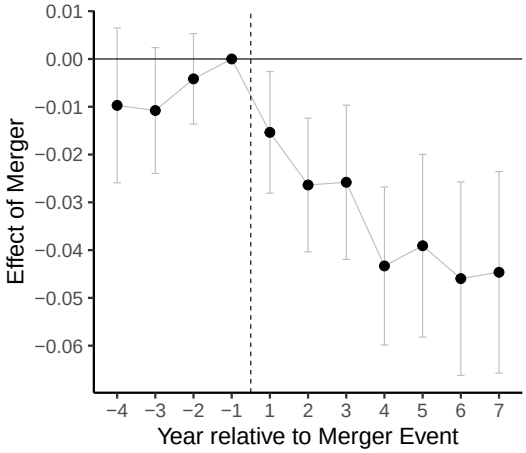
(a) Patient Care: Number of Workers, FTE (log)



(b) Patient Care: Hourly Wage (log)



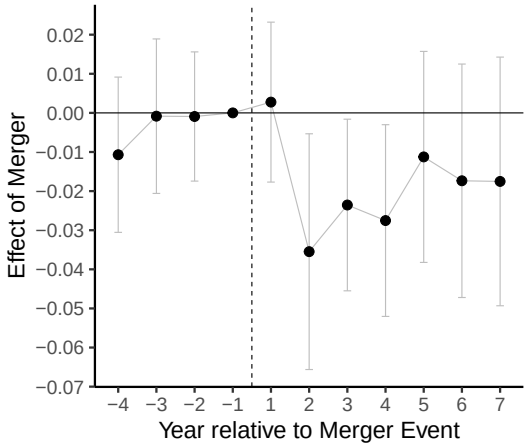
(c) Non-patient Care: Number of Workers, FTE (log)



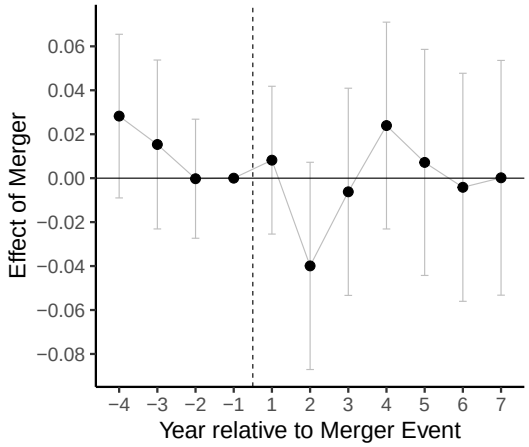
(d) Non-patient Care: Hourly Wage (log)

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

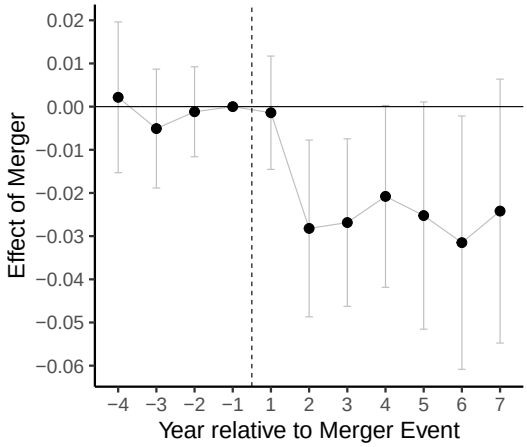
Notes: This figure presents difference-in-differences estimates that compare treated hospitals, defined as merging hospitals satisfying the "presumed anti-competitive" HHI thresholds, to merger-specific control groups of 10 hospitals from other markets matched to treated units by propensity score. The 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. 95% confidence intervals are displayed as brackets.



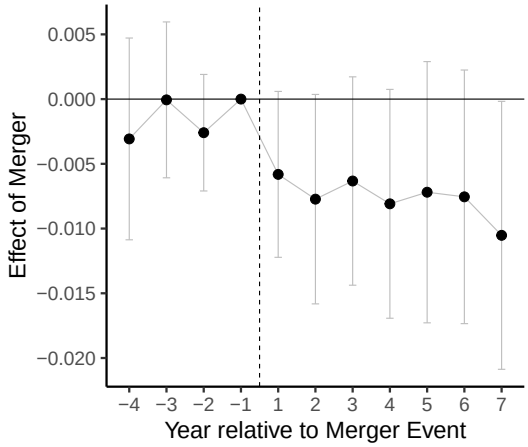
(a) Aggregate: Number of Patients (log)



(b) Aggregate: Price Index (log)



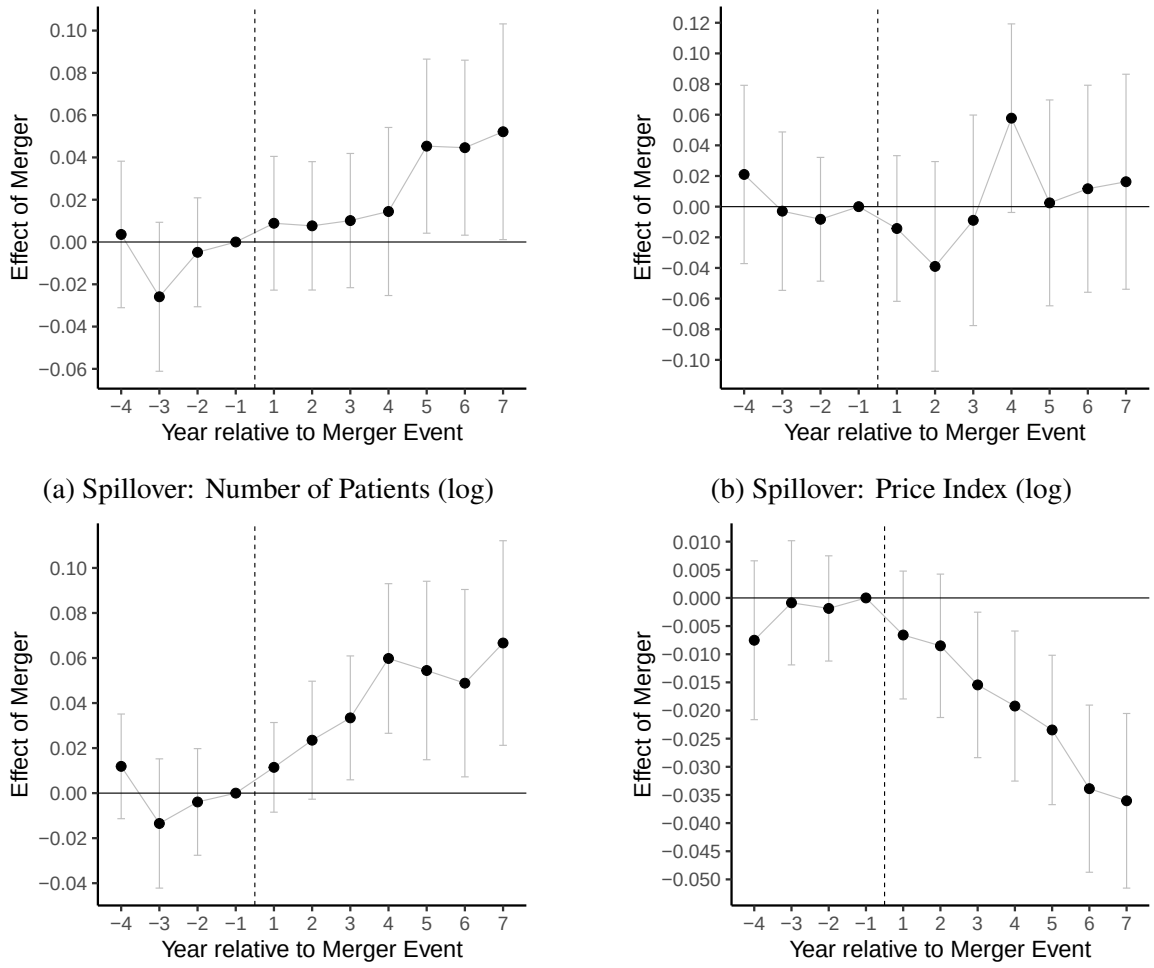
(c) Aggregate: Number of Workers (log)



(d) Aggregate: Hourly Wage (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. 95% confidence intervals are displayed as brackets.



(a) Spillover: Number of Patients (log)

(b) Spillover: Price Index (log)

(c) Spillover: Number of Workers (log)

(d) Spillover: Hourly Wage (log)

Figure A4: Spillover Effects of Mergers

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. 95% confidence intervals are displayed as brackets.