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,

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

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

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

where we denote the outside option share by $s^Q_{0t} \equiv 1 - \sum_h s^Q_{ht}$ 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,

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

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

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

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

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

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:

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

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}:

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

Using the chain rule on both sides:

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

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

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

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}:

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

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

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

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, $\frac{\partial \mathrm{MC}^L_{jt}}{\partial s^L_{0t}} < 0$, 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 $\mathrm{MR}^L_{ht} - \mathrm{MC}^L_{ht} = 0$ 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,

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

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

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

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

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

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

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

In sum, 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., $\tilde{\kappa}_h' = \tilde{\kappa}_h^* \exp(\tilde{\kappa}_\Delta^*)$. The simulation solves for the new equilibrium outcomes (new $P_h', Q_h', \ldots$) 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

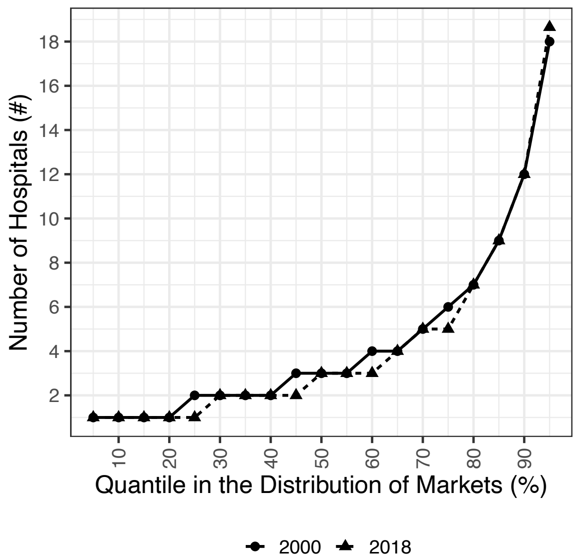

(a) Number of Hospitals per Market

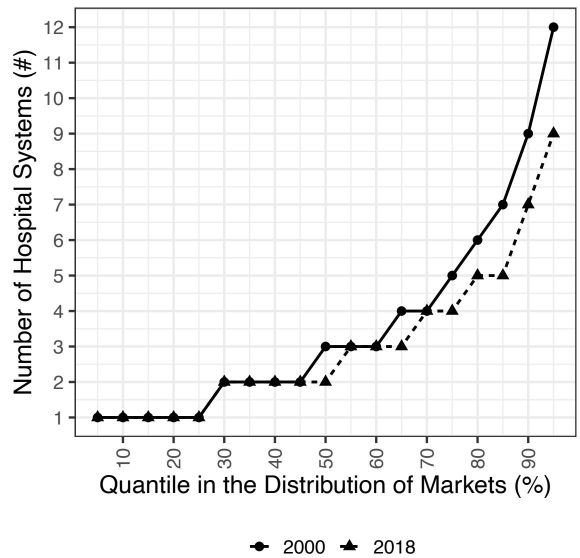

(b) Number of Hospital Systems per Market

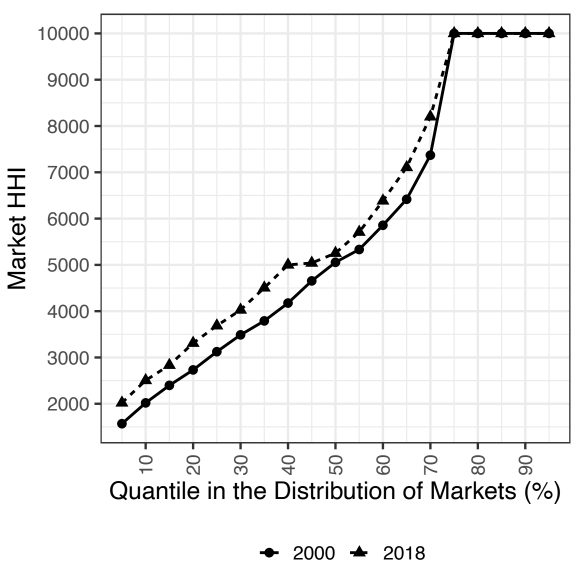

(c) HHI

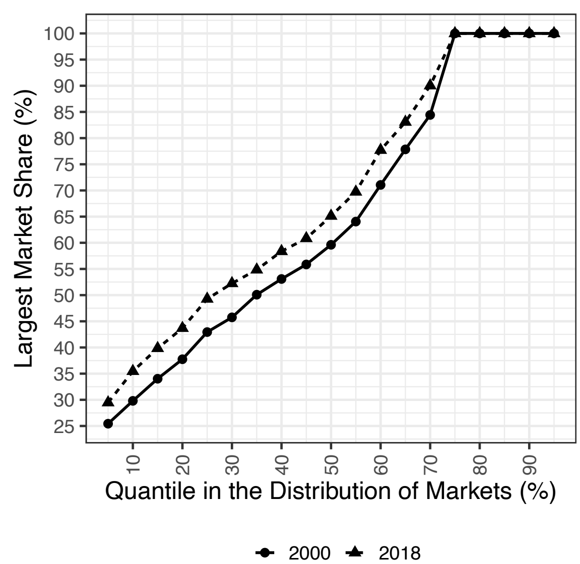

(d) Largest Market Share

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.

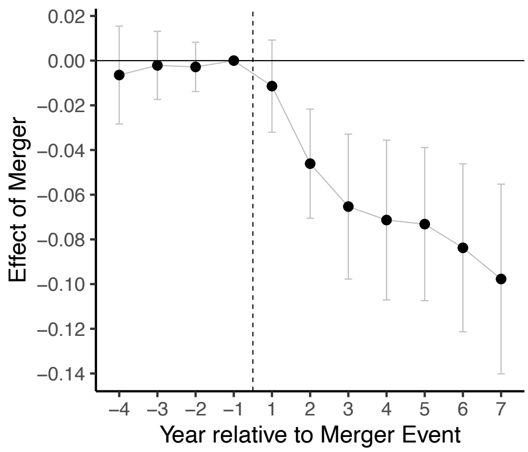

(a) Patient Care: Number of Workers, FTE (log)

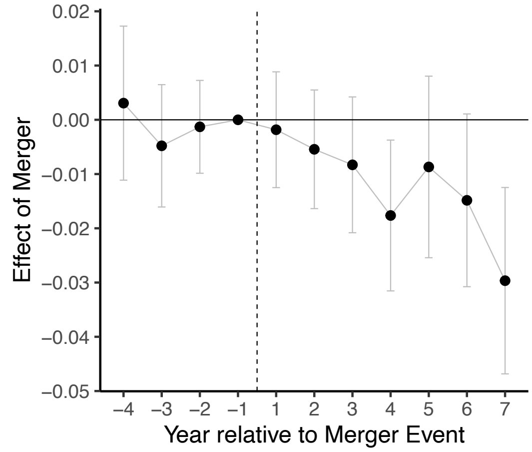

(b) Patient Care: Hourly Wage (log)

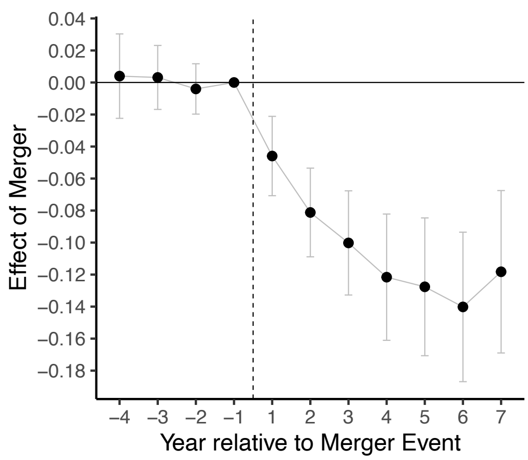

(c) Non-patient Care: Number of Workers, FTE (log)

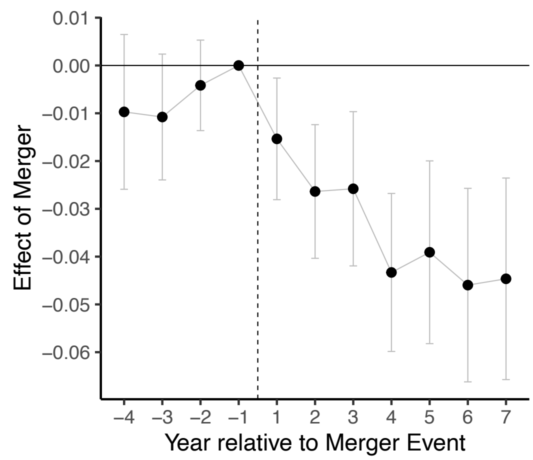

(d) Non-patient Care: Hourly Wage (log)

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.

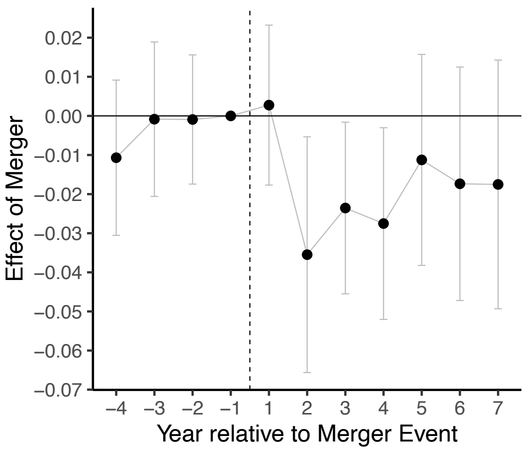

(a) Aggregate: Number of Patients (log)

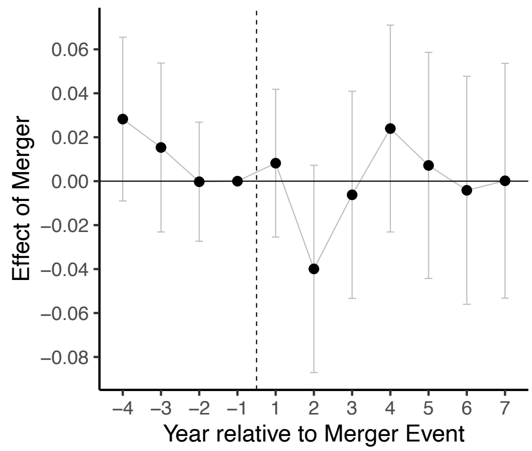

(b) Aggregate: Price Index (log)

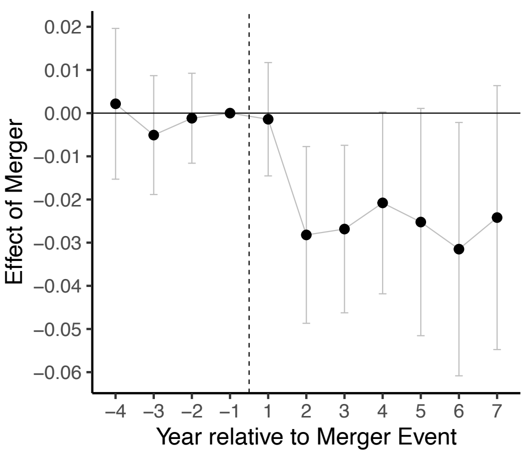

(c) Aggregate: Number of Workers (log)

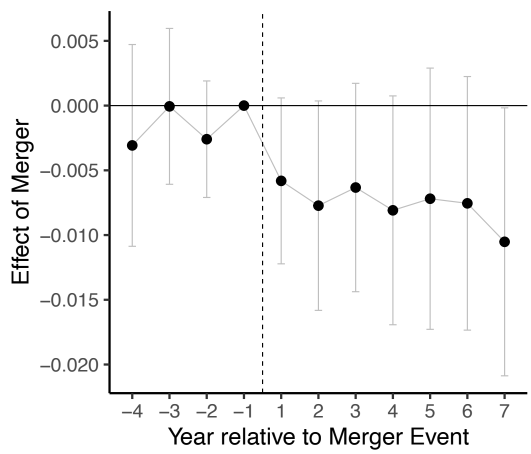

(d) Aggregate: Hourly Wage (log)

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.

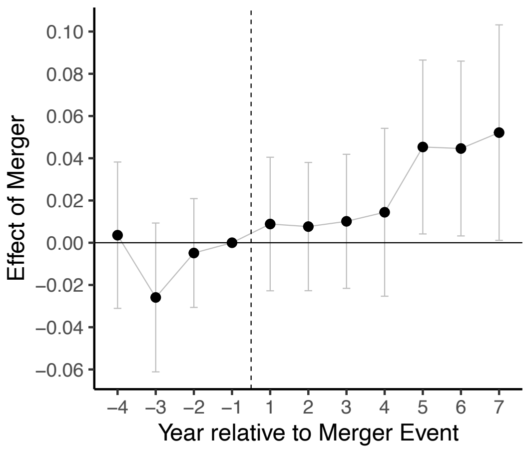

(a) Spillover: Number of Patients (log)

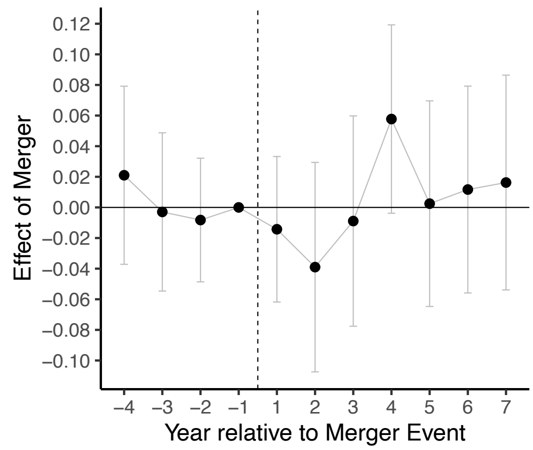

(b) Spillover: Price Index (log)

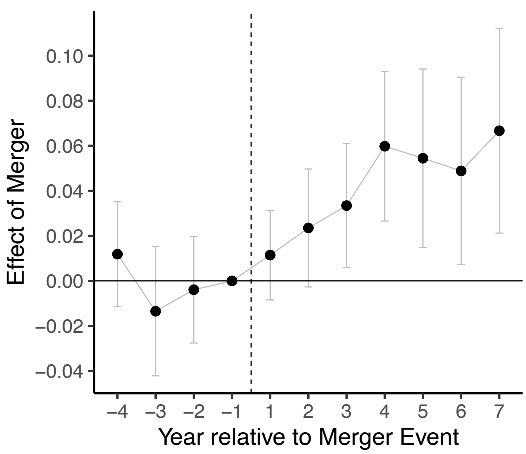

(c) Spillover: Number of Workers (log)

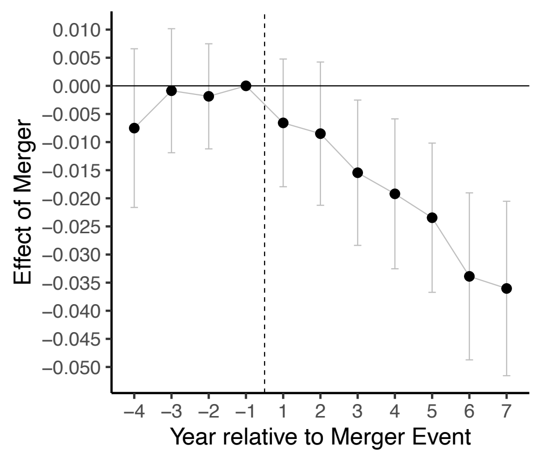

(d) Spillover: Hourly Wage (log)

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.